

Essays in Behavioral Game Theory Solution Concepts

Thesis by
Po-Hsuan Lin

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2024
Defended April 8, 2024

© 2024

Po-Hsuan Lin

ORCID: 0000-0003-3437-1734

All rights reserved

ACKNOWLEDGEMENTS

I am deeply grateful for the support of many people who have contributed to the completion of this thesis. While it is impossible to acknowledge everyone individually, I extend my sincerest appreciation to all those who have assisted me throughout my PhD journey.

First and foremost, I would like to express my deepest gratitude to my PhD advisor, Prof. Thomas Palfrey, for his invaluable guidance, patience, and advice throughout my PhD study at Caltech. His mentorship and continuous encouragement have profoundly impacted my academic journey, making it one of the most memorable and enjoyable experiences of my life.

I also owe great intellectual debts to three generous professors, Colin Camerer, Jean-Laurent Rosenthal, and Omer Tamuz. I am grateful to Prof. Colin Camerer for insightful feedback and constant support. Prof. Jean-Laurent Rosenthal has taught me how to be a better writer, and Prof. Omer Tamuz has taught me how to better present a paper. Without their guidance, I could not survive the job market.

I sincerely thank my undergraduate advisor and collaborator, Prof. Joseph Tao-yi Wang, for his support at various stages of my research career. He has never hesitated to offer his assistance since the first time I met him in his principles class. Additionally, I extend my appreciation to Dr. Walter Yuan, CEO of MobLab Inc., who hired me as a data science research intern before I started my PhD study at Caltech. This valuable experience helped prepare me for my academic career.

I am grateful for my friends in Caltech ACT, especially Yen-Yung, Yu-Chen, Meng-Jhang, Allen, Steve, Wen Chao, Chien-Ying, Michelle, Kenny, and Sean \times 2 (and my apologies if I missed anyone). I particularly want to thank Yen-Yung (aka Dr. Physics Bro) and Yu-Chen (aka Big Sis Hsu) for their unwavering support, ranging from emotional encouragement to choreographing my wedding dance. I am also grateful to have had Meng-Jhang (aka MJ), Allen, and Steve's company during the pandemic. I will never forget all the mahjong seminars we co-hosted. Lastly, I am thankful for playing softball with Wen Chao, Chien-Ying, Michelle, Kenny, and Sean \times 2. The ballgames and championships we won together are extremely memorable.

I am forever indebted to my family. Without their unconditional love and support, I could not have accomplished anything in life. I especially want to thank my wife,

Celine, for understanding and sharing everything we have been through. She has always been my cheerleader since day one of my PhD journey. While the past five years have not been easy, her love and companionship have made everything joyful. I also thank my brother. My childhood would never have been complete without you. Last but not least, I would like to convey my appreciation to my aunt, my parents and my grandparents, as this journey would not have been possible without their care and love. Their support has been instrumental in shaping who I am today.

ABSTRACT

This dissertation introduces two novel behavioral solution concepts for dynamic games: the cursed sequential equilibrium (CSE) and the dynamic cognitive hierarchy solution (DCH). Chapter 1 offers an overview of these theories and highlights their departure from standard equilibrium theory.

Chapter 2 develops the cursed sequential equilibrium, incorporating the bias where players neglect the correlation between other players' private types and actions into game theory. This framework extends the analysis of cursed equilibrium proposed by Eyster and Rabin (2005) from games in strategic form to multi-stage games, and applies it to various applications in economics and political economy.

Chapter 3 introduces the dynamic cognitive hierarchy solution, which relaxes the requirement of mutual consistency of beliefs by extending the cognitive hierarchy approach from games in strategic form to the extensive form. An important feature is that the solution can be dramatically different for games that are strategically equivalent from the perspective of standard equilibrium theory.

This property, which I call the "representation effect," has significant implications for experimental methodology and real-world phenomena. To test this effect, in Chapter 4, I design and conduct a laboratory experiment on the dirty-faces game, a simple multi-stage game of incomplete information. The experiment consists of two treatments, each implementing one of two strategically equivalent versions of the game. The dynamic cognitive hierarchy solution provides precise predictions about the differences in behavior between treatments, and the experimental results align with that prediction.

PUBLISHED CONTENT AND CONTRIBUTIONS

Lin, Po-Hsuan and Thomas R Palfrey (2024). “Cognitive hierarchies for games in extensive form”. In: *Social Science Working Paper 1460, California Institute of Technology*. DOI: <https://doi.org/10.7907/fv09c-16q24>.

Chapter 3 of the dissertation was previously circulated under the title “Cognitive Hierarchies in Extensive Form Games.” Po-Hsuan Lin contributed to the derivation of theoretical results, data analysis, and the writing of the manuscript.

Fong, Meng-Jhang, Po-Hsuan Lin, and Thomas R Palfrey (2023a). “A note on cursed sequential equilibrium and sequential cursed equilibrium”. In: *arXiv preprint arXiv:2304.05515*. DOI: <https://doi.org/10.48550/arXiv.2304.05515>.

Po-Hsuan Lin contributed to the derivation of theoretical results and applications, visualization, and the writing of the manuscript.

– (2023b). “Cursed sequential equilibrium”. In: *arXiv preprint arXiv:2301.11971*. DOI: <https://doi.org/10.48550/arXiv.2301.11971>.

Po-Hsuan Lin contributed to the derivation of theoretical results, visualization, and the writing of the manuscript.

Lin, Po-Hsuan (2023). “Cognitive hierarchies in multi-stage games of incomplete information: Theory and experiment”. In: *arXiv preprint arXiv:2208.11190v3*. DOI: <https://doi.org/10.48550/arXiv.2208.11190>.

Po-Hsuan Lin is the solo author of the manuscript.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	v
Published Content and Contributions	vi
Table of Contents	vi
List of Illustrations	ix
List of Tables	xi
Chapter I: Introduction	1
Chapter II: Cursed Sequential Equilibrium	7
2.1 Introduction	7
2.2 The Model	12
2.3 General Properties of χ -CSE	16
2.4 Applications	19
2.5 Concluding Remarks	42
Chapter III: Cognitive Hierarchies for Games in Extensive Form	49
3.1 Introduction	49
3.2 Related Literature	53
3.3 The Model	57
3.4 Properties of the DCH Solution	62
3.5 An Application: Centipede Games	73
3.6 Experimental Evidence	85
3.7 Discussion	88
3.8 Conclusions	96
Chapter IV: Cognitive Hierarchies in Multi-Stage Games of Incomplete In- formation: Theory and Experiment	104
4.1 Introduction	104
4.2 Related Literature	112
4.3 The Model	115
4.4 General Properties of the DCH Solution	118
4.5 DCH Analysis of the Dirty-Faces Games	122
4.6 Experimental Design, Hypotheses and Procedures	134
4.7 Experimental Results	142
4.8 Conclusions	153
Appendix A: Proofs for Chapter II	162
A.1 Proofs of General Properties	162
A.2 Proofs for Section 2.4	167
Appendix B: A Note on Cursed Sequential Equilibrium and Sequential Cursed Equilibrium	182
B.1 Introduction	182
B.2 Preliminary	183

B.3 Differences Between CSE and SCE	188
Appendix C: Proofs for Chapter III	210
C.1 Proofs for Section 3.4	210
C.2 Proofs for Section 3.5	213
Appendix D: Supplementary Information for Chapter IV	228
D.1 Proofs for General Properties	228
D.2 Proofs for Two-Person Dirty-Faces Games	233
D.3 Three-Person Three-Period Dirty-Faces Games	247
D.4 Detailed Analysis of Bayer and Chan (2007) Data	260
D.5 Supplementary Analysis for Experimental Data	269
D.6 Experimental Instructions	274

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
1.1 Overview of the dissertation.	2
2.1 Game Tree for Example 1.	21
2.2 Game Tree for <i>BH 3</i> and <i>BH 4</i> in Brandts and Holt (1993).	24
2.3 χ -CSE for Public Goods Game with Communication	29
2.4 Four-Stage Centipede Game.	30
2.5 χ -CSE for the Centipede Game with Altruistic Players	34
2.6 χ -CSE for the Sequential Voting Game	37
2.7 χ -CE and χ -CSE of the Dirty Faces Game	41
3.1 Game Tree of Example 3.4.2.1	66
3.2 Game Tree of Example 3.4.2.3	69
3.3 Game Tree of Γ	70
3.4 Game Tree of Γ'	71
3.5 2S-Move Centipede Game	73
3.6 Poisson-DCH of the Centipede Game	82
3.7 Games Trees in the Centipede Game Experiment	86
3.8 The Empirical CDF of the Centipede Game Experiment	87
3.9 Game Tree of the Chain-Store Game	90
3.10 Game Tree of Example 3.7.2	92
4.1 Illustration Game for Level- k Reasoning in Extensive Form	107
4.2 DCH Stopping Periods in Sequential Dirty-Faces Games	128
4.3 DCH Stopping Periods in Simultaneous Dirty-Faces Games	131
4.4 Illustration of the Representation Effect	133
4.5 Parameters for the Dirty-Faces Game Experiment	138
4.6 The CDFs of the Stopping Periods	143
4.7 The Empirical Frequencies of C in Period 1 and 2	144
4.8 The Empirical Frequencies of C in Each Payoff Configuration (1)	145
4.9 The Empirical Frequencies of C in Each Payoff Configuration (2)	146
4.10 Negative Log-Likelihood of Each Model	151
4.11 The Reaction Time in Both Treatments	152
B.1 Game Tree for the Illustrative Example.	197
B.2 The Solution Set of (χ_S, ψ_S) -SCE.	200

B.3	A Two-Player Game with Complete Information.	202
B.4	(χ_S, ψ_S) -SCE of an Unscrambled and Scrambled Signaling Game . .	206
D.1	Three-Person Dirty-Faces Games	259
D.2	The Choice Probabilities in Two-Person Games	266
D.3	The Choice Probabilities in Three-Person Games	267
D.4	The Decision Stage of the Sequential Treatment.	282
D.5	The Feedback Stage of the Sequential Treatment.	282
D.6	The Decision Stage of the Simultaneous Treatment.	283
D.7	The Feedback Stage of the Simultaneous Treatment (1).	283
D.8	The Feedback Stage of the Simultaneous Treatment (2).	284

LIST OF TABLES

<i>Number</i>	<i>Page</i>
2.1 Reduced Normal Form Centipede Game Payoff Matrix.	31
2.2 Player one's True Strategy.	32
2.3 Cursed Perception of Player one's Strategy.	32
2.4 Summary of Findings in Section 2.4.	43
3.1 Reduced Normal Form of Γ and Γ'	72
3.2 Level 2 Player 3's DCH Belief at Information Set I_3	93
4.1 Estimation Results for Two-Person Dirty-Faces Games	137
4.2 List of Game Parameters Implemented in the Experiment	141
4.3 Estimation Results for the Sequential and Simultaneous Treatments	150
D.1 Experimental Data from Bayer and Chan (2007).	261
D.2 Estimation Results for Treatment 1 and Treatment 2 Data.	265
D.3 Estimation Results for Pooled Data.	268
D.4 The Empirical Frequencies of C at Each Information Set.	269
D.5 The Empirical Frequencies of C at $(2, X)$ for Different Games.	270
D.6 The Magnitude of the Treatment Effect for Different Games.	270
D.7 Reaction Times When Seeing X in the Sequential Treatment	270
D.8 Reaction Times When Seeing X in the Simultaneous Treatment	271

Chapter 1

INTRODUCTION

This dissertation develops new behavioral solution concepts applicable to various economic scenarios where individuals interact with each other over time while possessing private information, and they seek to learn about others' private information from their past actions. Here are some examples: employers deciding whether to hire a job candidate based on their resume; in the used car market, buyers learning about the true condition of the car during negotiation; banks inspecting the loan applicant's credit report before approving the loan application; and legislators deciding whether to amend a policy based on the reports from interest groups.

In game theory, the standard approach to analyzing these situations is to model them as multi-stage games of incomplete information and solve for the equilibrium, either sequential equilibrium or perfect Bayesian equilibrium. Equilibrium analysis has been the cornerstone of economic theory since its introduction. However, there is overwhelming evidence from both the laboratory and the field showing that choices are often far away from equilibrium prediction.

To illustrate the extent to which our choices could deviate from the equilibrium prediction, let's consider a simple inductive puzzle known as the "dirty-faces game," a simple multi-stage game of incomplete information.

One day, two children, Ann and Bob, played in the park. While playing, they might accidentally get some dirt on their faces. They could only see each other's faces but not their own. When they got home, their mother (honestly) said, "At least one of you has a dirty face." Then she asked both children, "Do you have a dirty face?" If someone claims to have a dirty face, the mother will stop asking. Otherwise, the mother will keep asking the same question again and again until someone claims.

How will the children answer the question?

Standard equilibrium theory relies on three fundamental requirements: individuals must be capable of (1) forming **mutually consistent beliefs**, (2) following **Bayesian inference**, and (3) **best responding** to their beliefs at every information set. Under

these requirements, the equilibrium in this example can be solved through iterative rationality. If Ann were to see a clean face, she can infer that her own face must be dirty and claim immediately. However, if Ann sees a dirty face, no such inference is possible, so she will rationally wait first. In equilibrium, Ann believes that Bob would have claimed at the beginning if her own face were clean. Hence, if Bob waits at the beginning, due to mutual consistency and Bayesian inference, Ann will realize that her face is dirty and claim when the mother asks the question for the second time.

In laboratory experiments of this game, we find that the majority of participants (around 80% to 90%) claim immediately upon seeing a clean face. However, when seeing a dirty face, only about 50% of participants are able to follow the two-step equilibrium reasoning. This experimental finding has been replicated across different subject pools, suggesting that the equilibrium requirements may be empirically implausible.

To address this, in this dissertation, I develop two novel behavioral solution concepts: the **cursed sequential equilibrium (CSE)** and the **dynamic cognitive hierarchy solution (DCH)**, each of which relaxes the requirements of Bayesian inference and mutual consistency, respectively. This is summarized in the figure below.

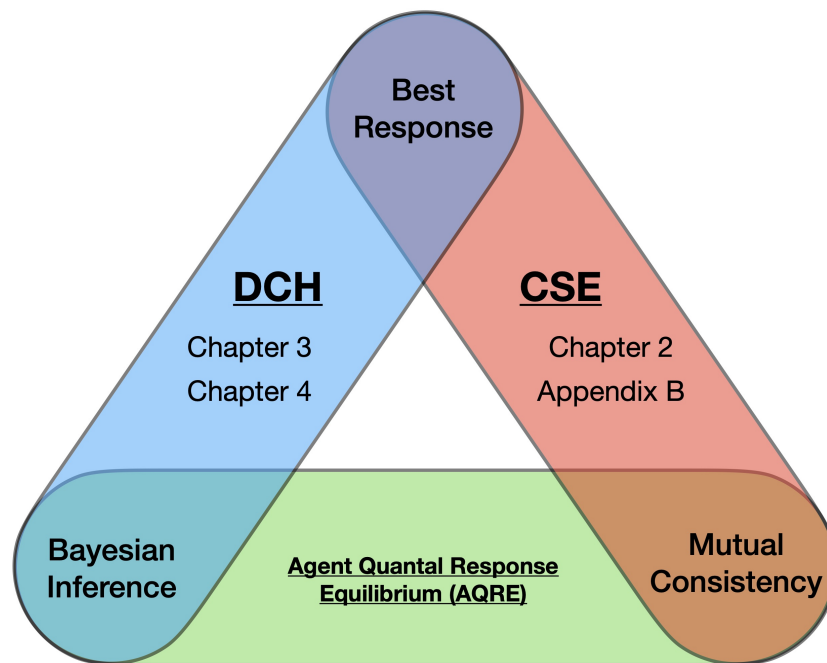


Figure 1.1: Overview of the dissertation.

Chapter 2 (written jointly with Meng-Jhang Fong and Thomas R Palfrey) develops the cursed sequential equilibrium (CSE) which maintains the requirements of best response and mutual consistency of the belief system while accounting for individuals' inability to make accurate Bayesian inferences. In CSE, "cursed" players neglect or partially neglect the correlation between other players' private types and actions while having accurate beliefs about others' average behavioral strategies. The incorrect updating of the beliefs about others' private information distorts their conjectures of how others will behave in the future. With this relaxation of Bayesian updating, CSE can generate more empirically plausible predictions than perfect Bayesian equilibrium and the standard cursed equilibrium proposed by Eyster and Rabin (2005) in both common-value and private-value games.

CSE provides new insights into many important applications. For instance, it offers new explanations for deviations from the intuitive criterion in signaling games and the failure of pivotal reasoning in strategic voting. Furthermore, it also yields clear qualitative and testable predictions in public goods games with communication, making CSE an experimentally falsifiable theory.

In chapter 3 (written jointly with Thomas R Palfrey), we propose the dynamic cognitive hierarchy solution (DCH) which relaxes the mutual consistency requirement by modeling individuals as heterogeneous in their beliefs about the strategic abilities of other players, while maintaining the requirements of Bayesian inference and best response. Conceptually, DCH extends the level- k model—a widely applied theory for organizing experimental data in games like the beauty contest game, coordination games, auctions, and more—from simultaneous games to dynamic games.

In the level- k model, each player is endowed with a "level" of sophistication, where a level k player incorrectly believes all others are level $(k-1)$ and best responds to this belief. To apply the standard level- k model to dynamic games, one needs to assume level- k players will choose an action that maximizes the continuation value of the game, while believing that all other players are level $(k-1)$ in the continuation game. Consequently, each player's belief about others' levels is fixed from the beginning, creating a logic conundrum as level k players can be "surprised" when an opponent's move that is inconsistent with the strategy of a level $(k-1)$ player.

To address this, DCH adopts the cognitive hierarchy approach, whereby a level k player believes all other players have lower levels distributed anywhere from level 0 to $k-1$. A player updates their beliefs about the distribution of each of the other players' levels as the history unfolds and best responds at every information set. In other

words, DCH provides us with the essential machinery to understand hierarchical reasoning behavior in dynamic games, such as learning and experimentation.

An important feature of DCH is the **representation effect**.¹ That is, the solution can be dramatically different when the same dynamic game is played according to its extensive form and reduced-normal form. Specifically, in the extensive form, players make decisions sequentially, acting only when the game reaches their own information sets. In contrast, in the reduced-normal form, all players simultaneously choose a reduced contingent strategy. Since a reduced contingent strategy specifies the decision at every possible situation, the reduced-normal form representation does not restrict the freedom of action. In other words, if a player is strategically rational in the sense of Nash equilibrium, they should make the same decision in both the extensive form and the corresponding reduced normal form. However, whether actual behavior under different representations is the same remains an open question, and DCH makes precise predictions about when and how different representations will induce different behavior.

The representation effect is important not only because it is a testable prediction but also because of its significant implications for experimental methodology and real-world phenomena. When conducting dynamic game experiments, researchers commonly employ two methods. The first method is the “direct-response method,” which implements the game in its extensive form. Alternatively, the second method is the “strategy method,” which implements the game in its reduced-normal form. The advantage of employing the strategy method is that one can collect much more data, particularly at information sets that are only occasionally reached during the course of play. From the perspective of DCH, these two methods have different extensive forms but share the same reduced-normal form, suggesting that implementing a dynamic game experiment using the strategy method could lead to behavioral distortions.

Furthermore, many real-world institutional settings, such as descending-clock auctions and runoff elections, are often implemented in their reduced-normal form due to simplicity or practical considerations. When implementing an auction, to avoid learning and potential collusion, rather than using the clock auction, auctioneers often opt for the simpler “sealed-bid” design in which bidders simply mail in their bids, and the auctioneer opens the envelopes and announces the winner, who pays their bids. Moreover, some countries, like France, adopt runoff elections to decide

¹In Chapter 3, the representation effect is also referred to as the “strategy-reduction effect.”

the new president. However, running two rounds of elections could be costly and politically chaotic. To avoid these issues, other countries, like Australia, adopt *contingent* runoff elections where voters submit their ranking of the candidates. This essentially means playing the voting game in reduced-normal form. These institutional designs implicitly rely on the insight of standard equilibrium theory that different representations of the game would not induce different behavior. Yet there is more and more evidence in different contexts showing that different representations could lead to different behavior, and DCH is a theory that can predict the occurrence and direction of the representation effect.

To test this prediction, in Chapter 3, we reanalyze a recent laboratory experiment on the centipede game, a classic game of perfect information where observed behavior is grossly inconsistent with standard equilibrium theory. This experiment compares behavior in four centipede games when played in extensive form and reduced-normal form. Standard equilibrium theory makes the same (outcome-equivalent) prediction under both representations for all four games. However, DCH predicts that a specific representation effect would occur in three games but not in one game. Surprisingly, the experimental findings align closely with the DCH prediction, providing the first empirical evidence of DCH in games of perfect information.

In Chapter 4, I delve deeper into the DCH representation effect in multi-stage games of incomplete information where the belief updating becomes much more intricate because players will learn about others' levels of sophistication and payoff-relevant types at the same time.

In addition to extending the theoretical characterization from games of perfect information to multi-stage games of incomplete information, I design and run a laboratory experiment on the dirty-faces game, aiming to assess the DCH representation effect. In this experiment, I manipulate both the timing structures (sequential vs. simultaneous) and the payoff parameters. The main challenge to designing the experiment is that the effect size, predicted by DCH, depends upon both the payoff parameters and the true distribution of levels, which is unknown before the experiment is run. To address this, I develop an optimal design approach where I first estimate the distribution of levels using an existing dirty-faces game experimental dataset from a published paper and then choose parameters that are maximally informative for this estimated distribution of levels. Employing this fine-tuned experimental design, we are able to detect a significant treatment effect. Importantly, both the direction and magnitude of the observed difference align with the predictions of DCH. This result

lays a strong empirical foundation for the DCH solution.

In summary, this dissertation introduces novel behavioral solution concepts for dynamic games. These new theories relax various requirements of standard equilibrium theory and produce predictions that are consistent with experimental evidence. They offer fresh insights into human behavior within diverse institutional settings, which were traditionally considered strategically equivalent according to standard equilibrium theory. In essence, this dissertation opens up new avenues for comprehending economic behavior in institutional designs.

CURSED SEQUENTIAL EQUILIBRIUM

2.1 Introduction

Cursed equilibrium (CE) proposed by Eyster and Rabin (2005) is a leading behavioral equilibrium concept that was developed to explain the “winner’s curse” and related anomalies in applied game theory. The basic idea behind CE is that individuals do not fully take account of the dependence of other players’ strategic actions on private information. Cursed behavior of this sort has been detected in a variety of contexts. Capen, Clapp, and Campbell (1971) first noted that in oil-lease auctions, “the winner tends to be the bidder who most overestimates the reserves potential” (Capen, Clapp, and Campbell, 1971, p. 641). Since then, this observation of overbidding relative to the Bayesian equilibrium benchmark, which can result in large losses for the winning bidder, has been widely documented in laboratory auction experiments (Bazerman and Samuelson, 1983; Kagel and Levin, 1986; Dyer, Kagel, and Levin, 1989; Forsythe, Isaac, and Palfrey, 1989; Kagel, Levin, et al., 1989; Lind and Plott, 1991; Kagel and Levin, 2009; Ivanov, Levin, and Niederle, 2010; Camerer, Nunnari, and Palfrey, 2016). In addition, the neglect of the connection between the opponents’ actions and private information is also found in non-auction environments, such as bilateral bargaining games (Samuelson and Bazerman, 1985; Holt and Sherman, 1994; Carrillo and Palfrey, 2009; Carrillo and Palfrey, 2011), zero-sum betting games with asymmetric information (Rogers, Palfrey, and Camerer, 2009; S¸ovik, 2009), and voting and jury decisions (Guarnaschelli, McKelvey, and Palfrey, 2000).

While CE provides a tractable alternative to Bayesian Nash equilibrium and can explain some anomalous behavior in games with a winner’s-curse structure, a significant limitation is that it is only developed as a strategic form concept for simultaneous-move Bayesian games. Thus, when applying the standard CE to dynamic games, the CE analysis is carried out on the strategic form representation of the game, implying that CE cannot distinguish behavior across dynamic games that differ in their timing of moves but have the same strategic form. That is, players are assumed to choose type-dependent contingent strategies simultaneously and *not* update their beliefs as the history of play unfolds. A further limitation implied

by the strategic form approach is that CE and standard Bayesian Nash equilibrium make identical predictions in games with a private-values information structure (Eyster and Rabin, 2005, Proposition 2). In this paper we extend the CE in a simple and natural way to multi-stage games of incomplete information. We call the new equilibrium concept *Cursed Sequential Equilibrium* (CSE).

In Section 2.2, we present the framework and our extension of cursed equilibrium to dynamic games. We consider the framework of multi-stage games with observed actions, introduced by Fudenberg and Tirole (1991b), where players' private information is represented by types, with the assumption that the set of available actions is independent of their types at each public history. Our new solution concept is in the same spirit of the cursed equilibrium—in our model, at each stage, players will (partially) neglect the dependence of the other players' behavioral strategies on their types, by placing some weight on the incorrect belief that all types adopt the average *behavioral* strategy. Specifically, at each public history, this corresponds to the average distribution of actions *given the current belief about others' types at that stage*. Therefore, as players update their beliefs about others' private information via Bayes' rule, but with incorrect beliefs about the other players' behavioral strategies, in later stages this can lead them to have *incorrect* beliefs about the other players' average distribution of actions.

Following Eyster and Rabin (2005)'s notion of cursedness, we parameterize the model by a single parameter $\chi \in [0, 1]$ which captures the degree of cursedness and define fully cursed ($\chi = 1$) CSE analogously to fully cursed ($\chi = 1$) CE. Recall that in a fully cursed ($\chi = 1$) CE, each type of each player chooses a best reply to expected (cursed) equilibrium distribution of other players' actions, averaged over the type-conditional strategies of the other players, with this average distribution calculated using the prior belief on types. Loosely speaking, a player best responds to the average CE strategy of the others. In a χ -CE, players are only partially cursed, in the sense that each player best responds to a χ -weighted linear combination of the *average* χ -CE strategy of the others and the *true* (type-dependent) χ -CE strategy of the others.

The extension of this definition to multi-stage games with observed actions is different from χ -CE in two essential ways: (1) the game is analyzed with behavioral strategies, and (2) we impose sequential rationality and Bayesian updating. In a fully cursed ($\chi = 1$) CSE, (1) implies at every stage t and each public history at t , each type of each player i chooses a best reply to the expected (cursed) equilibrium dis-

tribution of other players' stage- t actions, averaged over the type-conditional stage- t behavioral strategies of other players, with this average distribution calculated using i 's current belief about types at stage t . That is, player i best responds to the average stage- t CSE strategy of others. Moreover, (2) requires that each player's belief at each public history is derived by Bayes' rule wherever possible, and best replies are with respect to the continuation values computed by using the fully cursed beliefs about the behavioral strategies of the other players in current and future stages.

A χ -CSE, for $\chi < 1$, is then defined in analogously to χ -CE, except for using a χ -weighted linear combination of the *average* χ -CSE *behavioral strategies* of others and the *true* (type-dependent) χ -CSE *behavioral strategies* of others. Thus, similar to the fully cursed CE, in a fully cursed ($\chi = 1$) CSE, each player believes other players' actions at each history are *independent* of their private information. On the other hand, $\chi = 0$ corresponds to the standard sequential equilibrium where players have correct perceptions about other players' behavioral strategies and are able to make correct Bayesian inferences.¹

After defining the equilibrium concept, in Section 2.3 we explore some general properties of the model. We first prove the existence of a cursed sequential equilibrium in Proposition 2.1. Intuitively speaking, CSE mirrors the standard sequential equilibrium. The only difference is that players have incorrect beliefs about the other players' behavioral strategies at each stage since they fail to fully account for the correlation between others' actions and types at every history. We prove in Proposition 2.2 that the set of CSE is upper hemi-continuous with respect to χ . Consequently, every limit point of a sequence of χ -CSE points as χ converges to 0 is a sequential equilibrium. This result bridges our behavioral solution concept with the standard equilibrium theory. Finally, we also show in Proposition 2.4 that χ -CSE is equivalent to χ -CE for one-stage games, demonstrating the connection between the two behavioral solutions.

In multi-stage games, cursed beliefs about behavioral strategies will distort the evolution of a player's beliefs about the other players' types. As shown in Proposition 2.3, a direct consequence of the distortion is that in χ -CSE players tend to update their beliefs about others' types too passively. That is, there is some persistence in

¹For the off-path histories, similar to the idea of Kreps and Wilson (1982), we impose the χ -consistency requirement (see Definition 2.2) so the assessment is approachable by a sequence of totally mixed behavioral strategies. The only difference is that players' beliefs are incorrectly updated by assuming others play the χ -cursed behavioral strategies. Hence, in our approach if $\chi = 0$, a CSE is a sequential equilibrium.

beliefs in the sense that at each stage t , each χ -cursed player's belief about any type profile is at least χ times the belief about that type profile at stage $t - 1$. Among other things, this implies that if the *prior* belief about the types is full support and $\chi > 0$, the full support property will persist at all histories, and players will (possibly incorrectly) believe every profile of others' types is possible at every history.

This dampened updating property plays an important role in our framework. Not only does it contribute to the difference between CSE and the standard CE through the updating process, but it also implies additional restrictions on off-path beliefs. The effect of dampened updating is starkly illustrated in the pooling equilibria of signaling games where every type of sender behaves the same everywhere. In this case, Proposition 2.5 shows if an assessment associated with a pooling equilibrium is a χ -CSE, then it also a χ' -CSE for all $\chi' \leq \chi$, but it is not necessarily a pooling equilibrium for all $\chi' > \chi$. This contrasts with one of the main results about CE, that if a pooling equilibrium is a χ -CE for some χ , then it is a χ' -CE for all $\chi' \in [0, 1]$ (Eyster and Rabin, 2005, Proposition 3).

This suggests that perhaps the dampened updating property is an equilibrium selection device that eliminates some pooling equilibrium, but actually this is *not* a general property. As we demonstrate later, the χ -CE and χ -CSE sets can be non-overlapping, which we illustrate with a variety of applications. The intuition is that in CSE, players generally do not have correct beliefs about the opponents' average behavioral strategies. The pooling equilibrium is just a special case where players have correct beliefs.

In Section 2.4 we explore the implications of cursed sequential equilibrium with five applications in economics and political science. Section 2.4.1 analyzes the χ -CSE of signaling games. Besides studying the theoretical properties of pooling χ -CSE, we also analyze two simple signaling games that were studied in a laboratory experiment (Brandts and Holt, 1993). We show how varying the degree of cursedness can change the set of χ -CSE in these two signaling games in ways that are consistent with the reported experimental findings. Next, we turn to the exploration of how sequentially cursed reasoning can influence strategic communication. To this end, we analyze the χ -CSE for a public goods game with communication (Palfrey and Rosenthal, 1991; Palfrey, Rosenthal, and Roy, 2017) in Section 2.4.2, finding that χ -CSE predicts there will be less effective communication when players are more cursed.

Next, in Section 2.4.3 we apply χ -CSE to the centipede game studied experimentally by McKelvey and Palfrey (1992) where one of the players believes the other player

might be an “altruistic” player who always passes. This is a simple reputation-building game, where selfish types can gain by imitating altruistic types in early stages of the game. The public goods application and the centipede game are both private-values environments, so these two applications clearly demonstrate how CSE departs from CE and the Bayesian Nash equilibrium, and shows the interplay between sequentially cursed reasoning and the learning of types in private-value models.

In strategic voting applications, conditioning on “pivotality”—the event where your vote determines the final outcome—plays a crucial role in understanding equilibrium voting behavior. To illustrate how cursedness distorts the pivotal reasoning, in Section 2.4.4 we study the three-voter two-stage agenda voting game introduced by Ordeshook and Palfrey (1988). Since this is a private value game, the predictions of the χ -CE and the Bayesian Nash equilibrium coincide for all χ . That is, cursed equilibrium predicts no matter how cursed the voters are, they are able to correctly perform pivotal reasoning. On the contrary, our CSE predicts that cursedness will make the voters less likely to vote strategically. This is consistent with the empirical evidence about the prevalence of sincere voting over sequential agendas when inexperienced voters have incomplete information about other voters’ preferences (Levine and Plott, 1977; Plott and Levine, 1978; Eckel and Holt, 1989).

Finally, in Section 2.4.5 we study the relationship between cursedness and epistemic reasoning by considering the two-person dirty faces game previously studied by Weber (2001), Bayer and Chan (2007) and Lin (2023). In this game, χ -CSE predicts cursed players are, to some extent, playing a “coordination” game where they coordinate on a specific learning speed about their face types. Therefore, from the perspective of CSE, the non-equilibrium behavior observed in experiments can be interpreted as possibly due to a coordination failure resulting from cognitive limitations.

The cursed sequential equilibrium extends the concept of cursed equilibrium from static Bayesian games to multi-stage games with observed actions. This generalization preserves the spirit of the original cursed equilibrium in a simple and tractable way, and provides additional insights about the effect of cursedness in dynamic games. A contemporaneous working paper by Cohen and Li (2023) is closely related to our paper. That paper adopts an approach based on the coarsening of information sets to define sequential cursed equilibrium (SCE) for extensive form

games with perfect recall. The SCE model captures a different kind of cursedness² that arises if a player neglects the dependence of other players' unobserved (i.e., either future or simultaneous) actions on the history of play in the game, which is different from the dependence of other players' actions on their type (as in CE and CSE). In the terminology of Eyster and Rabin (2005), the cursedness is with respect to endogenous information, i.e., what players observe about the path of play (Eyster and Rabin, 2005, p. 1665). The idea is to treat the unobserved actions of other players in response to different histories (endogenous information) similarly to how cursed equilibrium treats players' types. A two-parameter model of partial cursedness is developed, and a series of examples demonstrate that for plausible parameter values, the model is consistent with some experimental findings related to the failure of subjects to fully take account of unobserved hypothetical events, whereas behavior is "more rational" if subjects make decisions after directly observing such events. A more detailed discussion of the differences between CSE and SCE can be found in Appendix B and Fong, Lin, and Palfrey (2023). At a more conceptual level, our paper is related to several other behavioral solution concepts developed for dynamic games, such as agent quantal response equilibrium (AQRE) (McKelvey and Palfrey, 1998), dynamic cognitive hierarchy solution (DCH) (Lin and Palfrey, 2022; Lin, 2023), and the analogy-based expectation equilibrium (ABEE) (Jehiel, 2005; Jehiel and Koessler, 2008), all of which modify the requirements of sequential equilibrium in different ways than cursed sequential equilibrium.

2.2 The Model

Since CSE is a solution concept for dynamic games of incomplete information, in this paper we will focus on the framework of multi-stage games with observed actions (Fudenberg and Tirole, 1991b). Section 2.2.1 defines the structure of multi-stage games with observed actions, followed by Section 2.2.2, where the χ -cursed sequential equilibrium is formally developed.

2.2.1 Multi-Stage Games with Observed Actions

Let $N = \{1, \dots, n\}$ be a finite set of players. Each player $i \in N$ has a *type* θ_i drawn from a finite set Θ_i . Let $\theta \in \Theta \equiv \times_{i=1}^n \Theta_i$ be the type profile and θ_{-i} be the type profile without player i . All players have the common (full support) prior distribution $F : \Theta \rightarrow (0, 1)$. At the beginning of the game, each player is told his own type, but is not informed anything about the types of others. Therefore, each

²We illustrate some implications in the application to signaling games in Section 2.4.1.

player i 's initial belief about the types of others when his type is θ_i is

$$F(\theta_{-i}|\theta_i) = \frac{F(\theta_{-i}, \theta_i)}{\sum_{\theta'_{-i} \in \Theta_{-i}} F(\theta'_{-i}, \theta_i)}.$$

If the types are independent across players, each player i 's initial belief about the types of others is $F_{-i}(\theta_{-i}) = \prod_{j \neq i} F_j(\theta_j)$ where $F_j(\theta_j)$ is the marginal distribution of player j 's type.

The game is played in stages $t = 1, 2, \dots, T$ where $T < \infty$. In each stage, players simultaneously choose their actions, which will be revealed at the end of the stage. The feasible set of actions can vary with histories, so games with alternating moves are also included. Let \mathcal{H}^{t-1} be the set of all available histories at stage t , where $\mathcal{H}^0 = \{h_0\}$ and \mathcal{H}^T is the set of terminal histories. Let $\mathcal{H} = \cup_{t=0}^T \mathcal{H}^t$ be the set of all available histories of the game, and let $\mathcal{H} \setminus \mathcal{H}^T$ be the set of non-terminal histories.

For every player i , the available information at stage t is in $\mathcal{H}^{t-1} \times \Theta_i$. Thus, each player i 's information sets can be specified as $\mathcal{I}_i \in \Pi_i = \{(h, \theta_i) : h \in \mathcal{H} \setminus \mathcal{H}^T, \theta_i \in \Theta_i\}$. That is, a type θ_i player i 's information set at the public history h^t can be defined as $\cup_{\theta_{-i} \in \Theta_{-i}} (h^t, \theta_{-i}, \theta_i)$. With a slight abuse of notation, it will be denoted as (h^t, θ_i) . For the sake of simplicity, the feasible set of actions for every player at every history is assumed to be type-independent. Let $A_i(h^{t-1})$ be the feasible set of actions for player i at history h^{t-1} and let $A_i = \times_{h \in \mathcal{H} \setminus \mathcal{H}^T} A_i(h)$ be the set of player i 's all feasible actions in the game. For each player i , A_i is assumed to be finite and $|A_i(h)| \geq 1$ for any $h \in \mathcal{H} \setminus \mathcal{H}^T$. Let $a_i^t \in A_i(h^{t-1})$ be player i 's action at history h^{t-1} , and let $a^t = (a_1^t, \dots, a_n^t) \in \times_{i=1}^n A_i(h^{t-1})$ denote the action profile at stage t . If a^t is the action profile chosen at stage t , then $h^t = (h^{t-1}, a^t)$.

A behavioral strategy for player i is a function $\sigma_i : \Pi_i \rightarrow \Delta(A_i)$ satisfying $\sigma_i(h^{t-1}, \theta_i) \in \Delta(A_i(h^{t-1}))$. Let $\sigma_i(a_i^t | h^{t-1}, \theta_i)$ denote the probability for player i to choose $a_i^t \in A_i(h^{t-1})$. A strategy profile $\sigma = (\sigma_i)_{i \in N}$ specifies a behavioral strategy for each player i . Lastly, each player i has a payoff function (in von Neumann-Morgenstern utilities) $u_i : \mathcal{H}^T \times \Theta \rightarrow \mathbb{R}$, and let $u = (u_1, \dots, u_n)$ be the profile of utility functions. A multi-stage game with observed actions, Γ , is defined by the tuple $\Gamma = \langle N, \mathcal{H}, \Theta, F, u \rangle$.

2.2.2 Cursed Sequential Equilibrium

In a multi-stage game with observed actions, a solution is defined by an ‘‘assessment,’’ which consists of a (behavioral) strategy profile σ , and a belief system μ . Since action profiles will be revealed to all players at the end of each stage, the belief

system specifies, for each player, a conditional distribution over the set of type profiles conditional on each history. Consider an assessment (μ, σ) . Following the spirit of the cursed equilibrium, for player i at stage t , we define the *average behavioral strategy profile of the other players* as

$$\bar{\sigma}_{-i}(a_{-i}^t | h^{t-1}, \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} \mu_i(\theta_{-i} | h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^t | h^{t-1}, \theta_{-i})$$

for any $i \in N$, $\theta_i \in \Theta_i$ and $h^{t-1} \in \mathcal{H}^{t-1}$.

In CSE, players have incorrect perceptions about other players' behavioral strategies. Instead of thinking they are using σ_{-i} , a χ -cursed³ type θ_i player i would believe the other players are using a χ -weighted average of the average behavioral strategy and the true behavioral strategy:⁴

$$\sigma_{-i}^\chi(a_{-i}^t | h^{t-1}, \theta_{-i}, \theta_i) = \chi \bar{\sigma}_{-i}(a_{-i}^t | h^{t-1}, \theta_i) + (1 - \chi) \sigma_{-i}(a_{-i}^t | h^{t-1}, \theta_{-i}).$$

The beliefs of player i about θ_{-i} are updated in the χ -CSE via Bayes' rule, whenever possible, assuming other players are using the χ -cursed behavioral strategy rather than the true behavioral strategy. We call this updating rule the *χ -cursed Bayes' rule*. Specifically, an assessment satisfies the χ -cursed Bayes' rule if the belief system is derived from the Bayes' rule while perceiving others are using σ_{-i}^χ rather than σ_{-i} .

Definition 2.1. (μ, σ) satisfies χ -cursed Bayes' rule if the following is applied to update the posterior beliefs as $\sum_{\theta'_{-i} \in \Theta_{-i}} \mu_i(\theta'_{-i} | h^{t-1}, \theta_i) \sigma_{-i}^\chi(a_{-i}^t | h^{t-1}, \theta'_{-i}, \theta_i) > 0$:

$$\mu_i(\theta_{-i} | h^t, \theta_i) = \frac{\mu_i(\theta_{-i} | h^{t-1}, \theta_i) \sigma_{-i}^\chi(a_{-i}^t | h^{t-1}, \theta_{-i}, \theta_i)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \mu_i(\theta'_{-i} | h^{t-1}, \theta_i) \sigma_{-i}^\chi(a_{-i}^t | h^{t-1}, \theta'_{-i}, \theta_i)}.$$

Let Σ^0 be the set of totally mixed behavioral strategy profiles, and let Ψ^χ be the set of assessments (μ, σ) such that $\sigma \in \Sigma^0$ and μ is derived from σ using χ -cursed Bayes' rule.⁵ Lemma 2.1 below shows that another interpretation of the χ -cursed Bayes' rule is that players have correct perceptions about σ_{-i} but are unable to make perfect Bayesian inference when updating beliefs. From this perspective, player i 's cursed

³We assume throughout the paper that all players are equally cursed, so there is no i subscript on χ . The framework is easily extended to allow for heterogeneous degrees of cursedness.

⁴If $\chi = 0$, players have correct beliefs about other players' behavioral strategies at every stage.

⁵In the following, we will use $\mu^\chi(\cdot)$ to denote the belief system derived under χ -cursed Bayes' Rule. Also, note that both σ_{-i}^χ and μ^χ are induced by σ ; that is, $\sigma_{-i}^\chi(\cdot) = \sigma_{-i}^\chi[\sigma](\cdot)$ and $\mu^\chi(\cdot) = \mu^\chi[\sigma](\cdot)$. For the ease of exposition, we drop $[\sigma]$ when it does not cause confusion.

belief is simply a linear combination of player i 's cursed belief at the beginning of that stage (with χ weight) and the Bayesian posterior belief (with $1 - \chi$ weight). Because σ is totally mixed, there are no off-path histories.

Lemma 2.1. *For any $(\mu, \sigma) \in \Psi^\chi$, $i \in N$, $h^t = (h^{t-1}, a^t) \in \mathcal{H} \setminus \mathcal{H}^T$ and $\theta \in \Theta$,*

$$\mu_i(\theta_{-i}|h^t, \theta_i) = \chi \mu_i(\theta_{-i}|h^{t-1}, \theta_i) + (1 - \chi) \left[\frac{\mu_i(\theta_{-i}|h^{t-1}, \theta_i) \sigma_{-i}(a^t_{-i}|h^{t-1}, \theta_{-i})}{\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \sigma_{-i}(a^t_{-i}|h^{t-1}, \theta'_{-i})} \right]$$

Proof: See Appendix A. ■

This is analogous to Lemma 1 of Eyster and Rabin (2005). Another insight provided by Lemma 2.1 is that even if player types are independently drawn, i.e., $F(\theta) = \prod_{i=1}^n F_i(\theta_i)$, players' cursed beliefs about other players' types are generally *not* independent across players. That is, in general, $\mu_i(\theta_{-i}|h^t, \theta_i) \neq \prod_{j \neq i} \mu_{ij}(\theta_j|h^t, \theta_i)$. The belief system will preserve the independence only when the players are either fully rational ($\chi = 0$) or fully cursed ($\chi = 1$).

Finally, we place a consistency restriction, analogous to consistent assessments in sequential equilibrium, on how χ -cursed beliefs are updated off the equilibrium path, i.e., when

$$\sum_{\theta'_{-i} \in \Theta_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \sigma_{-i}^\chi(a^t_{-i}|h^{t-1}, \theta'_{-i}, \theta_i) = 0.$$

An assessment satisfies χ -consistency if it is in the closure of Ψ^χ .

Definition 2.2. *(μ, σ) satisfies χ -consistency if there is a sequence of assessments $\{(\mu^k, \sigma^k)\} \subseteq \Psi^\chi$ such that $\lim_{k \rightarrow \infty} (\mu^k, \sigma^k) = (\mu, \sigma)$.*

For any $i \in N$, $\chi \in [0, 1]$, σ , and $\theta \in \Theta$, let $\rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi, \sigma_i)$ be i 's perceived conditional realization probability of terminal history $h^T \in \mathcal{H}^T$ at history $h^t \in \mathcal{H} \setminus \mathcal{H}^T$ if the type profile is θ and i uses the behavioral strategy σ_i whereas perceives other players' using the cursed behavioral strategy σ_{-i}^χ . At every non-terminal history h^t , a χ -cursed player in χ -CSE will use χ -cursed Bayes' rule (Definition 2.1) to derive the posterior belief about the other players' types. Accordingly, a type θ_i player i 's conditional expected payoff at history h^t is

$$\mathbb{E}u_i(\sigma|h^t, \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i(\theta_{-i}|h^t, \theta_i) \rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi, \sigma_i) u_i(h^T, \theta_{-i}, \theta_i).$$

Definition 2.3. *An assessment (μ^*, σ^*) is a χ -cursed sequential equilibrium if it satisfies χ -consistency and $\sigma_i^*(h^t, \theta_i)$ maximizes $\mathbb{E}u_i(\sigma^*|h^t, \theta_i)$ for all i , θ_i , $h^t \in \mathcal{H} \setminus \mathcal{H}^T$.*

2.3 General Properties of χ -CSE

In this section, we characterize some general theoretical properties of χ -CSE. The first result is the existence of the χ -CSE. The definition of χ -CSE mirrors the definition of the sequential equilibrium by Kreps and Wilson (1982)—the only difference is that players in χ -CSE update their beliefs by χ -cursed Bayes' rule and best respond to χ -cursed (behavioral) strategies. Therefore, one can prove the existence of χ -CSE in a similar way as in the standard argument of the existence of sequential equilibrium.

Proposition 2.1. *For any $\chi \in [0, 1]$ and any finite multi-stage game with observed actions, there is at least one χ -CSE.*

Proof: We briefly sketch the proof here, and the details can be found in Appendix A. Fix any $\chi \in [0, 1]$. For any $i \in N$ and any $\mathcal{I}_i = (h^{t-1}, \theta_i)$, player i has to choose every action $a_i^t \in A_i(h^{t-1})$ with probability at least ϵ . Since there are no off-path histories, the belief system is uniquely pinned down by χ -cursed Bayes' rule and a χ -CSE exists in this ϵ -constrained game. We denote this χ -CSE as $(\mu^\epsilon, \sigma^\epsilon)$. By compactness, there is a converging sub-sequence of assessments such that $(\mu^\epsilon, \sigma^\epsilon) \rightarrow (\mu^*, \sigma^*)$ as $\epsilon \rightarrow 0$, which is a χ -CSE. ■

Let $\Phi(\chi)$ be the correspondence that maps $\chi \in [0, 1]$ to the set of χ -CSE. Proposition 2.1 guarantees $\Phi(\chi)$ is non-empty for any $\chi \in [0, 1]$. Because χ -cursed Bayes' rule changes continuously in χ , we further prove that $\Phi(\chi)$ is an upper hemi-continuous correspondence.

Proposition 2.2. *$\Phi(\chi)$ is upper hemi-continuous with respect to χ .*

Proof: The proof follows a standard argument. See Appendix A for details. ■

As shown in Corollary 2.1, a direct consequence of upper hemi-continuity is that every limit point of a sequence of χ -CSE when $\chi \rightarrow 0$ is a sequential equilibrium. This result bridges our behavioral equilibrium concept with standard equilibrium theory.

Corollary 2.1. *Every limit point of a sequence of χ -CSE with χ converging to 0 is a sequential equilibrium.*

Proof: By Proposition 2.2, we know $\Phi(\chi)$ is upper hemi-continuous at 0. Consider of a sequence of χ -CSE. As $\chi \rightarrow 0$, the limit point remains a CSE, which is a sequential equilibrium at $\chi = 0$. This completes the proof. ■

Finally, by a similar argument to Kreps and Wilson (1982), for any $\chi \in [0, 1]$, χ -CSE is also upper hemi-continuous with respect to payoffs. In other words, our χ -CSE preserves the continuity property of sequential equilibrium.

The next result is the characterization of a necessary condition for χ -CSE. As seen from Lemma 2.1, players update their beliefs more passively in χ -CSE than in the standard equilibrium—they put χ -weight on their beliefs formed in previous stage. To formalize this, we define the χ -dampened updating property in Definition 2.4. An assessment satisfies this property if at *any* non-terminal history, the belief puts at least χ weight on the belief in previous stage—both on and off the equilibrium path. In Proposition 2.3, we show that χ -consistency implies the χ -dampened updating property.

Definition 2.4. *An assessment (μ, σ) satisfies the χ -dampened updating property if for any $i \in N$, $\theta \in \Theta$ and $h^t = (h^{t-1}, a^t) \in \mathcal{H} \setminus \mathcal{H}^T$,*

$$\mu_i(\theta_{-i}|h^t, \theta_i) \geq \chi \mu_i(\theta_{-i}|h^{t-1}, \theta_i).$$

Proposition 2.3. *χ -consistency implies χ -dampened updating for any $\chi \in [0, 1]$.*

Proof: See Appendix A. ■

It follows that if assessment (μ, σ) satisfies the χ -dampened updating property, then for any player i , any history h^t and any type profile θ , player i 's belief about θ_{-i} is bounded by

$$\chi \mu_i(\theta_{-i}|h^{t-1}, \theta_i) \leq \mu_i(\theta_{-i}|h^t, \theta_i) \leq 1 - \chi \sum_{\theta'_{-i} \neq \theta_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i).$$

One can see from this condition that when χ increases, the feasible range of $\mu_i(\theta_{-i}|h^t, \theta_i)$ shrinks, and the restriction on the belief system becomes more stringent. Moreover, if the history h^t is an off-path history of (μ, σ) , then this condition characterizes the feasible set of off-path beliefs, which shrinks as χ increases.

An important implication of this observation is that $\Phi(\chi)$ is not lower hemi-continuous with respect to χ . The intuition is that for some χ -CSE that contains

off-path histories, the off-path beliefs to support the equilibrium might not be χ -consistent for sufficiently large χ . In this case, the χ -CSE is not attainable by a sequence of χ_k -CSE where χ_k converges to χ from above, causing the lack of lower hemi-continuity.⁶

Lastly, another implication of χ -dampened updating property is that for each player i , history h^t and type profile θ , the belief $\mu_i(\theta_{-i}|h^t, \theta_i)$ has a lower bound that is *independent* of the strategy profile. The lower bound is characterized in Corollary 2.2. This result implies that when $\chi > 0$, $F(\theta_{-i}|\theta_i) > 0$ implies $\mu_i(\theta_{-i}|h^t, \theta_i) > 0$ for all h^t , so that if prior beliefs are bounded away from zero, beliefs are always bounded away from 0 as well. In other words, when $\chi > 0$, because of the χ -dampened updating, beliefs will always have full support even if at off-path histories.

Corollary 2.2. *For any χ -consistent assessment (μ, σ) , player $i \in N$, type profile $\theta \in \Theta$ and $h^t \in \mathcal{H} \setminus \mathcal{H}^T$,*

$$\mu_i(\theta_{-i}|h^t, \theta_i) \geq \chi^t F(\theta_{-i}|\theta_i)$$

Proof: See Appendix A. ■

If the game has only one stage, then the dampened updating property has no effect, in which case χ -CSE and χ -CE are equivalent solution concepts. This is formally stated and proved in Proposition 2.4.

Proposition 2.4. *For any one-stage game and for any $\chi \in [0, 1]$, χ -CSE and χ -CE are equivalent.*

Proof: For any one-stage game, the only public history is the initial history h_0 . Thus, in any χ -CSE, for each player $i \in N$ and type profile $\theta \in \Theta$, player i 's belief about other players' types at this history is $\mu_i(\theta_{-i}|h_0, \theta_i) = F(\theta_{-i}|\theta_i)$. Since the game has only one stage, the outcome is simply $a^1 = (a_1^1, \dots, a_n^1)$, the action profile at stage 1. Moreover, given any behavioral strategy profile σ , player i believes a^1 will be the outcome with probability

$$\sigma_i(a_i^1|h_0, \theta_i) \times [\chi \bar{\sigma}_{-i}(a_{-i}^1|h_0, \theta_i) + (1 - \chi)\sigma_{-i}(a_{-i}^1|h_0, \theta_{-i})].$$

Therefore, if σ is the behavioral strategy profile of a χ -CSE in an one-stage game, then for each player i , type $\theta_i \in \Theta_i$ and each $a_i^1 \in A_i(h_0)$ such that $\sigma_i(a_i^1|h_0, \theta_i) > 0$,

⁶An example is provided in Section 2.4.1 (see Footnote 7).

$$a_i^1 \in \operatorname{argmax}_{a_i^{1'} \in A_i(h_0)} \sum_{\theta_{-i} \in \Theta_{-i}} F(\theta_{-i} | \theta_i) \times \left\{ \sum_{a_{-i}^1 \in A_{-i}(h_0)} [\chi \bar{\sigma}_{-i}(a_{-i}^1 | h_0, \theta_i) + (1 - \chi) \sigma_{-i}(a_{-i}^1 | h_0, \theta_{-i})] \right\} u_i(a_{-i}^1, a_i^{1'}, \theta_{-i}, \theta_i),$$

which coincides with the maximization problem of χ -CE. ■

From the proof of Proposition 2.4, one can see that in one-stage games players have *correct* perceptions about the average strategy of others. Therefore, the maximization problem of χ -CSE coincides with the problem of χ -CE. For general multi-stage games, because of the χ -dampened updating property, players will update beliefs incorrectly and thus their perceptions about other players' future moves can also be distorted.

2.4 Applications

In this section, we will explore χ -CSE in five applications of multi-stage games with observed actions, in order to illustrate the range of effects it can have and to show how it is different from the χ -CE and sequential equilibrium. The omitted proofs in this section can be found in Appendix A.

Our first application is the sender-receiver signaling game, which is practically the simplest possible multi-stage game. From our analysis, we will see both the theoretical and empirical implications of our χ -CSE.

2.4.1 Pooling Equilibria in Signaling Games

We first make a general observation about pooling equilibria in multi-stage games. Player j follows a *pooling strategy* if for every non-terminal history, h^t , all types of player j take the same action $a_j^{t+1} \in A_j(h^t)$. Conceptually, since every type of player j takes the same action, players other than j cannot make any inference about j 's type from j 's actions. A *pooling χ -CSE* is a χ -CSE where every player follows a pooling strategy. Hence, every player has correct beliefs about any other player's future move because every type of every player chooses the same action.

Since in any pooling χ -CSE, players can correctly anticipate other players' future moves no matter how cursed they are, one may naturally conjecture that a pooling χ -CSE is also a χ' -CSE for any $\chi' \in [0, 1]$. As shown by Eyster and Rabin (2005),

this is true for one-stage Bayesian games: if a pooling strategy profile is a χ -cursed equilibrium, then it is also a χ' -cursed equilibrium for any $\chi' \in [0, 1]$. Surprisingly, this result does not extend to multi-stage games. Proposition 2.5 shows if a pooling behavioral strategy profile is a χ -CSE, then it remains a χ' -CSE only for $\chi' \leq \chi$, which is a weaker result than Eyster and Rabin (2005).

This result is driven by the χ -dampened updating property which restricts the set of off-path beliefs. As discussed above, when χ gets larger, the set of feasible off-path beliefs shrinks, eliminating some pooling χ -CSE.

Proposition 2.5. *A pooling χ -CSE is a χ' -CSE for $\chi' \leq \chi$.*

Proof: See Appendix A. ■

The proof strategy is similar to the one in Eyster and Rabin (2005) Proposition 3. Given a χ -CSE behavioral strategy profile, we can separate the histories into on-path and off-path histories. For on-path histories in a pooling equilibrium, since all types of players make the same decisions, players cannot make any inference about other players' types. Therefore, for on-path histories, their beliefs are the prior beliefs, which are independent of χ . On the other hand, for off-path histories, as shown in Proposition 2.3, a necessary condition for χ -CSE is that the belief system has to satisfy the χ -dampened updating property. As χ gets larger, this requirement becomes more stringent, and hence some pooling χ -CSE may break down.

Example 1 is a signaling game where the sender has only two types and two messages, and the receiver has only two actions. This example demonstrates the implication of Proposition 2.5 and shows the lack of lower hemi-continuity; i.e., it is possible for a pooling behavioral strategy profile to be a χ -CSE, but not a χ' -CSE for $\chi' > \chi$. We will also use this example to illustrate how the notion of cursedness in sequential cursed equilibrium proposed by Cohen and Li (2023) departs from CSE.

Example 1. The sender has two possible types drawn from the set $\Theta = \{\theta_1, \theta_2\}$ with $\Pr(\theta_1) = 1/4$. The receiver does not have any private information. After the sender's type is drawn, the sender observes his type and decides to send a message $m \in \{A, B\}$, or any mixture between the two. After that, the receiver decides between action $a \in \{L, R\}$ or any mixture between the two, and the game ends. The game tree is illustrated in Figure 2.1.




Figure 2.1: Game Tree for Example 1.

If we solve for the χ -CE of the game (or the sequential equilibria), we find that there are two pooling equilibria for every value of χ . In the first pooling χ -CE, both sender types choose A; the receiver chooses L in response to A and R at the off-path history B. In the second pooling χ -CE, both sender types pool at B and the receiver chooses R at both histories. By Proposition 3 of Eyster and Rabin (2005), these two equilibria are pooling χ -CE for all $\chi \in [0, 1]$. The intuition is that in a pooling χ -CE, players are not able to make any inference about other players' types from their actions because the average normal form strategy is the same as the type-conditional normal form strategy. Therefore, their beliefs are independent of χ , and hence a pooling χ -CE will still be an equilibrium for any $\chi \in [0, 1]$.

However, as summarized in Claim 2.1 below, the χ -CSE imposes *stronger* restrictions than χ -CE in this example, in the sense that when χ is sufficiently large, the second pooling equilibrium cannot be supported as a χ -CSE. The key reason is that when the game is analyzed in its normal form, the χ -dampened updating property shown in Proposition 2.3 does not have any bite, allowing both pooling equilibria to be supported as a χ -CE for any value of χ . Yet in the χ -CSE analysis, the additional restriction of χ -dampened updating property eliminates some extreme off-path beliefs, and hence, eliminates the second pooling χ -CSE equilibrium for sufficiently large χ . For simplicity, we use a four-tuple $[(m(\theta_1), m(\theta_2)); (a(A), a(B))]$ to denote a behavioral strategy profile.

Claim 2.1. *In this example, there are two pure pooling χ -CSE, which are:*

1. $[(A, A); (L, R)]$ is a pooling χ -CSE for any $\chi \in [0, 1]$.

2. $[(B, B); (R, R)]$ with $\mu_2(\theta_1|A) \in [\frac{1}{3}, 1 - \frac{3}{4}\chi]$ is a pooling χ -CSE if and only if $\chi \leq 8/9$.

From previous discussion, we know in general, the sets of χ -CSE and χ -CE are non-overlapping because of the nature of sequential distortion of beliefs in χ -CSE. Yet a pooling χ -CSE is an exception. In a pooling χ -CSE, players can correctly anticipate others' future moves, so a pooling χ -CSE will mechanically be a pooling χ -CE. In cases such as this, we can find that χ -CSE is a *refinement* of χ -CE.⁷

Remark 2.1. This game is useful for illustrating some of the differences between the notions of “cursedness” in χ -CSE and the sequential cursed equilibrium $((\chi_S, \psi_S)$ -SCE) proposed by Cohen and Li (2023).⁸ The first distinction is that the χ and χ_S parameters capture substantively different sources of distortion in a player's beliefs about the other players' strategies. In χ -CSE, the degree of cursedness, χ , captures how much a player neglects the dependence of the other players' behavioral strategies on those players' (exogenous) *private information*, i.e, types, drawn by nature, and as a result, mistakenly treats different types as behaving the same with probability χ . In contrast, in (χ_S, ψ_S) -SCE, the cursedness parameter, χ_S , captures how much a player neglects the dependence of the other players' strategies on future moves of the others, or current moves that are unobserved because of simultaneous play. Thus, it is a neglect related to endogenous information. If player i *observes* a previous move by some other player j , then player i correctly accounts for the dependence of player j 's chosen action on player j 's private type, as would be the case in χ -CSE only at the boundary where $\chi = 0$.

In the context of pooling equilibria in sender-receiver signaling games, if $\chi_S = 1$, then in SCE the sender believes the receiver will respond the same way both on and off the equilibrium path. This distorts how the sender perceives the receiver's future action in response to an off-equilibrium path message. In χ -CSE, cursedness does not hinder the sender from correctly perceiving the receiver's strategy since the receiver only has one type. Take the strategy profile $[(A, A); (L, R)]$ for example, which is a pooling χ -CSE equilibrium for all $\chi \in [0, 1]$. However, with (χ_S, ψ_S) -SCE, a sender misperceives that the receiver, upon receiving the off-path message

⁷Note that the χ -CSE correspondence $\Phi(\chi)$ is not lower hemi-continuous with respect to χ . To see this, we consider a sequence of $\{\chi_k\}$ where $\chi_k = \frac{8}{9} + \frac{1}{9k}$ for $k \geq 1$. From the analysis of Claim 2.1, we know $[(B, B); (R, R)] \notin \Phi(\chi_k)$ for any $k \geq 1$. However, in the limit where $\chi_k \rightarrow 8/9$, $[(B, B); (R, R)]$ with $\mu_2(\theta_1|A) = 1/3$ is indeed a CSE. That is, $[(B, B); (R, R)]$ is not approachable by this sequence of χ_k -CSE.

⁸To avoid confusion, we will henceforth add the subscript “S” to the parameters of SCE.

B , will, with probability χ_S , take the same action (L) as when receiving the on-path message A . If χ_S is sufficiently high, the sender will deviate to send B , which implies that $[(A, A); (L, R)]$ cannot be supported as an equilibrium when χ_S is sufficiently large ($\chi_S > 1/3$). The distortion induced by χ_S also creates an additional SCE if χ_S is sufficiently large: $[(B, B); (L, R)]$. To see this, if $\chi_S = 1$, then a sender incorrectly believes that the receiver will continue to choose R if the sender deviates to A , rather than switching to L , and hence B is optimal for both sender types. However, $[(B, B); (L, R)]$ is *not* a χ -CSE equilibrium for any $\chi \in [0, 1]$, or a χ -CE in the sense of Eyster and Rabin (2005), or a sequential equilibrium.

In the two possible pooling equilibria analyzed in the last paragraph, the second SCE parameter, ψ_S , does not have any effect, but the role of ψ_S can be illustrated in the context of the $[(B, B); (R, R)]$ sequential equilibrium. This second SCE parameter, ψ_S , is introduced to accommodate a player's possible failure to fully account for the informational content from *observed* events. The larger $(1 - \psi_S)$ is, the greater extent a player neglects the informational content of observed actions. Although the parameter ψ_S has a similar flavor to $1 - \chi$ in χ -CSE, it is different in a number of ways. In particular this parameter only has an effect via its interaction with χ_S and thus does not independently arise. In the two parameter model, the overall degree of cursedness is captured by the product, $\chi_S(1 - \psi_S)$, and thus any cursedness effect of ψ_S is shut down when $\chi_S = 0$. For instance, under our χ -CSE, the strategy profile $[(B, B); (R, R)]$ can only be supported as an equilibrium when χ is sufficiently small. However, $[(B, B); (R, R)]$ can be supported as a (χ_S, ψ_S) -SCE even when $(1 - \psi_S) = 1$ as long as χ_S is sufficiently small. In fact, when $\chi_S = 0$, a (χ_S, ψ_S) -SCE is equivalent to sequential equilibrium regardless of the value of ψ_S . See Appendix B and Fong, Lin, and Palfrey (2023) for a more detailed discussion.

Example 2. Here we analyze two signaling games that were studied experimentally by Brandts and Holt, 1993 (*BH 3* and *BH 4*) and show that χ -CSE can help explain some of their findings. In both Game *BH 3* and *BH 4*, the sender has two possible types $\{\theta_1, \theta_2\}$ which are equally likely. There are two messages $m \in \{I, S\}$ available to the sender.⁹ After seeing the message, the receiver chooses an action from $a \in \{C, D, E\}$. The game tree and payoffs for both games are plotted in Figure 2.2.

In both games, there are two pooling sequential equilibria. In the first equilibrium, both sender types send message I , and the receiver will choose C in response to I

⁹ I stands for “Intuitive” and S stands for “Sequential but not intuitive”, corresponding to the two pooling sequential equilibria of the two games.




Figure 2.2: Game Tree for *BH 3* and *BH 4* in Brandts and Holt (1993).

and *D* in response to *S*. In the second equilibrium, both sender types send message *S*, and the receiver will choose *D* in response to *I* while choose *C* in response to *S*. Both are sequential equilibria, in both games, but only the first equilibrium where the sender sends *I* satisfies the intuitive criterion proposed by Cho and Kreps (1987).

Since the equilibrium structure is similar in both games, the sequential equilibrium and the intuitive criterion predict the behavior should be the same in both games. However, this prediction is strikingly rejected by the data. Brandts and Holt (1993) report that in the later rounds of the experiment, almost all type θ_1 senders send *I* in Game *BH 3* (97%), and yet all type θ_1 senders send *S* in Game *BH 4* (100%). In contrast, type θ_2 senders behave similarly in both games—46.2% and 44.1% of type θ_2 senders send *I* in Games *BH 3* and *BH 4*, respectively. Qualitatively speaking, the empirical pattern reported by Brandts and Holt (1993) is that *sender type θ_1 is more likely to send I in Game BH 3 than Game BH 4 while sender type θ_2 's behavior is insensitive to the change of games.*

To explain this finding, Brandts and Holt (1993) propose a descriptive story based on *naive receivers*. A naive receiver will think both sender types are equally likely, regardless of which message is observed. This naive reasoning will lead the receiver to choose *C* in both games. Given this naive response, a type θ_1 sender has an incentive to send *I* in Game *BH 3* and choose *S* in Game *BH 4*. (Brandts and Holt, 1993, p. 284 – 285)

In fact, their story of naive reasoning echoes the logic of χ -CSE. When the receiver is fully cursed (or naive), he will ignore the correlation between the sender's action and type, causing him to not update the belief about the sender's type. Proposition 2.6 characterizes the set of χ -CSE of both games. Following the previous notation, we use a four-tuple $[(m(\theta_1), m(\theta_2)); (a(I), a(S))]$ to denote a behavioral strategy profile.

Proposition 2.6. *The set of χ -CSE of Game BH 3 and BH 4 are that:*

- *In Game BH 3, there are three pure χ -CSE:*
 1. $[(I, I); (C, D)]$ is a pooling χ -CSE if and only if $\chi \leq 4/7$.
 2. $[(S, S); (D, C)]$ is a pooling χ -CSE if and only if $\chi \leq 2/3$.
 3. $[(I, S); (C, C)]$ is a separating χ -CSE if and only if $\chi \geq 4/7$.
- *In Game BH 4, there are three pure χ -CSE:*
 1. $[(I, I); (C, D)]$ is a pooling χ -CSE if and only if $\chi \leq 4/7$.
 2. $[(S, S); (D, C)]$ is a pooling χ -CSE if and only if $\chi \leq 2/3$.
 3. $[(S, S); (C, C)]$ is a pooling χ -CSE for any $\chi \in [0, 1]$.

As noted earlier for Example 1, by Proposition 3 of Eyster and Rabin (2005), pooling equilibria (1) and (2) in games *BH 3* and *BH 4* survive as χ -CE for all $\chi \in [0, 1]$. Hence, Proposition 2.6 implies that χ -CSE refines the χ -CE pooling equilibria for larger values of χ . Moreover, χ -CSE actually eliminates all pooling equilibria in *BH 3* if $\chi > 2/3$. Proposition 2.6 also suggests that for any $\chi \in [0, 1]$, sender type θ_2 will behave similarly in both games, which is qualitatively consistent with the empirical pattern. In addition, χ -CSE predicts that a highly cursed ($\chi > 2/3$) type θ_1 sender will send different messages in different games—highly cursed type θ_1 senders will send *I* and *S* in Games *BH 3* and *BH 4*, respectively. This is consistent with the empirical data.

2.4.2 A Public Goods Game with Communication

Our second application is a threshold public goods game with private information and pre-play communication, variations of which have been studied in laboratory experiments (Palfrey and Rosenthal, 1991; Palfrey, Rosenthal, and Roy, 2017). Here

we consider the “unanimity” case where there are N players and the threshold is also N .

Each player i has a private cost parameter c_i , which is independently drawn from a uniform distribution on $[0, K]$ where $K > 1$. After each player’s c_i is drawn, each player observes their own cost, but not the others’ costs. Therefore, c_i is player i ’s private information and corresponds to θ_i in the general formulation.¹⁰ The game consists of two stages. After the profile of cost parameters is drawn, the game will proceed to stage 1 where each player simultaneously broadcasts a public message $m_i \in \{0, 1\}$ without any cost or commitment. After all players observe the message profile from this first stage, the game proceeds to stage 2 which is a unanimity threshold public goods game. Player i has to pay the cost c_i if he contributes, but the public good will be provided only if all players contribute. The public good is worth a unit of payoff for every player. Thus, if the public good is provided, each player’s payoff will be $1 - c_i$.

If there is no communication stage, the unique Bayesian Nash equilibrium is that no player contributes, which is also the unique χ -CE for any $\chi \in [0, 1]$. In contrast, with the communication stage, there exists an efficient sequential equilibrium where each player i sends $m_i = 1$ if and only if $c_i \leq 1$ and contributes if and only if all players send 1 in the first stage.¹¹ Since this is a private value game, the standard cursed equilibrium has no bite, and this efficient sequential equilibrium is also a χ -CE for all values of χ , by Proposition 2 of Eyster and Rabin (2005). In the following, we demonstrate that the prediction of χ -CSE is different from CE (and sequential equilibrium).

To analyze the χ -CSE, consider a set of “cutoff” costs, $\{C_c^\chi, C_0^\chi, C_1^\chi, \dots, C_N^\chi\}$. In the communication stage, each player communicates the message $m_i = 1$ if and only if $c_i \leq C_c^\chi$. In the second stage, if there are exactly $0 \leq k \leq N$ players sending $m_i = 1$ in the first stage, then such a player would contribute in the second stage if and only if $c_i \leq C_k^\chi$. A χ -CSE is a collection of these cost cutoffs such that the associated strategies are a χ -CSE for the public goods game with communication. The most efficient sequential equilibrium identified above for $\chi = 0$ corresponds to cutoffs with $C_0^0 = C_1^0 = \dots = C_{N-1}^0 = 0$ and $C_c^0 = C_N^0 = 1$.

¹⁰This application has a continuum of types. The framework of analysis developed for finite types is applied in the obvious way.

¹¹One can think of the first stage as a poll, where players are asked the following question: “Are you willing to contribute if everyone else says they are willing to contribute?” The message $m_i = 1$ corresponds to a “yes” answer and the message $m_i = 0$ corresponds to a “no” answer.

There are in fact multiple equilibria in this game with communication. In order to demonstrate how the cursed belief can distort players' behavior, here we will focus on the χ -CSE that is similar to the most efficient sequential equilibrium identified above, where $C_0^\chi = C_1^\chi = \dots = C_{N-1}^\chi = 0$ and $C_c^\chi = C_N^\chi$. The resulting χ -CSE is given in Proposition 2.7.

Proposition 2.7. *In the public goods game with communication, there is a χ -CSE where*

1. $C_0^\chi = C_1^\chi = \dots = C_{N-1}^\chi = 0$, and
2. there is a unique $C^*(N, K, \chi) \leq 1$ s.t. $C_c^\chi = C_N^\chi = C^*(N, K, \chi)$ that solves:

$$C^*(N, K, \chi) - \chi \left[\frac{C^*(N, K, \chi)}{K} \right]^{N-1} = 1 - \chi.$$

To provide some intuition, we sketch the proof by analyzing the two-person game, where the χ -CSE is characterized by four cutoffs $\{C_c^\chi, C_0^\chi, C_1^\chi, C_2^\chi\}$, with $C_0^\chi = C_1^\chi = 0$ and $C_c^\chi = C_2^\chi$. If players use the strategy that they would send message 1 if and only if the cost is less than C_c^χ , then by Lemma 2.1, at the history where both players send 1, player i 's cursed posterior belief density would be

$$\mu_i^\chi(c_{-i}|\{1, 1\}) = \begin{cases} \chi \cdot \left(\frac{1}{K}\right) + (1 - \chi) \cdot \left(\frac{1}{C_c^\chi}\right) & \text{if } c_{-i} \leq C_c^\chi \\ \chi \cdot \left(\frac{1}{K}\right) & \text{if } c_{-i} > C_c^\chi. \end{cases}$$

Notice that cursedness leads a player to put some probability weight on a type that is not compatible with the history. Namely, for χ -cursed players, when seeing another player sending 1, they still believe the other player might have $c_{-i} > C_c^\chi$. When χ converges to 1, the belief simply collapses to the prior belief as fully cursed players never update their beliefs. On the other hand, when χ converges to 0, the belief converges to $1/C_c^\chi$, which is the correct Bayesian inference.

Given this cursed belief density, the optimal cost cutoff to contribute, C_2^χ , solves

$$C_2^\chi = \int_0^{C_2^\chi} \mu_i^\chi(c_{-i}|\{1, 1\}) dc_{-i}.$$

Finally, at the first stage cutoff equilibrium, the C_c^χ type of player would be indifferent between sending 1 and 0 at the first stage. Therefore, C_c^χ satisfies

$$0 = \left(\frac{C_c^\chi}{K}\right) \left\{ -C_c^\chi + \int_0^{C_2^\chi} \mu_i^\chi(c_{-i}|\{1, 1\}) dc_{-i} \right\}.$$

After substituting $C_c^\chi = C_2^\chi$, we obtain the χ -CSE satisfies $C_c^\chi = C_2^\chi = (K - K\chi)/(K - \chi)$.

From this expression, one can see that the cutoff C_c^χ (as well as C_2^χ) is decreasing in χ and K . When $\chi \rightarrow 0$, C_c^χ converges to 1, which is the cutoff of the sequential equilibrium. On the other hand, when $\chi \rightarrow 1$, C_c^χ converges to 0, so there is no possibility for communication when players are fully cursed. Similarly, when $K \rightarrow 1$, C_c^χ converges to 1, which is the cutoff of the sequential equilibrium, while $\lim_{K \rightarrow \infty} C_c^\chi = 1 - \chi$.

These comparative statics results with respect to χ and K are not just a special property of the $N = 2$ case, but hold for all $N > 1$. Furthermore, there is a similar effect of increasing N that results in a lower cutoff (less effective communication). These properties of $C^*(N, K, \chi)$ are summarized in Corollary 2.3.

Corollary 2.3. *The efficient χ -CSE predicts for all $N \geq 2$ and $K > 1$:*

1. $C^*(N, K, 0) = 1$ and $C^*(N, K, 1) = 0$.
2. $C^*(N, K, \chi)$ is strictly decreasing in N , K , and χ for any $\chi \in (0, 1)$.
3. For all $\chi \in [0, 1]$, $\lim_{N \rightarrow \infty} C^*(N, K, \chi) = \lim_{K \rightarrow \infty} C^*(N, K, \chi) = 1 - \chi$.

These properties are illustrated in Figure 2.3. The left panel illustrates the equilibrium condition for C^* in a graph where the horizontal axis is $C \in [0, K]$. We can rewrite the characterization of $C^*(N, K, \chi)$ in Proposition 2.7 as a solution for C to the following equation:

$$\frac{1 - C}{\chi} = 1 - \left[\frac{C}{K} \right]^{N-1}.$$

The left panel displays the LHS of this equation, $\frac{1-C}{\chi}$, as the downward sloping line that connects the points $(0, \frac{1}{\chi})$ and $(1, 0)$. The RHS is displayed for $N = 2$ and $N = 3$ by the two curves that connect the points $(0, 1)$ and $(K, 0)$. The equilibrium, $C^*(N, K, \chi)$, is given by the (unique) intersection of the LHS and RHS curves. It is easy to see that $C^*(N, K, \chi)$ is strictly decreasing in N , K , and χ . When N increases, the RHS increases for all $C \in (0, K)$, resulting in an intersection at a lower value of C . When K increases, again the RHS increases for all $C \in (0, K)$, and also the intercept of the RHS on the horizontal axis increases, leading to a similar effect; and when χ increases, the intercept of the LHS on the horizontal axis decreases, resulting in an intersection at a lower value of C . In addition, when N grows without

bound, the RHS approaches to 1 for $C < K$, resulting in a limiting intersection at $C^*(\infty, K, \chi) = 1 - \chi$. This is illustrated in the middle panel of Figure 2.3, which graphs $C^*(2, 1.5, \cdot)$, $C^*(3, 1.5, \cdot)$, and $C^*(\infty, 1.5, \cdot)$. A similar effect occurs for $K \rightarrow \infty$, illustrated in the right panel of Figure 2.3, which displays $C^*(2, 1.25, \cdot)$, $C^*(2, 1.5, \cdot)$, and $C^*(2, \infty, \cdot)$.




Figure 2.3: (Left) Illustration of the χ -CSE equilibrium condition when $K = 1.5$ and $\chi = 0.5$. (Middle) The χ -CSE cutoff $C^*(N, K, \chi)$ for $N = 2, 3$ and for $N \rightarrow \infty$ when $K = 1.5$. (Right) The χ -CSE cutoff $C^*(N, K, \chi)$ for $K = 1.25, 1.5$ and for $K \rightarrow \infty$ when $N = 2$.

An interesting takeaway of this analysis is that in the public goods game with communication, *cursedness limits information transmission*: χ -CSE predicts when players are more cursed (higher χ), it will be harder for them to effectively communicate in the first stage for efficient coordination in the second stage. Moreover, Corollary 2.3 shows this χ -CSE varies systematically with *all three parameters of the model*: N, K , and χ . In contrast, in the standard χ -CE, players best respond to the *average type-contingent strategy* rather than the average behavioral strategy. Since it is a private value game, players do not care about the distribution of types, only the distribution of actions. Thus, the prediction of standard CE coincides with the equilibrium prediction *for all N, K , and χ* . This seems behaviorally implausible and is also suggestive of an experimental design that varies the two parameters N and K , since the qualitative effects of changing these parameters are identified.

2.4.3 Reputation Building: The Centipede Game with Altruists

In order to further demonstrate the difference between χ -CE and χ -CSE, in this section we consider a variation of the centipede game with private information, as analyzed in Kreps (1990) and McKelvey and Palfrey (1992). This game is an illustration of reputation-building, where a selfish player imitates an altruistic type in order to develop a reputation for passing, which in turn entices the opponent to pass and leads to higher payoffs.




Figure 2.4: Four-Stage Centipede Game.

There are two players and four stages, and the game tree is shown in Figure 2.4. In stage 1, player one can choose either Take ($T1$) or Pass ($P1$). If she chooses $T1$, the game ends and the payoffs to players one and two are 4 and 1, respectively. If she chooses the action $P1$, the game continues and player two has a choice between take ($T2$) and pass ($P2$). If he chooses $T2$, the game ends and the payoffs to player one and two are 2 and 8, respectively. If he chooses $P2$, the game continues to the third stage where player one chooses between $T3$ and $P3$. Similar to the previous stages, if she chooses $T3$, the payoffs to player one and two are 16 and 4, respectively. If she chooses $P3$, the game proceeds to the last stage where player two chooses between $T4$ and $P4$. If player two chooses $T4$ the payoffs are 8 and 32, respectively. If player two alternatively chooses $P4$, the payoffs are 64 and 16, respectively.

Player one has two types, selfish and altruistic. Selfish types are assumed to have a utility function that is linear in their own payoff. Altruistic types instead have a utility function that is linear in the sum of the two payoffs. For the sake of simplicity, we assume that player two has only one type: selfish. The common probability that player one is altruistic is α . Player one knows her own type, but player two does not. Thus, player one's type is her private information. In the following, we focus on the interesting case where $\alpha \leq 1/7$.¹²

¹²If $\alpha > 1/7$, player two always chooses $P2$ in the second stage since the probability of encountering altruistic player one is sufficiently high. Selfish player one would thus choose $P1$ in the first stage and choose $T3$ in the third stage.

Because this is a game of incomplete information with *private values*, the standard χ -CE is equivalent to the Bayesian Nash equilibrium of the game for all $\chi \in [0, 1]$, and yields the same take probabilities as the Bayesian equilibrium. Since altruistic player one wants to maximize the sum of the payoffs, it is optimal for her to always pass. The equilibrium behavior is summarized in Claim 2.2.

Claim 2.2. *In the Bayesian Nash equilibrium, selfish player one will choose P_1 with probability $\frac{6\alpha}{1-\alpha}$ and choose T_3 with probability 1; player two will choose P_2 with probability $\frac{1}{7}$ and choose T_4 with probability 1.*

It is useful to see exactly why, in this example (and more generally) the standard χ -CE is the same as the perfect Bayesian equilibrium. In particular, why it is not the case that cursed beliefs will change player two's updating process after observing P_1 at stage one. Belief updating is *not* a property of the standard χ -CE as the analysis is in the strategic form, and thus is solved as a BNE of the game in the reduced normal form.¹³ Table 2.1 summarizes the payoff matrices in the reduced normal form of centipede game for selfish and altruistic type.

Table 2.1: Reduced Normal Form Centipede Game Payoff Matrix.

selfish ($1 - \alpha$)	T_2	P_2T_4	P_2P_4	altruistic (α)	T_2	P_2T_4	P_2P_4
T_1	4, 1	4, 1	4, 1	T_1	5, 1	5, 1	5, 1
P_1T_3	2, 8	16, 4	16, 4	P_1T_3	10, 8	20, 4	20, 4
P_1P_3	2, 8	8, 32	64, 16	P_1P_3	10, 8	40, 32	80, 16

It is easily verified that at the Bayesian Nash equilibrium, selfish player one would choose T_1 with probability $(1 - 7\alpha)/(1 - \alpha)$ and choose P_1T_3 with probability $6\alpha/(1 - \alpha)$, while player two would choose T_2 with probability $6/7$.

To solve the standard χ -CE, let selfish player one choose T_1 with probability p and P_1T_3 with probability $1 - p$. Let player two choose T_2 with probability q and P_2T_4 with probability $1 - q$. Notice that for player two, P_2P_4 is a dominated strategy and given this, it is also sub-optimal for selfish player one to choose P_1P_3 . In this case, selfish player one would choose T_1 if and only if $4 \geq 2q + 16(1 - q) \iff q \geq 6/7$, implying that selfish player one's best response correspondence in the standard cursed analysis coincides with the Bayesian Nash equilibrium analysis. On the other hand, to solve for player two's best responses we need to first solve for the *perceived* strategy. When player two is χ -cursed, he would think that player one

¹³The analysis is similar for the non-reduced normal form.

is using $\sigma_1^\chi(a|\theta)$ where $a \in \{T_1, P_1T_3, P_1P_3\}$ and $\theta \in \{\text{selfish, altruistic}\}$. Player one's true strategy is given in Table 2.2.

Table 2.2: Player one's True Strategy.

$\sigma_1(a \theta)$	player one's type	
	selfish	altruistic
T_1	p	0
P_1T_3	$1 - p$	0
P_1P_3	0	1

In this case, player one's *average strategy* is simply

$$\bar{\sigma}_1(T_1) = (1 - \alpha)p, \quad \bar{\sigma}_1(P_1T_3) = (1 - \alpha)(1 - p), \quad \bar{\sigma}_1(P_1P_3) = \alpha.$$

By definition, $\sigma_1^\chi(a|\theta) = \chi\bar{\sigma}_1(a) + (1 - \chi)\sigma_1(a|s)$ and hence $\sigma_1^\chi(a|\theta)$ is given in Table 2.3.

Table 2.3: Cursed Perception of Player one's Strategy.

$\sigma_1^\chi(a \theta)$	player one's type	
	selfish	altruistic
T_1	$p(1 - \chi\alpha)$	$p\chi(1 - \alpha)$
P_1T_3	$(1 - p)(1 - \chi\alpha)$	$(1 - p)\chi(1 - \alpha)$
P_1P_3	$\chi\alpha$	$1 - \chi + \chi\alpha$

From player two's perspective, given any action profile, player two's expected payoff is not affected by whether player one is selfish or altruistic. Hence, player two only cares about the *marginal* distribution of player one's actions. In this case, χ -cursed player two believes player one will choose $a \in \{T_1, P_1T_3, P_1P_3\}$ with probability $\bar{\sigma}_1(a)$. Therefore, it is optimal for player two to choose T_2 if and only if

$$\bar{\sigma}_1(T_1) + 8[1 - \bar{\sigma}_1(T_1)] \geq \bar{\sigma}_1(T_1) + 4\bar{\sigma}_1(P_1T_3) + 32\bar{\sigma}_1(P_1P_3) \iff p \leq \frac{1 - 7\alpha}{1 - \alpha},$$

implying player two's best responses in the standard cursed analysis also coincides with the Nash best responses. As a result, one concludes that standard χ -CE would make exactly the same prediction as the Bayesian Nash equilibrium regardless how cursed the players are.

In contrast, the χ -CSE will exhibit distortions to the *conditional beliefs* of player two, given that player one has passed, because player two incorrectly takes into account

how player one's choice to pass depended on player one's private information. In particular, it is harder to build a reputation, since a selfish type will have to imitate altruists in such a way that the true posterior on altruistic type conditional on a pass is higher than in the perfect Bayesian equilibrium, because the updating by player two about player one's type is dampened relative to this true posterior due to cursedness. This distorted belief updating will result in *less* passing by player one compared to the Bayesian equilibrium. Formally, the χ -CSE is described in Proposition 2.8.

Proposition 2.8. *In the χ -CSE, selfish player one will choose P1 with probability q_1^χ and choose T3 with probability 1; player two will choose P2 with probability q_2^χ and choose T4 with probability 1 where*

$$q_1^\chi = \begin{cases} \left[\frac{7\alpha - 7\alpha\chi}{1 - 7\alpha\chi} - \alpha \right] / (1 - \alpha) & \text{if } \chi \leq \frac{6}{7(1-\alpha)} \\ 0 & \text{if } \chi > \frac{6}{7(1-\alpha)} \end{cases} \quad \text{and} \quad q_2^\chi = \begin{cases} 1/7 & \text{if } \chi \leq \frac{6}{7(1-\alpha)} \\ 0 & \text{if } \chi > \frac{6}{7(1-\alpha)}. \end{cases}$$

In order to see how the cursedness affects the equilibrium behavior, here we focus on the case of $\chi \leq \frac{6}{7(1-\alpha)}$ where selfish player one and player two will both mix at stage one and two. Given selfish player one chooses P1 with probability q_1^χ , by Lemma 2.1, we know when the game reaches stage two, player two's belief about player one being altruistic becomes

$$\mu^\chi = \chi\alpha + (1 - \chi) \left[\frac{\alpha}{\alpha + (1 - \alpha)q_1^\chi} \right].$$

Here we see that when χ is larger, player two will update his belief more slowly. Therefore, in order to maintain indifference at the mixed equilibrium, selfish player one has to pass with *lower* probability so that P1 is a more informative signal to player two. Thus, to make player two indifferent between T2 and P2, the following condition must hold at the equilibrium:

$$\mu^\chi = \frac{1}{7} \iff q_1^\chi = \left[\frac{7\alpha - 7\alpha\chi}{1 - 7\alpha\chi} - \alpha \right] / (1 - \alpha).$$

To conclude this section, in Figure 2.5, we plot the probabilities of choosing P1 and P2 at χ -CSE when there is a five percent chance that player one is an altruist (i.e., $\alpha = 0.05$). From our analysis above, we can find that both the standard equilibrium theory and χ -CE predict selfish player one chooses P1 with probability 0.32 and player two chooses P2 with probability 0.14. Moreover, these probabilities are independent of χ . However, χ -CSE predicts when players are more cursed,




Figure 2.5: χ -CSE for the centipede game with altruistic players ($\alpha = 0.05$).

selfish player one is less likely to choose $P1$. When players are sufficiently cursed ($\chi \geq 0.91$), selfish player one and player two will never pass, i.e., behave as if there were no altruistic players.

2.4.4 Sequential Voting over Binary Agendas

In this section, we apply the concept of χ -CSE to the model of strategic binary amendment voting with incomplete information studied by Ordeshook and Palfrey (1988). Let $N = \{1, 2, 3\}$ denote the set of voters. These three voters will vote over three possible alternatives in $X = \{a, b, c\}$. Voting takes place in a two-stage agenda. In the first stage, voters vote between a and b . In the second stage, voters vote between c and the majority rule winner of the first stage. The majority rule winner of the second stage is the outcome.

Each voter i has three possible private-value types where $\Theta \in \{\theta_1, \theta_2, \theta_3\}$ is the set of possible types. Each voter's type is independently drawn from a common prior distribution of types, p . In other words, the probability of a voter being type θ_k is p_k . Each voter's type is their own private information. Each voter has the same type-dependent payoff function, which is denoted by $u(x|\theta)$ for any $x \in X$ and $\theta \in \Theta$. We summarize the payoff function with the following table.

Notice that $v \in (0, 1)$ is a parameter that measures the *intensity* of the second ranked outcome relative to the top ranked outcome. This intensity parameter, v , is assumed to be the same for all types of all voters. Because this is a game of private values,

		x		
		a	b	c
θ	$u(x \theta)$			
	θ_1	1	v	0
	θ_2	0	1	v
θ_3	v	0	1	

the standard χ -CE and the Bayesian Nash equilibrium coincide.

We use $a_i^1(\theta)$ to denote type θ voter i 's action at stage 1. As is standard in majority voting games we will focus on the analysis of symmetric pure-strategy equilibria where voters do not use weakly dominated strategies. In other words, we will consider $a_i^t(\cdot) = a_j^t(\cdot)$ for all $i, j \in N$, and will drop the subscript.

In this PBE (and χ -CE) all voters will vote sincerely in equilibrium *except for* type θ_1 voters at stage 1. To see this, first note that voting insincerely in the last stage is dominated and thus eliminated, so all types of voters vote for their preferred alternative on the last ballot. Second, voting sincerely in both stages is a dominant strategy for a type θ_2 voter, who prefers any lottery between b and c to either a or c . Third, voting sincerely in both stages is also dominant for a type θ_3 voter in the sense that, in the event that neither of the other two voters are type θ_3 , then any lottery between a and c is better than a vote between b and c since b (i.e., type θ_3 's least preferred alternative) will win.¹⁴ The PBE (and χ -CE) prediction about a type θ_1 voter's strategy at stage 1 is summarized in the following claim.

Claim 2.3. *The symmetric (undominated pure) PBE strategy for type θ_1 voters in the first stage can be characterized as follows.*

1. $a^1(\theta_1) = b$ is a PBE strategy if and only if $v \geq \frac{p_1}{p_1+p_2}$.
2. $a^1(\theta_1) = a$ is a PBE strategy if and only if $v \leq \frac{p_1}{p_1+p_3}$.

Proof: See Ordeshook and Palfrey (1988). ■

Claim 2.3 shows that, if v is relatively large, only type θ_1 voting sophisticatedly for b instead of sincerely for a can be supported by a PBE. Conditional on being pivotal, voting for b in the first stage guarantees an outcome of b and thus guarantees getting

¹⁴When there is another type θ_3 voter, the first ballot does not matter since their most preferred alternative c will always win in the second stage.

v , while voting for a leads to a lottery between a and c . Thus, when v is sufficiently high, a type θ_1 voter will have an incentive to strategically vote for b to avoid the risk of having c elected stage 2.

The analysis of a cursed sequential equilibrium is different from the standard cursed equilibrium in strategic form because the cursedness affects belief updating over the stages of the game, and players anticipate future play of the game. Due to the dynamics and the anticipation of future cursed behavior, such cursed behavior at later stages of a game can feedback and affect strategic behavior earlier in the game.

In the context of the two-stage binary amendment strategic voting model, cursed behavior and belief updating mean that voters in the first stage use the expected cursed beliefs in the second stage to compute the continuation values in the two continuation games of the second stage, either a vote between a and c or a vote between b and c . Because they have a cursed understanding about the relationship between types and the voting behavior in the first stage, this affects their predictions about which alternative wins in the second stage, conditional on which alternative wins in the first stage.

It is noteworthy that, given any $\chi \in [0, 1]$, all voters will still vote sincerely in χ -CSE *except for* type θ_1 voters at stage 1. As implied by Proposition 2.4, a voter in the last stage would act as if solving a maximization problem of χ -CE but under an (incorrectly) updated belief. Therefore, we can follow the same arguments as solving for the undominated Bayesian equilibrium and conclude that type θ_2 and θ_3 voters as well as type θ_1 voters at stage 2 will vote sincerely under a χ -CSE.

Proposition 2.9 establishes that the set of parameters v and p that can support a χ -CSE in which type θ_1 voters vote sophisticatedly for b shrinks as χ increases.

Proposition 2.9. *If $a^1(\theta_1) = b$ can be supported by a symmetric χ -CSE, then it can also be supported by a symmetric χ' -CSE for all $\chi' \leq \chi$.*

The intuition behind strategic voting over agendas mainly comes from the information content of *hypothetical* pivotal events. However, a cursed voter does not (fully) take such information into consideration, and thus becomes overly optimistic about his favorite alternative a being elected in the second stage. Therefore, a type θ_1 voter has a stronger incentive to deviate from sophisticated voting to sincere voting in stage 1 as χ increases.

Interestingly, the set of v and p that can support a χ -CSE in which type θ_1 voters vote sincerely for a does not necessarily expand as the level of cursedness becomes higher, as characterized in Proposition 2.10.

Proposition 2.10. *Given p and $v \in (0, 1)$, there exists $\tilde{\chi}(p, v)$ such that*

1. *If $v > \frac{p_1}{p_1+p_3}$, then $a^1(\theta_1) = a$ is a χ -CSE strategy if and only if $\chi \geq \tilde{\chi}(p, v)$;*
2. *If $v < \frac{p_1}{p_1+p_3}$, then $a^1(\theta_1) = a$ is a χ -CSE strategy if and only if $\chi \leq \tilde{\chi}(p, v)$.*

Thus, Proposition 2.10 shows that, when χ is sufficiently large, there are some values of (v, p) that cannot support sincere voting for type θ_1 voters under PBE (and χ -CE) but can support it under χ -CSE. Alternatively, there also exist some values of (v, p) that can support sincere voting under PBE but fail to support it under χ -CSE when χ is large.




Figure 2.6: χ -CSE for Sophisticated (left) and Sincere (right) Voting When $v = 0.7$.

To illustrate this, Figure 2.6 plots the set of p (fixing $v = 0.7$) that can support a χ -CSE for type θ_1 voters at stage 1 to vote sophisticatedly for b and sincerely for a . The left panel of Figure 2.6 shows that a sophisticated voting χ -CSE becomes harder to be supported as χ increases, as indicated by Proposition 2.9. For example, when $p \equiv (p_1, p_2, p_3) = (0.6, 0.3, 0.1)$, type θ_1 voters will not vote for second preferred alternative b if $\chi > 0.18$.

On the other hand, the right panel of Figure 2.6 shows that, while type θ_1 voters who sincerely vote for a at stage 1 cannot be supported under PBE when p_3 is large, they may emerge in a χ -CSE with sufficiently high χ . Also note that when p_2 is large, sincere voting by type θ_1 voters is no longer a χ -CSE with high χ . In such a sincere voting equilibrium, a fully rational type θ_1 voter knows there will be only one type θ_2 voter among the other two voters when being pivotal. As a result, whether to sincerely vote for a is determined by the ratio of p_1 to p_3 . When p_3 is large, sincere voting at stage 1 will likely lead to zero payoff for type θ_1 voters and thus cannot be a PBE strategy. However, cursed type θ_1 voters will take the possibility of having two type θ_2 voters into account since they are not correctly conditioning on pivotality. As a result, when p_2 is large, sincere voting at stage 1 will likely lead to zero payoff for type θ_1 voters, and thus cannot be a χ -CSE strategy with high χ , while voting sophisticatedly for b can likely secure a payoff of v .

2.4.5 The Dirty Faces Game

The dirty faces game was first described by Littlewood (1953) to study the relationship between common knowledge and behavior.¹⁵ There are several different variants of this game, but here we focus on a simplified version, the two-person dirty faces game, which was theoretically analyzed by Fudenberg and Tirole (1991a) and Lin (2023) and was experimentally studied by Weber (2001), Bayer and Chan (2007) and Lin (2023).

Let $N = \{1, 2\}$ be the set of players. For each $i \in N$, let $x_i \in \{O, X\}$ represent whether player i has a clean face (O) or a dirty face (X). Each player's face type is independently and identically determined by a commonly known probability $p = \Pr(x_i = X) = 1 - \Pr(x_i = O)$. Once the face types are drawn, each player i can observe the other player's face x_{-i} but not their own face.¹⁶ If there is at least one player with a dirty face, a public announcement of this fact is broadcast to both players at the beginning of the game. Let $\omega \in \{0, 1\}$ denote whether there is an announcement or not. If there is an announcement ($\omega = 1$), all players are informed there is at least one dirty face but not the identities. When $\omega = 0$, it is common knowledge to both players that their faces are clean and the game becomes trivial.

¹⁵The dirty faces game has also been reframed as the “cheating wives puzzle” (Gamow and Stern, 1958), the “cheating husbands puzzle” (Moses, Dolev, and Halpern, 1986), the “muddy children puzzle” (Barwise, 1981) and (Halpern and Moses, 1990), and the “red hat puzzle” (Hardin and Taylor, 2008).

¹⁶To fit into the framework, each player's “type” (their own private information) can be specified as “other players' faces.” That is, $\theta_i = x_{-i}$.

Hence, in the following, we will focus only on the interesting case where $\omega = 1$.

There are a finite number of $T \geq 2$ stages. In each stage, each player i simultaneously chooses $s_i \in \{U, D\}$. The game ends as soon as either player (or both) chooses D , or at the end of stage T in case neither player has chosen D . Actions are revealed at the end of each stage. Payoffs depend on own face types and action. If a player chooses D , he will get $\alpha > 0$ if he has a dirty face while receive -1 if he has a clean face. We assume that

$$p\alpha - (1 - p) < 0 \iff 0 < \bar{\alpha} \equiv \frac{\alpha}{(1 - p)(1 + \alpha)} < 1, \quad (2.1)$$

where $p\alpha - (1 - p)$ is the expected payoff of D when the belief of having a dirty face is p . Thus, Assumption (2.1) guarantees it is strictly dominated to choose D at stage 1 when observing a dirty face. In other words, players will be rewarded when correctly inferring the dirty face but penalized when wrongly claiming the dirty face.

The payoffs are discounted with a common discount factor $\delta \in (0, 1)$. To summarize, conditional on reaching stage t , each player's payoff function (which depends on their own face and action) can be written as

$$u_i(s_i|t, x_i = X) = \begin{cases} \delta^{t-1}\alpha & \text{if } s_i = D \\ 0 & \text{if } s_i = U \end{cases} \quad \text{and} \quad u_i(s_i|t, x_i = O) = \begin{cases} -\delta^{t-1} & \text{if } s_i = D \\ 0 & \text{if } s_i = U. \end{cases}$$

Therefore, a two-person dirty faces game is defined by a tuple $\langle p, T, \alpha, \delta \rangle$.

Since the game ends as soon as some player chooses D , the information sets of the game can be specified by the face type the player observes and the stage number. Thus a behavioral strategy can be represented as

$$\sigma : \{1, \dots, T\} \times \{O, X\} \rightarrow [0, 1],$$

which is a mapping from information sets to the probability of choosing D , where $\{O, X\}$ corresponds to a player's observation of the other player's face.

There is a unique sequential equilibrium. When observing a clean face, a player would immediately know his face is dirty. Hence, it is strictly dominant to choose D at stage 1 in this case. On the other hand, when observing a dirty face, because of Assumption (2.1), it is optimal for the player to choose U at stage 1. However, if the game proceeds to stage 2, the player would know his face is dirty because the other player would have chosen D at stage 1 if his face were clean and the

game would not have reached stage 2. This result is independent of the payoffs, the timing, the discount factor, and the (prior) probability of having a dirty face. The only assumption for this argument is common knowledge of rationality.

Alternatively, when players are “cursed,” they are not able to make perfect inferences from the other player’s actions. Specifically, since a cursed player has incorrect perceptions about the relationship between the other player’s actions and their private information after seeing the other player choose U in stage 1, a cursed player does not believe they have a dirty face for sure. At the extreme when $\chi = 1$, fully cursed players never update their beliefs. In the following, we will compare the predictions of the standard χ -CE and the χ -CSE. A surprising result is that there is always a *unique* χ -CE, but there can be multiple χ -CSE.

For the sake of simplicity, we focus on the characterization of pure strategy equilibrium in the following analysis. Since the game ends when some player chooses D , we can equivalently characterize a stopping strategy as a mapping from the observed face type to a stage in $\{1, 2, \dots, T, T + 1\}$ where $T + 1$ corresponds to the strategy of never stopping. Furthermore, both χ -CE and χ -CSE will be symmetric because if players were to stop at different stages, least one of the players would have a profitable deviation. Finally, we use $\hat{\sigma}^\chi(x_{-i})$ and $\tilde{\sigma}^\chi(x_{-i})$ to denote the equilibrium stopping strategies of χ -CE and χ -CSE, respectively.

We characterize the χ -CE in Proposition 2.11. Since χ -CE is defined for simultaneous move Bayesian games, to solve for the χ -CE, we need to look at the corresponding normal form where players simultaneously choose $\{1, 2, \dots, T, T + 1\}$ given the observed face type.

Proposition 2.11. *The χ -cursed equilibrium can be characterized as follows.*

1. *If $\chi > \bar{\alpha}$, the only χ -CE is that both players choose:*

$$\hat{\sigma}^\chi(O) = 1 \quad \text{and} \quad \hat{\sigma}^\chi(X) = T + 1.$$

2. *If $\chi < \bar{\alpha}$, the only χ -CE is that both players choosing*

$$\hat{\sigma}^\chi(O) = 1 \quad \text{and} \quad \hat{\sigma}^\chi(X) = 2.$$

Proposition 2.11 shows that χ -CE makes an extreme prediction, when observing a dirty face, players would either choose D at stage 2 (the equilibrium prediction)

or never choose D . In addition, the prediction of χ -CE is unique for $\chi \neq \bar{\alpha}$. As characterized in Proposition 2.12, for extreme values of χ , the prediction of χ -CSE coincides with χ -CE. But for intermediate values of χ , there can be multiple χ -CSE.

Proposition 2.12. *The pure strategy χ -CSE can be characterized as follows.*

1. $\tilde{\sigma}^\chi(O) = 1$ for all $\chi \in [0, 1]$.
2. Both players choosing $\tilde{\sigma}^\chi(X) = T + 1$ is a χ -CSE if and only if $\chi \geq \bar{\alpha}^{\frac{1}{T+1}}$.
3. Both players choosing $\tilde{\sigma}^\chi(X) = 2$ is a χ -CSE if and only if $\chi \leq \bar{\alpha}$.
4. For any $3 \leq t \leq T$, both players choosing $\tilde{\sigma}^\chi(X) = t$ is a χ -CSE if and only if

$$\left(\frac{1 - \kappa(\chi)}{1 - p} \right)^{\frac{1}{t-2}} \leq \chi \leq \bar{\alpha}^{\frac{1}{t-1}} \quad \text{where}$$

$$\kappa(\chi) \equiv \frac{[(1 + \alpha)(1 + \delta\chi) - \alpha\delta] - \sqrt{[(1 + \alpha)(1 + \delta\chi) - \alpha\delta]^2 - 4\delta\chi(1 + \alpha)}}{2\delta\chi(1 + \alpha)}.$$

Illustrative Example




Figure 2.7: χ -CE vs. χ -CSE When $(\alpha, \delta, p, T) = \left(\frac{1}{4}, \frac{4}{5}, \frac{2}{3}, 5\right)$.

To illustrate the sharp contrast between the predictions of χ -CE and χ -CSE, here we consider an illustrative example where $\alpha = 1/4$, $\delta = 4/5$, $p = 2/3$ and the horizon

of the game is $T = 5$. As characterized by Proposition 2.11, χ -CE predicts players will choose $\hat{\sigma}^\chi(X) = 2$ if $\chi \leq \bar{\alpha} = 0.6$; otherwise, they will choose $\hat{\sigma}^\chi(X) = 6$, i.e., they never choose D when observing a dirty face. As demonstrated in the left panel of Figure 2.7, χ -CE is (generically) unique and it predicts players will either behave extremely sophisticated or unresponsive to the other's action at all.

In contrast, as characterized by Proposition 2.12, there can be multiple χ -CSE. As shown in the right panel of Figure 2.7, when $\chi \leq \bar{\alpha} = 0.6$, both players stopping at stage 2 is still an equilibrium, but it is not unique except for very low values of χ . For $0.168 \leq \chi \leq 0.505$, both players stopping at stage 3 is also a χ -CSE, and for $0.505 \leq \chi \leq 0.6$, there are three pure strategy χ -CSE where both players stop at stage 2, 3, or 4, respectively.

The existence of multiple χ -CSE in which both players stop at $t > 2$ highlights a player's learning process in a multi-stage game, which does not happen in strategic form cursed equilibrium. In the strategic form, a player has no opportunity to learn about the other player's type in middle stages. Thus, when level of cursedness is not low enough to support a χ -CE with stopping at stage 2, both players would never stop. However, in a χ -CSE of the multi-stage game, a cursed player would still learn about his own face being dirty as the game proceeds, even though he might not be confident enough to choose D at stage 2. If χ is not too large, the expected payoff of choosing D would eventually become positive at some stage before the last stage T .¹⁷ For some intermediate values of χ , there might be multiple stopping stages which yield positive expected payoffs. In this case, the dirty faces game becomes a special type of coordination games where both players coordinate on stopping strategies, resulting in the existence of multiple χ -CSE.¹⁸

2.5 Concluding Remarks

In this paper, we formally developed Cursed Sequential Equilibrium, which extends the strategic form cursed equilibrium (Eyster and Rabin, 2005) to multi-stage games, and illustrated the new equilibrium concept with a series of applications. While the standard CE has no bite in private value games, we show that cursed beliefs can actually have significant consequences for dynamic private value games. In the private value games we consider, our cursed sequential equilibrium predicts (1)

¹⁷The upper bound of the inequality in Proposition 2.12 characterizes the stages at which stopping yields positive expected payoffs.

¹⁸Note that players with low levels of cursedness would not coordinate on stopping at late stages since the discount factor shrinks the informative value of waiting (i.e., both choosing U). This result is characterized by the lower bound of the inequality in Proposition 2.12.

under-contribution caused by under-communication in the public goods game with communication, (2) low passing rate in the presence of altruistic players in the centipede game, and (3) less sophisticated voting in the sequential two-stage binary agenda game. We also illustrate the distinction between CE and CSE in some non-private value games. In simple signaling games, χ -CSE implies refinements of pooling equilibria that are not captured by traditional belief-based refinements (or χ -CE), and are qualitatively consistent with some experimental evidence. Lastly, we examine the dirty face game, showing that the CSE further expands the set of equilibrium and predicts stopping in middle stages of the game. Our findings are summarized in Table 2.4.

Table 2.4: Summary of Findings in Section 2.4.

	Private-Value Game	χ -CE vs. BNE	χ -CSE vs. χ -CE
Signaling Games with Pooling Equilibrium	No	\neq	χ -CSE \subset χ -CE
Public Goods Game with Communication	Yes	=	\neq
Centipede Game with Altruists	Yes	=	\neq
Sequential Voting Game	Yes	=	\neq
Dirty Faces Game	No	\neq	\neq

The applications we consider are only a small sample of the possible dynamic games where CSE could be usefully applied. One prominent class of problems where it would be interesting to study the dynamic effects of cursedness is social learning. For example, in the standard information cascade model of Bikhchandani, Hirshleifer, and Welch (1992), we conjecture that the effect would be to delay the formation of an information cascade because players will partially neglect the information content of prior decision makers. Laboratory experiments report evidence that subjects underweight the information contained in prior actions relative to their own signal (Goeree et al., 2007). A related class of problems involves information

aggregation through sequential voting and bandwagon effects (Callander, 2007; Ali, Goeree, et al., 2008; Ali and Kartik, 2012). A natural conjecture is that CSE will impede information transmission in committees and juries as later voters will underappreciate the information content of the decisions by early voters. This would dampen bandwagon effects. The centipede example we studied suggests that CSE might have broader implications for behavior in reputation-building games, such as the finitely repeated prisoner's dilemma or entry deterrence games such as the chain store paradox.

The generalization of CE to dynamic games presented in this paper is limited in several ways. First, the CSE framework is formally developed for finite multi-stage games with observed actions. We do not extend CSE for games with continuous types but we do provide one application that shows how such an extension is possible. However, a complete generalization to continuous types (or continuous actions) would require more technical development and assumptions. We also assume that the number of stages is finite, and extending this to infinite horizon multi-stage games would be a useful exercise. Extending CSE to allow for imperfect monitoring in the form of private histories is another interesting direction to pursue. The SCE approach in Cohen and Li (2023) allows for cursedness with respect to both public and private endogenous information, which leads to some important differences from our CSE approach. In CSE, we find that subjects are limited in their ability to make correct inferences about hypothetical events, but the mechanism is different from SCE, which introduces a second free parameter that modulates cursedness with respect to hypothetical events. For a more detailed discussion of these and other differences and overlaps between CSE and SCE, see Appendix B and Fong, Lin, and Palfrey (2023).

As a final remark, our analysis of applications of χ -CSE suggests some interesting experiments. For instance, χ -CSE predicts in the public goods game with communication, when either the number of players (N) or the largest possible contribution cost (K) increases, pre-play communication will be less effective, while the prediction of sequential equilibrium and χ -CE is independent of N and K . In other words, in an experiment where N and K are manipulated, significant treatment effects in this direction would provide evidence supporting χ -CSE over χ -CE. Also, χ -CSE makes qualitatively testable predictions in the sequential voting games and the dirty faces games, which have not been extensively studied in laboratory experiments. In the sequential voting game, it would be interesting to test how sensitive strategic

(vs. sincere) voting behavior is to preference intensity (v) and the type distribution. In the dirty faces game, it would be interesting to design an experiment to identify the extent to which deviations from sequential equilibrium are related to the coordination problem that arises in χ -CSE.

References

- Ali, Nageeb, Jacob K Goeree, et al. (2008). “Information aggregation in standing and ad hoc committees”. In: *American Economic Review* 98.2, pp. 181–186.
- Ali, Nageeb and Navin Kartik (2012). “Herding with collective preferences”. In: *Economic Theory* 51, pp. 601–626.
- Barwise, Jon (1981). “Scenes and other situations”. In: *Journal of Philosophy* 78.7, pp. 369–397.
- Bayer, Ralph C and Mickey Chan (2007). *The dirty faces game revisited*. Tech. rep. University of Adelaide, School of Economics.
- Bazerman, Max H and William F Samuelson (1983). “I won the auction but don’t want the prize”. In: *Journal of Conflict Resolution* 27.4, pp. 618–634.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992). “A theory of fads, fashion, custom, and cultural change as informational cascades”. In: *Journal of Political Economy* 100.5, pp. 992–1026.
- Brandts, Jordi and Charles A Holt (1993). “Adjustment patterns and equilibrium selection in experimental signaling games”. In: *International Journal of Game Theory* 22.3, pp. 279–302.
- Callander, Steven (2007). “Bandwagons and momentum in sequential voting”. In: *Review of Economic Studies* 74, pp. 653–684.
- Camerer, Colin F, Salvatore Nunnari, and Thomas R Palfrey (2016). “Quantal response and non-equilibrium beliefs explain overbidding in maximum value auctions”. In: *Games and Economic Behavior* 98, pp. 243–263.
- Capen, Edward C, Robert V Clapp, and William M Campbell (1971). “Competitive bidding in high-risk situations”. In: *Journal of Petroleum Technology* 23.06, pp. 641–653.
- Carrillo, Juan and Thomas R Palfrey (2009). “The compromise game: Two-sided adverse selection in the laboratory”. In: *American Economic Journal: Microeconomics* 1.1, pp. 151–181.
- (2011). “No trade”. In: *Games and Economic Behavior* 71.1, pp. 66–87.
- Cho, In-Koo and David M Kreps (1987). “Signaling games and stable equilibria”. In: *Quarterly Journal of Economics* 102.2, pp. 179–221.
- Cohen, Shani and Shengwu Li (2023). “Sequential cursed equilibrium”. In: *arXiv preprint arXiv:2212.06025v3*.

- Dyer, Douglas, John H Kagel, and Dan Levin (1989). “A comparison of naive and experienced bidders in common value offer auctions: A laboratory analysis”. In: *Economic Journal* 99.394, pp. 108–115.
- Eckel, Catherine C and Charles A Holt (1989). “Strategic voting in agenda controlled committee experiments”. In: *American Economic Review* 79.3, pp. 763–773.
- Eyster, Erik and Matthew Rabin (2005). “Cursed equilibrium”. In: *Econometrica* 73.5, pp. 1623–1672.
- Fong, Meng-Jhang, Po-Hsuan Lin, and Thomas R Palfrey (2023). “A note on cursed sequential equilibrium and sequential cursed equilibrium”. In: *arXiv preprint arXiv:2304.05515*.
- Forsythe, Robert E, R Mark Isaac, and Thomas R Palfrey (1989). “Theories and tests of blind bidding in sealed-bid auctions”. In: *Rand Journal of Economics* 20.2, pp. 214–238.
- Fudenberg, Drew and Jean Tirole (1991a). *Game theory*. MIT press.
- (1991b). “Perfect Bayesian equilibrium and sequential equilibrium”. In: *Journal of Economic Theory* 53.2, pp. 236–260.
- Gamow, G and M Stern (1958). “Forty unfaithful wives”. In: *Puzzle Math*, pp. 20–23.
- Goeree, Jacob K et al. (2007). “Self-correcting information cascades”. In: *Review of Economic Studies* 74.3, pp. 733–762.
- Guarnaschelli, Serena, Richard D McKelvey, and Thomas R Palfrey (2000). “An experimental study of jury decision rules”. In: *American Political Science Review* 94.2, pp. 407–423.
- Halpern, Joseph Y and Yoram Moses (1990). “Knowledge and common knowledge in a distributed environment”. In: *Journal of the ACM (JACM)* 37.3, pp. 549–587.
- Hardin, Christopher S and Alan D Taylor (2008). “An introduction to infinite hat problems”. In: *The Mathematical Intelligencer* 30.4, pp. 20–25.
- Holt, Charles A and Roger Sherman (1994). “The loser’s curse”. In: *American Economic Review* 84.3, pp. 642–652.
- Ivanov, Asen, Dan Levin, and Muriel Niederle (2010). “Can relaxation of beliefs rationalize the winner’s curse”. In: *Econometrica* 78.4, pp. 1435–1452.
- Jehiel, Philippe (2005). “Analogy-based expectation equilibrium”. In: *Journal of Economic Theory* 123.2, pp. 81–104.
- Jehiel, Philippe and Frédéric Koessler (2008). “Revisiting games of incomplete information with analogy-based expectations”. In: *Games and Economic Behavior* 62.2, pp. 533–557.
- Kagel, John H and Dan Levin (1986). “The winner’s curse and public information in common value auctions”. In: *American Economic Review*, pp. 894–920.

- Kagel, John H and Dan Levin (2009). “Common value auctions and the winner’s curse”. In: *Common Value Auctions and the Winner’s Curse*. Princeton University Press.
- Kagel, John H, Dan Levin, et al. (1989). “First-price common value auctions: Bidder behavior and the “winner’s curse””. In: *Economic Inquiry* 27.2, pp. 241–258.
- Kreps, David M (1990). *A course in microeconomic theory*. Princeton university press.
- Kreps, David M and Robert Wilson (1982). “Sequential equilibria”. In: *Econometrica* 50.4, pp. 863–894.
- Levine, Michael E and Charles R Plott (1977). “Agenda influence and its implications”. In: *Virginia Law Review*, pp. 561–604.
- Lin, Po-Hsuan (2023). “Cognitive hierarchies in multi-stage games of incomplete information: Theory and experiment”. In: *arXiv preprint arXiv:2208.11190v3*.
- Lin, Po-Hsuan and Thomas R Palfrey (2022). “Cognitive hierarchies in extensive form games”. In: *Social Science Working Paper 1460, California Institute of Technology*.
- Lind, Barry and Charles R Plott (1991). “The winner’s curse: Experiments with buyers and with sellers”. In: *American Economic Review* 81.1, pp. 335–346.
- Littlewood, John E (1953). *A mathematician’s miscellany*. London, England: Meuthen & Co. Ltd.
- McKelvey, Richard D and Thomas R Palfrey (1992). “An experimental study of the centipede game”. In: *Econometrica*, pp. 803–836.
- (1998). “Quantal response equilibria for extensive form games”. In: *Experimental Economics* 1.1, pp. 9–41.
- Moses, Yoram, Danny Dolev, and Joseph Y Halpern (1986). “Cheating husbands and other stories: A case study of knowledge, action, and communication”. In: *Distributed Computing* 1.3, pp. 167–176.
- Ordeshook, Peter C and Thomas R Palfrey (1988). “Agendas, strategic voting, and signaling with incomplete information”. In: *American Journal of Political Science*, pp. 441–466.
- Palfrey, Thomas R and Howard Rosenthal (1991). “Testing for effects of cheap talk in a public goods game with private information”. In: *Games and Economic Behavior* 3.2, pp. 183–220.
- Palfrey, Thomas R, Howard Rosenthal, and Nilanjan Roy (2017). “How cheap talk enhances efficiency in threshold public goods games”. In: *Games and Economic Behavior* 101, pp. 234–259.
- Plott, Charles R and Michael E Levine (1978). “A model of agenda influence on committee decisions”. In: *American Economic Review* 68.1, pp. 146–160.

- Rogers, Brian W, Thomas R Palfrey, and Colin F Camerer (2009). "Heterogeneous quantal response equilibrium and cognitive hierarchies". In: *Journal of Economic Theory* 144.4, pp. 1440–1467.
- Samuelson, William F and Max H Bazerman (1985). "Negotiation under the winner's curse". In: *Research in Experimental Economics* 3, pp. 105–138.
- Søvik, Ylva (2009). "Strength of dominance and depths of reasoning: An experimental study". In: *Journal of Economic Behavior & Organization* 70.1-2, pp. 196–205.
- Weber, Roberto A (2001). "Behavior and learning in the "dirty faces" game". In: *Experimental Economics* 4.3, pp. 229–242.

*Chapter 3*COGNITIVE HIERARCHIES FOR GAMES IN EXTENSIVE
FORM**3.1 Introduction**

“What surprised me was how bad they played.”

—Beth Harmon, *The Queen’s Gambit*

In many situations, people interact with one another over time, in a multi-stage environment, such as playing chess or bargaining with alternating offers. The standard approach to studying these situations is to model them as games in extensive form where equilibrium theory is applied, often with refinements such as subgame perfection or other notions of sequential rationality. However, the standard equilibrium concepts, such as sequential equilibrium and its refinements, impose strong assumptions about the strategic sophistication of the players—perhaps implausibly strong from an empirical standpoint, as behavior in many laboratory experiments has suggested (see, for example, Camerer 2003).

In response to these anomalous findings, researchers have proposed a variety of models that relax the requirement of mutual consistency of beliefs embodied in standard equilibrium concepts. The focus of this paper is the “level- k ” family of models, which assume a hierarchical structure of strategic sophistication among the players. In this family of models, each player is endowed with a specific level of sophistication. Level-0 players are non-strategic and choose their actions randomly. Level- k players, on the other hand, can think k strategic steps and believe everyone else is less sophisticated in the sense that they think fewer than k strategic steps. The *standard* level- k model assumes level- k players believe all other players are level- $(k-1)$ (see Nagel 1995).

However, applications of the level- k approach have been limited almost exclusively to the analysis of games in strategic form, where all players make their moves simultaneously, and the theory has not been formally developed for the analysis of general games in extensive form. To apply the standard level- k model to games in extensive form, one would assume that at each decision node, a level- k player will choose an action that maximizes the continuation value of the game, assuming all

other players are level- $(k-1)$ players in the continuation game. As a result, each player's belief about other players' level is *fixed* at the beginning. However, as the game proceeds, this fixed belief can lead to a logical conundrum, as a level- k player can be “surprised” by an opponent's previous move that is not consistent with the strategy of a level- $(k-1)$ player.¹

If one closely examines this problem, the incompatibility derives from two sources that imply players cannot learn: (1) each level of player's prior belief about the other players' levels is degenerate, i.e., a singleton; and (2) players ignore the information contained in the history of the game. To solve both of these problems at the same time, as an alternative to the standard level- k approach, we use the cognitive hierarchy (CH) version of level- k , as proposed by Camerer, Ho, and Chong (2004), and extend it to games in extensive form. Like the standard level- k model, the CH framework posits that players are heterogeneous with respect to levels of strategic sophistication and believe that other players are less sophisticated. However, their beliefs are not degenerate. A level- k player believes all other players have lower levels distributed anywhere from level 0 to $k-1$.

Furthermore, the CH framework imposes a partial consistency requirement that ties the players' prior beliefs on the level-type space to the true underlying distribution of levels. Specifically, a level- k player's beliefs are specified as the truncated true distribution of levels, conditional on levels ranging from 0 to $k-1$, i.e., players have “truncated rational expectations.” This specification has the important added feature, relative to the standard level- k model, that more sophisticated players also have beliefs that are closer to the true distribution of levels, and very high level types have approximate rational expectations about the behavior of the other players. Thus, the CH approach blends aspects of purely behavioral models and equilibrium theory.

In our extension of CH to games in extensive form, a player's prior beliefs over lower levels are updated as the history of play in the game unfolds, revealing information about the distribution of other players' levels of sophistication. These learning effects can be quite substantial as we illustrate later in the paper. Hence, the main contribution of this paper is to propose a new CH framework—dynamic cognitive

¹An exception is a recent paper by Schipper and Zhou (2023) that was developed independently and contemporaneously with ours. They take an alternative approach to resolving this conundrum for extending the standard level- k analysis to games in extensive form, which we discuss in section 3.2.

hierarchy (DCH)—for the analysis of general games in extensive form and, in doing so, provide new insights beyond those offered by the original CH model.

Our first result establishes that in games with perfect information, every player will update their belief about each of the opponent's levels independently (Proposition 3.1). Second, we show that when the history of play in the game unfolds, players become more certain about the opponents' levels of sophistication, in a specific way. Formally, the support of their beliefs shrinks as the history gets longer (Proposition 3.2). Third, we show that the probability of paths with strictly dominated strategies being realized converges to zero as the distribution of levels increases (Proposition 3.3). Nonetheless, solution concepts based on iterated dominance, such as forward induction, can be inconsistent with DCH even at the limit when the average level of sophistication converges to infinity. Relatedly, even though the players fully exploit the information from the history, it is not guaranteed that high-level players will use strategies that are consistent with the subgame perfect equilibrium of the game. In fact, behavior of the most sophisticated players can be inconsistent with backward induction, even at the limit when the level of sophistication of all players is arbitrarily high.

Another important property of the DCH solution is that it is not reduced normal form invariant. In many games, players have multiple strategies that are realization-equivalent. The (structurally) reduced normal form of the game consolidates such strategies into a single strategy, thus reducing the cardinality of the set of pure strategies. In DCH, level-0 players uniformly randomize over the set of actions at each information set, which is realization-equivalent to uniformly randomizing over the set of pure strategies in the normal form representation of the game. However, it is not equivalent to uniform randomization over the set of pure strategies in the reduced normal form, because of the reduced cardinality of the set of strategies. For this reason, we refer to this property of DCH as the *strategy-reduction effect*.²

The strategy-reduction effect has important implications in experimental economics, as sequential game experiments are often implemented by having subjects simultaneously choose reduced strategies, a procedure known as *the strategy method*. There is considerable evidence that qualitatively different data is observed in experimental studies of games depending on whether the strategy method or the direct-response (sequential choice) method is used to elicit choices. How and why such differences

²In a special class of games that satisfy *balancedness*, the consolidation of realization equivalent strategies is innocuous and does not lead to a strategy reduction effect in DCH. Battigalli (2023) presents a formal statement and proof of this result.

arise is not well understood.³ Because of the strategy-reduction effect, the DCH solution has the potential to provide new insights into this issue.

With this goal in mind, we explore the implications of the strategy-reduction effect of DCH in an application to a prominent class of games where observed behavior is grossly inconsistent with the standard equilibrium theory, the increasing-pie “centipede game.” In this class of games, the DCH solution makes clear predictions about the direction of the strategy-reduction effect.⁴

Specifically, we consider two common implementations of centipede games studied in experiments. The first (and the usual) implementation is to play the game in its extensive form, as an alternating-move sequential game that terminates as soon as one player takes (*the direct-response method*). That is, first player 1 decides to take or pass. If they take, the game ends; if they pass, it is player 2’s turn to pass or take, and so on. The second implementation asks both players to decide at which node to take if the game gets that far, or always pass (*the strategy method*). That is, players simultaneously choose from their respective set of reduced pure strategies. The benefit of doing so, from a methodological standpoint, is that one can gather more experimental data, particularly at histories that are only occasionally reached when the game is played out sequentially.

However, from the perspective of DCH, implementing the strategy method reduces the number of pure strategies, causing level-0 players to behave differently in the two protocols. Consequently, the behavior of higher-level players (who always believe in the existence of level-0 players) is also affected. In Theorem 3.1, we show that this strategy-reduction effect implies that increasing-sum centipede games will end *earlier, with lower payoffs to the players*, when implemented by the direct-response method than when implemented by the strategy method.

While the direct-response method is the most commonly used method to implement centipede game experiments, there are a few exceptions. Nagel and Tang (1998) is the first paper to report the results from a centipede game experiment conducted as a simultaneous move game, the reduced normal form. In their 12-node centipede games, each player has seven available strategies that correspond to an intended

³The survey by Brandts and Charness (2011) provides an extensive account of the evidence and discussion of possible explanations for these differences.

⁴These are alternating-move finite-horizon two-person games, where, in turn, each player can either “take” the larger of two pieces of a growing pie, which terminates the game and leaves the other player with the smaller piece of the current pie, or “pass,” which increases the size of the pie and it is the other player’s turn to take or pass. Payoffs are such that the subgame perfect equilibrium of this game is solved by backward induction with the game ending immediately.

“take-node” or always passing, and they make their decisions simultaneously. Pooling the data over many repetitions, they find that only 0.5% of the outcomes coincide with the equilibrium prediction, suggesting that the non-equilibrium behavior in the centipede games cannot solely be attributed to the violation of backward induction. However, as the authors remarked, the results may be confounded with the effect of reduced normal form: “...There might be substantial differences in behavior in the extensive form game and in the normal form game...” (Nagel and Tang, 1998, p. 357). One of our contributions is to show that DCH provides a theoretical rationale for the existence of such effects.

A recent experiment by García-Pola, Iriberry, and Kovářik (2020a) specifically studies how the direct-response method and the strategy method would affect the behavior in four different centipede games. The DCH solution makes a sharp prediction about the strategy-reduction effect—in three of the four games, the distribution of terminal nodes under the strategy method will first order stochastically dominate the distribution under the direct-response method, but not in the fourth game. We revisit the data from that experiment, and show that the results from all four of their games are consistent with the strategy-reduction effect predicted by DCH.

The paper is organized as follows. The related literature is discussed in the next section. Section 3.3 sets up the model. Section 3.4 establishes properties of the DCH belief-updating process and explores the relationship between our model and subgame perfect equilibrium with several examples. In Section 3.5, the strategy-reduction effect is explored in a detailed analysis of centipede games with a linearly increasing pie. Experimental data that provide a test the strategy-reduction effect hypothesis is presented and discussed in Section 3.6. We discuss several additional features of our model in Section 3.7 and conclude in Section 3.8.

3.2 Related Literature

The idea of limited depth of reasoning in games of strategy has been proposed and studied by economists and game theorists for at least thirty years (see, for example, Binmore, 1987; Binmore, 1988; Selten, 1991; Aumann, 1992; Stahl, 1993; Selten, 1998; Alaoui and Penta, 2016; Alaoui and Penta, 2022). On the empirical side, Nagel (1995) conducts the first laboratory experiment explicitly designed to study hierarchical reasoning in simultaneous move games, using the “beauty contest” game. Each player chooses a number between 0 and 100. The winner is the player whose choice is closest to the average of all the chosen numbers discounted by a

parameter $p \in (0, 1)$. To analyze the data, Nagel (1995) assumes level-0 players choose randomly. Level-1 players believe all other players are level-0 and best respond to them by choosing $50p$. Following the same logic, level- k players believe all other players are level- $(k-1)$ and best respond to them with $50p^k$.

This iterative definition of hierarchies has been applied to a range of different environments. For instance, Ho, Camerer, and Weigelt (1998) also analyze the beauty contest game while Costa-Gomes, Crawford, and Broseta (2001) and Crawford and Iriberri (2007a) consider the strategic levels in a variety of simultaneous move games. Costa-Gomes and Crawford (2006) study the “two-person guessing game,” a variant of the beauty contest game. Finally, the level- k approach has also been applied to games of incomplete information. Crawford and Iriberri (2007b) apply this approach to reanalyze auction data, and Cai and Wang (2006) and Wang, Spezio, and Camerer (2010) use the level- k model to organize empirical patterns in experimental sender-receiver games. All these studies assume level- k players best respond to degenerate beliefs of level- $(k-1)$ players.

This standard level- k model has been extended in a number of ways. One such approach is that each level of player best responds to a mixture of all lower levels. Stahl and P. W. Wilson (1995) are the first to construct and estimate a specific mixture model of bounded rationality in games where each level of player best responds to a mixture between lower levels and equilibrium players.⁵ Camerer, Ho, and Chong (2004) develop the CH framework, where level- k players best respond to a *mixture* of the behavior of all lower level behavioral types from 0 to $k-1$. In addition, players have correct beliefs about the relative proportions of these lower levels, so it includes a consistency restriction on beliefs in the form of truncated rational expectations.

A second direction is to endogenize the strategic levels of players, using a cost benefit approach. Alaoui and Penta (2016) develop a model of endogenous depth of reasoning, where each player trades off the benefit of additional levels of sophistication against the cost of doing so. Players can have different benefit and cost functions, depending on their beliefs and strategic abilities, respectively. A model is developed for two-person games with complete information and calibrated against experimen-

⁵In the same spirit of Stahl and P. W. Wilson (1995), Levin and Zhang (2022) propose the NLK solution which is an equilibrium model where each player best responds to a mixture of an exogenously determined naive strategy (with probability λ) and the equilibrium strategy (with probability $1 - \lambda$). The main difference between NLK and our DCH solution is that NLK requires the belief system to be mutually consistent (as NLK is a solution of a fixed problem) while DCH relaxes this requirement.

tal data. Alaoui, Janezic, and Penta (2020) provide some additional analysis and a laboratory experiment that further explores the implications of this model.

Third, De Clippel, Saran, and Serrano (2019) study the implications of the standard level- k approach to mechanism design. They establish a form of the revelation principle for level- k implementation and obtain conditions on the implementability of social choice functions under a range of assumptions about level-0 behavior.

Fourth, there have been several papers that model how an individual's strategic level evolves when the same game is repeated multiple times.⁶ The standard level- k model is ideally suited to understanding how naive individuals behave when they encounter a game for the first time. This is a limitation since, in most laboratory experiments in economics and game theory, subjects play the same game with multiple repetitions, in order to gain experience and to facilitate convergence to equilibrium behavior. It is also a limitation since many games studied by economists and other social scientists are aimed at understanding strategic interactions between highly experienced players (oligopoly, procurement auctions, legislative bargaining, for example), where some convergence to equilibrium would be natural to expect.

In this vein, Ho and Su (2013) and Ho, Park, and Su (2021) propose a modification of CH that allows for learning across repeated plays of the same sequential game, in a different way than in Stahl (1996), but in the same spirit. In their setting, an individual player repeatedly plays the same game (such as the guessing game) and updates his or her beliefs about the distribution of levels *after* observing past outcomes of earlier games, but holding fixed beliefs during each play of the game. In addition to updating beliefs about other players' levels, a player also *endogenously* chooses a new level of strategic sophistication for themselves, in the spirit of Stahl (1996), for the next iteration of the game. This is different from our DCH framework where each player updates their beliefs about the levels of other players after each move within a single game. Moreover, because players are forward-looking in DCH, they are *strategic learners*—i.e., they correctly anticipate the evolution of their posterior beliefs in later stages of the game—which leads to a much different learning dynamic compared with naive adaptive learning models.

All of these extensions add significantly to the literature on level- k behavior for games in strategic form, by allowing for a richer set of heterogeneous beliefs, by incorporating cognitive costs into the model, and by showing how the model can be

⁶Nagel (1995) proposes a model for the evolution of levels based on changing reference points to explain unravelling in guessing games.

used to address classic mechanism design problems, but all of them are limited by a restriction to simultaneous move games. In extensive games, the timing structure is crucial, and beliefs evolve as the game unfolds and players have an opportunity to adjust their beliefs in response to past actions, which is the focus of this paper. Our DCH provides an extension in this direction, under the assumption that players are forward looking about the actions of their opponents in the entire game tree.⁷ Rampal (2022) develops an alternative approach for multi-stage games of perfect information. He models levels of sophistication by assuming that players have limited foresight in the sense of a rolling horizon; that is, players only look forward a fixed number of stages. This creates a hierarchy of strategic sophistication that depends on the length of a player’s rolling horizon. Players are uncertain about their opponents’ foresight. The baseline game of perfect information is then transformed into a game of incomplete information, with specific assumptions about players’ beliefs about payoffs at non-terminal nodes that correspond to the current limit of their horizon in the game.

At a more conceptual level, our dynamic generalization of CH is related to other behavioral models in game theory. There is a connection between DCH and misspecified learning models (see, for example, Hauser and Bohren, 2021) in the sense that level- k players wrongly believe all other players are less sophisticated. However, in contrast to categorical types of players in misspecified learning models, DCH provides added structure to the set of types in a systematic way, such that higher-level types have a more accurate belief about opponents’ rationality at the aggregate level. In the context of social learning, application of our model to the investment game is related to Eyster and Rabin (2010) and Bohren (2016) who model the updating process when there exist some behavioral types of players in the population.

Our DCH solution is also related to two other behavioral solution concepts of the dynamic games—the Agent Quantal Response Equilibrium (AQRE) by McKelvey and Palfrey (1998) and the Cursed Sequential Equilibrium (CSE) by Fong, Lin, and Palfrey (2023)—in the sense that each of these solution concepts relaxes different requirements of the standard equilibrium theory. DCH is a non-equilibrium model

⁷The recent paper by Schipper and Zhou (2023) also identifies the logical conundrum when applying the standard level- k model to games in extensive form, and resolve the conundrum in a different way. In contrast to the DCH solution, Schipper and Zhou (2023) adopts a non-Bayesian approach to updating, called strong level- k thinking, which posits that for every information set, level- k players attach the maximum level- l thinking for $l < k$ to their opponents consistent with the information set.

which allows different levels of players to best respond to different conjectures about other players' strategies while AQRE is an equilibrium model where players make stochastic choices. Both DCH and AQRE depict that players are able to perform Bayesian inferences. In contrast, CSE is an equilibrium model where players are able to make best response but fail to perform Bayesian inferences. Both AQRE and CSE predict a strategy-reduction effect, but in different manners from DCH.

The key application in this paper is about the centipede game, which has been the subject of many theoretical and experimental studies. Rosenthal (1981) first introduced the centipede game to demonstrate how backward induction can be challenging and implausible to hold in some environments due to logical issues about updating off-path beliefs. His example is a ten-node game with a linearly increasing pie. Later on a shorter variant with an exponentially increasing pie, called "share or quit," is studied by Megiddo (1986) and Aumann (1988). The name centipede was coined by Binmore (1987), and named for a 100-node variant.

Starting with McKelvey and Palfrey (1992), many laboratory and field experiments have reported experimental data from centipede games in a range of environments, such as different lengths of the game (see McKelvey and Palfrey, 1992 and Fey, McKelvey, and Palfrey 1996), different subject pools (see Palacios-Huerta and Volij, 2009; Levitt, List, and Sadoff, 2011, and Li et al., 2021), different payoff configurations (see Fey, McKelvey, and Palfrey, 1996; Zauner, 1999; Kawagoe and Takizawa, 2012; Healy, 2017, and García-Pola, Iriberry, and Kovářik, 2020b) and different experimental methods (Nagel and Tang, 1998; Rapoport et al., 2003; Bornstein, Kugler, and Ziegelmeyer, 2004, and García-Pola, Iriberry, and Kovářik, 2020a). Although standard game theory predicts the game should end in the first stage, such behavior is rarely observed.

3.3 The Model

This section formally develops the dynamic cognitive hierarchy (DCH) solution for games in extensive form. In Section 3.3.1, we introduce the notation for (finite) games in extensive form by following Osborne and Rubinstein (1994). Next, we define the DCH updating process in Section 3.3.2, specifying how players' beliefs about other players' levels evolve from the history of play. This leads to a definition of the *DCH solution* of a game.

3.3.1 Games in Extensive Form

Let $N_C = \{C, 1, \dots, n\} \equiv \{C\} \cup N$ be a finite set of *players*, where player C is called “Chance.” Let H be a finite set of *sequences of actions* that satisfies the following two properties.

1. The empty sequence ϕ (the initial history) is a member of H .
2. If $(a^j)_{j=1, \dots, J} \in H$ and $L < J$, then $(a^j)_{j=1, \dots, L} \in H$.

Each $h \in H$ is a *history* and each component of a history is an *action* taken by a player. In addition, for any non-initial history $h = (a^j)_{j=1, \dots, L}$, we use $\alpha(h) = (a^j)_{j=1, \dots, L-1}$ to denote the immediate predecessor of h .⁸ A history $(a^j)_{j=1, \dots, J} \in H$ is a *terminal history* if there is no a^{J+1} such that $(a^j)_{j=1, \dots, J+1} \in H$. The set of terminal histories is denoted as Z and $H \setminus Z$ is the set of non-terminal histories. Moreover, for every non-terminal history $h \in H \setminus Z$, $A(h) = \{a : (h, a) \in H\}$ is the set of available actions after the history h , and Z_h is the set of terminal histories after h .

The *player function* $P : H \setminus Z \rightarrow N_C$ assigns to each non-terminal history a player of N_C . In other words, $P(h)$ is the player who takes an action after history h . With this, for any player $i \in N_C$, $H_i = \{h \in H \setminus Z : P(h) = i\}$ is the set of histories where player $i \in N_C$ is the active player. Therefore, H_C is the set of non-terminal histories where Chance determines the action taken after history h . A function σ_C specifies a probability measure $\sigma_C(\cdot|h)$ on $A(h)$ for every $h \in H_C$. That is, $\sigma_C(a|h)$ is the probability that a occurs after the history h .

For each $i \in N$, a partition \mathcal{I}_i of H_i defines i 's information sets. Information set $I_i \in \mathcal{I}_i$ specifies a subset of histories contained in H_i that i cannot distinguish from one other, where for any $h \in H_i$, $I_i(h)$ is the element of \mathcal{I}_i that contains h . Furthermore, i 's available actions are the same for all histories in the same information set. Formally, for any history $h' \in I_i(h) \in \mathcal{I}_i$, $A(h') = A(h)$.⁹ Therefore, we use $A(I_i)$ to denote the set of available actions at information set I_i . In addition, each player $i \in N$ has a *payoff function* (in von Neumann-Morgenstern utilities) $u_i : Z \rightarrow \mathbb{R}$. Finally, a game in extensive form, Γ , is defined by the tuple $\Gamma = \langle N_C, H, P, \sigma_C, (\mathcal{I}_i)_{i \in N}, (u_i)_{i \in N} \rangle$.

In a game in extensive form Γ , for each player $i \in N$, the set of *behavioral strategies* for player i is $\Sigma_i \equiv \times_{I_i \in \mathcal{I}_i} \Delta(A(I_i))$ and $\sigma_i \in \Sigma_i$ is a behavioral strategy of player i .

⁸If $L = 1$, then $\alpha(h) = \phi$.

⁹We assume that all players in the game have perfect recall. See Osborne and Rubinstein (1994) Chapter 11 for a definition.

For notational convenience, we use $\Sigma \equiv \times_{i \in N} \Sigma_i$ to denote the set of all profiles of behavioral strategies, and we use the notation $\Sigma_{-i} = \times_{j \neq i} \Sigma_j$ and write elements of Σ as $\sigma = (\sigma_i, \sigma_{-i})$ when focusing on a particular player $i \in N$.

3.3.2 Cognitive Hierarchies and Belief Updating

Prior Beliefs about Levels of Sophistication

Each player i is endowed with a level of sophistication, $\tau_i \in \mathbb{N}_0$, where $\Pr(\tau_i = k) = p_{ik}$ for all $i \in N$ and $k \in \mathbb{N}_0$, and the distribution is *independent* across players.¹⁰ Without loss of generality, we assume $p_{ik} > 0$ for all $i \in N$ and $k \in \mathbb{N}_0$. Let $\tau = (\tau_1, \dots, \tau_n)$ be the profile of levels and τ_{-i} be the profile of levels without player i . Each level $k > 0$ of each player i has a prior belief about all other players' levels and these prior beliefs satisfy *truncated rational expectations*. That is, for each i and $k > 0$, a level- k player i believes all other players in the game are at most level- $(k-1)$. For each i , $j \neq i$ and k , let $\mu_{ij}^k(\tau_j)$ be level- k player i 's prior belief about player j 's level, and $\mu_i^k(\tau_{-i}) = \left(\mu_{ij}^k(\tau_j) \right)_{j \neq i}$ be level- k player i 's profile of prior beliefs. Furthermore, for each i and k , level- k player i believes any other player j 's level is independently distributed according to the lower truncated probability distribution function:

$$\mu_{ij}^k(\kappa) = \begin{cases} \frac{p_{j\kappa}}{\sum_{m=0}^{k-1} p_{jm}} & \text{if } \kappa < k \\ 0 & \text{if } \kappa \geq k. \end{cases} \quad (3.1)$$

The assumption underlying μ_{ij}^k is that level- k types of each player have a correct belief about the relative proportions of players who are less sophisticated than they are, but maintain the incorrect belief that other players of level $\kappa \geq k$ do not exist. The j subscript indicates that different players can have different level distributions.

Level-Dependent Profile of Behavioral Strategies

A profile of behavioral strategies is now a *level-dependent* profile of behavioral strategies specifying the behavioral strategy for each level of each player. We denote this profile as $\sigma = (\sigma_i^k)_{i \in N, k \in \mathbb{N}_0}$ where σ_i^k is the behavioral strategy adopted by level- k player i . Level-0 players are assumed to uniformly randomize at each

¹⁰For the sake of simplicity, we assume that the distribution of levels is *independent* of the probability distribution of Chance's moves σ_C .

information set.¹¹ That is, for all $i \in N$, $I \in \mathcal{I}_i$ and $a \in A(I)$, $\sigma_i^0(a | I) = \frac{1}{|A(I)|}$. In the following we may interchangeably call level-0 players *non-strategic players* and level $k \geq 1$ players *strategic players*.

In the DCH solution, strategic level- k players believe all other players are strictly less sophisticated than level- k . Therefore, in the following, we use $\sigma_j^{<k} = (\sigma_j^0, \dots, \sigma_j^{k-1})$ to denote the profile of behavioral strategies adopted by the levels below k of player j . In addition, let $\sigma_{-i}^{<k} = (\sigma_1^{<k}, \dots, \sigma_{i-1}^{<k}, \sigma_{i+1}^{<k}, \dots, \sigma_n^{<k})$ denote the profile of behavioral strategies of the levels below k of all players other than player i .

Level-Dependent Profile of Posterior Beliefs

For each player i with level $k \geq 1$, level k player i forms a posterior belief about the joint distribution of other players' levels of sophistication and the distribution of histories in any information set $I \in \mathcal{I}_i$. These posterior beliefs depend on the level-dependent profile of behavioral strategies of the other players, and their prior beliefs about the distribution of levels, $\mu_i^k(\tau_{-i})$.¹²

Specifically, for any $i \in N$ and $k \geq 1$, level- k player i forms posterior beliefs about the joint distribution of other players' levels of sophistication, τ_{-i} (lower than k) and histories $h \in I$ for any $I \in \mathcal{I}_i$. That is, level- k player i 's *personal system of beliefs* is defined as a function $v_i^k : \mathcal{I}_i \rightarrow \Delta(\{0, \dots, k-1\} \times H \setminus Z)$ such that

$$\sum_{h \in I} \sum_{\{\tau_{-i} : \tau_j < k \forall j \neq i\}} v_i^k(\tau_{-i}, h | I) = 1 \quad \text{for any } I \in \mathcal{I}_i.$$

DCH imposes a consistency restriction that the posterior beliefs of any level- k player i are derived from Bayes' rule, conditioned on the level-dependent profile of

¹¹Uniform randomization is not the only way to model non-strategic level-0 players, but there are several justifications for doing so. Firstly, one compelling reason is that it is universally applicable to all games, in exactly the same way, unlike almost any other specification since the cardinality of the set of available actions will typically vary across players and information sets. Probably for this reason, it is the most commonly used approach in applications of CH, including the original specification in Camerer, Ho, and Chong (2004). By taking the same approach in our generalization of CH to games in extensive form, it allows for clear comparisons to the original CH framework. Secondly, the choice of uniform randomization also appeals to the principle of insufficient reason, not only from the agnostic standpoint of the researcher, but also in the sense of being a neutral assumption about strategic players' beliefs about the behavior of nonstrategic players. Because uniform randomization is nondegenerate, a notable implication is that there is no off-path event in the DCH solution. Lastly, it is simple and parsimonious, as uniform randomization is not based on ad hoc assumptions tailored to specific games. In principle, the number of specifications of level-0 behavior is enormous, especially for games with many actions and information sets. Some alternatives to uniform randomization are noted in Section 3.7.4.

¹²Level-1 players do not update, since they have a degenerate prior belief that all other players are level-0.

behavioral strategies of lower levels being $\sigma_{-i}^{<k}$ and the game being at information set I . Because $\mu_{ij}^k(0) > 0$ for all $i, j \in N, k \geq 1$ and $\sigma_j^0(a|I) > 0$ for all $j, I \in \mathcal{I}_j$ and $a \in A(I)$, all strategic players believe every *history* (and hence information set) is reached with positive probability. Therefore, these posterior beliefs derived using Bayes' rule are well-defined everywhere. In the following, we use $v_i^k(\tau_{-i}, h | I; \sigma_{-i}^{<k})$ to denote the posterior belief of level- k player i , induced by $\sigma_{-i}^{<k}$, conditional on being at information set $I \in \mathcal{I}_i$. We call $v = (v_i^k)_{i \in N, k \geq 1}$ the *level-dependent profile of personal systems of beliefs*.

The DCH Solution

In the DCH solution, the posterior distribution of levels of other players will generally be different for different levels of the same player at the same information set, since the supports of those distributions will generally differ.¹³ This, in turn, induces different levels of the same player to have different beliefs about the probability distribution over the terminal payoffs that can be reached from that information set. For each $i \in N, k \geq 1, \sigma$, and τ_{-i} such that $\tau_j < k$ for all $j \neq i$, let $\tilde{\rho}_i^k(z | h, \tau_{-i}, \sigma_{-i}^{<k}, \sigma_i^k)$ be level- k player i 's belief about the conditional realization probability of $z \in Z_h$ at history $h \in H \setminus Z$, if the profile of levels of the other players is τ_{-i} and i is using strategy σ_i^k . Finally, level- k player i 's conditional expected payoff at information set I is given by

$$\mathbb{E}u_i^k(\sigma|I) = \sum_{h' \in I} \sum_{\{\tau_{-i}: \tau_j < k \forall j \neq i\}} \sum_{z \in Z_{h'}} v_i^k(\tau_{-i}, h' | I; \sigma_{-i}^{<k}) \tilde{\rho}_i^k(z | \tau_{-i}, h', \sigma_{-i}^{<k}, \sigma_i^k) u_i(z). \quad (3.2)$$

The *DCH solution* of the game is defined as the level-dependent profile of behavioral strategies, σ^* , such that $\sigma_i^{k*}(\cdot|I)$ maximizes $\mathbb{E}u_i^k(\sigma^*|I)$ for all i, k and $I \in \mathcal{I}_i$.¹⁴ Moreover, the *DCH belief system* is the level-dependent profile of personal belief systems induced by σ^* .

Remark 3.1. In the DCH solution, each level- k player i 's personal beliefs are defined as a joint distribution over histories and other players' levels of sophistication. An

¹³However, the support of the beliefs of all levels of all players will always include the type profile τ_{-i}^0 , in which all other players are level-0. That is, $v_i^k(\tau_{-i}^0, h | I; \sigma_{-i}^{<k}) > 0$ for all $i \in N, k \geq 1$, information set I and $h \in I$.

¹⁴We assume (as is typical in level- k models) that players randomize uniformly over optimal actions when indifferent. This assumption is convenient because it ensures a unique DCH solution to every game, so we assume it here. Note that while the DCH solution is defined as a fixed point, it can be solved for recursively, starting with the lowest level and iteratively working up to higher levels.

alternative formulation of personal beliefs would be to define level- k player i 's updating process on τ_{-i} history-by-history, i.e., $v_i^k(\tau_{-i}|h; \sigma_{-i}^{\leq k})$, and then derive the conditional posterior beliefs, $\pi_i^k(h)$ over the histories in every information set of player i in the game, for all i, k , and $I \in \mathcal{I}_i$, $\sum_{h' \in I} \pi_i^k(h') = 1$. See Lin and Palfrey (2022). For games with perfect information, these two formulations are identical since every information set is a singleton. Therefore, in games of perfect information, for every level- k player i and any $h \in H_i$, the DCH belief system simply reduces to a profile of personal belief systems $v = (v_i^k)_{i \in N, k \geq 1}$ where $v_i^k : H_i \rightarrow \Delta(\{0, \dots, k-1\})$ and $v_i^k(\tau_{-i}|h; \sigma_{-i}^{\leq k})$ is the posterior belief about the level profile of other players at history h .

To simplify notation and exposition, most of the remainder of the paper studies DCH in games of perfect information. Some the properties established in the next section for games of perfect information apply more generally, and these cases are noted. We discuss additional extensions to games of imperfect information in Section 3.7.2 and 3.7.3.

3.4 Properties of the DCH Solution

Section 3.4.1 first establishes the general properties of the belief-updating process. Section 3.4.2 explores the relationship between the DCH solution and subgame perfect equilibrium. Finally, we illustrate the *strategy-reduction effect* predicted by DCH in Section 3.4.3. Specifically, we show that the DCH solution might be dramatically different for two different games in extensive form that share the same reduced normal form.

3.4.1 Properties of the Belief-Updating Process

The first result shows that for games of perfect information, the updating process satisfies an independence property. Specifically, the following proposition establishes that all levels of all players will update their posterior beliefs about other players' levels independently.

Proposition 3.1. *For any finite game of perfect information Γ , any $i \in N$, any $h \in H_i$, and for any $k \in \mathbb{N}$, level- k player i 's posterior belief about other players' levels at history h is independent across players. That is, $v_i^k(\tau_{-i} | h; \sigma_{-i}^{\leq k}) = \prod_{j \neq i} v_{ij}^k(\tau_j | h; \sigma_{-i}^{\leq k})$ where $v_{ij}^k(\tau_j | h; \sigma_{-i}^{\leq k})$ is level- k player i 's marginal posterior belief about player j being level τ_j at history h .*

Proof: We prove this proposition by induction on the length of the sequence of h , which we denote by $|h|$. Let σ be any level-dependent strategy profile and p be any prior distribution over types. First we can notice that at the initial history, i.e., $|h| = 0$, $i = P(\phi)$, and any level $k > 0$, $v_i^k(\tau_{-i} | \phi; \sigma_{-i}^{<k}) = \prod_{j \neq i} \mu_{ij}^k(\tau_j)$ as players' levels are independently determined. For any history h' and h'' , we define a partial order $<$ on H such that $h' < h''$ if and only if h' is a subsequence of h'' . In the following, for any $h', h'' \in H$ where $h' < h''$, we use $\alpha(h', h'')$ to denote the unique action at h' that leads to h'' .

To establish the base case, consider any player i with any level $k > 0$ and any $h \in H_i$ such that $|h| = 1$ and $j = P(\phi)$ where $j \neq i$. Because player j has made the only move in the game so far, and the prior distribution of types is assumed to be independent across players, we have, for any τ_{-i} such that $\tau_{i'} < k \forall i' \neq i, j$:

$$\begin{aligned} v_i^k(\tau_{-i} | h; \sigma_{-i}^{<k}) &= \frac{\sigma_j^{\tau_j}(\alpha(\phi, h) | \phi) \mu_{ij}^k(\tau_j)}{\sum_{l=0}^{k-1} \sigma_j^l(\alpha(\phi, h) | \phi) \mu_{ij}^k(l)} \prod_{i' \neq i, j} \mu_{ii'}^k(\tau_{i'}) \\ \implies v_i^k(\tau_{-i} | h; \sigma_{-i}^{<k}) &= \prod_{j \neq i} v_{ij}^k(\tau_j | h; \sigma_{-i}^{<k}), \end{aligned}$$

where we know $\sum_{l=0}^{k-1} \sigma_j^l(\alpha(\phi, h) | \phi) > 0$ because $\sigma_j^0(\alpha(\phi, h) | \phi) = \frac{1}{|A(\phi)|} > 0$. Therefore, the result is true for $|h| = 1$.

Next, consider any player i with any level $k > 0$ and suppose that $v_i^k(\tau_{-i} | h; \sigma_{-i}^{<k}) = \prod_{j \neq i} v_{ij}^k(\tau_j | h; \sigma_{-i}^{<k})$ for all $h \in H_i$ such that $|h| = 1, 2, \dots, t-1$. It suffices to complete the proof by considering any $h \in H_i$ such that $|h| = t$. Because in games of perfect information, actions are perfectly observed, level k player i 's belief about the level profile of others being τ_{-i} such that $\tau_j < k$ for any $j \neq i$ is:

$$\begin{aligned} v_i^k(\tau_{-i} | h; \sigma_{-i}^{<k}) &= \frac{\prod_{\{\tilde{h}: \tilde{h} < h, P(\tilde{h}) \neq i\}} \sigma_{P(\tilde{h})}^{\tau_{P(\tilde{h})}}(\alpha(\tilde{h}, h) | \tilde{h})}{\sum_{\{\tau'_{-i}: \tau'_j < k \forall j \neq i\}} \prod_{\{\tilde{h}: \tilde{h} < h, P(\tilde{h}) \neq i\}} \sigma_{P(\tilde{h})}^{\tau'_{P(\tilde{h})}}(\alpha(\tilde{h}, h) | \tilde{h})} \\ &= \frac{\prod_{j \neq i} f_j(h | \sigma_j^{\tau_j})}{\sum_{\{\tau'_{-i}: \tau'_j < k \forall j \neq i\}} \prod_{j \neq i} f_j(h | \sigma_j^{\tau'_j})} = \prod_{j \neq i} \left[\frac{f_j(h | \sigma_j^{\tau_j})}{\sum_{l=0}^{k-1} f_j(h | \sigma_j^l)} \right] \\ \implies v_i^k(\tau_{-i} | h; \sigma_{-i}^{<k}) &= \prod_{j \neq i} v_{ij}^k(\tau_j | h; \sigma_{-i}^{<k}), \end{aligned}$$

where $f_j(h | \sigma_j^l)$ is the probability that player j moves along the path to reach h

given player j is using σ_j^l . That is,

$$f_j(h \mid \sigma_j^l) = \begin{cases} \prod_{\{\tilde{h}: \tilde{h} < h, P(\tilde{h})=j\}} \sigma_j^l(\alpha(\tilde{h}, h) \mid \tilde{h}) & \text{if } \{\tilde{h} : \tilde{h} < h, P(\tilde{h}) = j\} \neq \emptyset \\ 1 & \text{otherwise.} \end{cases}$$

This completes the proof. ■

What drives this result is that in games of perfect information, when player j moves, all players perfectly observe this history. As a result, all players other than j only update their beliefs about the level of player j , and do not update their beliefs about any of the other players. In addition, the assumption of independence of the distribution of player levels is used. If levels are correlated across players then it's possible that player i can update their beliefs about the level of player j based on actions taken by player l . From Proposition 3.1, we can see that the marginal posterior belief of level- k player i to player j 's belief only depends on player j 's moves along the history. Therefore, we can obtain that $v_{ij}^k(\kappa \mid h; \sigma_{-i}^{<k}) = v_{ij}^k(\kappa \mid h; \sigma_j^{<k})$.

The second property of the DCH solution is that in the later histories, the support of the posterior beliefs is (weakly) shrinking. In this sense, the players would have a more precise posterior belief when the history gets longer. For any player $i, j \in N$ such that $i \neq j$, for any $h \in H_i$, and for any $k \in \mathbb{N}$, we denote the support of level- k player i 's belief about player j 's level as

$$\text{supp}_{ij}^k(h) \equiv \{\tau_j \in \{0, 1, \dots, k-1\} \mid v_{ij}^k(\tau_j \mid h; \sigma_j^{<k}) > 0\}.$$

This property is formally stated in the following proposition.

Proposition 3.2. *In any finite game of perfect information Γ , for all $i, j \in N$, $k \in \mathbb{N}$, and any $h, h' \in H_i$, if $h' < h$, then $\text{supp}_{ij}^k(h) \subseteq \text{supp}_{ij}^k(h')$.*

Proof: To prove the statement, it suffices to show that $\kappa \notin \text{supp}_{ij}^k(h') \implies \kappa \notin \text{supp}_{ij}^k(h)$ for all $\kappa = 0, 1, \dots, k-1$. From the proof of Proposition 3.1, we can obtain that

$$\begin{aligned} v_{ij}^k(\kappa \mid h; \sigma_j^{<k}) &= \frac{f_j(h \mid \sigma_j^\kappa)}{\sum_{l=0}^{k-1} f_j(h \mid \sigma_j^l)} \\ &= \begin{cases} \frac{\prod_{\{\tilde{h}: h' < \tilde{h} < h, P(\tilde{h})=j\}} \sigma_j^\kappa(\alpha(\tilde{h}, h) \mid \tilde{h}) v_{ij}^k(\kappa \mid h'; \sigma_j^{<k})}{\sum_{l=0}^{k-1} \prod_{\{\tilde{h}: h' < \tilde{h} < h, P(\tilde{h})=j\}} \sigma_j^l(\alpha(\tilde{h}, h) \mid \tilde{h}) v_{ij}^k(l \mid h'; \sigma_j^{<k})} & \text{if } \{\tilde{h} : h' < \tilde{h} < h, P(\tilde{h}) = j\} \neq \emptyset \\ v_{ij}^k(\kappa \mid h'; \sigma_j^{<k}) & \text{otherwise.} \end{cases} \end{aligned}$$

Hence, $v_{ij}^k(\kappa | h'; \sigma_j^{<k}) = 0 \implies v_{ij}^k(\kappa | h; \sigma_j^{<k}) = 0$, so $\text{supp}_{ij}^k(h) \subseteq \text{supp}_{ij}^k(h')$. This completes the proof. ■

Notice that although players are unable to perfectly observe all previous actions in games of imperfect information, Proposition 3.2 still holds—the support of marginal beliefs about others' levels in later information sets is (weakly) shrinking. See Lin and Palfrey (2022). Besides, the assumption of independence of the distribution of levels is not used in the proof. In other words, the shrinkage of the support is irreversible, regardless of how the levels are distributed.

In addition, there are a few additional remarks about the properties of the updating process in the DCH solution that are worth highlighting. First, there is a second source of learning, besides the shrinking support property, which is that after each move by an opponent, each strategic player with level $k \geq 2$ updates the probability that the opponent is level-0.¹⁵ This in turn leads to updating of the relative likelihood of the higher strategic types of the opponent, since the probabilities have to sum to 1. Second, as the game unfolds, the beliefs of higher level players about their opponents can be updated in either direction, in the sense of believing an opponent is either more or less sophisticated. Examples in the next section will illustrate this. Third, while players' belief-updating process is adaptive, nonetheless all players are *strategically forward-looking* (rather than myopic) in the sense that players take into account and correctly anticipate how all players in the game will update beliefs at each future history.

Since the players are forward-looking and have truncated rational expectations, it is natural to ask if there is any connection between our model and perfect or sequential equilibrium. We explore this relationship in the next section.

3.4.2 DCH and Subgame Perfect Equilibrium

In this section, we study the relationship between the DCH solution and subgame perfect equilibrium through two simple examples. One question we address is whether sufficiently high-level players always behave consistently with rational backward induction. As it turns out, this is *not* generally true. In the following series of simple two-person games in extensive form, we demonstrate how high-level players could violate backward induction either on or off the equilibrium path, suggesting the DCH solution is fundamentally different from subgame perfection. For the sake

¹⁵The updating by strategic players' beliefs about level-0 opponents can be either increasing or decreasing.

of simplicity, in this section and for the rest of the paper (except for Section 3.4.2.2) we assume every player's level distribution is identical.

3.4.2.1 Violating Backward Induction at Some Subgame




Figure 3.1: Game Tree of Example 3.4.2.1. A “+” sign indicates a move is chosen by the specified level and all higher levels. The subgame perfect equilibrium moves are marked with arrows.

Example 3.4.2.1 demonstrates how backward induction could be violated by every level of player at some subgame. The game tree for this two-person game of perfect information is shown in Figure 3.1. Suppose every player's level is independently drawn from $\text{Poisson}(1.5)$, which has been suggested by Camerer, Ho, and Chong (2004) as an empirically plausible distribution. Every level of players' move choices are labelled in the figure, with a “+” sign indicating a move is chosen by the specified level type and all higher levels. For instance, level-1 player 1 chooses r_{1a} at the beginning while level-2 and above choose l_{1a} . Calculations can be found in Appendix C.

To illustrate the mechanics of the DCH solution in this example, it is useful to begin by focusing on subgame 2a. In this subgame, level-2 and higher-level of player 2 would update from the information that player 1 is not a level-1 player, leading a level-2 player 2 to choose l_{2a} because the updated belief puts all weight on player 1 being level-0. However, a level-3 player 2 places positive posterior probability

on player 1 being level-2, and as long as this posterior probability is high enough, it is optimal for level-3 player 2 to choose r_{2a} —as if player 2 were engaged in the same backward induction reasoning used to justify the subgame perfect equilibrium. Following a similar logic, all high-level players would behave this way in the left branch of the game, where player 1 chooses l_{1a} at the beginning.

However, this is not the case for the right branch of the game after player 1 chooses r_{1a} at the beginning. At subgame $1c$, the move predicted by the subgame perfect equilibrium is *never* chosen by any strategic player 1. Hence, in the DCH solution for this example, high-level player 1's behavior is consistent with subgame perfect equilibrium on the left branch but not on the right branch. \square

3.4.2.2 Dominated Actions

As we examine this example carefully, we can find the key of this phenomenon is that player 1 knows the subgame $h = 1c$ can be reached only if player 2 chooses a strictly dominated action¹⁶ in the previous stage. One can think of player 2's decision at subgame $h = 2b$ as a rationality check in the following sense. Whenever player 2 chooses r_{2b} , the support of strategic player 1's posterior belief will shrink to a singleton—he will believe player 2 is level-0. This extreme posterior belief would lead a strategic player 1 to deviate from subgame perfect strategy.

Generally speaking, if a history contains some player's strictly dominated action, then all other players will immediately believe this player is non-strategic and best respond to such strategy. As a result, it is possible that the strategy profile will not be the subgame perfect equilibrium for every strategic level. This argument holds as long as level-0 player's strategy is the beginning of the hierarchical reasoning process—no matter how small the proportion of level-0 players is. However, since paths with strictly dominated actions can be realized only if some player is level-0, paths containing strictly dominated actions occur with vanishing probability as the proportion of level-0 players converges to 0. Proposition 3.3 formally shows this conclusion.

¹⁶Formally speaking, at any history $h \in H \setminus Z$ with $i = P(h)$, an action $a' \in A(h)$ is strictly dominated if there is another action $a'' \in A(h)$ such that

$$\min_{z \in Z_{h''}} u_i(z) > \max_{z \in Z_{h'}} u_i(z),$$

where $h' = (h, a')$ and $h'' = (h, a'')$.

Proposition 3.3. *Consider any finite game of perfect information where each player i 's level is drawn from the distribution $p_i = (p_{ik})_{k=0}^{\infty}$. If some history h can occur only if some player chooses a strictly dominated action, then the probability for such history being realized converges to 0 as $p_0 = (p_{i0})_{i \in N} \rightarrow (0, \dots, 0)$.*

Proof: Consider any h that can occur only if some player chooses a strictly dominated action. That is, there is h' that is a subsequence of h with $i = P(h')$ such that there is a strictly dominated action $a' \in A(h')$ and (h', a') is a subsequence of h . Since this is a strictly dominated action, it can only be chosen by a level-0 player. Therefore, the ex ante probability for player i to choose a' at h' is

$$\Pr(a' | h') = \sum_{j=0}^{\infty} \sigma_i^j(a' | h') p_{ij} = \sigma_i^0(a' | h') p_{i0} = \frac{1}{|A(h')|} p_{i0}.$$

Lastly, the ex ante probability for h to be realized, $\Pr(h)$, is smaller than $\Pr(a'|h')$ and hence

$$\lim_{p_0 \rightarrow (0, \dots, 0)} \Pr(h) \leq \lim_{p_0 \rightarrow (0, \dots, 0)} \Pr(a' | h') = \lim_{p_0 \rightarrow (0, \dots, 0)} \frac{1}{|A(h')|} p_{i0} = 0.$$

This completes the proof. ■

It is worth noticing that the independence of the distribution of the levels is not required in this proposition as strictly dominated actions will only be chosen by level-0 players. Moreover, one can see this principle in play in Example 3.4.2.1 where player 1's anomalous behavior only happens when player 2 chooses a strictly dominated action, which is only chosen by level-0. For other parts of the game, if both players are at least level 3, DCH predicts the game will follow the subgame perfect equilibrium path.

Since the subgame perfect equilibrium path never contains strictly dominated actions, one might be tempted to conjecture that the equilibrium path is always followed by sufficiently sophisticated players. The next example demonstrates that this is not true. In fact, it is possible that the subgame perfect equilibrium path is *never* chosen by strategic players, so high-level players in our model do not necessarily converge to the subgame perfect equilibrium.

3.4.2.3 Violating Backward Induction on the Equilibrium Path

Example 3.4.2.3 is modified from the previous example by changing player 1's payoff from 4 to $\frac{3}{2}$ as he chooses r_{1b} at history $1b$. Decreasing the payoff does




Figure 3.2: Game Tree of Example 3.4.2.3. A “+” sign indicates a move is chosen by the specified level and all higher levels. The subgame perfect equilibrium moves are marked with arrows.

not affect the subgame perfect equilibrium. However, this change makes low-level players think the subgame perfect equilibrium actions are not profitable, causing a domino effect that high-level players think the equilibrium actions are not optimal as well. Here we consider an arbitrary prior distribution $p = (p_k)_{k=0}^{\infty}$. The game tree is shown in Figure 3.2 with every level of players’ decisions. The calculations can be found in Appendix C.

Level-1 players will behave the same as in the previous example. However, the change of payoffs makes l_{1a} not profitable for level-2 player 1 at the initial history. Hence, player 2 would believe player 1 is certainly level-0 whenever the game proceeds to the left branch. Moreover, every level of players would behave the same by the same logic. As a result, the subgame perfect equilibrium path is *never* chosen by strategic players. If p_0 is close to 0, the subgame perfect equilibrium outcome will almost never be reached.

Instead, there is an imperfect Nash equilibrium that can be supported by the strategy profile of every strategic level of both players. Loosely speaking, the belief updating process gets “stuck” at this equilibrium, causing all higher-level players behave in

the same way.¹⁷ □

3.4.3 The Strategy-Reduction Effect of DCH




Figure 3.3: Game Tree of Γ . A “+” sign indicates a move is chosen by the specified level and all higher levels. The subgame perfect equilibrium moves are marked with arrows.

An interesting feature of the DCH solution is *the strategy-reduction effect*. That is, the DCH solution can differ for two games that share the same reduced normal form. To illustrate this effect, we first consider a toy example in this section and provide a detailed analysis of the strategy-reduction effect for a class of increasing-sum centipede games in the subsequent section.

Consider the game in extensive form Γ , whose game tree is shown in Figure 3.3. This example is almost exactly the same as Example 3.4.2.3, with the single exception being that player 1’s payoff changes from 3 to 8 after choosing r_{2a} at subgame $2a$. This change does not affect the subgame perfect equilibrium, but makes choosing l_{1a} profitable again for high-level player 1. (Here we again assume the prior distribution follows $\text{Poisson}(1.5)$.) Consequently, higher levels ($k \geq 3$) of DCH players in this

¹⁷The following strategy profile defines this imperfect equilibrium: player 1 chooses r_{1a} at the beginning, r_{1b} at subgame $h = 1b$, and chooses l_{1c} at subgame $h = 1c$; player 2 chooses l_{2a} at subgame $h = 2a$, l_{2b} at subgame $h = 2b$, and chooses r_{2c} at subgame $h = 2c$. Therefore, $(2, 5)$ is an equilibrium outcome.

game will choose actions that lead to the subgame perfect equilibrium outcome, (8, 4).

This switch to the subgame perfect outcome is a direct consequence of the belief-updating process of the DCH solution. Although the payoff 10 is really attractive to player 1, strategic player 1 will realize he can get it only if player 2 is level-0. Therefore, if there is a high enough probability of higher levels of player 2, player 1 will realize he is likely to get the lower payoff of 3 at node $2c$. Hence, a high-level player 1 will choose l_{1a} at the beginning (as if conducting backward induction). As long as there are enough strategic types of player 1 choosing l_{1a} , higher levels of player 2 will update accordingly and choose the subgame perfect equilibrium action r_{2a} . The calculations can be found in Appendix C.

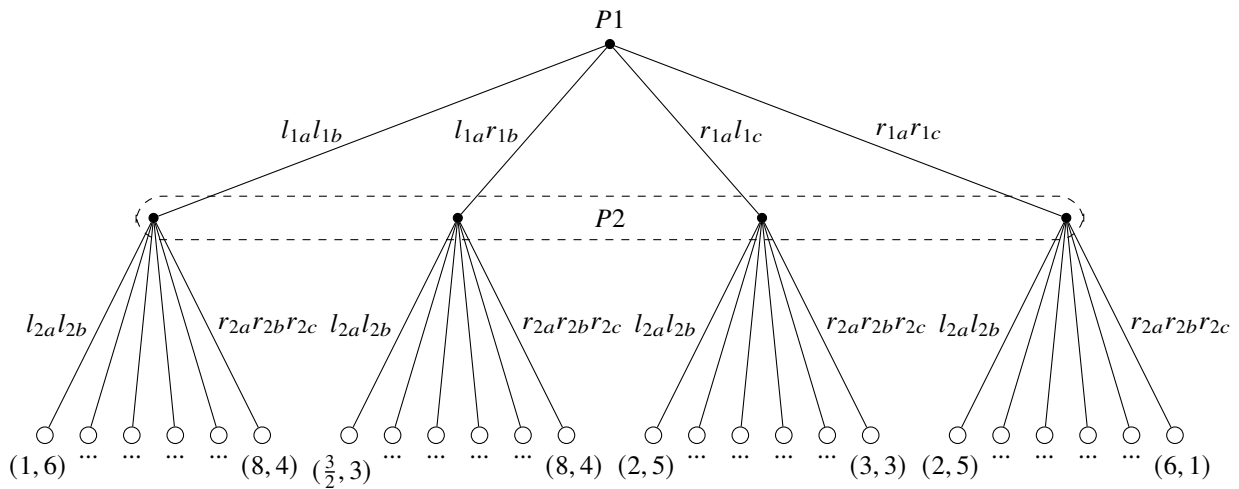


Figure 3.4: Game Tree of Γ' where player 1 first chooses a reduced strategy of Γ and then player 2 chooses a reduced strategy of Γ without observing player 1's action.

To illustrate the strategy-reduction effect, consider another game in extensive form Γ' , whose game tree is shown in Figure 3.4. In this game, player 1 moves first and chooses one of the reduced strategies in Γ . That is, player 1's action set is $A_1 = \{l_{1a}l_{1b}, l_{1a}r_{1b}, r_{1a}l_{1c}, r_{1a}r_{1c}\}$. After that, player 2 chooses one of the reduced strategies in Γ' without observing player 1's action. Γ and Γ' share the same reduced normal form as shown in Table 3.1.

Since Γ' is essentially a simultaneous-move game, the DCH solution of Γ' coincides with the standard CH solution of the 4×6 matrix game displayed in Table 3.1. It is easy to see that level-1 and higher-level of player 1 will choose the strategy $r_{1a}l_{1c}$ and

level-1 and higher-level of player 2 will choose the strategy $l_{2a}l_{2b}$, as indicated in the table. This result illustrates how DCH could dramatically differ between different games that share the same reduced normal form. Specifically, DCH predicts that all strategic levels of player 1 and 2 in Γ' will choose $r_{1a}l_{1c}$ and $l_{2a}l_{2b}$, respectively, leading to a different outcome compared to the DCH solution of Γ .

Table 3.1: Reduced Normal Form of Γ and Γ' .

Player 1		Player 2					
		$l_{2a}l_{2b}$	$r_{2a}l_{2b}$	$l_{2a}r_{2b}l_{2c}$	$l_{2a}r_{2b}r_{2c}$	$r_{2a}r_{2b}l_{2c}$	$r_{2a}r_{2b}r_{2c}$
		L1+					
$l_{1a}l_{1b}$		1,6	8,4	1,6	1,6	8,4	8,4
$l_{1a}r_{1b}$		3/2,3	8,4	3/2,3	3/2,3	8,4	8,4
$r_{1a}l_{1c}$	L1+	2,5	2,5	10,2	3,3	10,2	3,3
$r_{1a}r_{1c}$		2,5	2,5	6,1	6,1	6,1	6,1

When examining this result carefully, one can realize that the driving force behind this non-equivalence between Γ and Γ' is the difference in the numbers of available strategies, which alters the behavior of level-0. In particular, level-0 player 2 uniformly randomizes between six reduced strategies in Γ' , whereas level-0 player 2 uniformly randomizes among eight non-reduced strategies in Γ . This difference changes the behavior of level-1 player 1 (who best responds to level-0 player 2), which in turn alters the behavior of level-2 player 2, and so on for all higher levels. \square

Remark 3.2. Battigalli, Leonetti, and Maccheroni (2020) prove that two extensive game structures with imperfect information share the same mapping from profiles of reduced strategies to induced terminal paths if and only if one can be transformed into the other using two elementary transformations: *Interchanging of simultaneous moves (INT)* and *Coalescing move/sequential agent splitting (COA)*. The DCH solution is invariant under INT, but is not invariant under COA unless a balancedness property is satisfied. For a more detailed discussion, see Battigalli (2023).

In addition, one property of the standard CH model identified by Camerer, Ho, and Chong (2004) is that if a level- k player plays a (pure) equilibrium strategy, then all higher levels of that player will play that strategy too. One may wonder if an analogous property holds in the DCH solution. That is, if some level of a player chooses on the equilibrium path, do all higher-levels of that player choose that action too? The game Γ provides a counterexample for this conjecture. At the

initial history, level-1 player 1 chooses the equilibrium path l_{1a} . However, level-2 player 1 switches to r_{1a} , and level-3 (and above) player 1 switches back to l_{1a} .

The underlying reason is that even if a level- k player chooses the equilibrium path, a higher-level player could still deviate from the equilibrium path if other players do not move along the equilibrium path in later subgames. In this example, level-1 player 1 chooses l_{1a} at the beginning to best respond to level-0 player 2. Yet level-1 player 2 does not choose the equilibrium path at the subgame $h = 2a$, causing level-2 player 1 to choose r_{1a} at the beginning. Level-2 (and above) player 2 switches to the equilibrium path at the subgame $h = 2a$, and this information can only be updated by level-3 (and above) player 1. Finally, as long as there are enough level-2 (and above) players, high-level player 1 would switch back to the equilibrium path, creating a non-monotonicity.

3.5 An Application: Centipede Games




Figure 3.5: $2S$ -Move Centipede Game

In this section, we explore the strategy-reduction effect in much more detail, focusing on the class of “linear centipede games,” which is illustrated in Figure 3.5. The games in this class are described in the following way. Player 1, the first-mover, and player 2, the second-mover, alternate over a sequence of moves. At each move, the player whose turn it is can either end the game (“take”) and receive the larger of two payoffs or allow the game to continue (“pass”), in which case both the large and the small payoffs are incremented by an amount $c > 0$. The difference between the large and the small payoffs equals 1 and does not change. The game continues for at most $2S$ decision nodes (stages) where $S \geq 2$, and we label the decision nodes by $\{1, 2, \dots, 2S\}$. Player 1 moves at odd nodes and player 2 moves at even nodes. If the game is ended by a player at stage $j \leq 2S$, the payoffs are $(1 + (j - 1)c, (j - 1)c)$ if j is odd and $((j - 1)c, 1 + (j - 1)c)$ if j is even. If no player ever takes, the payoffs

are $(1 + 2Sc, 2Sc)$. Thus, a linear centipede game has two parameters: (S, c) . To avoid trivial cases, we assume $\frac{1}{3} < c < 1$.¹⁸

Specifically, we will compare each level of players' behavior when the centipede game is played in two different representations given the same prior distribution. In Section 3.5.1, we first characterize the DCH solution of the centipede game when it is played in its original extensive form (as shown in Figure 3.5), by which we mean that the game is played as an alternating-move sequential game. In Section 3.5.2, we then characterize the DCH solution of the centipede game when it is played in reduced normal form, by which we mean that both players simultaneously choose a reduced strategy.¹⁹ From the perspective of the standard equilibrium theory, these two implementations of the game would not induce different outcomes because they share the same reduced normal form. However, whether playing the game in extensive form (using the direct-response method) or in reduced normal form (using the strategy method) will induce the same behavior is still an open question that is under debate in experimental methodology.

From the perspective of DCH, the key difference between the direct-response method and the strategy method is that the cardinalities of action sets are different, which implies different behavior for level-0 players. Since the DCH solution is solved recursively from the bottom of the hierarchy, this non-equivalence of level-0 behavior triggers a chain reaction that affects the behavior of all higher levels.

3.5.1 DCH for the Centipede Game in Extensive Form

When the centipede game is played in extensive form, players take turns moving in an alternating-move sequential game, where each player can move at (most) S stages. Therefore, a (behavioral) strategy for player i is an S -tuple where each element is the probability to take at the corresponding decision node. That is, $\sigma_1 = (\sigma_{1,1}, \dots, \sigma_{1,S})$ and $\sigma_2 = (\sigma_{2,1}, \dots, \sigma_{2,S})$ are player 1 and 2's strategies, respectively. For every $1 \leq j \leq S$, $\sigma_{1,j}$ is the probability that player 1 would take at stage $2j - 1$ and $\sigma_{2,j}$ is the probability that player 2 would take at stage $2j$.

Following the notation introduced earlier, we use σ_1^k and σ_2^k to denote level- k players' strategies. Level-0 players uniformly randomize at each stage. That is,

¹⁸If $c > 1$, the unique equilibrium is for every player to pass at every node. If $c < \frac{1}{3}$, all strategic players will always take, so CH behavior is the same as subgame perfect equilibrium behavior.

¹⁹When we say a centipede game is played in its reduced normal form, we mean that a player's strategy corresponds to the node at which they will stop the game by taking (or always pass). Therefore, each player has $S + 1$ available strategies in the reduced normal form.

$\sigma_1^0 = \sigma_2^0 = \left(\frac{1}{2}, \dots, \frac{1}{2}\right)$. Finally, to simplify the notation, for every stage $1 \leq j \leq 2S$, we let $v_j^k(\cdot)$ be level- k stage j -mover's belief about the opponent's level at stage j where $v_j^k(\tau_{-P(j)}) \equiv v_{P(j), -P(j)}^k(\tau_{-P(j)} \mid j; \sigma_{-P(j)}^{<k})$.

To fully characterize every level of players' strategies, we need to compute every level of players' best responses at every subgame. In principle, we have to solve the behavior of each level recursively. However, since each level of players' strategy is *monotonic*—when the player decides to take at some stage, he will take in all of his later subgames—we can alternatively characterize the solution by identifying the lowest level of player to take at every subgame.

In Lemma 3.1, we characterize level-1 players' behavior and establish the monotonicity result. These results are straightforward and follow from the assumption that $\frac{1}{3} < c < 1$.

Lemma 3.1. *In linear centipede games with an extensive form shown in Figure 3.5, if $\frac{1}{3} < c < 1$, then*

1. $\sigma_{2,S}^k = 1$ for all $k \geq 1$.
2. $\sigma_1^1 = (0, \dots, 0)$ and $\sigma_2^1 = (0, \dots, 0, 1)$.
3. For every $k \geq 2$ and every $1 \leq j \leq S - 1$,
 - (i) $\sigma_{1,j}^k = 0$ if $\sigma_{2,j}^m = 0$ for every $1 \leq m \leq k - 1$;
 - (ii) $\sigma_{2,j}^k = 0$ if $\sigma_{1,j+1}^m = 0$ for every $1 \leq m \leq k - 1$.

Proof: See Appendix C. ■

Lemma 3.1 has three parts: (1) every strategic player 2 takes at the last stage; (2) completely characterizes level-1 strategies—player 1 passes at every stage and player 2 passes at every stage except for the last stage; (3) provides necessary conditions for higher levels to take at some stage. For any level $k \geq 2$ and any stage $1 \leq j \leq 2S - 1$, a level- k player would take at stage j only if there is some lower level player that would take at the next stage. Otherwise, it is optimal for a level- k player to pass at stage j .

The general characterization of level- k optimal strategies is in terms of the following **cutoffs**, specifying, for each stage, the lowest level type to take at that stage.

Definition 3.1. For every stage j where $1 \leq j \leq 2S$, define the **cutoff**, K_j^* be the lowest level of player that would take at this stage. In other words,

$$K_j^* = \begin{cases} \arg \min_k \left\{ \sigma_{1, \frac{j+1}{2}}^k = 1 \right\}, & \text{if } j \text{ is odd} \\ \arg \min_k \left\{ \sigma_{2, \frac{j}{2}}^k = 1 \right\}, & \text{if } j \text{ is even} \\ \infty, & \text{if } \nexists k \text{ s.t. } \sigma_{1, \frac{j+1}{2}}^k = 1 \text{ or } \sigma_{2, \frac{j}{2}}^k = 1. \end{cases}$$

Based on Definition 3.1, the monotonicity obtained in part (3) of Lemma 3.1 implies the following two results about cutoffs and strategies. Together they show that for any stage, a player's strategy will be to take at that stage if and only if his level is greater or equal to the cutoff.

Proposition 3.4. For every $1 \leq j \leq 2S - 1$,

1. $K_j^* \geq K_{j+1}^* + 1$ if $K_{j+1}^* < \infty$;
2. $K_j^* = \infty$ if $K_{j+1}^* = \infty$.

Proof: See Appendix C. ■

Proposition 3.5. For every $1 \leq j \leq 2S - 1$,

1. $\sigma_{1, \frac{j+1}{2}}^k = 1$ for all $k \geq K_j^*$ if j is odd and $K_j^* < \infty$;
2. $\sigma_{2, \frac{j}{2}}^k = 1$ for all $k \geq K_j^*$ if j is even and $K_j^* < \infty$.

Proof: See Appendix C. ■

Hence, cutoffs characterize optimal strategies of each level of each player, with a cutoff defining the lowest level that would take at each stage and all higher levels of that player would also take at that stage. The next two propositions establish recursive necessary and sufficient conditions for the existence of some level of some player to take at each stage. The proofs of these propositions provide a recipe for computing cutoffs.

Proposition 3.6. $K_{2S-1}^* < \infty \iff p_0 < \frac{2^S}{2^{S+} \left(\frac{3c-1}{1-c} \right)}$

First, we note that the proofs are simplified somewhat by observing the following identity:

$$p_0 < \frac{2^S}{2^S + \left(\frac{3c-1}{1-c}\right)} \iff \frac{p_0 \left(\frac{1}{2}\right)^S}{p_0 \left(\frac{1}{2}\right)^{S-1} + (1-p_0)} < \frac{1-c}{1+c}.$$

Proof: Only if: Suppose $K_{2S-1}^* < \infty$. By Proposition 3.4, $K_j^* \geq K_{2S-1}^*$ for all $j < 2S - 1$. Hence, the belief of level K_{2S-1}^* of player 1 that player 2 is level-0 at stage $2S - 1$ equals

$$v_{2S-1}^{K_{2S-1}^*}(0) = \frac{p_0 \left(\frac{1}{2}\right)^{S-1}}{p_0 \left(\frac{1}{2}\right)^{S-1} + \sum_{l=1}^{K_{2S-1}^*-1} p_l},$$

since it is optimal for $K_{2S-1}^* < \infty$ to take at $2S - 1$. This implies $\frac{p_0 \left(\frac{1}{2}\right)^S}{p_0 \left(\frac{1}{2}\right)^{S-1} + \sum_{l=1}^{K_{2S-1}^*-1} p_l} < \frac{1-c}{1+c}$ and hence

$$\frac{p_0 \left(\frac{1}{2}\right)^S}{p_0 \left(\frac{1}{2}\right)^{S-1} + (1-p_0)} < \frac{1-c}{1+c} \iff p_0 < \frac{2^S}{2^S + \left(\frac{3c-1}{1-c}\right)}.$$

If: Suppose $K_{2S-1}^* = \infty$. Then from Proposition 3.4, $K_j^* = \infty$ for all $j < 2S - 1$. That is, all levels of both players pass at every stage up to and including $2S - 1$. Hence, the belief of level $k \geq 1$ of player 1 that player 2 is level-0 at stage $2S - 1$ equals to

$$v_{2S-1}^k(0) = \frac{p_0 \left(\frac{1}{2}\right)^{S-1}}{p_0 \left(\frac{1}{2}\right)^{S-1} + \sum_{l=1}^{k-1} p_l} > \frac{p_0 \left(\frac{1}{2}\right)^{S-1}}{p_0 \left(\frac{1}{2}\right)^{S-1} + (1-p_0)}.$$

Since $K_{2S-1}^* = \infty$, it is optimal to pass at $2S - 1$ for all levels $k \geq 1$ of player 1, implying

$$\frac{p_0 \left(\frac{1}{2}\right)^S}{p_0 \left(\frac{1}{2}\right)^{S-1} + (1-p_0)} \geq \frac{1-c}{1+c}.$$

This completes the proof. ■

Thus, p_0 must be sufficiently small, and the condition is easier to satisfy the smaller c is (the potential gains to passing) and the larger is S (the horizon). If this condition

holds, there exists some strategic player 1 that takes at stage $2S - 1$. The proof also provides an insight for how the cutoffs can be computed. Specifically, the K_{2S-1}^* cutoff is computed as

$$K_{2S-1}^* = \arg \min_k \left\{ \frac{p_0 \left(\frac{1}{2}\right)^S}{p_0 \left(\frac{1}{2}\right)^{S-1} + \sum_{l=1}^{k-1} p_l} < \frac{1-c}{1+c} \right\}.$$

Cutoffs for earlier stages can be derived recursively as the following proposition establishes.

Proposition 3.7. *For every $1 \leq j \leq 2S - 2$,*

$$K_j^* < \infty \iff \frac{p_0 \left(\frac{1}{2}\right)^{\lfloor \frac{j}{2} \rfloor + 1} + \sum_{l=1}^{K_{j+1}^* - 1} p_l}{p_0 \left(\frac{1}{2}\right)^{\lfloor \frac{j}{2} \rfloor} + (1 - p_0)} < \frac{1-c}{1+c}. \quad (3.3)$$

Proof: The logic of the proof is similar to the proof of Proposition 3.6. See Appendix C for details. ■

A simple economic interpretation of the conditions obtained in Proposition 3.6 and 3.7 is as follows. At any stage s , if the other player will take at the next stage, the net gain to taking at s is $[1 + (s - 1)c] - sc = 1 - c$. On the other hand, if the other player passes at the next stage, the net gain to taking at stage $s + 2$ is $[1 + (s + 1)c] - sc = 1 + c$. Hence, the right-hand side is simply the ratio of payoffs to the current player depending on the opponent taking or passing at the next stage, assuming the current player will take in the subsequent stage. Thus, a player will take in the current stage if and only if the posterior probability the opponent will take in the next stage is less than this ratio.

The information contained in the history is that *if the game proceeds to later stages, the opponent is less likely to be a level-0 player*. If the game reaches stage j , the player would know the opponent has passed $\lfloor \frac{j}{2} \rfloor$ times, which would occur with probability (conditional on the opponent being level-0) $1/2^{\lfloor \frac{j}{2} \rfloor}$ which rapidly approaches 0.

3.5.2 DCH for the Centipede Game in Reduced Normal Form

In contrast, when the $2S$ -move centipede game is played in reduced normal form, both players simultaneously choose the node at which they will stop the game or

always pass. Therefore, $A_1 = A_2 = \{1, \dots, S + 1\}$ is the set of actions for each player. Action $s \leq S$ represents a plan to pass at the first $s - 1$ opportunities and take at the s -th opportunity. Strategy $S + 1$ is the plan to always pass. Player 1 and 2's strategies are denoted by a_1 and a_2 , respectively. If $a_1 \leq a_2$, then the payoffs are $(1 + (2a_1 - 2)c, (2a_1 - 2)c)$; if $a_1 > a_2$, then the payoffs are $((2a_2 - 1)c, 1 + (2a_2 - 1)c)$.

To characterize the DCH solution for the linear centipede game in reduced normal form, we let a_i^k denote level- k player i 's strategy. A level-0 player uniformly randomizes across all available strategies. With a minor abuse of notation, denote $a_i^0 = \frac{1}{S+1}$ for $i \in \{1, 2\}$. Lemma 3.2 establishes level-1 players' behavior and the monotonicity, similarly to Lemma 3.1.

Lemma 3.2. *In linear centipede games of reduced normal form, if $\frac{1}{3} < c < 1$, then*

1. $a_1^1 = S + 1$ and $a_2^1 = S$.
2. For every $k \geq 2$,
 - (i) $a_1^k \geq \min\{a_2^m : 1 \leq m \leq k - 1\}$;
 - (ii) $a_2^k \geq \min\{a_1^m : 1 \leq m \leq k - 1\} - 1$.
3. $a_i^{k+1} \leq a_i^k$ for all $k \geq 1$ and for all $i \in \{1, 2\}$.

Proof: See Appendix C. ■

Lemma 3.2 has essentially the same three parts as Lemma 3.1, but stated in terms of the stopping point strategies rather than behavioral strategies. Therefore, as in the extensive form, optimal strategies are given by **cutoffs**, defined analogously to Definition 3.1.

Definition 3.2. *For every stage s where $1 \leq j \leq 2S$, define the **cutoff** \tilde{K}_j^* to be the lowest level of player that would take no later than this stage. In other words,*

$$\tilde{K}_j^* = \begin{cases} \arg \min_k \left\{ a_1^k \leq \frac{j+1}{2} \right\}, & \text{if } j \text{ is odd} \\ \arg \min_k \left\{ a_2^k \leq \frac{j}{2} \right\}, & \text{if } j \text{ is even} \\ \infty, & \text{if } \nexists k \text{ s.t. } a_1^k \leq \frac{j+1}{2} \text{ or } a_2^k \leq \frac{j}{2}. \end{cases}$$

By Lemma 3.2, we know $a_2^1 = S$. Therefore, $\tilde{K}_{2S}^* = 1$. Proposition 3.8 and 3.9 are parallel to Proposition 3.6 and 3.7, providing necessary and sufficient conditions for the existence of some strategic players to take before a particular stage.

Proposition 3.8. $\tilde{K}_{2S-1}^* < \infty \iff p_0 < \frac{S+1}{(S+1)+(\frac{3c-1}{1-c})}$.

Proof: See Appendix C. ■

Proposition 3.9. For every $1 \leq j \leq 2S - 2$,

$$\tilde{K}_j^* < \infty \iff p_0 \left(\frac{S}{S+1} - \frac{2 \lfloor \frac{j}{2} \rfloor c}{(S+1)(1+c)} \right) + \sum_{k=1}^{\tilde{K}_{j+1}^*-1} p_k < \frac{1-c}{1+c}. \quad (3.4)$$

Proof: The logic of the proof is similar to the proof of Proposition 3.8. See Appendix C for details. ■

Propositions 3.6 and 3.8 identify conditions on p such that there is some level $k > 0$ of player 1 who would take at some stage when the centipede game is played in extensive form while every strategic level of player 1 would choose “always pass” when it is played in reduced normal form.

Corollary 3.1. If $\frac{S+1}{(S+1)+(\frac{3c-1}{1-c})} \leq p_0 < \frac{2^S}{2^S+(\frac{3c-1}{1-c})}$, then $K_{2S-1}^* < \infty$ and $\tilde{K}_{2S-1}^* = \infty$.

Proof: Since $2^S > S + 1$ for all $S \geq 2$, this is a direct consequence of Propositions 3.6 and 3.8. ■

An implication of propositions 3.7 and 3.9 is that if p_0 is small, then the difference in behavior under the two different representations of the game will also be small, since the left hand side of inequalities (3.3) and (3.4) both converge to $\sum_{k=1}^{\tilde{K}_{j+1}^*-1} p_k$. This result is intuitive. If there is no level-0 in the population, the difference in the behavior of level-0 will not trigger the chain reaction that affects higher-level players' behavior, as the behavior of level-1 players does not differ when the game is played in extensive form or reduced normal form.

However, regardless of how small p_0 is (as long as it is positive), DCH predicts the extensive form and the reduced form representations lead to systematically different behavioral predictions. These differences lead to the main result of this section, Theorem 3.1, which establishes that players are more likely to take at every stage when the game is played in extensive form.

Theorem 3.1 (Strategy-reduction effect). *For every stage $1 \leq j \leq 2S$,*

$$K_j^* \leq \tilde{K}_j^*.$$

Proof: See Appendix C. ■

This result provides a testable prediction that these centipede games will end earlier if played in the extensive form rather than in the reduced normal form. Moreover, this result is robust to the prior distributions of levels. Therefore, DCH predicts players will exhibit more sophisticated behavior in the extensive form since the information from the history is that the opponent is less likely to be a level-0 player.

3.5.3 Results for the Poisson-DCH Model

In previous applications of the CH model, it has been useful to assume the distribution of levels is given by a Poisson distribution (Camerer, Ho, and Chong, 2004). We obtain some additional results here for this one-parameter family of distributions that allow us to further pin down the differences between the centipede game when played in extensive form and when played in reduced normal form. The Poisson-DCH model assumes

$$p_k \equiv \frac{e^{-\lambda} \lambda^k}{k!}, \text{ for all } k = 0, 1, 2, \dots$$

where $\lambda > 0$ is the mean of the Poisson distribution.

Finally, we write the cutoffs as functions of λ . In the extensive form, the cutoff function for stage j is $K_j^*(\lambda)$. In the reduced normal form, the cutoff function for stage j is $\tilde{K}_j^*(\lambda)$.

As previously discussed, due to the realization-nonequivalence of level-0 players in different representations, beliefs of level $k \geq 2$ about the opponent being level-0 shrink much faster when the game is played in extensive form compared to when it is played in reduced normal form. To quantify the effect, Proposition 3.10 demonstrates the difference between two representations at stage $2S - 1$, where player 1 has the best information.

Proposition 3.10. *As the prior distribution follows Poisson(λ), then*

$$(i) \ K_{2S-1}^*(\lambda) < \infty \iff \lambda > \ln \left[1 + \left(\frac{1}{2} \right)^S \left(\frac{3c-1}{1-c} \right) \right];$$

$$(ii) \tilde{K}_{2S-1}^*(\lambda) < \infty \iff \lambda > \ln \left[1 + \left(\frac{1}{S+1} \right) \left(\frac{3c-1}{1-c} \right) \right].$$

Proof: The result is obtained by substituting $p_0 = e^\lambda$ in the formulas given by Propositions 3.6 and 3.8, and with some algebra. See Appendix C for details. ■

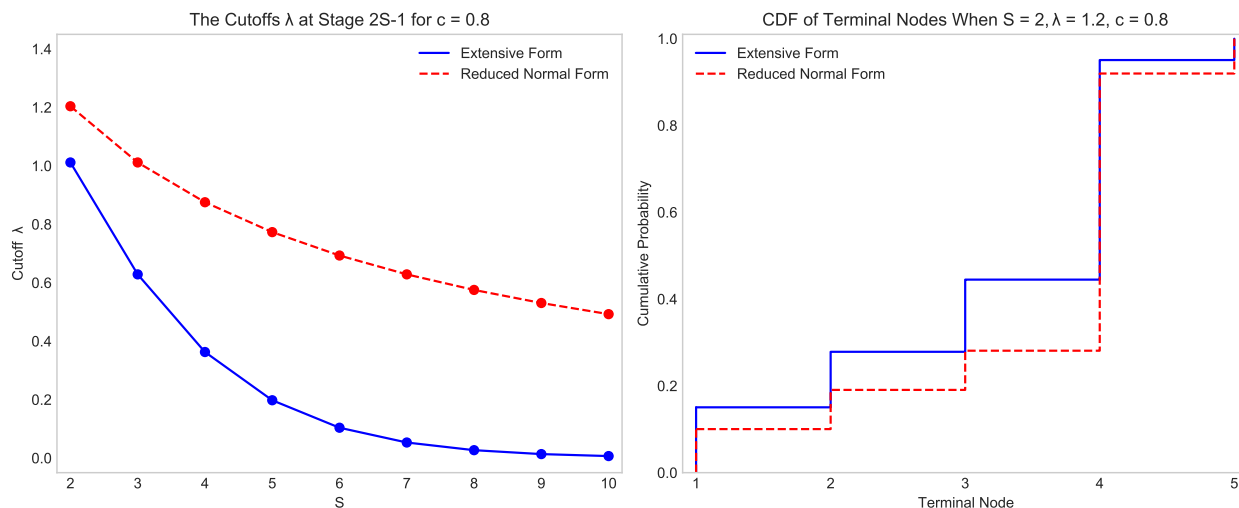


Figure 3.6: (Left) The minimum value of λ needed to support taking at stage $2S - 1$ when the centipede game is played in extensive form (solid) and when played in reduced normal form (dashed) for $c = 0.8$, with S on the horizontal axis and λ on the vertical axis. (Right) The CDFs of terminal nodes in four-node centipede games when the game is played in different representations predicted by DCH.

Proposition 3.10 provides a closed form solution for the minimum λ to support some level of player 1 to take at stage $2S - 1$ in both the extensive form and the reduced normal form. The left panel of Figure 3.6 plots the lowest λ . From the figure, we can notice that at stage $2S - 1$, the minimum value of λ to start unraveling is much smaller in the extensive form than in the reduced normal form. Moreover, the minimum λ converges to 0 much faster in the extensive form than in the reduced normal form as S gets higher, which is derived from the belief updating of DCH.

On the other hand, in the right panel of Figure 3.6, we focus on the four-move centipede game ($S = 2$) and plot the CDF of terminal nodes when the game is played in extensive form and reduced normal form predicted by DCH. First of all, we can observe the distribution of terminal nodes under the reduced normal form first order stochastically dominates the distribution under the extensive form. In fact, the FOSD relationship holds for any S, c, λ . This leads to a second interpretation

of Theorem 3.1; since the cutoffs when the game is played in extensive form are uniformly smaller than when played in reduced normal form, there are more levels of players that would take at every stage, thus generating the FOSD relationship.

When λ gets larger, the distribution of levels will shift to the right and players tend to be more sophisticated at the aggregate level. Proposition 3.11 shows that for sufficiently large λ , highly sophisticated players would take at every stage in both the extensive form and reduced normal form of the centipede game.

Proposition 3.11. *In both extensive form and reduced normal form linear centipede games, there exists sufficiently high λ such that unraveling occurs, i.e., for each S :*

- (i) $\exists \lambda^* < \infty$ such that $K_1^*(\lambda) < \infty$ for all $\lambda > \lambda^*$;
- (ii) $\exists \tilde{\lambda}^* < \infty$ such that $\tilde{K}_1^*(\lambda) < \infty$ for all $\lambda > \tilde{\lambda}^*$.

Proof: See Appendix C. ■

This result shows that DCH predicts unravelling occurs if λ is sufficiently high, in both representations. However, it leaves open questions about how this unravelling differs between the two representations. To this end, Proposition 3.12 provides some insight on this issue, in particular that the reduced normal form requires strictly more “density shift” (higher λ) in order to completely unravel for high-level players.

Proposition 3.12. *For any j where $1 \leq j \leq 2S - 1$, let λ_{2S-j}^{**} be the lowest λ such that $K_{2S-j}^*(\lambda) = j + 1$ for all $\lambda > \lambda_{2S-j}^{**}$, and let $\tilde{\lambda}_{2S-j}^{**}$ be the lowest λ such that $\tilde{K}_{2S-j}^*(\lambda) = j + 1$ for all $\lambda > \tilde{\lambda}_{2S-j}^{**}$. Then $\lambda_{2S-j}^{**} < \tilde{\lambda}_{2S-j}^{**}$ for all $1 \leq j \leq 2S - 1$.*

Proof: See Appendix C. ■

In other words, we can view the difference of density shifts between two representations (so that every level of players completely unravels) as a measure of the strategy-reduction effect. As shown in Proposition 3.12, we can always find a non-trivial set of λ such that players have already unravelled in the extensive form but not in the reduced normal form.

Finally, in the Poisson family, we can obtain an unambiguous comparative static result on the change of λ . Proposition 3.13 shows that when λ increases, the cutoff level of each stage is weakly decreasing. That is, when the average sophistication

of the players increases, play is closer to the fully rational model because strategic players believe the opponent is less likely to be level-0.

Proposition 3.13. *For every $1 \leq j \leq 2S$, $K_j^*(\lambda)$ and $\tilde{K}_j^*(\lambda)$ are weakly decreasing in $\lambda > 0$.*

Proof: See Appendix C. ■

3.5.4 Non-Linear Centipede Games

The results of this section about the exact characterization of behavior in extensive form and reduced normal form centipede games only consider games with a linearly increasing pie. A natural robustness question is whether the qualitative findings apply more generally to other families of centipede games. The key assumption in our analysis is that the increment of pie is not too fast or too slow. If the increment is too fast (i.e., $c > 1$), then it is optimal to pass everywhere. On the other hand, if the increment is too slow (i.e., $c < \frac{1}{3}$), even the lowest level of players would take at the first stage. In all cases within this range, the DCH strategy-reduction effect occurs, resulting in earlier taking if the centipede game is played in extensive form. This would seem to be a general property of increasing-pie centipede games. That is, unless the pie sizes grow so fast that all positive levels of players will always pass, or so slowly that positive levels will always take, then the realization-nonequivalence of level-0 will lead to different behavior of higher-level players when the game is played in different representations. Moreover, since beliefs of level $k \geq 2$ about the opponent being level-0 shrink much faster when the game is played in extensive form, DCH predicts that playing the game sequentially in extensive form will result in earlier taking compared to playing the game simultaneously in reduced normal form, under mild conditions.

For example, the analysis can be extended to the class of centipede games with an exponentially increasing pie, as studied in the McKelvey and Palfrey (1992) experiment. Similar to the previous analysis, two players alternate over a sequence of moves in an exponential centipede game with $2S$ nodes. At each node, if a player passes, both the large and small (positive) payoffs would be *multiplied* by $c > 1$. In addition, the ratio between the large and the small payoff is equal to $\pi > 1$ and does not change as the game progresses. Therefore, an exponential centipede game is parameterized by (S, π, c) : if the game is terminated by a player at stage $j \leq 2S$, the payoffs are $(c^{j-1}\pi, c^{j-1})$ if j is odd and $(c^{j-1}, c^{j-1}\pi)$ if j is even. If no one ever takes, then the payoffs will be $(c^{2S}\pi, c^{2S})$. In this class of centipede games, the

multiplier c governs the growth rate of pie, and the logic of the proofs of propositions for the linear games is similar for exponential games as long as

$$\frac{-1 + \sqrt{1 + 8\pi^2}}{2\pi} < c < \pi,$$

which rules out trivial cases, in the same way as the assumption of $\frac{1}{3} < c < 1$ rules out trivial cases in linear centipede games.

All of the qualitative results for linearly increasing centipede games also hold for exponential centipede games, with the only difference being the analytical expression of the cutoffs. In particular, Theorem 3.1, the strategy-reduction effect, continues to hold.

3.6 Experimental Evidence

Since DCH is a solution concept developed for games in extensive form, the solutions could appear dramatically different for different games that share the same reduced normal form. Specifically, DCH predicts that players would behave differently if centipede games of a certain class are implemented under the direct-response method (extensive form) and the strategy method (reduced normal form). To empirically test the strategy-reduction effect predicted by DCH, we revisit a recent experiment conducted by García-Pola, Iriberri, and Kovářík (2020a) which compared the behavior in four centipede games under the direct response method and the strategy method.

The game trees of the four centipede games (CG 1 to CG 4) studied in the experiment are plotted in Figure 3.7. CG 1 is a centipede game with an exponentially-increasing pie while CG 2 is a centipede game with a constant-sum pie. By contrast, the change of the pie size in CG 3 and CG 4 is not monotonic and player 1's payoff is always greater than player 2's payoff. Among these four centipede games, DCH predicts that in CG 1, CG 2 and CG 4, the distribution of terminal nodes under the strategy method will *first order stochastically dominate* the distribution of terminal nodes under the direct response method. Yet DCH predicts the FOSD relationship does not necessarily hold in CG 3.²⁰

²⁰Since CG 1 is an exponential centipede game, as discussed in section 3.5.4, we can use the same argument as Theorem 3.1 to show the FOSD relationship. By a similar argument, we can show that DCH also predicts the FOSD relationship in CG 2 and CG 4. We can also prove that the FOSD relationship in CG 3 does not hold for all distributions of levels, unlike CG 1, CG 2, and CG 4, where DCH makes unambiguous predictions. In particular, the FOSD relationship is *violated* for empirically plausible prior distributions of levels, for example if levels are distributed Poisson with mean equal to 1.

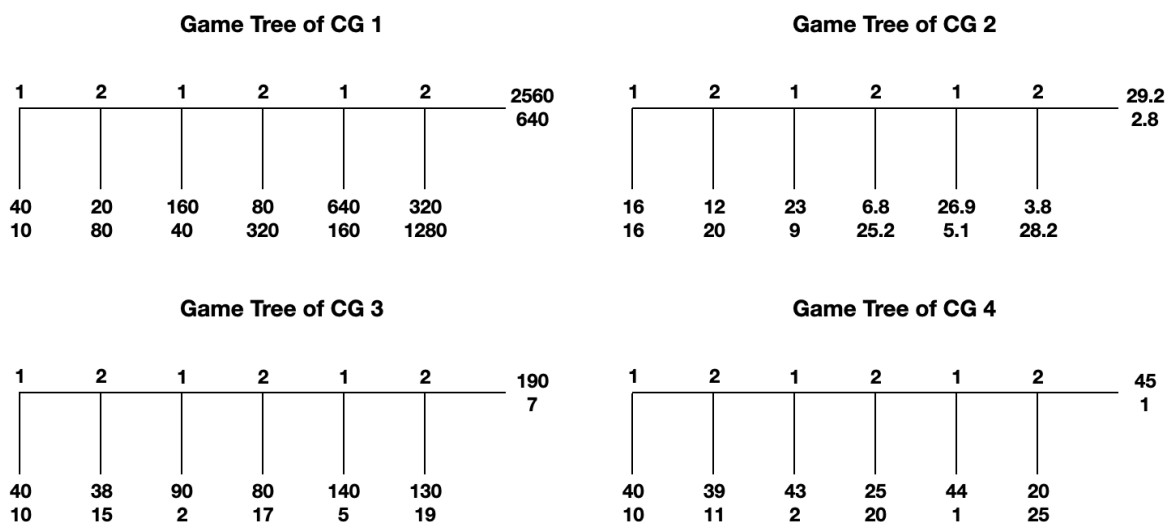


Figure 3.7: The game trees of CG 1 to CG 4 studied in García-Pola, Iriberry, and Kovářík (2020a).

DCH Prediction: *In CG 1, CG 2 and CG 4, the distribution of terminal nodes under the strategy method will first order stochastically dominate the distribution of terminal nodes under the direct response method. However, the FOSD relationship can be violated in CG 3, depending on the prior distribution of levels.*

This experiment consists of two treatments—the direct response method (the hot treatment) and the strategy method (the cold treatment)—with between-subject design. That is, each subject only participates in one of the two treatments. There are 151 subjects in the cold treatment, 76 in the role of first-mover and 75 in the role of second-mover. Each subject chose a stopping point for 16 different centipede games (including the four games in the figure) without feedback between games, and was subsequently matched with a random player of the other role to determine payment. There were 352 subjects in the hot treatment, and each subject only played only one of the four centipede games in the figure.²¹

The two treatments share the following features. The subjects are given identical instructions in both treatments, except for the specific way subject decisions are

²¹In the hot treatment, there were 90 subjects (45 in each role) participating in each of CG 1, CG 2, and CG 4, while there were 82 subjects (41 in each role) participating in CG 3. Every subject in a session played 10 repetitions of the same game in the experiment, with feedback, using a matching protocol that was designed to minimize reputation effects. To avoid the analysis from being confounded with learning effects from repetition, our analysis only uses data from the first match of each game.

elicited. In particular, the “frame” of the game is explained and presented on subject computer interfaces in game-tree form to reflect the timing of the centipede game in both treatments. In the cold (strategy-method) treatment, each subject was instructed to click on the first node at which they wanted to stop, if the game got that far. Thus, the task for a subject in the cold treatment was to choose one of four pure strategies of the reduced normal form of the game: take at the first opportunity (T); pass at the first opportunity and take at the second opportunity (PT); pass at the first two opportunities and take at the third opportunity (PPT); or never take (PPP). The instructions provide subjects with explanations of the decision screens, the matching protocols, the payment method, etc. Thus, there is no possibility of a “framing” effect as could happen, for example if the cold treatment were presented as a 4×4 matrix game.²²

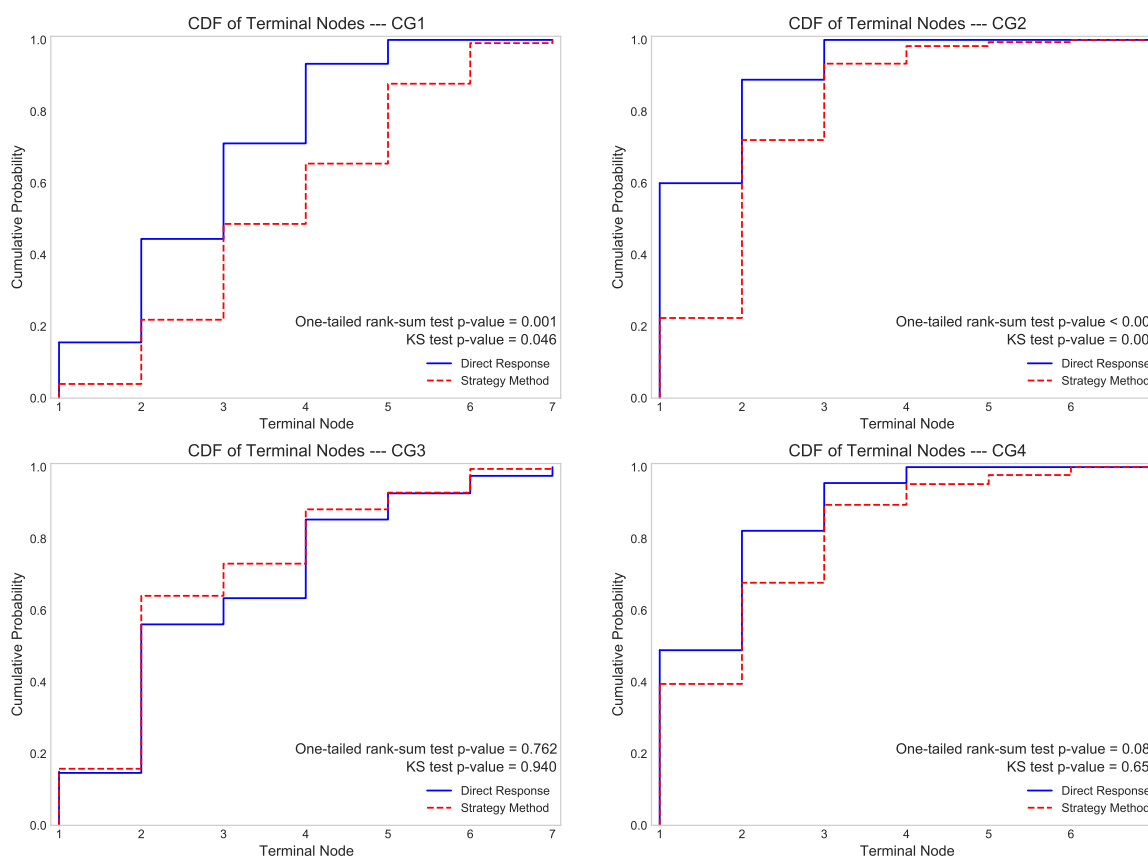


Figure 3.8: The empirical CDF of CG 1 to CG 4 in García-Pola, Iriberry, and Kovářík (2020a) under the direct response (solid) and the strategy method (dashed).

²²See García-Pola, Iriberry, and Kovářík (2020a) and García-Pola, Iriberry, and Kovářík (2020b) for copies of the instructions and exact details of the experimental procedures.

Figure 3.8 plots the empirical CDFs of the terminal nodes under two different methods in all four games. In the hot treatment, since players are randomly paired at the beginning of the game, we plot the observed distribution of terminal nodes. On the other hand, in the cold (reduced normal form) treatment, players are not paired into groups before each game, and hence a distribution of terminal nodes is not directly observed. However, the empirical conditional take probabilities of each stage is directly observed, and from this one can easily compute the implied distribution of terminal nodes.²³

Focusing on CG 1, CG 2, and CG 4, we can see that the distribution of terminal nodes under the strategy method indeed first order stochastically dominates the distribution under the direct response method although the strategy-reduction effect is weaker in CG 4. The effect is only marginally significant for CG 4 using a one-tailed Mann-Whitney rank-sum test,²⁴ and is insignificant using the low-powered Kolmogorov-Smirnov test.²⁵ In stark contrast, the *opposite* FOSD relationship is observed in CG 3—earlier *taking* in the cold treatment, and it is not significant (one-tailed Mann-Whitney rank-sum test $p = 0.762$; Kolmogorov-Smirnov test $p = 0.940$). This experimental evidence supports DCH at the aggregate level.

3.7 Discussion

In this section, we briefly discuss several additional features and potential applications of DCH. Section 3.7.1 illustrates how reputation effects can arise with DCH, and in fact are a built-in feature of the solution concept. This follows from the fact that strategic players are not myopic, but are *forward looking* and take into account how their current actions will affect other players' beliefs and actions. Thus, in

²³For instance, the probability that the game ends at the first stage is equal to the fraction of subjects in the first-mover role who chose the “T” strategy. The probability that the game ends at the second stage is equal to one minus the fraction of subjects in the first-mover role who chose the “T” strategy times the fraction of subjects in the second-mover role who chose the “T” strategy. The probabilities that the game ends at later stages are computed in the similar way.

²⁴A one-tailed Mann-Whitney rank-sum test is performed because DCH makes a clear directional prediction between the two treatments.

²⁵One-tailed Mann-Whitney rank test: CG 1 $p = 0.001$, CG 2 $p = 0.003$, CG 4 $p = 0.084$; Kolmogorov-Smirnov test: CG 1 $p = 0.046$, CG 2 $p = 0.003$, CG 4 $p = 0.653$. To compute the p-values, we follow the approach of García-Pola, Iriberry, and Kovářík (2020a) and generate 100,000 random sub-samples of subjects from the cold treatment to match the number of subjects in each role in each game in the corresponding hot treatment. Next, in each random sample, we randomly match subjects into pairs to obtain the distribution of terminal nodes. This process yields 100,000 simulated CDFs of the terminal nodes from the cold treatment. Then, we perform the KS test and one-tailed Mann-Whitney rank-sum test on each simulated CDF of the terminal nodes against the CDF from the hot treatment and report the median p-value. The KS p-values computed using this process are almost identical to the KS p-values reported by García-Pola, Iriberry, and Kovářík (2020a).

DCH higher-level players can mimic lower-level types in order to affect lower levels of other players' beliefs and hence their future play. In Sections 3.7.2 and 3.7.3, we highlight some complications of the DCH belief system that arise in games of imperfect or incomplete information. Finally, Section 3.7.4 discusses general issues related to the equivalence or non-equivalence of DCH when analyzed in the *non-reduced* normal form.

3.7.1 Reputation Formation

In addition to the strategy-reduction effect, another interesting phenomenon that can arise in DCH is *reputation building* by higher-level players. Since in DCH, players will update their beliefs about others' levels as the history unfolds, it is possible for higher-level players to *mimic* some lower-level players' strategies in order to maintain the reputation of being some lower levels and benefit from this reputation.

It is worth noticing that the reputation concerns predicted by DCH do not only appear in games of incomplete information, but also in games of *complete information*. In other words, the reputation formation in DCH is driven by manipulating the beliefs about levels of sophistication rather than the beliefs about exogenous types. We illustrate this point with the "chain-store game" introduced by Selten (1978).

Illustrative Example

In this game, there are $N + 1$ players: one chain-store (CS) and N competitors, numbered $1, \dots, N$. In each period, one of the potential competitors decides whether to compete with CS or not ("In" or "Out"); in period i , it is competitor i 's turn to decide. If competitor i chooses "In," then CS decides either to fight ("F") or cooperate ("C"). CS responds to competitor i before competitor $(i + 1)$'s turn. Hence, in each period, there are three possible outcomes $\{Out, (In, F), (In, C)\}$. The game tree of each period i is plotted in Figure 3.9. In addition, at every point of the game, all players know all actions taken previously, which makes this game an extensive game of perfect information. Finally, the payoff of CS in this game is the sum of its payoffs in N periods.

There are multiple Nash equilibria where the outcome in any period is either *Out* or (In, C) . Specifically, in any equilibrium where competitor i chooses "Out," CS's strategy is to fight if the competitor chooses "In." However, these equilibria are *imperfect*. The unique subgame perfect equilibrium is that all competitors will choose "In" and CS will always choose "C," which fails to capture the reality that




Figure 3.9: Game tree of period i in the chain-store game. The first number of each pair is competitor i 's payoff and the second number is CS's payoff.

CS may choose “F” to deter entrance.

Unlike Kreps and R. Wilson (1982)'s approach to rationalize the deterrence by introducing payoff-relevant private types, DCH predicts that higher-level CS might purposely choose to fight in early periods to make potential competitors think they are facing a level-0 CS. We demonstrate this by considering the following distribution of levels

$$(p_0, p_1, p_2, p_3) = (0.10, 0.15, 0.60, 0.15).$$

Under this distribution, it suffices to characterize the DCH solution by analyzing level-1 to level-3 players' behavior.

Level-1 players believe all other players are level-0 and best respond to this belief. Therefore, level-1 competitors will choose “Out” and level-1 CS will choose “C.” Next, from level-2 CS's perspective, all competitors are either level-0 or level-1 and the only strategic level of competitors will choose “Out,” the most preferred outcome of CS. Therefore, level-2 CS does not have incentive to “F,” and will behave like level-1 to always choose “C.”

On the other hand, level-2 competitors will update their beliefs about CS's level based on the history. If level-2 competitors have observed that CS has cooperated for T times and never fought, the belief about CS being level-0 is

$$\nu \equiv \frac{p_0 \left(\frac{1}{2}\right)^T}{p_0 \left(\frac{1}{2}\right)^T + p_1} = \frac{2 \left(\frac{1}{2}\right)^T}{2 \left(\frac{1}{2}\right)^T + 3}.$$

The expected payoff of choosing “In” is $2\nu + 4(1 - \nu)$ and therefore, it's optimal for a level-2 competitor to choose “In” if and only if $2\nu + 4(1 - \nu) < 3 \iff \nu < \frac{1}{2}$.

In this case, if CS has never fought, level-2 competitors will always choose “In.” However, if CS has ever fought,²⁶ level-2 competitors will believe CS is level-0 and always choose “Out.” Besides, because level-2 CS behaves the same as level-1, level-3 competitors will therefore behave the same as level-2—they will choose “In” if CS has never chosen “F” but choose “Out” if CS has ever chosen “F.”

Finally, in the early periods,²⁷ if the relative proportion of level-2 players is sufficiently high, it would be profitable for level-3 CS to purposely choose “F” to make all competitors believe the chain store is level-0 and choose “Out” in later periods. Specifically, level-3 CS would choose to fight if

$$4 + \left[\frac{p_0}{\sum_{k=0}^2 p_k} \times 7 + \frac{p_1}{\sum_{k=0}^2 p_k} \times 10 + \frac{p_2}{\sum_{k=0}^2 p_k} \times 4 \right] < 0 + \left[\frac{p_0}{\sum_{k=0}^2 p_k} \times 7 + \frac{p_1}{\sum_{k=0}^2 p_k} \times 10 + \frac{p_2}{\sum_{k=0}^2 p_k} \times 10 \right] \iff \frac{p_2}{p_0 + p_1 + p_2} > \frac{2}{3}. \quad \square$$

3.7.2 Correlated Beliefs in Games of Imperfect Information

There is a wide range of applications of games in extensive form in economics and political science where players have private information, either due to privately known preferences and beliefs about other players, or from imperfect observability of the histories of play in the game. These applications would include many workhorse models, such as signaling, information transmission, information design, social learning, entry deterrence, reputation building, crisis bargaining, and so forth. Hence the natural next step is to investigate more deeply our approach to extensive games of imperfect information. In such environments, one complication is that players not only learn about the opponents’ levels of sophistication but also about more basic elements of the game structure, such as the opponents’ private information, payoff types, and prior moves.

One observation is that allowing for imperfect information in the DCH approach does not introduce any problems of off-path beliefs. The reason is that at every information set of the game, all levels of all players have posterior beliefs over the opponents’ levels that include a positive probability they are facing level-0 players. Hence, there is no issue of specifying off-path beliefs in an ad hoc fashion and therefore we avoid the complications of belief-based refinements.

²⁶By Proposition 3.2, we know if CS has fought once, then all competitors will always believe CS is level-0, regardless of how many time CS has cooperated.

²⁷In the last period, there is no reputation concern and therefore, level-3 CS will choose “C” if he has a chance to move.

In games of imperfect information, the DCH belief system is a level-dependent profile of posterior beliefs that assigns to every information set a *joint distribution* of other players' levels and the histories in the information set. When information sets are not singleton, at some information set, the marginal beliefs about other players' levels can be correlated across players. In other words, Proposition 3.1 might break down in games of imperfect information. We illustrate this by using the following three-person game where each player moves once. The game tree is shown in Figure 3.10.

Illustrative Example




Figure 3.10: Game Tree of Example 3.7.2. Dashed lines are the paths selected by level-1 players.

Player 1 chooses first whether to go left or right. After that, player 2 chooses to go left or right. If player 1 and 2 make the same decision, the game ends. Otherwise, player 3 makes the final decision. However, at that stage, player 3 only knows that one of the previous players chose l and the other chose r , but does not know which one chose l .

Level-1 players believe all other players are level-0. As we compute the expected payoff of each action, level-1 player 1 will choose r at the initial node. Level-1 player 2 will choose r at subgame $h = l$ and $h = r$. At player 3's information set, since level-1 player 3 thinks both players are level-0, he would believe both histories are equally likely, and choose l .

Conditional on the game reaching player 3's information set I_3 , level-2 player 3's DCH belief $\nu_3^2(\tau_{-3}, h|I_3)$ is a joint distribution of the histories $h \in \{lr, rl\}$ and the profile of levels of other players $\tau_{-3} = (\tau_1, \tau_2) \in \{(0, 0), (0, 1), (1, 0), (1, 1)\}$. Let $(p_k)_{k \in \mathbb{N}_0}$ be the true distribution of levels of all players, and level-2 player 3's DCH belief can be summarized in Table 3.2.

Table 3.2: Level 2 player 3's DCH belief at information set I_3 .

$\nu_3^2(\tau_{-3}, h I_3)$	(0, 0)	(0, 1)	(1, 0)	(1, 1)
$h = lr$	$\frac{0.25p_0}{0.5p_0+p_1}$	$\frac{0.5p_1}{0.5p_0+p_1}$	0	0
$h = rl$	$\frac{0.25p_0}{0.5p_0+p_1}$	0	$\frac{0.5p_1}{0.5p_0+p_1}$	0

We can first observe that given level-0 and level-1 players' strategies, at player 3's information set, level-2 player 3 would think player 1 and player 2 cannot both be level-1 players; otherwise, the game will not reach this information set. In other words, level-2 player 3's marginal belief $\nu_3^2(\tau_{-3} = (1, 1)|I_3) = 0$. Nonetheless, the marginal belief of *each* player is $\nu_3^2(\tau_1 = 1|I_3) = \nu_3^2(\tau_2 = 1|I_3) = \frac{0.5p_1}{0.5p_0+p_1}$, suggesting that the marginal beliefs about levels are correlated across players as

$$\nu_3^2(\tau_1 = 1|I_3) \times \nu_3^2(\tau_2 = 1|I_3) \neq 0 = \nu_3^2(\tau_{-3} = (1, 1)|I_3). \quad \square$$

Remark 3.3. The underlying reason why the DCH beliefs are correlated across players in this example is that player 3 is unable to distinguish the actions of player 1 and 2. Therefore, this game does not belong to the class of games with *observable deviators* (see Fudenberg and Levine (1993), Battigalli (1996), and Battigalli (1997) for the definition). Battigalli (2023) shows that the DCH beliefs remain product measures in all games with observable deviators.

3.7.3 DCH in Multi-Stage Games with Observed Actions

In line with the observation of Battigalli (2023), Lin (2023) demonstrates that the DCH beliefs indeed conform to product measures across players in the framework of *multi-stage games with observed actions*, as introduced by Fudenberg and Tirole (1991). This framework captures situations where every player observes the actions of every other player—the only uncertainty is about other players' payoff-relevant private information which is determined by an initial chance move.

In a multi-stage game with observed action, each player $i \in N$ has a payoff-relevant *type* θ_i drawn from a finite type set Θ_i according to the distribution F_i .²⁸ After the types are assigned, each player will learn about their own type but not others' types. The game is played in periods $t = 1, \dots, T$. In each period, every players will simultaneously choose an action and the action profile will be revealed to all players at the end of the period. Therefore, in a multi-stage game with observed actions, each player i 's information sets can be specified as (h, θ_i) where h is a non-terminal public history.²⁹

For any $k \in \mathbb{N}$ and $i \in N$, level- k player i 's DCH belief at information set (h, θ_i) is a joint distribution of other players' types and levels, denoted as $\nu_i^k(\tau_{-i}, \theta_{-i}|h, \theta_i)$. Proposition 1 of Lin (2023) shows that $\nu_i^k(\tau_{-i}, \theta_{-i}|h, \theta_i)$ is a product measure across players. That is,

$$\nu_i^k(\tau_{-i}, \theta_{-i}|h, \theta_i) = \prod_{j \neq i} \nu_{ij}^k(\tau_j, \theta_j|h, \theta_i).$$

From the comparison between Example 3.7.2 and multi-stage games with observed actions, we can find that the observability of actions plays a crucial role in the independence property of the DCH beliefs.

3.7.4 (Non-)Equivalence on the Normal Form

The analysis in Section 3.4.3 and Section 3.5 shows that the DCH solution is *not reduced-normal-form invariant*—the DCH solution can look dramatically different for two extensive games that share the same reduced normal form. To this end, one may naturally wonder whether DCH is *normal-form invariant*.³⁰

The intuition behind such invariance, which is explained in more detail in Battigalli (2023), is as follows. First, observe that a level-0 player's *behavioral strategy* (uniform randomization at every subgame) in extensive form and a level-0 player's *mixed strategy* (uniform randomization over all contingent strategies) in normal form are *realization-equivalent*.³¹ Second, either in extensive form or in normal

²⁸Without loss of generality, we assume that $F_i(\theta_i) > 0$ for all $\theta_i \in \Theta_i$ and F_i is independent of F_j for any $i, j \in N$ and $i \neq j$.

²⁹See Lin (2023) for a detailed description of the framework of multi-stage games with observed actions and the DCH solution for this framework.

³⁰We are grateful to Pierpaolo Battigali for observing the normal-form invariance of DCH. The formal analysis for games of perfect information can be found in Battigalli (2023).

³¹Two (mixed or behavioral) strategies are realization-equivalent if for every collection of pure strategies of the other players the two strategies induce the same distribution of outcomes (see Osborne and Rubinstein (1994) Chapter 11).

form, all strategic players (with level $k \geq 1$) *best respond* to totally mixed strategies due to the ever-presence of level-0 players. Since expected payoff maximization is dynamically consistent, the *ex ante* best response (the optimal contingent strategy in normal form) must be realization-equivalent to the *sequential* best response (the optimal behavioral strategy in extensive form). Hence, the DCH solution will be realization-equivalent between two extensive games that share the same normal form.

The equivalence of DCH solution in extensive form and non-reduced normal form relies heavily on two assumptions: (1) *uniform randomization* by level-0 players and (2) the dynamic consistency of best response.

In standard level- k and CH models, level-0 is assumed to uniformly randomize across all available actions³², and this approach has carried over to most applications. The specification of uniform randomization by level-0 players has several advantages: it is well-defined (and applied equally and in the same way) for all games; it is nondegenerate, so all paths of play can be rationalized by all strategic players; and it is simple and parsimonious. Some non-uniform specifications of level-0 behavior have been tailored to specific games of interest in particular applications. For example, alternative approaches include modeling level-0 players as choosing (or avoiding) a *salient* action (e.g. Crawford and Iriberri 2007a), an *instinctive* action (e.g. Rubinstein 2007) or a *minimum-payoff averse* action (e.g. Chong, Ho, and Camerer 2016) from the action set. Regardless of the specification of level-0 behavior, if level-0 players behave differently in different games sharing the same reduced normal form, the DCH strategy-reduction effect would still occur.

A second approach that is often used relaxes the perfect best response assumption of strategic types. In this alternative approach, strategic levels of players are typically assumed to make *better responses*, whereby players choose actions at each information set stochastically (with full support), and the choice probabilities are increasing in the continuation values, usually specified by a quantal response function such as the logit choice rule (e.g. Camerer, Nunnari, and Palfrey, 2016, Stahl and P. W. Wilson 1995). In this case, for any strategic level of player, the quantal response behavioral strategy in extensive form will generally not be realization-equivalent to the quantal response mixed strategy in non-reduced normal form. This relaxation of perfect best responses could be useful for estimating DCH in experimental data

³²Uniform randomization is specifically assumed in the original formulation of CH in Camerer, Ho, and Chong (2004).

sets.³³ It results in a smoother updating process, and implies full support of beliefs of about other players' (lower) levels at every information set of the game.

3.8 Conclusions

We conclude by emphasizing the key motivation for this paper: to provide a theoretical framework that characterizes hierarchical reasoning in sequential games. As documented in the literature, sequential equilibrium based on fully rational backward induction is not only mathematically fragile but also empirically implausible to hold. To narrow the gap between the theory and empirical patterns in sequential games, it is natural to extend the level- k approach to such games, as it has already demonstrated considerable success in narrowing the gap for games played simultaneously. However, the conundrum for directly applying the standard level- k approach is that players may observe actions that are incompatible with their beliefs, which leads to the widely known problem of specifying off-path beliefs. The DCH solution avoids this issue with a simple structure that allows players with heterogeneous levels of sophistication to update their beliefs everywhere as history unfolds, using Bayes' rule.

We characterize properties of the belief-updating process and explore how it can affect players' strategic behavior. The key of our framework is that the history of play contains substantial information about other players' levels of sophistication, and therefore as play unfolds, players learn about their opponents' strategic sophistication and update their beliefs about the continuation play in the game accordingly. In this way our DCH solution departs from the standard level- k approach and generates new insights, including experimentally testable implications.

We obtain two main results that apply generally to all finite games in extensive form. Proposition 3.1 establishes that a player's updating process is independent across the other players. That is, for every player and every non-terminal history, the joint distribution of the beliefs of the levels of the other players is the product of the personal posterior distribution of the levels of each of those other players. In games of imperfect information, the information sets are non-singleton and the beliefs could be correlated across the histories at some information set.

In addition, Proposition 3.2 establishes that the updating process filters out possible level types of opponents as the game proceeds, and it is irreversible. That is, over the course of play, it is possible that a player eliminates some levels of another player

³³See Lin (2023) for an example of estimating DCH with quantal responses.

from the support of his beliefs, and as the game continues, these levels can never be added back to the support. Hence, in addition to updating posterior beliefs over the support of level types, the support also shrinks over time. However, the level-0 players always remain in the support of beliefs, and hence every player believes every future information set can be reached with positive probability.

The second half of the paper provides a rigorous analysis of a class of increasing-pie centipede games and generates testable predictions about how play depends on whether the game is played in its original extensive form or in its reduced normal form. One direct implication is the strategy-reduction effect given by Theorem 3.1. The theorem states that playing a centipede game in its extensive form representation, i.e., as a sequential move game, would lead to more taking than the reduced normal form representation, where the two players simultaneously announce the stage at which they will take.

This result provides a prediction that may be useful for experimental testing, since the claim is independent of the length of the centipede and the increment of the pie. Moreover, the statement is true for any prior belief about the strategic levels. García-Pola, Iriberri, and Kovářík (2020a) recently reported the results of an experiment that is ideally suited to test the strategy-reduction effect implied by DCH. That experiment explored whether there were differences in behavior in four different centipede games, depending on whether they were played sequentially or (simultaneously) in the reduced normal form. DCH predicts a strategy-reduction effect with earlier taking in the sequential treatment, in three of the four games, and this is exactly what they find, and the effect is statistically significant in two of the three games. DCH does not predict this strategy-reduction effect in the fourth game, which is also what their experiment finds. This provides empirical support for the strategy-reduction effect we identify, and suggests additional experimental studies would be valuable to establish robustness of these effects, and to see if the findings extend the linear centipede games that we focused on in Section 3.5.

Another direction worth pursuing would be to incorporate some salient features of alternative behavioral models of learning in extensive games into our approach. In the approach taken here, the learning process is “extreme” in the sense that players will completely rule out some levels from their beliefs whenever they observe incompatible actions. For example, players will believe the opponent is level-0 with certainty if a strictly dominated action is taken. Yet it is possible that the player is strategic and the action is taken by mistake. In this sense, one could incorporate

some elements of the extensive form QRE, where players choose actions at each information set stochastically, and the choice probabilities are increasing in the continuation values. In fact, this approach has been used with some success in simultaneous move games (Crawford and Iriberry, 2007a). As shown in Proposition 3.2, in the current version of DCH, there is no way to *expand* the support of a player's belief about the other players' types. However, if players choose stochastically, then no level type is ever ruled out from the support, which smoothes out the updating process. Because players' beliefs maintain full support on lower types throughout the game, a natural conjecture is that arbitrarily high-level players will approach backward induction when the error is sufficiently small.

As a final remark, while the main point of this paper is to develop a general theoretical foundation for applying CH to games in extensive form, the ultimate hope is that this framework can be usefully applied to gain insight into specific economic models. There are a number of possible such applications one might imagine, where some or all agents in the model have opportunities to learn about the strategic sophistication of the other agents in ways that could significantly affect their choices in the game. One such application is reputation building, which we briefly examined in section 3.7.1 and deserves more extensive study. As another possible application, Chamley and Gale (1994) analyze a dynamic investment game with social learning, where investments are valuable only if enough other agents are able to invest, and learning occurs as investment decisions are observed over time. The DCH solution, which combines learning and updating, but without common knowledge of rationality or fully rational expectations, might be a useful alternative approach to this problem. Models of sequential voting on agendas (McKelvey and Niemi, 1978; Banks, 1985), limit pricing and entry deterrence (Selten, 1978; Milgrom and Roberts, 1982), and dynamic public good provision (Marx and Matthews, 2000; Duffy, Ochs, and Vesterlund, 2007; Choi, Gale, and Kariv, 2008) are some additional areas of applied interest where the DCH approach could be useful.

References

- Alaoui, Larbi, Katharina A Janezic, and Antonio Penta (2020). "Reasoning about others' reasoning". In: *Journal of Economic Theory* 189, p. 105091.
- Alaoui, Larbi and Antonio Penta (2016). "Endogenous depth of reasoning". In: *Review of Economic Studies* 83.4, pp. 1297–1333.
- (2022). "Cost-benefit analysis in reasoning". In: *Journal of Political Economy* 130.4, pp. 881–925.

- Aumann, Robert (1988). “Preliminary notes on integrating irrationality into game theory”. In: *International Conference on Economic Theories of Politics, Haifa*.
- (1992). “Irrationality in game theory”. In: *Economic Analysis of Markets and Games*, pp. 214–227.
- Banks, Jeffrey S (1985). “Sophisticated voting outcomes and agenda control”. In: *Social Choice and Welfare* 1.4, pp. 295–306.
- Battigalli, Pierpaolo (1996). “Strategic independence and perfect Bayesian equilibria”. In: *Journal of Economic Theory* 70.1, pp. 201–234.
- (1997). “Games with observable deviators”. In: *Decisions, Games and Markets*, pp. 57–96.
- (2023). “A note on reduced strategies and cognitive hierarchies in the extensive and normal form”. In: *IGIER Working Paper n.706*.
- Battigalli, Pierpaolo, Paolo Leonetti, and Fabio Maccheroni (2020). “Behavioral equivalence of extensive game structures”. In: *Games and Economic Behavior* 121, pp. 533–547.
- Binmore, Ken (1987). “Modeling rational players: Part I”. In: *Economics & Philosophy* 3.2, pp. 179–214.
- (1988). “Modeling rational players: Part II”. In: *Economics & Philosophy* 4.1, pp. 9–55.
- Bohren, J Aislinn (2016). “Informational herding with model misspecification”. In: *Journal of Economic Theory* 163, pp. 222–247.
- Bornstein, Gary, Tamar Kugler, and Anthony Ziegelmeyer (2004). “Individual and group decisions in the centipede game: Are groups more “rational” players?” In: *Journal of Experimental Social Psychology* 40.5, pp. 599–605.
- Brandts, Jordi and Gary Charness (2011). “The strategy versus the direct-response method: A first survey of experimental comparisons”. In: *Experimental Economics* 14.3, pp. 375–398.
- Cai, Hongbin and Joseph Tao-Yi Wang (2006). “Overcommunication in strategic information transmission games”. In: *Games and Economic Behavior* 56.1, pp. 7–36.
- Camerer, Colin F (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.
- Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong (2004). “A cognitive hierarchy model of games”. In: *Quarterly Journal of Economics* 119.3, pp. 861–898.
- Camerer, Colin F, Salvatore Nunnari, and Thomas R Palfrey (2016). “Quantal response and nonequilibrium beliefs explain overbidding in maximum-value auctions”. In: *Games and Economic Behavior* 98, pp. 243–263.

- Chamley, Christophe and Douglas Gale (1994). “Information revelation and strategic delay in a model of investment”. In: *Econometrica*, pp. 1065–1085.
- Choi, Syngjoo, Douglas Gale, and Shachar Kariv (2008). “Sequential equilibrium in monotone games: A theory-based analysis of experimental data”. In: *Journal of Economic Theory* 143.1, pp. 302–330.
- Chong, Juin-Kuan, Teck-Hua Ho, and Colin F Camerer (2016). “A generalized cognitive hierarchy model of games”. In: *Games and Economic Behavior* 99, pp. 257–274.
- Costa-Gomes, Miguel and Vincent P Crawford (2006). “Cognition and behavior in two-person guessing games: An experimental study”. In: *American Economic Review* 96.5, pp. 1737–1768.
- Costa-Gomes, Miguel, Vincent P Crawford, and Bruno Broseta (2001). “Cognition and behavior in normal-form games: An experimental study”. In: *Econometrica* 69.5, pp. 1193–1235.
- Crawford, Vincent P and Nagore Iriberri (2007a). “Fatal attraction: Saliency, naivete, and sophistication in experimental “hide-and-seek” games”. In: *American Economic Review* 97.5, pp. 1731–1750.
- (2007b). “Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner’s curse and overbidding in private-value auctions?” In: *Econometrica* 75.6, pp. 1721–1770.
- De Clippel, Geoffroy, Rene Saran, and Roberto Serrano (2019). “Level-k mechanism design”. In: *Review of Economic Studies* 86.3, pp. 1207–1227.
- Duffy, John, Jack Ochs, and Lise Vesterlund (2007). “Giving little by little: Dynamic voluntary contribution games”. In: *Journal of Public Economics* 91.9, pp. 1708–1730.
- Eyster, Erik and Matthew Rabin (2010). “Naive herding in rich-information settings”. In: *American Economic Journal: Microeconomics* 2.4, pp. 221–43.
- Fey, Mark, Richard D McKelvey, and Thomas R Palfrey (1996). “An experimental study of constant-sum centipede games”. In: *International Journal of Game Theory* 25.3, pp. 269–287.
- Fong, Meng-Jhang, Po-Hsuan Lin, and Thomas R Palfrey (2023). “Cursed sequential equilibrium”. In: *arXiv preprint arXiv:2301.11971*.
- Fudenberg, Drew and David K Levine (1993). “Self-confirming equilibrium”. In: *Econometrica*, pp. 523–545.
- Fudenberg, Drew and Jean Tirole (1991). “Perfect Bayesian equilibrium and sequential equilibrium”. In: *Journal of Economic Theory* 53.2, pp. 236–260.
- García-Pola, Bernardo, Nagore Iriberri, and Jaromír Kovářík (2020a). “Hot versus cold behavior in centipede games”. In: *Journal of the Economic Science Association* 6.2, pp. 226–238.

- García-Pola, Bernardo, Nagore Iriberry, and Jaromír Kovářik (2020b). “Non-equilibrium play in centipede games”. In: *Games and Economic Behavior* 120, pp. 391–433.
- Hauser, Daniel N and J Aislinn Bohren (2021). “Learning with heterogeneous misspecified models: Characterization and robustness”. In: *PIER Working Paper*.
- Healy, Paul J (2017). “Epistemic experiments: Utilities, beliefs, and irrational play”. In: *mimeo, Ohio State University*.
- Ho, Teck-Hua, Colin F Camerer, and Keith Weigelt (1998). “Iterated dominance and iterated best response in experimental “p-beauty contests””. In: *American Economic Review* 88.4, pp. 947–969.
- Ho, Teck-Hua, So-Eun Park, and Xuanming Su (2021). “A Bayesian level-k model in n-person games”. In: *Management Science* 67.3, pp. 1622–1638.
- Ho, Teck-Hua and Xuanming Su (2013). “A dynamic level-k model in sequential games”. In: *Management Science* 59.2, pp. 452–469.
- Kawagoe, Toshiji and Hirokazu Takizawa (2012). “Level-k analysis of experimental centipede games”. In: *Journal of Economic Behavior & Organization* 82.2-3, pp. 548–566.
- Kreps, David M and Robert Wilson (1982). “Reputation and imperfect information”. In: *Journal of Economic Theory* 27.2, pp. 253–279.
- Levin, Dan and Luyao Zhang (2022). “Bridging level-k to Nash equilibrium”. In: *Review of Economics and Statistics* 104.6, pp. 1329–1340.
- Levitt, Steven D, John A List, and Sally E Sadoff (2011). “Checkmate: Exploring backward induction among chess players”. In: *American Economic Review* 101.2, pp. 975–90.
- Li, Zhi et al. (2021). “Conducting large, repeated, multi-game economic experiments using mobile platforms”. In: *PloS one* 16.4, e0250668.
- Lin, Po-Hsuan (2023). “Cognitive hierarchies in multi-stage games of incomplete information: Theory and experiment”. In: *arXiv preprint arXiv:2208.11190v3*.
- Lin, Po-Hsuan and Thomas R Palfrey (2022). “Cognitive hierarchies in extensive form games”. In: *Social Science Working Paper 1460, California Institute of Technology*.
- Marx, Leslie M and Steven A Matthews (2000). “Dynamic voluntary contribution to a public project”. In: *Review of Economic Studies* 67.2, pp. 327–358.
- McKelvey, Richard D and Richard G Niemi (1978). “A multistage game representation of sophisticated voting for binary procedures”. In: *Journal of Economic Theory* 18.1, pp. 1–22.
- McKelvey, Richard D and Thomas R Palfrey (1992). “An experimental study of the centipede game”. In: *Econometrica*, pp. 803–836.

- McKelvey, Richard D and Thomas R Palfrey (1998). “Quantal response equilibria for extensive form games”. In: *Experimental Economics* 1.1, pp. 9–41.
- Megiddo, Nimrod (1986). *Remarks on bounded rationality*. Citeseer.
- Milgrom, Paul and John Roberts (1982). “Limit pricing and entry under incomplete information: An equilibrium analysis”. In: *Econometrica*, pp. 443–459.
- Nagel, Rosemarie (1995). “Unraveling in guessing games: An experimental study”. In: *American Economic Review* 85.5, pp. 1313–1326.
- Nagel, Rosemarie and Fang Fang Tang (1998). “Experimental results on the centipede game in normal form: an investigation on learning”. In: *Journal of Mathematical Psychology* 42.2-3, pp. 356–384.
- Osborne, Martin J and Ariel Rubinstein (1994). *A course in game theory*. MIT press.
- Palacios-Huerta, Ignacio and Oscar Volij (2009). “Field centipedes”. In: *American Economic Review* 99.4, pp. 1619–35.
- Rampal, Jeevant (2022). “Limited foresight equilibrium”. In: *Games and Economic Behavior* 132, pp. 166–188.
- Rapoport, Amnon et al. (2003). “Equilibrium play and adaptive learning in a three-person centipede game”. In: *Games and Economic Behavior* 43.2, pp. 239–265.
- Rosenthal, Robert W (1981). “Games of perfect information, predatory pricing and the chain-store paradox”. In: *Journal of Economic Theory* 25.1, pp. 92–100.
- Rubinstein, Ariel (2007). “Instinctive and cognitive reasoning: A study of response times”. In: *Economic Journal* 117.523, pp. 1243–1259.
- Schipper, Burkhard C and Hang Zhou (2023). “Level-k thinking in the extensive form”. In: *arXiv preprint arXiv:2311.07789*.
- Selten, Reinhard (1978). “The chain store paradox”. In: *Theory and Decision* 9.2, pp. 127–159.
- (1991). “Anticipatory learning in two-person games”. In: *Game Equilibrium Models I*. Springer, pp. 98–154.
- (1998). “Features of experimentally observed bounded rationality”. In: *European Economic Review* 42.3-5, pp. 413–436.
- Stahl, Dale O (1993). “Evolution of smartn players”. In: *Games and Economic Behavior* 5.4, pp. 604–617.
- (1996). “Boundedly rational rule learning in a guessing game”. In: *Games and Economic Behavior* 16.2, pp. 303–330.
- Stahl, Dale O and Paul W Wilson (1995). “On players’ models of other players: Theory and experimental evidence”. In: *Games and Economic Behavior* 10.1, pp. 218–254.

Wang, Joseph Tao-Yi, Michael Spezio, and Colin F Camerer (2010). “Pinocchio’s pupil: using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games”. In: *American Economic Review* 100.3, pp. 984–1007.

Zauner, Klaus G (1999). “A payoff uncertainty explanation of results in experimental centipede games”. In: *Games and Economic Behavior* 26.1, pp. 157–185.

Chapter 4

COGNITIVE HIERARCHIES IN MULTI-STAGE GAMES OF INCOMPLETE INFORMATION: THEORY AND EXPERIMENT

4.1 Introduction

Multi-stage games of incomplete information are important workhorse models in economics, political science, finance, social networks, and even biology. These games tend to be more intricate than games of perfect information due to the potentially large number of information sets, regardless of how simple the game rules are.¹ The standard approach for analyzing these games is to solve for the sequential equilibrium, wherein players are assumed to form *mutually consistent* beliefs at each information set. In other words, each player’s conjecture about the behavioral strategies of others aligns with the actual strategies of those other players.

The assumption of mutual consistency of the belief system is crucial in standard equilibrium theory, as it, along with the best response requirement, pins down precise predictions of equilibrium outcomes. However, this requirement may be implausibly strong from an empirical standpoint, especially for complicated multi-stage games of incomplete information, as indicated by behavior observed in many laboratory experiments (e.g., Camerer, 2003).

In response to these findings, I develop a new tool for analyzing multi-stage games of incomplete information without relying on mutual consistency of beliefs: the “Dynamic Cognitive Hierarchy (DCH) Solution.” The contribution of this paper encompasses both theoretical and experimental aspects. Theoretical contributions involve extending the DCH solution from games of perfect information, as characterized by Lin and Palfrey (2022), to multi-stage games of incomplete information. On the experimental front, I design and conduct a laboratory experiment to test a key implication of DCH—the *violation of invariance under strategic equivalence*²—in the context of the dirty-faces game, a classic game for studying iterative rationality.

Two extensive games are strategically equivalent if they share the same reduced nor-

¹For example, Johanson (2013) estimates that in a two-person Texas Hold’em game, the number of information sets is around 10^{162} , which is 10^{82} times larger than the number of atoms in the observed universe.

²The violation of invariance under strategic equivalence is sometimes referred to as “the representation effect” in the literature.

mal form. It has been argued that any good equilibrium should exhibit invariance in strategically equivalent games (e.g., Kohlberg and Mertens, 1986). This invariance property is also appealing from the standpoint of experimental design as it suggests that implementing a dynamic game with its reduced normal form does not distort behavior. By doing so, one can gather more experimental data, particularly at information sets that are only occasionally reached. This approach to experimental design is commonly referred to as the “strategy method” (Selten, 1967).

In experimental methodology, an ongoing debate surrounds whether the use of the strategy method distorts behavior (Brandts and Charness, 2011). DCH sheds light on this debate by indicating a potential violation of invariance under strategic equivalence, suggesting that the strategy method could theoretically create distortions in behavior. To empirically test this prediction, I implement a dirty-faces game experiment, which consists of two treatments: the sequential and the simultaneous treatments. In the sequential treatment, the game is played period-by-period. Players can observe the history and are asked to make decisions at realized information sets. In contrast, in the simultaneous treatment, the game is played in reduced normal form, where players choose their contingent strategies.

Furthermore, the game parameters used in the experiment are selected using an “optimal design approach,” where I first calibrate DCH using data from previous dirty-faces game experiments reported by other studies, and then select the game parameters to maximize the expected treatment effect. The utilization of an optimal design approach offers a systematic method for experimenters to choose game parameters, which is not commonly applied in economic experiments. To some extent, the experimental design in this paper serves as a proof-of-concept illustrating how this approach can help experimenters in designing future theory-testing experiments.

By employing this fine-tuned experimental design, significant treatment effects involving the violations of invariance are detected in the data. Both the direction and magnitudes of the observed differences align with the predictions of DCH. Furthermore, when comparing DCH with alternative behavioral solution concepts that relax the best response requirement and the ability to make Bayesian inferences, we find that DCH significantly outperforms other models in both treatments. While there is evidence of the failure of best responses and Bayesian inferences, the observed violation of invariance in the data is primarily attributed to the relaxation of mutual consistency.

To offer readers a better intuitive understanding of the paper, I will next provide an

overview of the DCH solution, illustrate its application in the dirty-faces game, and discuss the experimental design and findings.

Overview of the DCH Solution

The DCH solution is akin to the “level- k model,” a non-equilibrium framework introduced by Nagel (1995). That model relaxes the mutual consistency requirement in simultaneous-move games by assuming a hierarchical structure of strategic sophistication among the players. In the level- k model, each player is endowed with a specific level of sophistication. Level 0 players are non-strategic and choose their actions randomly. Level k players, on the other hand, incorrectly believe that all other players are level $k - 1$ and best respond to this belief.³

The level- k model has been widely applied to organize experimental data in simultaneous move games like the beauty contest game, coordination games, sender-receiver games, auction games, and more. Nevertheless, when applying the standard level- k model to dynamic games, a logical conundrum arises: level k players are assumed to choose actions that maximize the continuation value of the game, while believing that all other players are level $k - 1$ in the continuation game. Consequently, each player’s belief about the levels of others remains fixed from the beginning, potentially leading to situations where level k players are “surprised” by an opponent’s move that contradicts the strategy of a level $k - 1$ player.

To illustrate this conundrum, consider the extensive game shown in Figure 4.1. Suppose the level- k model predicts that level 1 player 1 will choose A , while level 2 player 1 will choose B . Since level 3 player 2 thinks player 1 is level 2, he believes that player 1 will certainly choose B . From the perspective of level 3 player 2, A is an “off-path event” of player 1. If A is chosen, level 3 player 2’s belief about player 1’s level is incompatible with the history.

To avoid this issue, the DCH solution assumes that level k players believe all other players have lower levels distributed anywhere from level 0 to $k - 1$, and update their beliefs about others’ levels as the history unfolds. Specifically, suppose each player’s level is drawn from the distribution $p = (p_k)_{k=0}^{\infty}$. Level 0 players uniformly randomize at every information set. Level k players’ prior belief about any other player being level $j < k$ is $p_j / \sum_{l=0}^{k-1} p_l$, which follows the Cognitive Hierarchy

³In this paper, level 0 players may be interchangeably referred to as “non-strategic players,” and level $k \geq 1$ players as “strategic players” since they best respond to their beliefs.

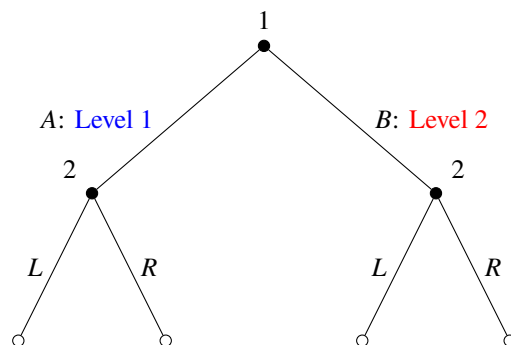


Figure 4.1: An illustration game for the level- k model in extensive games.

(CH) specification proposed by Camerer, Ho, and Chong (2004). Level k players correctly perceive lower-level players' behavioral strategies and update their beliefs about the other players' levels using Bayes' rule as the game progresses.

The previous illustrative example clearly demonstrates how DCH solves the conundrum. Suppose DCH level 1 player 1 chooses A , while DCH level 2 player 1 selects B . In contrast to the level- k model, DCH assumes that level 3 player 2's prior belief about player 1 being level $j = 0, 1, 2$ is $p_j / \sum_{l=0}^2 p_l$. Moreover, after observing player 1's choice of A , level 3 player 2 eliminates the possibility of player 1 being level 2 (otherwise, B would have been chosen), and the beliefs about player 1 being level 0 and level 1 are $\frac{0.5p_0}{0.5p_0+p_1}$ and $\frac{p_1}{0.5p_0+p_1}$, respectively. Because level 0 players uniformly randomize at every information set, strategic players' beliefs are always well-defined, effectively resolving the conundrum.

When extending DCH from games of perfect information to multi-stage games of incomplete information, players will update their beliefs about others' levels *and* payoff-relevant private types at the same time. That is, the DCH belief system is a joint measure about the types and levels of other players.

Proposition 4.1 establishes that if the private types are independently drawn across players, every level of every player's posterior belief remains independent across players at every information set. It's important to note that this property of DCH holds only in multi-stage games with observed actions⁴ but not in general extensive games. Furthermore, when private types are correlated across players, Proposition 4.2 demonstrates that the original game (with correlated types) can be transformed into another game with independent types, and the DCH behavioral strategy profiles remain invariant in both games. These two propositions provide the recipe for

⁴Games of perfect information belong to multi-stage games with observed actions.

solving DCH in multi-stage games of incomplete information. Lastly, because level 0 players uniformly randomize at every information set, when the sizes of action sets differ, level 0 players' behavioral strategies might not be outcome-equivalent. This, in turn, affects all higher-level players, as DCH is solved recursively from the bottom of the hierarchy, causing DCH solutions to differ across strategically equivalent games. This violation of invariance under strategic equivalence is then illustrated in the dirty-faces game.

The Dirty-Faces Game

The dirty-faces game is a diagnostic game to study iterative rationality. It was originally introduced as a mathematical puzzle by (Littlewood, 1953, p. 3-4):

Three ladies, A,B,C, in a railway carriage all have dirty faces and are all laughing. It suddenly flashes on A: why doesn't B realize C is laughing at her? Heavens! I must be laughable.

This game is frequently discussed in understanding iterated reasoning and plays a central role in many studies of common knowledge (see, for example, Binmore and Brandeburger, 1988; Fudenberg and Tirole, 1993 & Geanakoplos, 1994). To illustrate the violation of invariance under strategic equivalence, I consider two versions of the game: the *sequential* and *simultaneous* versions.

The *sequential* (two-person) dirty-faces game was previously studied experimentally by Weber (2001) and Bayer and Chan (2007).⁵ In this game, each player is randomly assigned a face type: either "dirty" or "clean." Once the face types are determined, players can see the other player's face, but not their own. Additionally, if at least one player has a dirty face, a public announcement is made, ensuring that this information becomes common knowledge. After observing the other's face and the announcement, players take actions in a series of periods. In each period, players simultaneously choose between "wait" and "claim" (to have a dirty face). The actions are revealed to both players at the end of each period. If both players decide to wait, the game proceeds to the next period. Otherwise, the game ends after the

⁵Weber (2001) and Bayer and Chan (2007) studied both two-person and three-person games. Since the experimental focus of this paper is on the two-person game, I will primarily focus on the discussion of two-person games in this paper and provide the discussion of three-person games in Appendix D.

period in which at least one player chooses to claim. Players receive rewards for correctly claiming to have a dirty face but are penalized for making false claims.

The second strategically equivalent version is the *simultaneous* (two-person) dirty-faces game, which, to the best of my knowledge, has never been studied experimentally before. In this version, the information structure remains identical to the sequential version. The only difference is that after observing the other's face and announcement, both players simultaneously decide the earliest period to claim as if the game were played in the sequential version. The payoffs are then determined accordingly.

Since the two versions are strategically equivalent, the sequential equilibrium is outcome-equivalent. Consider two players, Ann and Bob. When a player, let's say Ann, sees a clean face (along with the announcement), she should recognize that her own face is dirty and claim in period 1. Yet if Ann sees a dirty face and is rational, she will wait in period 1 since she has no knowledge of her own face type. In equilibrium, Ann believes that Bob would have claimed in the first period if her own face were clean. If the game does indeed reach the second period, Ann will realize that her face is dirty and claim.

In contrast, due to the presence of non-strategic level 0 players, the DCH solution lacks common knowledge of rationality, causing it to differ dramatically from the standard equilibrium solution. Furthermore, DCH makes distinct predictions between the two versions, demonstrating the violation of invariance under strategic equivalence.

To understand the intuition of DCH, we first focus on the sequential version. In period 1, all strategic players behave like rational players, claiming immediately upon seeing a clean face and the announcement, while waiting when they see a dirty face. In contrast, level 0 players choose randomly regardless of their observations. Unlike strategic players, the actions of level 0 players convey no information about the true face types. Therefore, when observing a dirty face and the game proceeds beyond period 1, strategic players will believe that there are two possible situations:

1. If the other player is level 0, then their action is randomly determined and provides no information about my own face.
2. If the other player is level $k \geq 1$, then my own face is certainly dirty because the other player would have claimed in period 1 if my own face were clean.

Consequently, after period 1, a strategic player faces a dynamic tradeoff. If she waits and the game proceeds to the next period, she will become more certain about having a dirty face because the other player is less likely to be a level 0 player. However, the risk of waiting is that the game might be randomly terminated (by a level 0 opponent) and the payoff is further discounted due to impatience. As a result, the DCH solution is characterized by level-dependent stopping periods, which depend on the prior distribution of levels and the payoffs. Because lower-level players are more likely to believe the other is level 0, they need to wait longer to become certain enough about having a dirty face.

The intuition of DCH in the simultaneous version is the same as in the sequential version. However, in the simultaneous version, the number of available strategies changes, causing level 0 players' behavioral strategies to differ between the two versions. Consequently, all higher-level players behave differently between the two versions, demonstrating the violation of invariance under strategic equivalence in DCH.⁶ Furthermore, the magnitude of the difference in behavior predicted by DCH depends on the game parameters. Specifically, there exists two disjoint sets of game parameters, where strategic players tend to claim earlier when observing a dirty face in one set in the sequential version and later in the other set.

Experimental Design and Findings

To assess the violation of invariance under strategic equivalence in the dirty-faces game, I design and conduct a laboratory experiment that manipulates the timing structures (sequential vs. simultaneous) using a between-subject design. The main challenge in designing the experiment, as suggested by DCH, is the selection of game parameters. Specifically, the magnitude of the treatment effect predicted by DCH depends on both the game parameters and the true distribution of levels, which is unknown before the experiment is conducted.

To address this, I develop an “optimal design approach” where I first estimate the distribution of levels using data from an experimental dirty-faces game reported by Bayer and Chan (2007). Then, I select game parameters to maximize diagnosticity by considering a mix of parameters expected to yield various magnitudes of the treatment effect.⁷

⁶See Section 4.5.3 for a detailed discussion of the violation of invariance in dirty-faces games.

⁷It is worth remarking that the optimal design approach is not referring to the one in the statistical literature, whose goal is to maximize the determinant of the information matrix. For more information

This optimal design approach has two advantages. First, it provides experimenters with a systematic method for selecting game parameters when designing experiments. This is important to experimenters since, as noted by Moffatt (2020), when choosing the parameters, “most experimenters have followed an informal approach” (Moffatt, 2020, pp. 335). Second, it serves as a stress test for DCH. After calibrating the prior distribution of levels, DCH offers precise predictions regarding the magnitudes of violations of invariance. Instead of fitting the model *ex post*, this approach provides a benchmark prediction before the experiment is conducted, enabling us to assess the predictive power of DCH.

A significant violation of invariance under strategic equivalence is detected in the data, and more importantly, *how* it is violated aligns with DCH. Additionally, to analyze whether the difference found in the data can be attributed to the relaxation of other equilibrium requirements, I compare DCH with two alternative models: the Agent Quantal Response Equilibrium by McKelvey and Palfrey (1998) and the Cursed Sequential Equilibrium by Fong, Lin, and Palfrey (2023b). These alternatives relax the best response requirement and Bayesian inferences, respectively. The estimation results indicate that for both treatments, DCH fits the data significantly better than the other two solutions. Although relaxing other requirements could improve the fitness, the observed difference is primarily attributed to the relaxation of mutual consistency.

This experimental result highlights how implementing a dynamic game experiment in reduced normal form (using the “strategy method”) can lead to significant distortions in behavior. In addition, DCH provides a better explanation for how these distortions arise compared to other behavioral solution concepts. This suggests that if using the strategy method is necessary, DCH can offer a more reasonable assessment of behavioral distortions compared to the natural approach, which allows subjects to make decisions when it’s their turn, as if the game tree were being fully implemented.⁸

The paper is organized as follows. Section 4.2 discusses the related literature. Section 4.3 sets up the model, and general properties of the DCH solution are established in Section 4.4. In Section 4.5, I demonstrate the violation of invariance under strategic equivalence of DCH in a class of two-person dirty-faces games. Section

on optimal design in the statistical literature and its applications in risky lottery experiments, see Moffatt (2020) Chapter 14, and Bland (2023).

⁸This approach is commonly referred to as the “direct-response method.”

4.6 describes the experimental design, and Section 4.7 reports the experimental results. Finally, Section 4.8 concludes the paper.

4.2 Related Literature

The DCH solution is closely related to a number of behavioral models of games. Over the past thirty years, the idea of limited depth of reasoning has been theoretically studied by various researchers, including Selten (1991), Aumann (1992), Stahl (1993), Selten (1998), Alaoui and Penta (2016), and Alaoui and Penta (2022) & Lin and Palfrey (2022). In addition to theoretical work, Nagel (1995) conducted the first experiment on the “beauty contest game” to study people’s iterative reasoning process. In this game, each player simultaneously chooses an integer between 0 and 100. The winner is the player whose choice is closest to the average of all numbers multiplied by $p \in (0, 1)$. The unique equilibrium predicts that all players should choose 0. However, empirical observations show that almost no player chooses the equilibrium action. Instead, players seem to behave as if they are performing some finite number of iterative best responses.⁹

To explain the data, Nagel (1995) proposed the “level- k model,” which assumes that each player is endowed with a specific “level” of reasoning. Level 0 players randomly select actions from their action sets. For every $k \geq 1$, level k players believe that they are one level of reasoning higher than the rest and best respond accordingly. The level- k model has been applied to various environments, including simultaneous-move games (Costa-Gomes, Crawford, and Broseta, 2001; Crawford and Iriberri, 2007a), two-person guessing games (Costa-Gomes and Crawford, 2006), auctions (Crawford and Iriberri, 2007b), and sender-receiver games (Cai and Wang, 2006; Wang, Spezio, and Camerer, 2010).

The standard level- k model has been successful in explaining the data and has been extended in various ways. One such approach is the CH framework proposed by Camerer, Ho, and Chong (2004), which assumes that players best respond to a mixture of lower-level players.¹⁰ In this framework, level k players best respond to a mixture of lower levels, ranging from level 0 to $k - 1$. Furthermore, players hold accurate beliefs about the relative proportions of the lower levels. However, this

⁹This empirical pattern has been robustly replicated in different environments. For instance, Ho, Camerer, and Weigelt (1998) and Bosch-Domenech et al. (2002) have found similar results in both laboratory and field experiments.

¹⁰In a similar vein, Stahl and Wilson (1995) and Levin and Zhang (2022) allow each level of players to best respond to a mixture of lower-level players and equilibrium players.

approach is primarily developed for simultaneous-move games, and DCH extends it to dynamic games.

Another direction is to endogenize the levels of players. Alaoui and Penta (2016) considered a cost-benefit analysis approach, where players decide their levels of sophistication by weighing the benefits of additional levels against the costs of doing so. The implications of this model were further explored in Alaoui, Janezic, and Penta (2020). In the same spirit as Stahl (1996), Ho and Su (2013), and Ho, Park, and Su (2021) considered a canonical laboratory environment where players repeatedly play the same game and endogenously choose a new level of sophistication for the next iteration of the game. This approach is different from DCH, where players update their beliefs about other players' levels after each move within a single game. Furthermore, DCH players are strategic learners as they can correctly anticipate the evolution of posterior beliefs in later information sets. This leads to a much different learning dynamic compared to naive adaptive learning models.

At a more conceptual level, the DCH solution is related to other solution concepts for dynamic games that relax the requirements of sequential equilibrium. DCH is a non-equilibrium model that allows players at different levels to best respond to different conjectures about other players' strategies, while Agent Quantal Response Equilibrium (AQRE) by McKelvey and Palfrey (1998) is an equilibrium model in which players make stochastic choices. Both DCH and AQRE assume that players follow Bayes' rule to make inferences. In contrast, Cursed Sequential Equilibrium (CSE) by Fong, Lin, and Palfrey (2023b) and Sequential Cursed Equilibrium (SCE) by Cohen and S. Li (2023) are two different equilibrium models in which players are able to make best responses but are unable to make correct Bayesian inferences.¹¹

One common theoretical property of these behavioral solution concepts is the violation of invariance under strategic equivalence, albeit in different ways. However, whether implementing the same dynamic game with different methods creates any behavioral distortion is an empirical question. Brandts and Charness (2011) surveyed 29 experiments that compared the behavior under different elicitation methods and found that the invariance may be violated under certain conditions.¹² More recently, S. Li (2017) compared the second-price auction and the ascending clock

¹¹The CSE proposed by Fong, Lin, and Palfrey (2023b) captures the situation where players fail to understand how other players' actions depend on their own private information. On the other hand, the SCE introduced by Cohen and S. Li (2023) depicts the bias where people fail to realize how others' action depend on their information set partitions. See Appendix B and Fong, Lin, and Palfrey (2023a) for a detailed comparison of the two solution concepts.

¹²In their survey, they observed no difference in 16 studies, systematic differences in four studies,

auction and found that players are more likely to follow dominant strategies in the ascending clock auction. Additionally, García-Pola, Iriberry, and Kovářík (2020) experimentally studied the invariance in four centipede games and reported that in three of the four games, players tend to terminate the game earlier in the sequential version of the game, which is consistent with DCH (Lin and Palfrey, 2022).¹³ Finally, it is worth noting that D. L. Chen and Schonger (2023a) and D. L. Chen and Schonger (2023b) point out that the violation of invariance is linked to the emotional salience induced during the experiment.

This paper also contributes to the literature on dirty-faces games. The concept of dirty-faces game was originally introduced by Littlewood (1953) as a means to illustrate the transmission of common knowledge. Binmore and Brandeburger (1988) and Fudenberg and Tirole (1993), and Geanakoplos (1994) were the first to theoretically study the dirty faces games with the knowledge operator. Furthermore, Liu (2008) demonstrated that if players are unaware of other players' face types, they might incorrectly claim their face types, and hence influence the transmission of knowledge among the players.

Weber (2001) and Bayer and Chan (2007) conducted the first two experiments on dirty-faces games and found that many subjects fail to perform such iterative reasoning. More recent experiments have further demonstrated the persistence of failure in iterative reasoning even when playing against fully rational robot players (Grehl and Tutić, 2015; W. J. Chen, Fong, and Lin, 2023). This failure has also been found to be correlated with cognitive abilities (Devetag and Warglien, 2003; Bayer and Renou, 2016a; Bayer and Renou, 2016b), while the deviations from the equilibrium significantly decrease when the participants are selected through a market mechanism (Choo and Zhou, 2022). Overall, these experimental findings provide support for the existence of non-strategic types of players who are not sequentially rational, highlighting the heterogeneity in strategic sophistication within the population.

and mixed evidence in nine of them. In particular, they found suggestive evidence that the frequency of violation of invariance is related to the number of available actions.

¹³In the three centipede games where termination occurs earlier in the sequential version, DCH predicts that the distribution of terminal nodes from the simultaneous version (strategy method) will first-order stochastically dominate the distribution from the sequential version (direct response method). In the fourth centipede game where FOSD is not predicted, the empirical distributions of terminal nodes from the two versions are almost identical.

4.3 The Model

Section 4.3.1 introduces the multi-stage games with observed actions, as proposed by Fudenberg and Levine (1983) and Fudenberg and Tirole (1991). This framework provides a tractable approach to studying how players learn about the types and levels of others. Next, the DCH solution for this family of games is defined in section 4.3.2.

4.3.1 Multi-Stage Games with Observed Actions

Let $N = \{1, \dots, n\}$ be a finite set of players. Each player $i \in N$ has a *type* θ_i drawn from a finite set Θ_i . Let $\theta \in \Theta \equiv \times_{i=1}^n \Theta_i$ be the type profile and θ_{-i} be the type profile without player i . All players have the common (full support) prior distribution $F : \Theta \rightarrow (0, 1)$. At the beginning of the game, each player is told his own type, but is not informed anything about the types of others. Therefore, each player i 's initial belief about the types of others when his type is θ_i is

$$F(\theta_{-i}|\theta_i) = \frac{F(\theta_{-i}, \theta_i)}{\sum_{\theta'_{-i} \in \Theta_{-i}} F(\theta'_{-i}, \theta_i)}.$$

If the types are independent across players, each player i 's initial belief about the types of others is $F_{-i}(\theta_{-i}) = \prod_{j \neq i} F_j(\theta_j)$ where $F_j(\theta_j)$ is the marginal distribution of player j 's type.

The game is played in “periods” $t = 1, 2, \dots, T$ where $T < \infty$. In each period, players simultaneously choose their actions, which will be revealed at the end of the period. The feasible set of actions can vary with histories, so games with alternating moves are also included. Let \mathcal{H}^{t-1} be the set of all available histories at period t , where $\mathcal{H}^0 = \{h_0\}$ and \mathcal{H}^T is the set of terminal histories. Let $\mathcal{H} = \cup_{t=0}^T \mathcal{H}^t$ be the set of all available histories of the game, and let $\mathcal{H} \setminus \mathcal{H}^T$ be the set of non-terminal histories.

For every player i , the available information at period t is in $\mathcal{H}^{t-1} \times \Theta_i$. Thus, each player i 's information sets can be specified as $\mathcal{I}_i \in \Pi_i = \{(h, \theta_i) : h \in \mathcal{H} \setminus \mathcal{H}^T, \theta_i \in \Theta_i\}$. That is, a type θ_i player i 's information set at the public history h^t can be defined as $\cup_{\theta_{-i} \in \Theta_{-i}} (h^t, \theta_{-i}, \theta_i)$. With a slight abuse of notation, it will be denoted as (h^t, θ_i) . For the sake of simplicity, the feasible set of actions for every player at every history is assumed to be type-independent. Let $A_i(h^{t-1})$ be the feasible set of actions for player i at history h^{t-1} and let $A_i = \times_{h \in \mathcal{H} \setminus \mathcal{H}^T} A_i(h)$ be the set of player i 's all feasible actions in the game. For each player i , A_i is assumed to be finite and $|A_i(h)| \geq 1$ for any $h \in \mathcal{H} \setminus \mathcal{H}^T$. Let $a_i^t \in A_i(h^{t-1})$ be player i 's action at history

h^{t-1} , and let $a^t = (a_1^t, \dots, a_n^t) \in \times_{i=1}^n A_i(h^{t-1})$ denote the action profile at stage t . If a^t is the action profile chosen at stage t , then $h^t = (h^{t-1}, a^t)$.

A behavioral strategy for player i is a function $\sigma_i : \Pi_i \rightarrow \Delta(A_i)$ satisfying $\sigma_i(h^{t-1}, \theta_i) \in \Delta(A_i(h^{t-1}))$. Let $\sigma_i(a_i^t | h^{t-1}, \theta_i)$ denote the probability for player i to choose $a_i^t \in A_i(h^{t-1})$. A strategy profile $\sigma = (\sigma_i)_{i \in N}$ specifies a behavioral strategy for each player i . Lastly, each player i has a payoff function (in von Neumann-Morgenstern utilities) $u_i : \mathcal{H}^T \times \Theta \rightarrow \mathbb{R}$, and let $u = (u_1, \dots, u_n)$ be the profile of utility functions. A multi-stage game with observed actions, Γ , is defined by the tuple $\Gamma = \langle N, \mathcal{H}, \Theta, F, u \rangle$.

4.3.2 Dynamic Cognitive Hierarchy Solution

Each player i is endowed with a level of sophistication $\tau_i \in \mathbb{N}_0$ which is independently drawn from the distribution $P_i(\tau_i)$. Without loss of generality, I assume $P_i(\tau_i) > 0$ for all $i \in N$ and $\tau_i \in \mathbb{N}_0$. Let $\tau = (\tau_1, \dots, \tau_n)$ be the level profile and τ_{-i} be the level profile without player i . Due to the independence, the level profile is drawn from a distribution $P : \mathbb{N}_0^{|N|} \rightarrow (0, 1)$ such that $P(\tau) = \prod_{i=1}^n P_i(\tau_i)$. Following Lin and Palfrey (2022), I assume that players' "types" and "levels" are drawn independently.

Assumption 4.1. *F and P are independent distributions.*

Each player i has a prior belief about the opponents' levels which satisfies the property of *truncated rational expectations*. That is, for each $i, j \neq i$, and k , let $\hat{P}_{ij}^k(\tau_j)$ be level k player i 's prior belief about player j 's level, and $\hat{P}_{ij}^k(\tau_j)$ satisfies

$$\hat{P}_{ij}^k(\tau_j) = \begin{cases} \frac{P_j(\tau_j)}{\sum_{m=0}^{k-1} P_j(m)} & \text{if } \tau_j < k \\ 0 & \text{if } \tau_j \geq k. \end{cases} \quad (4.1)$$

The intuition of (4.1) is that despite mistakenly believing all other players are at most level $(k-1)$,¹⁴ each level of players have a correct belief about the relative proportions of players who are less sophisticated than they are.

In the DCH solution, a strategy profile is a level-dependent profile of behavioral strategy of each level of each player. Let σ_i^k be level k player i 's behavioral strategy,

¹⁴The cognitive hierarchy specification is in line with behavioral and psychological evidence of overconfidence across various domains (see, for instance, Camerer and Lovallo, 1999; Moore and Healy, 2008 & Enke, Graeber, and Oprea, 2023). While recent findings by Halevy, Hoelzemann, and Kneeland (2021) suggest the possibility of players believing others to be more sophisticated than themselves, this behavior falls beyond the scope of this paper, as DCH is developed within the confines of the standard CH framework.

where level 0 players uniformly randomize at every information set.¹⁵ That is, for every $i \in N$, $\theta_i \in \Theta_i$, $h \in \mathcal{H} \setminus \mathcal{H}^T$, and for all $a \in A_i(h)$,

$$\sigma_i^0(a \mid \theta_i, h) = \frac{1}{|A_i(h)|}.$$

At every history h^t , every strategic level k player i forms a *joint belief* about all other players' types and levels.¹⁶ Their posterior beliefs at history h^t depend on the level-dependent strategy profile and the prior beliefs. To formalize the belief updating process, let $\sigma_j^{<k} = (\sigma_j^0, \dots, \sigma_j^{k-1})$ be the profile of strategies adopted by the levels below k of player j . Furthermore, let $\sigma_{-i}^{<k} = (\sigma_1^{<k}, \dots, \sigma_{i-1}^{<k}, \sigma_{i+1}^{<k}, \dots, \sigma_n^{<k})$ be the profile of behavioral strategies of the levels below k of all players other than player i . It is worth noticing that all strategic players believe every history is possible because $\hat{P}_{ij}^k(0) > 0$ for all $i, j \in N$ and $k > 0$, and $\sigma_j^0(a \mid h, \theta_j) > 0$ for all j, θ_j, h and $a \in A_j(h)$. Consequently, Bayes' rule can be applied to derive every level of players' posterior belief about other players' types and levels. Specifically, for any $i \in N$, $k \geq 1$ and $\theta_i \in \Theta_i$, a level-dependent strategy profile will induce the posterior belief $\mu_i^k(\tau_{-i}, \theta_{-i} \mid h, \theta_i)$ at every $h \in \mathcal{H} \setminus \mathcal{H}^T$ with $\mu_i^k(\theta_{-i} \mid h^{t-1}, \theta_i)$ and $\mu_i^k(\tau_{-i} \mid h^{t-1}, \theta_i)$ being level k player i 's marginal beliefs of other players' types and levels at history h^{t-1} , respectively. Lastly, for any $j \neq i$, let $\mu_{ij}^k(\tau_j, \theta_j \mid h^{t-1}, \theta_i)$ denote level k player i 's belief about player j 's type and level at history h^{t-1} .

In the DCH solution, players correctly anticipate how they will update their posterior beliefs at all future histories of the game, i.e., players are strategic learners. Therefore, for any i, k, θ_i and any level-dependent strategy-profile of others $\sigma_{-i}^{<k}$, type θ_i level k player i believes the probability of $a_{-i}^t \in A_{-i}(h^{t-1})$ being chosen is

$$\tilde{\sigma}_{-i}^{-k}(a_{-i}^t \mid h^{t-1}, \theta_i) \equiv \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{\{\tau_{-i}: \tau_j < k \forall j \neq i\}} \mu_i^k(\tau_{-i}, \theta_{-i} \mid h^{t-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau_j}(a_j^t \mid h^{t-1}, \theta_j).$$

Furthermore, for every level of players, given lower-level players' strategies, they can compute the probability of any outcome being realized at any non-terminal history. In particular, for any $i \in N$, $\tau_i > 0$, $\theta \in \Theta$, σ , and τ_{-i} such that $\tau_j < \tau_i$ for any $j \neq i$, let $P_i^{\tau_i}(h^T \mid h^{t-1}, \theta, \tau_{-i}, \sigma_{-i}^{<\tau_i}, \sigma_i^{\tau_i})$ be level τ_i player i 's belief about the conditional realization probability of $h^T \in \mathcal{H}^T$ at history $h^{t-1} \in \mathcal{H} \setminus \mathcal{H}^T$ if the type

¹⁵Uniform randomization is not the only way to model level 0 players' behavior; however, one compelling justification for its use is its universal applicability to all games in the same manner. In fact, the DCH solution is well-defined as long as level 0 players' behavioral strategy is fully mixed at every information set.

¹⁶Level 1 players always believe other players are level 0 whose actions are uninformative about their types. Therefore, they don't update their beliefs about the levels and types of others.

profile is θ , the level profile is τ , and player i uses $\sigma_i^{\tau_i}$. Finally, level τ_i player i 's expected payoff at any $h^t \in \mathcal{H} \setminus \mathcal{H}^T$ is

$$\begin{aligned} \mathbb{E}u_i^{\tau_i}(\sigma \mid h^t, \theta_i) = \\ \sum_{h^T \in \mathcal{H}^T} \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{\{\tau_{-i}: \tau_j < k \ \forall j \neq i\}} \mu_i^{\tau_i}(\tau_{-i}, \theta_{-i} \mid h^t, \theta_i) P_i^{\tau_i}(h^T \mid h^t, \theta, \tau_{-i}, \sigma_{-i}^{< \tau_i}, \sigma_i^{\tau_i}) u_i(h^T, \theta_{-i}, \theta_i). \end{aligned} \quad (4.2)$$

The *DCH solution* of the game is defined as the level-dependent assessment (σ^*, μ^*) , such that $\sigma_i^{k*}(\cdot \mid h^t, \theta_i)$ maximizes (4.2) for all i, k, θ_i and $h^t \in \mathcal{H} \setminus \mathcal{H}^T$ and the *DCH belief system* μ^* is induced by σ^* . Moreover, players are assumed to uniformly randomize over optimal actions when they are indifferent. This is a typical assumption in level- k models, and it is convenient because it ensures a unique DCH solution.

Lemma 4.1. *The DCH solution is unique.*

Proof: See Appendix D. ■

Remark 4.1. *For one-stage games, the DCH solution reduces to the standard CH solution because one-stage games are essentially static games.*

4.4 General Properties of the DCH Solution

In this section, I first characterize some general properties of the belief updating process of DCH. Assume for now that players' types are independently drawn, i.e., $F(\theta) = \prod_{i \in N} F_i(\theta_i)$. With this assumption, Proposition 4.1 shows that at every information set, the posterior beliefs are independent across players. In other words, the DCH belief system is a product measure.

Proposition 4.1. *For any multi-stage game with observed actions Γ , any $h \in \mathcal{H} \setminus \mathcal{H}^T$, any $i \in N$, $\theta_i \in \Theta_i$, and for any $k \in \mathbb{N}$, if the prior distribution of types is independent across players, i.e., $F(\theta) = \prod_{i=1}^n F_i(\theta_i)$, then level k player i 's posterior belief about other players' types and levels at h is independent across players. That is,*

$$\mu_i^k(\tau_{-i}, \theta_{-i} \mid h, \theta_i) = \prod_{j \neq i} \mu_{ij}^k(\tau_j, \theta_j \mid h, \theta_i).$$

Proof: See Appendix D. ■

Proposition 4.1 extends the independence property shown by Lin and Palfrey (2022) from games of perfect information to multi-stage games with observed actions. This

generalization relies on that (1) the actions are perfectly observed and (2) players are able to perform Bayesian inferences.

In multi-stage games with observed actions, as the prior distribution of types and levels is independent across players,¹⁷ because every player's action is perfectly monitored, strategic players understand that each player's action does not convey any information about other players' private information. In this case, Proposition 4.1 shows the belief system remains to be a product measure in any information set. However, as pointed out by Lin and Palfrey (2022), this is not true for general dynamic games of imperfect information. When the actions are not perfectly observed, the marginal beliefs about others' levels could be correlated across players (see section 7.2 of Lin and Palfrey, 2022).

Besides, the ability to perform Bayesian inferences plays a crucial role in maintaining the independence property. In other behavioral solution concepts, such as the Cursed Sequential Equilibrium proposed by Fong, Lin, and Palfrey (2023b), where players are unable to perform Bayesian inferences, players may mistakenly believe that others' actions are informative about another player's private information, even though the actions are perfectly observed and the prior distribution of types is independent across players.

It is worth noticing that the independence property is useful for solving the DCH solution. When there are more players or when the game structure becomes more complex, computing the posterior belief can become challenging, as it involves level-dependent probability measures. Yet Proposition 4.1 guarantees that the DCH belief system can be computed player-wise rather than information-set-wise, which simplifies the computation process.

Next, I consider the case where the prior distribution of types is not independent across players. When the types are correlated across players, their actions are informative about not only their own private information but also the private information of players whose types are correlated with them. Similar to the observations of Myerson (1985) and Fudenberg and Tirole (1991), to deal with correlated types, the original game (with correlated types) can be simply transformed into one game with independent types with a specific transformation.

For any multi-stage game with observed actions Γ , consider a corresponding transformed game $\hat{\Gamma}$ where the prior distribution of types is the product of independent

¹⁷The independence property does not rely on Assumption 4.1. It holds as long as the priors distributions of types and levels are both independent across players.

uniform marginal distributions. Namely,

$$\hat{F}(\theta) = \frac{1}{\prod_{i=1}^n |\Theta_i|} \quad \forall \theta \in \Theta.$$

In addition, the utility functions are transformed to be

$$\hat{u}_i(h^T, \theta_{-i}, \theta_i) = F(\theta_{-i}|\theta_i)u_i(h^T, \theta_{-i}, \theta_i).$$

Proposition 4.2 shows that the DCH level-dependent behavioral strategy profile is invariant under the transformation between the transformed and original game, suggesting that the independence assumption of the types is without loss of generality. Moreover, it is surprising that the transformation is *level-independent*, given that the best response of each level is determined iteratively. The intuition behind this is that, due to the independence of types and levels (Assumption 4.1), players cannot make inferences about others' types based on their knowledge of others' levels.

Proposition 4.2. *The level-dependent assessment $(\hat{\sigma}, \hat{\mu})$ is the DCH solution of the transformed game (with independent types) if and only if the level-dependent assessment (σ, μ) is the DCH solution of the original game (with correlated types) where $\sigma = \hat{\sigma}$ and for any $i \in N$, $\theta_i \in \Theta_i$, $k > 0$, and $h^t \in \mathcal{H} \setminus \mathcal{H}^T$,*

$$\mu_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i) = \frac{F(\theta_{-i}|\theta_i)\hat{\mu}_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}; \tau'_j < k \forall j \neq i\}} F(\theta'_{-i}|\theta_i)\hat{\mu}_i^k(\tau'_{-i}, \theta'_{-i}|h^t, \theta_i)}.$$

Proof: See Appendix D. ■

Proposition 4.2 shows that at any (h^t, θ_i) , player i 's belief of a specific type-level profile (τ_{-i}, θ_{-i}) is proportional to prior belief of θ_{-i} conditional on θ_i . In addition, if $F(\theta_{-i}|\theta_i) \rightarrow 1$, i.e., θ_i is almost perfectly correlated with θ_{-i} , then

$$\begin{aligned} \mu_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i) &= \frac{F(\theta_{-i}|\theta_i)\hat{\mu}_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}; \tau'_j < k \forall j \neq i\}} F(\theta'_{-i}|\theta_i)\hat{\mu}_i^k(\tau'_{-i}, \theta'_{-i}|h^t, \theta_i)} \\ &\rightarrow \frac{\hat{\mu}_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i)}{\sum_{\{\tau'_{-i}; \tau'_j < k \forall j \neq i\}} \hat{\mu}_i^k(\tau'_{-i}, \theta_{-i}|h^t, \theta_i)} = \hat{\mu}_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i), \end{aligned}$$

implying that the belief in the transformed game aligns with the belief in the original game. Intuitively speaking, if the types are almost perfectly correlated, then the remaining information to be learned is solely others' levels. Since the DCH behavioral strategy profile is invariant under the transformation, the belief about others' levels will also be invariant under the transformation.

The next property is about the evolution of the support of the beliefs. Lin and Palfrey (2022) have shown that in games of perfect information, the support of beliefs about the levels weakly shrinks as the history unfolds. Proposition 4.3 extends this result to multi-stage games with observed actions, indicating that the support of the marginal beliefs about the levels weakly shrinks in later periods. Additionally, Proposition 4.3 demonstrates that the marginal beliefs about other players' types always maintain full support. In other words, in the DCH solution, players will become more certain about the levels of others in later periods while never completely ruling out any possibility of a type profile. To formally state the proposition, I first define the support of the marginal beliefs.

Definition 4.1 (Support). *For any multi-stage game with observed actions Γ , any $i \in N$, any $\tau_i \in \mathbb{N}$, any $\theta_i \in \Theta_i$, and any history $h \in \mathcal{H} \setminus \mathcal{H}^T$, let $\text{supp}_i(\theta_{-i}|h, \theta_i, \tau_i)$ and $\text{supp}_i(\tau_{-i}|h, \theta_i, \tau_i)$ be the support of level τ_i player i 's marginal belief about other players' types and levels at information set (h, θ_i) , respectively. In other words, for any θ'_{-i} and τ'_{-i} ,*

$$\begin{aligned}\theta'_{-i} \in \text{supp}_i(\theta_{-i}|h, \theta_i, \tau_i) &\iff \mu_i^{\tau_i}(\theta'_{-i} | h, \theta_i) > 0, \\ \tau'_{-i} \in \text{supp}_i(\tau_{-i}|h, \theta_i, \tau_i) &\iff \mu_i^{\tau_i}(\tau'_{-i} | h, \theta_i) > 0.\end{aligned}$$

Proposition 4.3. *Consider any multi-stage game with observed actions Γ , any $i \in N$, any $\tau_i \in \mathbb{N}$, and any $\theta_i \in \Theta_i$. The following two statements hold.*

1. *For any $h^t = (h^{t-1}, a^t) \in \mathcal{H}^t \setminus \mathcal{H}^T$, $\text{supp}_i(\tau_{-i}|h^t, \theta_i, \tau_i) \subseteq \text{supp}_i(\tau_{-i}|h^{t-1}, \theta_i, \tau_i)$.*
2. *For any $h \in \mathcal{H} \setminus \mathcal{H}^T$, $\text{supp}_i(\theta_{-i}|h, \theta_i, \tau_i) = \Theta_{-i}$.*

Proof: See Appendix D. ■

The intuition behind Proposition 4.3 is that since it is always possible for other players to be level 0, players can always rationalize any type profile by assuming all other players are level 0.¹⁸ This implies that no matter how sophisticated the players are, common knowledge of rationality is never reached in DCH, which suggests that DCH and the equilibrium theory are fundamentally different solution concepts.

Furthermore, another feature of DCH that sharply contrasts with the equilibrium theory is the violation of invariance under strategic equivalence. In DCH, since level

¹⁸If the action sets vary not only with histories but also with types, players may rule out the possibility of certain type profiles when specific actions are chosen.

0 players uniformly randomize at every information set, their behavioral strategies might not be outcome-equivalent if the cardinality of action sets changes. All higher-level players are subsequently affected, as DCH is solved recursively from the bottom of the hierarchy.¹⁹ In the remaining part of the paper, I will demonstrate theoretically and experimentally how the invariance property is violated in a family of dirty-faces games.

Finally, it is worth remarking that the DCH posterior beliefs about the types and levels are generally *correlated* despite of that the types and levels are determined independently (Assumption 4.1). This will also be illustrated in the next section.

4.5 DCH Analysis of the Dirty-Faces Games

The dirty-faces game was originally developed by Littlewood (1953) to study the role of common knowledge.²⁰ In this section, I focus on a particular specification of the game that has been theoretically and experimentally studied in the literature (see e.g., Fudenberg and Tirole, 1993, Weber, 2001 and Bayer and Chan, 2007).

There are two players $N = \{1, 2\}$ and there are up to $2 \leq T < \infty$ periods. At the beginning of the game, each player i is randomly assigned a face type, denoted as x_i , which can be either $x_i = O$ (representing a clean face) or $x_i = X$ (representing a dirty face). The face types are i.i.d. drawn from the distribution $p = \Pr(x_i = X) = 1 - \Pr(x_i = O)$ where $p > 0$ represents the probability of having a dirty face. After the face types are determined, each player i can observe the other player's face type x_{-i} but not their own face. Hence, player i 's private information is the other player's face type x_{-i} . Furthermore, if at least one player has a dirty face, a public announcement is made, informing both players of this fact. If there is no announcement, it is common knowledge to both players that both faces are clean. To avoid triviality, I will focus on the case where an announcement is made.

After seeing the other player's face type and the announcement, in each period, every player i simultaneously chooses to "Wait" (W) or "Claim" (to have a dirty face, C) and their actions are revealed at the end of each period. The game will end

¹⁹According to Thompson (1952) and Elmes and Reny (1994), two extensive games share the same reduced normal form if and only if they can be transformed into each other using a small set of elementary transformations. Specifically, Elmes and Reny (1994) propose three such transformations: INT, COA, and ADD, which preserve perfect recall. Because DCH is sensitive to the cardinality of action sets, it varies under COA while remaining invariant under INT and ADD.

²⁰It has also been referred to as the "cheating wives puzzle" (Gamow and Stern, 1958), the "cheating husbands puzzle" (Moses, Dolev, and Halpern, 1986), the "muddy children puzzle" (Barwise, 1981; Halpern and Moses, 1990), and the "red hat puzzle" (Hardin and Taylor, 2008).

after any period where some player chooses C or after period T . The last period of the game is called the “terminal period,” and both players’ payoffs are determined by their own face types and their actions in the terminal period.

Suppose period t is the terminal period. If player i chooses W in the terminal period, his payoff for this game is 0 regardless of his face type. On the other hand, if player i chooses C to terminate the game, he will receive $\alpha > 0$ if his face is dirty and -1 if his face is clean. Besides, payoffs are discounted with a common discount factor $\delta \in (0, 1)$ per period. That is, if player i claims in period t and $x_i = X$, player i will receive $\delta^{t-1}\alpha$, but player i will receive $-\delta^{t-1}$ if $x_i = O$. To make players unattractive to gamble if they do not have additional information except for the prior, following Weber (2001) and Bayer and Chan (2007), I assume that

$$p\alpha - (1 - p) < 0 \iff 0 < \bar{\alpha} \equiv \frac{p\alpha}{1 - p} < 1, \quad (4.3)$$

which guarantees it is strictly dominated to choose C in period 1 when seeing a dirty face. Thus, a two-person dirty-faces game is defined by a tuple $\langle T, \delta, \alpha, p \rangle$ where $(\delta, \bar{\alpha}) \in (0, 1)^2$.

With common knowledge of rationality, the unique equilibrium can be solved through the following iterative reasoning: When player i sees a clean face, the public announcement will lead him to realize that his own face is dirty and claim in period 1. On the other hand, when player i sees a dirty face, he will wait in period 1 because of the uncertainty about his own face. However, if player $-i$ also waits in period 1, player i will then recognize that his own face is dirty and claim in period 2, as player i knows that if his own face were clean, player $-i$ would have claimed in period 1.

When implementing this game in a laboratory experiment, the natural approach is to specify this game as a *sequential* dirty-faces game and allow subjects to make decisions period-by-period, following the rules described above, using the direct-response method. Alternatively, the other approach is the strategy method which specifies this game as a *simultaneous* dirty-faces game—after seeing the other’s face and the announcement, players simultaneously decide a “plan” which specifies the period to claim or always wait. From the standard game-theoretic perspective, the sequential and simultaneous dirty-faces game are strategically equivalent as they share the same reduced normal form. In the following, I will demonstrate that the DCH solution varies in these two versions of the game, illustrating the violation of invariance under strategic equivalence of DCH.

4.5.1 DCH Solution for the Sequential Dirty-Faces Games

In the sequential dirty-faces game, since there are (at most) T periods, a behavioral strategy for player i is a mapping from the period and the observed face type ($x_{-i} \in \{O, X\}$) to the probability of C . The behavioral strategy is denoted by

$$\sigma_i : \{1, \dots, T\} \times \{O, X\} \rightarrow [0, 1].$$

For the sake of simplicity, I assume that each player i 's level is i.i.d. drawn from the distribution $p = (p_k)_{k=0}^{\infty}$ where $p_k > 0$ for all k . In the DCH solution, each player's optimal behavioral strategy is level-dependent. Let the behavioral strategy of level k player i be σ_i^k . Following previous notations, let $\mu_i^k(x_i, \tau_{-i}|t, x_{-i})$ be level k player i 's belief about their own face and the level of the other player, conditional on observing x_{-i} and being at period t . Level 0 players will uniformly randomize everywhere, so $\sigma_i^0(t, x_{-i}) = 1/2$ for all t and x_{-i} .

Proposition 4.4 fully characterizes the DCH solution for the sequential dirty-faces games. When observing a clean face, a player can immediately figure out that his face is dirty. Therefore, DCH coincides with the equilibrium prediction when $x_{-i} = O$. However, if a player sees a dirty face and the other player waits in period 1, he cannot tell his face type for sure, no matter how sophisticated he is. Instead, he will believe that he is more likely to have a dirty face as the game continues. As a result, conditional on observing a dirty face, level $k \geq 2$ players will claim as long as the reward $\bar{\alpha}$ is high enough or the discount rate δ is sufficiently low. Otherwise, they will wait for more evidence.

Proposition 4.4. *For any sequential two-person dirty-faces game, the level-dependent strategy profile of the DCH solution satisfies that for any $i \in N$,*

1. $\sigma_i^k(t, O) = 1$ for any $k \geq 1$ and $1 \leq t \leq T$.
2. $\sigma_i^1(t, X) = 0$ for any $1 \leq t \leq T$. Moreover, for any $k \geq 2$,
 - (1) $\sigma_i^k(1, X) = 0$,
 - (2) for any $2 \leq t \leq T - 1$, $\sigma_i^k(t, X) = 1$ if and only if

$$\bar{\alpha} \geq \frac{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta \right] p_0 + (1 - \delta) \sum_{j=1}^{k-1} p_j},$$

(3) $\sigma_i^k(T, X) = 1$ if and only if

$$\bar{\alpha} \geq \frac{\left(\frac{1}{2}\right)^{T-1} p_0}{\left(\frac{1}{2}\right)^{T-1} p_0 + \sum_{j=1}^{k-1} p_j}.$$

Proof: See Appendix D. ■

To gain insights into the mechanics of the model, I analyze the behavior of level 1 and 2 players. Level 1 players believe the other player is non-strategic, implying that the other player's actions do not provide any information about their face type. Consequently, for level 1 players, the announcement and their own observations are the only relevant sources of information. As a result, level 1 players in each period will behave exactly the same as in period 1—they will claim when seeing a clean face, and wait when seeing a dirty face. This is because they cannot gather any additional information from the other's actions, and their belief about their own face type remains unchanged in every period.

For level 2 players, they will also claim to have a dirty face immediately upon seeing a clean face. Furthermore, level 2 players are aware that level 1 players will claim in period 1 if they observe a clean face. In contrast, when observing a dirty face, level 2 players will wait in period 1 (which is a strictly dominant strategy) and form a joint belief about the other player's level and their own face type if the game proceeds to period 2. Due to the presence of level 0 players, even if the game proceeds to period 2, level 2 players are still uncertain about their face types. However, they will know it is *impossible* that the other player is level 1 and their own face is clean. Specifically, level 2 players' posterior belief about the their own face x_i and the other player's level τ_{-i} is $\mu_i^2(x_i, \tau_{-i} | 2, X)$ where

$$\begin{aligned} \mu_i^2(X, 0|2, X) &= \frac{\left(\frac{1}{2}\right) p_0 p}{\left(\frac{1}{2}\right) p_0 + p p_1}, & \mu_i^2(O, 0|2, X) &= \frac{\left(\frac{1}{2}\right) p_0 (1-p)}{\left(\frac{1}{2}\right) p_0 + p p_1}, \\ \mu_i^2(X, 1|2, X) &= \frac{p p_1}{\left(\frac{1}{2}\right) p_0 + p p_1}, & \mu_i^2(O, 1|2, X) &= 0. \end{aligned}$$

As the game proceeds beyond period 2, level 2 players will make the inference that if the other player is level 1, then their own face is dirty; otherwise, the other player's actions are uninformative about their face types. Moreover, at any period $2 \leq t \leq T$,

level 2 players' marginal belief about having a dirty face is

$$\mu_i^2(X|t, X) = \underbrace{\frac{\left(\frac{1}{2}\right)^{t-1} p_0 p}{\left(\frac{1}{2}\right)^{t-1} p_0 + p p_1}}_{= \mu_i^2(X,0|t,X)} + \underbrace{\frac{p p_1}{\left(\frac{1}{2}\right)^{t-1} p_0 + p p_1}}_{= \mu_i^2(X,1|t,X)} = \frac{p \left[\left(\frac{1}{2}\right)^{t-1} p_0 + p_1 \right]}{\left(\frac{1}{2}\right)^{t-1} p_0 + p p_1},$$

which is increasing in t , suggesting that level 2 players are more certain about having a dirty face in later periods. This is level 2 players' benefit of waiting when seeing a dirty face. However, the cost of waiting is that the other player may randomly end the game (if the other is level 0) and the payoff is discounted. Therefore, level 2 players' tradeoff is analogous to the sequential sampling problem of Wald (1947)—they decide the *optimal stopping period* to claim. The optimal stopping period depends on the parameters $\bar{\alpha}$ and δ , as well as the distribution of levels. This is in sharp contrast with the equilibrium prediction that the equilibrium prediction is independent of the parameters.

In particular, for any period $2 \leq t \leq T$, level 2 player i 's expected payoff of claiming to have a dirty face is

$$\mathbb{E}u_i^2(C|t) := \delta^{t-1} [\alpha \mu_i^2(X|t, X) - \mu_i^2(O|t, X)].$$

At period T , the last period of the game, it is optimal to claim if and only if

$$\mathbb{E}u_i^2(C|T) \geq 0 \iff \alpha \geq \frac{\mu_i^2(O|T, X)}{\mu_i^2(X|T, X)} \iff \bar{\alpha} \geq \frac{\left(\frac{1}{2}\right)^{T-1} p_0}{\left(\frac{1}{2}\right)^{T-1} p_0 + p_1}.$$

For any other period $2 \leq t' \leq T - 1$, it is optimal to claim at period t' only if

$$\mathbb{E}u_i^2(C|t') \geq \Pr(t' + 1|t', X) \mathbb{E}u_i^2(C|t' + 1),$$

where $\Pr(t' + 1|t', X)$ is level 2 player i 's belief about the probability that player $-i$ would wait in period t' .²¹ Rearranging the inequality yields the condition stated in Proposition 4.4. Furthermore, the proof in Appendix D shows that these conditions

²¹ $\Pr(t' + 1|t', X)$ is the probability that the other player chooses to wait in period t' , which is

$$\Pr(t' + 1|t', X) = 0.5 \cdot \mu_i^2(O|t', X) + 1 \cdot \mu_i^2(1|t', X) = \frac{\left(\frac{1}{2}\right)^{t'} p_0 + p p_1}{\left(\frac{1}{2}\right)^{t'-1} p_0 + p p_1}.$$

are not only necessary but also sufficient to pin down level 2 players' optimal stopping periods. By induction on the levels, it can be shown that no matter how sophisticated the players are, the behavior is characterized by the solution of a sequential sampling problem.

Proposition 4.4 characterizes the level-dependent behavioral strategies. Alternatively, the DCH solution can be characterized by the level-dependent stopping period (given observing x_{-i}), which is formally defined in Definition 4.2.

Definition 4.2 (Stopping Period). *For any sequential two-person dirty-faces game and its DCH level-dependent strategy profile σ , let $\hat{\sigma}_i^k(x_{-i})$ be level k player i 's earliest period to claim to have a dirty face conditional on observing x_{-i} for any $k \geq 1$ and $i \in N$. Specifically,*

$$\hat{\sigma}_i^k(x_{-i}) = \begin{cases} \min \{t' : \sigma_i^k(t', x_{-i}) = 1\}, & \text{if } \exists t \text{ s.t. } \sigma_i^k(t, x_{-i}) = 1 \\ T + 1, & \text{otherwise.} \end{cases}$$

With Definition 4.2, Corollary 4.1 is a direct consequence of Proposition 4.4. If $x_{-i} = O$, every strategic level of players will know their face is dirty and claim to have a dirty face in period 1, viz. $\hat{\sigma}_i^k(O) = 1$ for every $k \geq 1$. In contrast, if $x_{-i} = X$, Corollary 4.1 shows that the optimal stopping period is monotonically decreasing in k , implying that higher-level players tend to claim in fewer periods.

Corollary 4.1. *For any sequential two-person dirty-faces game, the DCH level-dependent strategy profile σ can be equivalently characterized by level-dependent stopping periods. Moreover, for any $i \in N$, we know*

1. $\hat{\sigma}_i^k(O) = 1$ for any $k \geq 1$,
2. $\hat{\sigma}_i^1(X) = T + 1$, and $\hat{\sigma}_i^k(X) \geq 2$ for all $k \geq 2$.
3. $\hat{\sigma}_i^k(X)$ is weakly decreasing in k .

Proof: See Appendix D. ■

To summarize, I illustrate the DCH optimal stopping periods of level 2 and level infinity players when seeing a dirty face, i.e., $\hat{\sigma}_i^2(X)$ and $\hat{\sigma}_i^\infty(X)$. Because the set of dirty-faces games is described by $(\delta, \bar{\alpha})$, it is simply the unit square on the $(\delta, \bar{\alpha})$ -plane. For the illustrative purpose, I consider $T = 5$ and the distribution of levels

follows Poisson(1.5), which is an empirically regular prior according to Camerer, Ho, and Chong (2004).

The DCH stopping periods can be solved according to Proposition 4.4 and are plotted in Figure 4.2. From the figure, we can find that DCH predicts it is possible for strategic players to choose any stopping period in $\{2, 3, 4, 5, 6\}$, depending on the parameters $\bar{\alpha}$ and δ . For instance, level 2 players will claim in period 2 (red area) if and only if

$$\bar{\alpha} \geq \frac{\left(\frac{1}{2} - \frac{1}{4}\delta\right) p_0}{\left(\frac{1}{2} - \frac{1}{4}\delta\right) p_0 + (1 - \delta)p_1} = \frac{\left(\frac{1}{2} - \frac{1}{4}\delta\right) e^{-1.5}}{\left(\frac{1}{2} - \frac{1}{4}\delta\right) e^{-1.5} + (1 - \delta)1.5e^{-1.5}} = \frac{2 - \delta}{8 - 7\delta}.$$

In addition, DCH predicts the comparative statics that the optimal stopping period is weakly decreasing in $\bar{\alpha}$ and weakly increasing in δ for any level $k \geq 2$. The intuition is that when $\bar{\alpha}$ is larger or δ is smaller, waiting becomes more costly, which causes the players to claim earlier with a less certain belief about their own face type.

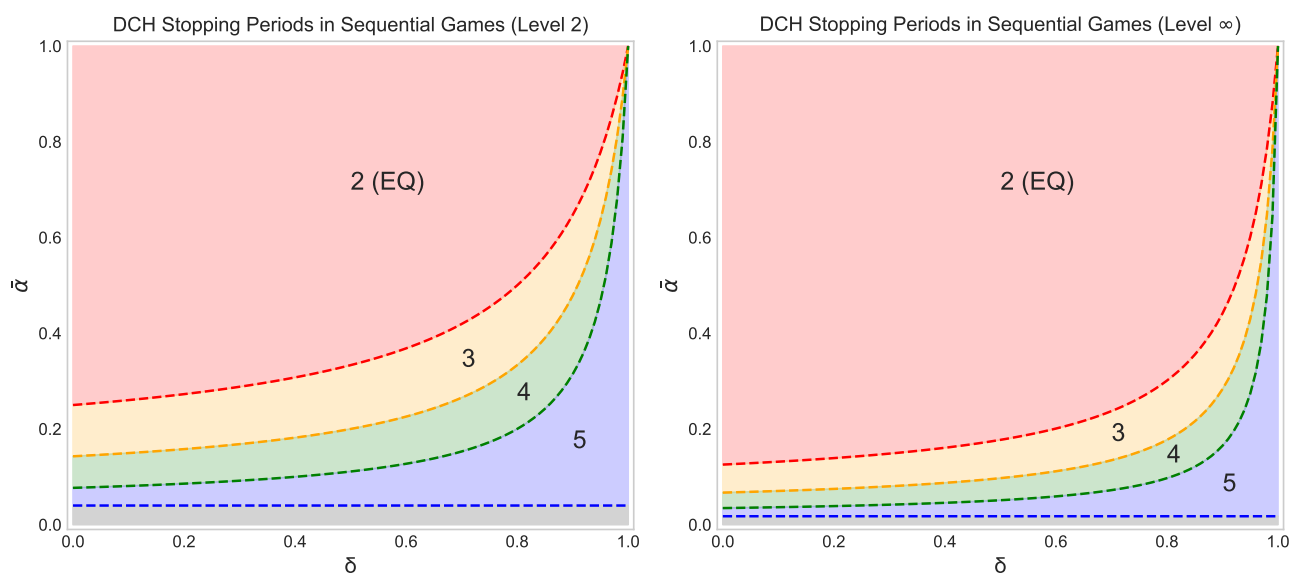


Figure 4.2: DCH stopping periods in sequential dirty-faces games for level 2 (left) and level ∞ players (right) as $x_{-i} = X$ where $T = 5$ and the distribution of levels follows Poisson(1.5).

4.5.2 DCH Solution for the Simultaneous Dirty-Faces Games

In contrast, the strategically equivalent simultaneous dirty-faces game is essentially a one-period game where players simultaneously choose an action from the set

$S = \{1, \dots, T + 1\}$. Action $t \leq T$ represents the plan to wait from period 1 to $t - 1$ and claim in period t . Action $T + 1$ is the plan to always wait. In the simultaneous dirty-faces game, a mixed strategy for player i is a mapping from the observed face type ($x_{-i} \in \{O, X\}$) to a probability distribution over the action set. The mixed strategy is denoted by

$$\tilde{\sigma}_i : \{O, X\} \rightarrow \Delta(S).$$

Suppose (s_i, s_{-i}) is the action profile. If $s_i \leq s_{-i}$, then the payoff for player i is computed as the case where player i claims in period s_i ; if $s_i > s_{-i}$, then player i 's payoff is 0.

The DCH solution for one-stage games coincides with the standard CH solution. For the sake of simplicity, I again assume that each player's level is i.i.d. drawn from the distribution $p = (p_k)_{k=0}^{\infty}$ where $p_k > 0$ for all k . Level 0 players will uniformly randomize, regardless of what they observe, so $\tilde{\sigma}_i^0(x_{-i}) = \frac{1}{T+1}$ for all i, x_{-i} . Since level $k \geq 1$ players will generically choose pure strategies, I slightly abuse the notation to use $\tilde{\sigma}_i^k(x_{-i})$ to denote the pure strategies.²²

Proposition 4.5 is parallel to Proposition 4.4 that characterizes the DCH solution for the simultaneous dirty-faces games. The intuition is similar to the analysis of sequential dirty-faces games. When observing a clean face, players can figure out their face types immediately. Hence, they will choose the strictly dominant strategy $\tilde{\sigma}_i^k(O) = 1$ for all $k \geq 1$. On the other hand, when observing a dirty face, players have to make hypothetical inferences about their face types and the other player's level of sophistication.

Proposition 4.5. *For any simultaneous two-person dirty-faces game, the level-dependent strategy profile of the DCH solution satisfies that for any $i \in N$,*

1. $\tilde{\sigma}_i^k(O) = 1$ for any $k \geq 1$.
2. $\tilde{\sigma}_i^1(X) = T + 1$. Moreover, for any $k \geq 2$,

$$(1) \tilde{\sigma}_i^k(X) \geq 2,$$

²²Specifically, for any $t \in \{1, 2, \dots, T, T + 1\}$, we use $\tilde{\sigma}_i^k(x_{-i}) = t$ to denote the degenerated distribution:

$$\tilde{\sigma}_i^k(x_{-i})(t) = 1, \quad \text{and} \quad \tilde{\sigma}_i^k(x_{-i})(t') = 0 \quad \forall t' \neq t.$$

(2) for any $2 \leq t \leq T - 1$, $\tilde{\sigma}_i^k(X) \leq t$ if and only if

$$\bar{\alpha} \geq \frac{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1} \delta \right] p_0}{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1} \delta \right] p_0 + (1 - \delta) \sum_{j=1}^{k-1} p_j},$$

(3) $\tilde{\sigma}_i^k(X) \leq T$ if and only if

$$\bar{\alpha} \geq \frac{\frac{2}{T+1} p_0}{\frac{2}{T+1} p_0 + \sum_{j=1}^{k-1} p_j}.$$

Proof: See Appendix D. ■

The characterization is similar to Proposition 4.4. When seeing a dirty face, strategic players will not choose 1, i.e., wait in period 1. Instead, they will claim in period t if and only if the reward $\bar{\alpha}$ is sufficiently high or the discount rate δ is low enough. However, the critical level of $\bar{\alpha}$ is different, indicating a violation of invariance under strategic equivalence. Although DCH makes different quantitative predictions in the sequential and simultaneous dirty-faces games, it makes a similar qualitative prediction that higher-level players tend to claim earlier than lower-level players. This is proven in Corollary 4.2.

Corollary 4.2. *For any simultaneous two-person dirty-faces game, the DCH level-dependent strategy profile $\tilde{\sigma}$ satisfies for any $i \in N$ and any $k \geq 2$, $\tilde{\sigma}_i^k(X)$ is weakly decreasing in k .*

Proof: See Appendix D. ■

To conclude, I illustrate the DCH strategies for the simultaneous dirty-faces games of level 2 and level infinity players when seeing a dirty face, i.e., $\tilde{\sigma}_i^2(X)$ and $\tilde{\sigma}_i^\infty(X)$, with $T = 5$ and the prior distribution of levels being Poisson(1.5). The DCH strategies can be solved by Proposition 4.5 and plotted in the unit square on the $(\delta, \bar{\alpha})$ -plane.

From Figure 4.3, we can observe that the DCH solution for the simultaneous games is similar to the DCH solution for the sequential games. In both games, DCH predicts that the stopping periods are weakly decreasing in $\bar{\alpha}$ and weakly increasing in δ for any level $k \geq 2$, although the boundaries of the areas are different.




Figure 4.3: DCH stopping periods in simultaneous dirty-faces games for level 2 (left) and level ∞ players (right) as $x_{-i} = X$ where $T = 5$, and the distribution of levels follows Poisson(1.5).

4.5.3 The Violation of Invariance under Strategic Equivalence

As discussed in the previous sections, DCH predicts that players might behave differently in two strategically equivalent dirty-faces games. This result is driven by the fact that when the cardinalities of the action sets differ, the behavioral strategies of level 0 players may not be outcome-equivalent. This initiates a chain reaction that affects the behavior of higher-level players because the DCH solution is solved recursively. In this subsection, I will examine *how* changes in the cardinalities of the action sets influence behavior in dirty-faces games.

In the sequential game, the cardinality of each player i 's action set (conditional on each x_{-i}) is 2^T , while in the simultaneous game, the cardinality is $T + 1$. This difference in cardinalities leads to different behavior among level 0 players. For example, in the first period, level 0 players will claim to have a dirty face with a probability of $1/2$ in the sequential game, while in the simultaneous game, they will claim with a probability of $1/(T + 1)$. Although this difference does not impact level 1 players, it significantly influences how level 2 (and more sophisticated) players update their beliefs regarding their own face types.

Remark 4.2. *The standard CH solutions for the sequential and simultaneous dirty-*

faces games coincide. Therefore, the difference of DCH between two versions of the game can alternatively be interpreted as the difference between the DCH and the standard CH on sequential games.

To characterize this distinction, we can begin by partitioning the set of dirty-faces games based on the stopping rules of each level in the sequential games, i.e., $\hat{\sigma}_i^k(X)$. Specifically, for any level $k \geq 1$, we define \mathcal{E}_t^k as the set of games in which level k players will claim *no later than* period t when they observe a dirty face.²³ That is,

$$(\delta, \bar{\alpha}) \in \mathcal{E}_t^k \iff \hat{\sigma}_i^k(X) \leq t \text{ under } (\delta, \bar{\alpha}).$$

The partition is visualized in Figure 4.2. For instance, \mathcal{E}_2^2 corresponds to the “2 (EQ)” area in the left panel.²⁴ Second, the set of dirty-faces games can be alternatively partitioned by the stopping rules of each level in the simultaneous games, i.e., $\tilde{\sigma}_i^k(X)$. For any $t \geq 1$ and $k \geq 1$, let \mathcal{S}_t^k be the set of dirty-faces games where $\tilde{\sigma}_i^k(X) \leq t$, which is illustrated in Figure 4.3. Proposition 4.6 compares the DCH solutions in different versions of the game with the set inclusions of \mathcal{E}_t^k and \mathcal{S}_t^k .

Proposition 4.6. *Consider any $T \geq 2$ and the set of two-person dirty-faces games. For any level $k \geq 2$, the following relationships hold.*

1. $\mathcal{S}_T^k \subset \mathcal{E}_T^k$.
2. $\mathcal{S}_t^k \subset \mathcal{E}_t^k$ for any $[\ln(T+1)/\ln 2] \leq t \leq T-1$.
3. *There is no set inclusion relationship between \mathcal{S}_t^k and \mathcal{E}_t^k for $2 \leq t < [\ln(T+1)/\ln 2]$. Moreover, for any $i \in N$, there exists $\bar{\delta}(T, t) \in (0, 1)$ such that $t = \hat{\sigma}_i^k(X) \leq \tilde{\sigma}_i^k(X)$ if $\delta \leq \bar{\delta}(T, t)$ and $\hat{\sigma}_i^k(X) \geq \tilde{\sigma}_i^k(X) = t$ if $\delta > \bar{\delta}(T, t)$. Specifically,*

$$\bar{\delta}(T, t) = \frac{(2^t - 2)(T + 1) - (t - 1)2^t}{(2^t - 1)(T + 1) - t2^t}.$$

Proof: See Appendix D. ■

²³Therefore, Corollary 4.1 implies for level 1 players, $\mathcal{E}_t^1 = \emptyset$ for all $t = 1, \dots, T$ and $\mathcal{E}_{T+1}^1 = (0, 1)^2$; for higher-level players, $\mathcal{E}_1^k = \emptyset$ for all $k \geq 1$.

²⁴Formally, when the distribution of levels follows Poisson(1.5), \mathcal{E}_2^2 is characterized by: $(\delta, \bar{\alpha}) \in \mathcal{E}_2^2 \iff (2 - \delta)/(8 - 7\delta) \leq \bar{\alpha} < 1$, and $0 < \delta < 1$.

Proposition 4.6 formally compares the DCH solutions of the sequential games and the simultaneous games. First, $\mathcal{S}_T^k \subset \mathcal{E}_T^k$ for any $k \geq 2$ implies that when seeing a dirty face, players are more likely to claim before the game ends in the sequential game than the simultaneous game. Yet this does not imply players will always claim earlier in the sequential games than the simultaneous games. The second and third results show that when the horizon is long enough and the players are sufficiently patient, i.e., $\delta > \bar{\delta}(T, t)$, it is possible for players to claim *later* in the sequential games. More surprisingly, the cutoff $\bar{\delta}(T, t)$ is independent of the level of sophistication and the prior distributions of levels, suggesting that the differences between the two versions have the same impact on each level of players.

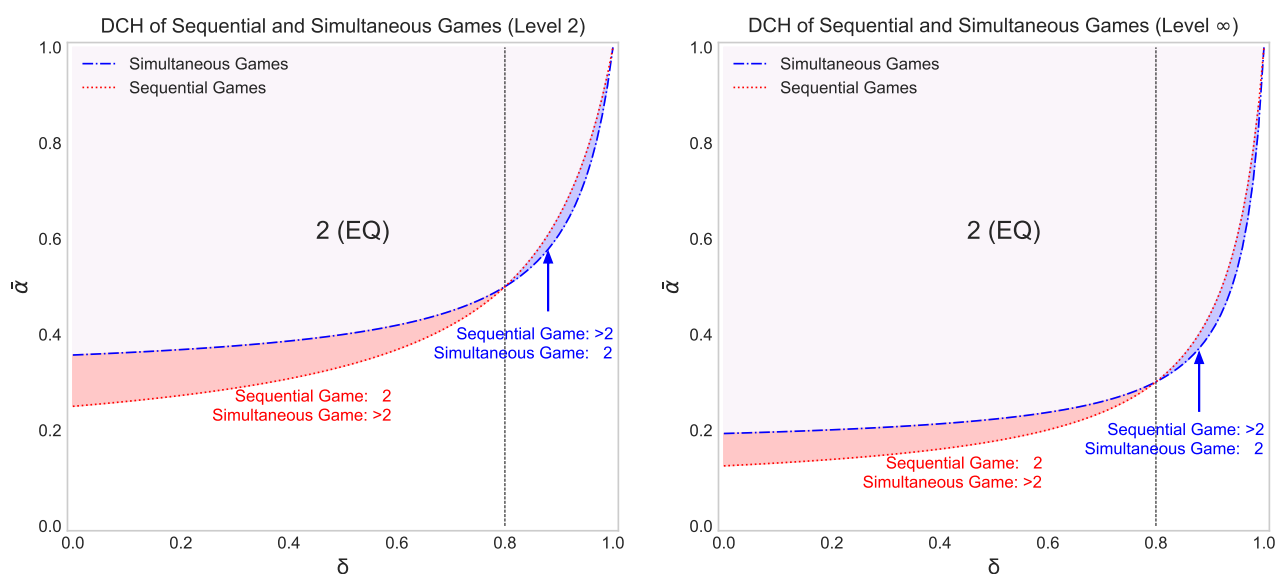


Figure 4.4: The set of dirty-faces games where $\tilde{\sigma}_i^k(X) = 2$ or $\hat{\sigma}_i^k(X) = 2$ for $k = 2$ (left) and $k = \infty$ (right) where $T = 5$ and the distribution of levels follows Poisson(1.5).

To illustrate this proposition, consider the case where $T = 5$ and the distribution of levels follows Poisson(1.5). By Proposition 4.6, we can find that $\mathcal{S}_t^k \subset \mathcal{E}_t^k$ for any $k \geq 2$ and $3 \leq t \leq 5$, while there is no set inclusion relation between \mathcal{S}_2^k and \mathcal{E}_2^k . These two sets for $k = 2$ and ∞ are plotted in Figure 4.4. By Proposition 4.4 and

Proposition 4.5, \mathcal{E}_2^2 and \mathcal{S}_2^2 can be characterized by

$$(\delta, \bar{\alpha}) \in \mathcal{E}_2^2 \iff \bar{\alpha} \geq \frac{\left(\frac{1}{2} - \frac{1}{4}\delta\right) e^{-1.5}}{\left(\frac{1}{2} - \frac{1}{4}\delta\right) e^{-1.5} + (1 - \delta)1.5e^{-1.5}} = \frac{2 - \delta}{8 - 7\delta},$$

$$(\delta, \bar{\alpha}) \in \mathcal{S}_2^2 \iff \bar{\alpha} \geq \frac{\left(\frac{5}{6} - \frac{2}{3}\delta\right) e^{-1.5}}{\left(\frac{5}{6} - \frac{2}{3}\delta\right) e^{-1.5} + (1 - \delta)1.5e^{-1.5}} = \frac{5 - 4\delta}{14 - 13\delta}.$$

The boundaries of \mathcal{E}_2^2 and \mathcal{S}_2^2 intersect at $\delta = 0.8$, suggesting that when $\delta < 0.8$, level 2 players tend to claim earlier in the sequential games, and vice versa. By similar calculations, it can be shown that the boundaries of \mathcal{E}_2^∞ and \mathcal{S}_2^∞ also intersect at $\delta = 0.8$, illustrating that the cutoff $\bar{\delta}(5, 2)$ is the same for every level.

Lastly, as the maximum horizon T increases, the cardinalities of the action sets in the sequential and simultaneous games become increasingly distinct. Consequently, the behavior of level 0 players diverges even further between the sequential and simultaneous games. As $T \rightarrow \infty$, Proposition 4.6 implies for any period $t \geq 2$ and level $k \geq 2$, \mathcal{S}_t^k and \mathcal{E}_t^k do not have set inclusion relationship, suggesting that higher-level players do not definitely learn their face types earlier in one game or another. Their behavior depends on the parameters $(\delta, \bar{\alpha})$. The result is formally presented in Corollary 4.3.

Corollary 4.3. *When $T \rightarrow \infty$, for any $t \geq 2$ and $k \geq 2$, there is no set inclusion relationship between \mathcal{S}_t^k and \mathcal{E}_t^k . Specifically, if $\delta < \bar{\delta}^*(t)$, then $t = \hat{\sigma}_i^k(X) \leq \tilde{\sigma}_i^k(X)$; and if $\delta > \bar{\delta}^*(t)$, then $\hat{\sigma}_i^k(X) \geq \tilde{\sigma}_i^k(X) = t$ where*

$$\bar{\delta}^*(t) = [2^t - 2]/[2^t - 1].$$

Proof: See Appendix D. ■

When there are more than two players, DCH predicts a bigger difference between the two versions in the sense that the boundaries between the sequential and simultaneous games are further apart. For the purpose of illustration, in Appendix D, I characterize the DCH solutions of three-person three-period games and find that players tend to learn their face types earlier in the sequential games.

4.6 Experimental Design, Hypotheses and Procedures

As demonstrated in the previous section, DCH makes various predictions about how people's behavior would vary with the timing (sequential vs. simultaneous) and the

payoff structures of the dirty-faces games. To test these predictions, I conduct a laboratory experiment on two-person dirty-faces games tailored to evaluate the DCH solution.

Specifically, the primary goal of the experiment is to measure the violation of invariance under strategic equivalence and understand how it interacts with the payoff structures. Furthermore, the variation of the payoff structures provides the opportunity to explore the sensitivity of behavior to payoffs in both the sequential and simultaneous versions of the game. Lastly, the stylized facts found in this experiment will help identify the strengths and the weaknesses of the DCH solution and alternative theories.

The theoretical analysis of DCH suggests that the main challenge in designing the experiment lies in the fact that the magnitude of the difference between the two versions depends on the payoff structure and the distribution of levels, which remains unknown before the experiment is run. To address this, I first estimate the distribution of levels using the dirty-faces game experimental data collected by Bayer and Chan (2007), and then choose the game parameters to maximize the diagnosticity based on the calibration results.

4.6.1 Calibration

The dirty-faces game experiment by Bayer and Chan (2007) is implemented under the direct response method with two treatments: two-person two-period games and three-person three-period games. In both treatments, the prior probability of having a dirty face is $2/3$, the discount factor δ is $4/5$, and the reward α is $1/4$.²⁵ I will focus on the data from two-person games because this environment is the closest to my experiment.²⁶ A detailed analysis of the data can be found in Appendix D.

There are 42 subjects (from two sessions) in the two-person treatment of Bayer and Chan (2007). At the beginning of the experiment, the computer randomly matches two subjects into a group. Subjects play 14 rounds of dirty-faces games against the same opponent, with the face types in each round being independently drawn according to the prior probabilities. In each round, an announcement is made on

²⁵In Bayer and Chan (2007), the payoff of correctly claiming a dirty face is 100 ECU (experimental currency unit) and the penalty of wrongly claiming a dirty face is -400 ECU. Therefore, the relative reward of correctly claiming a dirty face $\alpha = 1/4$ can be obtained by normalizing the payoffs.

²⁶Weber (2001)'s dataset consists of two experiments where experiment 2 is comparable with Bayer and Chan (2007)'s design. However, there are much fewer observations in this experiment than Bayer and Chan (2007) and there is no discount factor, making this dataset less ideal for the purpose of calibration.

the screen to both subjects if there is at least one person having a dirty face (type X). At the end of each round, subjects are told their own payoffs from that round and they are paid with the sum of the earnings of all 14 rounds.

In the calibration exercise, I exclude the data from the situation where there is no public announcement²⁷, resulting in 690 observations at the information set level. Following previous notations, I use (t, x_{-i}) to denote the situation where subject i sees type x_{-i} at period t . Table 4.1 reports the empirical frequency of choosing *claim* at each information set, revealing that the behavior is inconsistent with the prediction of standard equilibrium theory, particularly when observing a dirty face.

Following the literature on the cognitive hierarchy theory, I assume the prior distribution of levels follows a Poisson distribution. In the Poisson-DCH model, each individual i 's level is identically and independently drawn from $(p_k)_{k=0}^{\infty}$ where $p_k = e^{-\tau} \tau^k / k!$ for all $k \in \mathbb{N}_0$ and $\tau > 0$. Once the distribution of levels is specified, DCH makes a precise prediction about the aggregate choice frequency at each information set.²⁸ The rationale for estimating the Poisson-DCH model is to find τ , estimated using the maximum likelihood method, which minimizes the difference between the choice frequencies predicted by DCH and the empirical frequencies. See Appendix D for the details on the construction of the likelihood function.

It is worth remarking that since τ is the mean (and variance) of the Poisson distribution, the economic interpretation of τ is as the average level of sophistication among the population. Additionally, another property of the Poisson-DCH model is that as $\tau \rightarrow \infty$, the aggregate choice frequencies predicted by DCH converge to the equilibrium predictions. This provides a second interpretation of τ : the higher the value of τ , the closer the predictions are to the equilibrium. See Proposition D.1 in Appendix D for the proof.

Table 4.1 reports the estimation results of the Poisson-DCH model. Additionally, I estimate the standard Poisson-CH model by Camerer, Ho, and Chong (2004)

²⁷If there is no public announcement, it is common knowledge that both subjects' faces are clean.

²⁸The aggregate choice frequency can be constructed as follows. Consider any game, any player i , any information set \mathcal{I}_i , and any available action c_i at this information set. Let $P_k(c_i|\mathcal{I}_i)$ represent the probability of level k player i choosing c_i at \mathcal{I}_i . Additionally, let $f(k|\mathcal{I}_i, \tau)$ be the posterior distribution of levels at information set \mathcal{I}_i . The choice frequency predicted by DCH for action c_i at information set \mathcal{I}_i is the aggregation of choice probabilities from all levels, weighted by the proportion $f(k|\mathcal{I}_i, \tau)$:

$$\mathcal{D}(c_i|\mathcal{I}_i, \tau) \equiv \sum_{k=0}^{\infty} f(k|\mathcal{I}_i, \tau) P_k(c_i|\mathcal{I}_i, \tau).$$

Table 4.1: Estimation Results for Two-Person Dirty-Faces Games

	(t, x_{-i})	N	$\sigma_i^*(t, x_{-i})$	$\hat{\sigma}_i(t, x_{-i})$	DCH	Standard CH
$\sigma_i(t, x_{-i})$	(1, O)	123	1.000	0.943	0.859	0.791
	(2, O)	6	1.000	0.500	0.500	0.500
	(1, X)	391	0.000	0.210	0.141	0.104
	(2, X)	170	1.000	0.618	0.503	0.477
Parameter	τ				1.269	1.161
	S.E.				(0.090)	(0.095)
Fitness	LL				-360.75	-381.46
	AIC				723.50	764.91
	BIC				728.04	769.45
Vuong Test						6.517
p-value						< 0.001

Note: The equilibrium and the empirical frequencies of C at each information set are denoted as σ_i^* and $\hat{\sigma}_i$, respectively. There are 294 games (rounds \times groups).

as a benchmark.²⁹ Comparing the fitness of these models, I find that the log-likelihood of DCH is significantly higher than standard CH (Vuong test p-value < 0.001), suggesting that DCH outperforms the standard CH in capturing the empirical pattern. Besides, the estimated τ of Poisson-DCH falls within the range of commonly observed τ in various environments, with a value of 1.269. In the following, I will design the experiment by treating Poisson(1.269) as the true prior distribution of levels.

4.6.2 Games and Hypotheses

In this experiment, I employ a between-subject design where each participant is assigned to either the “sequential treatment” (using the direct-response method) or the “simultaneous treatment” (using the strategy method). To observe potential heterogeneity in stopping periods, I set the maximum length to be $T = 5$ for both treatments.

Assessing whether the difference between the two treatments is challenging because level 1 players—the most common types of players according to the calibration

²⁹The logit-AQRE proposed by McKelvey and Palfrey (1998) is also estimated. The likelihood scores between Poisson-DCH and logit-AQRE are not significantly different. See Appendix D for the details.

result—will behave the same under the two treatments. When observing a dirty face, they will always wait in both the sequential and simultaneous games. Therefore, to diagnose the predictivity of DCH, the game parameters are chosen to make level 2 players behave differently under different representations. The behavioral change of level 2 players (around 22.6% based on the calibration) is anticipated to yield a sizable treatment effect.




Figure 4.5: (Left) The set of dirty-faces games where at information set $(2, X)$, level 2 players behave differently in the two versions when $T = 5$ and the distribution of levels follows Poisson(1.269). (Right) The two diagnostic games and the four control games in the experiment.

According to Proposition 4.6, DCH predicts the existence of a set of dirty-faces games in which level 2 players exhibit different behavior in the sequential and simultaneous games at information set $(2, X)$. The left panel of Figure 4.5 illustrates this set of games when the distribution of levels follows Poisson(1.269). From this figure, we can observe the following:

- (1) For $\delta < 0.8$, there is a range of games (red area) where level 2 players choose to claim at $(2, X)$ in the sequential games but not in the simultaneous games.
- (2) For $\delta = 0.8$, level 2 (and more sophisticated) players behave the same in the sequential and simultaneous games.

- (3) For $\delta > 0.8$, there is a range of games (blue area) where level 2 players choose to claim at $(2, X)$ in the simultaneous games but not in the sequential games.

Guided by DCH, I consider the following six dirty-faces games $(\delta, \bar{\alpha})$ as depicted in the right panel of Figure 4.5.

The set of games consists of two *diagnostic games* where $(\delta, \bar{\alpha}) = (0.6, 0.45)$ and $(0.95, 0.8)$ and four *control games* where $(\delta, \bar{\alpha}) = (0.6, 0.8)$, $(0.8, 0.45)$, $(0.8, 0.8)$ and $(0.95, 0.45)$. DCH predicts in the diagnostic games, level 2 players will behave differently in two treatments, but not in the control games. This variation allows us to examine the interplay between the violation of invariance under strategic equivalence varies with the payoff structures.

DCH makes several predictions about the comparative statics. First, by Proposition 4.4 and 4.5, DCH predicts that no matter in sequential games or in simultaneous games, when observing a dirty face, players will choose to claim *earlier* when δ is smaller or $\bar{\alpha}$ is higher. An implication is that in both treatments, at information set $(2, X)$, players are more likely to claim when δ decreases or $\bar{\alpha}$ increases.

Hypothesis 4.1. *In both the sequential and simultaneous treatments, at information set $(2, X)$, the empirical frequency of choosing C is higher when δ decreases or $\bar{\alpha}$ increases.*

Besides, DCH makes a specific prediction regarding the relative magnitude of the treatment effect among these six games. First, in the DCH solution, part of the treatment effect is attributed to the difference in level 0 players' strategies between the two treatments. In the sequential games, level 0 players uniformly randomize at every information set, resulting in a conditional probability to claim at $(2, X)$ is $1/2$. Yet in the simultaneous games, level 0 players uniformly randomize across all reduced contingent strategies, leading to a conditional probability to claim at $(2, X)$ is $1/5$. In other words, the difference in level 0 players' strategies generates a mechanical effect that increases the likelihood of players choosing to claim at $(2, X)$ in the sequential games. Because in all four control games, strategic players behave the same at $(2, X)$ under two representations, DCH predicts the magnitude of the violation of invariance under strategic equivalence will be similar in the control games. Particularly, in the game $(\delta, \bar{\alpha})$, the treatment effect can be quantified by computing the difference between the conditional probabilities of choosing to claim

at $(2, X)$ in the sequential version and the simultaneous version. This difference is denoted by $\Delta(\delta, \bar{\alpha})$.³⁰

Second, in the game where $(\delta, \bar{\alpha}) = (0.6, 0.45)$, level 2 players will claim at $(2, X)$ in the sequential version but not in the simultaneous version. As a result, DCH predicts that the difference between the two treatments in this diagnostic game will be stronger compared to the effect observed in the control games. On the contrary, in the game where $(\delta, \bar{\alpha}) = (0.95, 0.8)$, level 2 players will claim at $(2, X)$ in the simultaneous version but not in the sequential version. This offsets the mechanical effect caused by level 0 players. The expected differences based on the calibration results are summarized below.

Hypothesis 4.2. *Based on the calibration results, the expected differences are:*

$$\Delta(0.6, 0.45) > \Delta(0.6, 0.8) = \Delta(0.8, 0.8) \approx \Delta(0.8, 0.45) \approx \Delta(0.9, 0.45) > \Delta(0.95, 0.8).$$

31.15%	7.4%	7.4%	4.82%	3.26%	-18.93%

4.6.3 Experimental Procedures

The experimental sessions were conducted at the Experimental Social Science Laboratory (ESSL) located on the campus of the University of California, Irvine. Subjects were recruited from the general undergraduate population, from all majors. Experiments were conducted through oTree software (D. L. Chen, Schonger, and Wickens, 2016). I conducted 10 sessions with a total of 118 subjects. No subject participated in more than one session. Each session lasted around 45 minutes, and the average earnings was \$33.36, including the \$10 show-up fee (max \$52 and min \$10).

Subjects were given instructions at the beginning and the instructions were read aloud. Subjects were allowed to ask any questions during the whole instruction process. The questions were answered so that every one can hear. Afterwards, they had to answer several comprehension questions on the computer screen in order to proceed. The instructions for both the sequential treatment and the simultaneous treatments are identical except for the instructions about the choices and the feedback after each game. The instructions for both treatments can be found in Appendix D.

³⁰In the sequential version, the observed conditional probability of claiming at $(2, X)$ is simply the empirical $\sigma_i(2, X)$. For the simultaneous version, the conditional probability can be computed from the empirical $\tilde{\sigma}_i(X)$ by $\tilde{\sigma}_i(2, X) \equiv \Pr(\tilde{\sigma}_i(X) = 2) / \sum_{t=2}^6 \Pr(\tilde{\sigma}_i(X) = t)$. Therefore, the treatment effect is quantified by the (empirical) difference between $\sigma_i(2, X)$ and $\tilde{\sigma}_i(2, X)$, i.e., $\Delta \equiv \sigma_i(2, X) - \tilde{\sigma}_i(2, X)$.

Table 4.2: List of Game Parameters Implemented in the Experiment

	Game Parameters				Normalized Probabilities	
	δ	α	p	$\bar{\alpha}$	one X, one O	two X
Diagnostic Game 1	0.60	0.225	0.67	0.45	0.25	0.50
Diagnostic Game 1'	0.60	0.150	0.75	0.45	0.20	0.60
Diagnostic Game 2	0.95	0.400	0.67	0.80	0.25	0.50
Diagnostic Game 2'	0.95	0.267	0.75	0.80	0.20	0.60
Control Game 1	0.60	0.400	0.67	0.80	0.25	0.50
Control Game 1'	0.60	0.267	0.75	0.80	0.20	0.60
Control Game 2	0.80	0.225	0.67	0.45	0.25	0.50
Control Game 2'	0.80	0.150	0.75	0.45	0.20	0.60
Control Game 3	0.80	0.400	0.67	0.80	0.25	0.50
Control Game 3'	0.80	0.267	0.75	0.80	0.20	0.60
Control Game 4	0.95	0.225	0.67	0.45	0.25	0.50
Control Game 4'	0.95	0.150	0.75	0.45	0.20	0.60

Each session comprised 12 games with different (δ, α, p) configurations, as summarized in Table 4.2.³¹ The sequence of these games was randomized, and in each game, subjects were randomly paired into groups. The draws for player types were independent, and the protocol was common knowledge. To prevent any framing effect, the “dirty face” and the “clean face” were labelled as “red” and “white” in the instruction, respectively. Besides, the actions were labelled as “I’m red” and “wait.” Finally, to avoid situations where both faces are clean, the probabilities were normalized to ensure that having two clean faces was impossible.³²

After observing the other’s face type, subjects were asked to simultaneously choose their actions. In the sequential treatment, subjects simultaneously chose either “I’m red” or “wait.” If both subjects chose to wait, the game would proceed to the next period and they were asked to choose again. The game ended after the period where some one chose “I’m red” or after period 5. On the other hand, in the simultaneous treatment, subjects simultaneously chose one of the six plans (the period to choose “I’m red” or always wait) and the plans were implemented by the computer. At the end of each match, the subjects were informed of their own payoffs, the true types

³¹Notice that the parameters are selected such that each $(\delta, \bar{\alpha})$ is played twice.

³²For example, in the game with $p = 2/3$, subjects were informed that the probability of one dirty face and one clean face was $1/4$, and the probability of two dirty faces was $1/2$. Therefore, if the other’s face was clean, the subject could infer that his own face was dirty. Conversely, if the other’s face was dirty, the subject’s belief about his own face being dirty was $2/3$.

and the histories of the game.³³

Lastly, subjects were paid in cash based on their total points earned from the 12 games. The highest possible earnings of each game was 100 points.³⁴ The conversion rate was two US dollars for every 100 points. Following previous dirty face game experiments (Weber, 2001; Bayer and Chan, 2007), each subject was provided an endowment of 900 points at the beginning of the experiment to prevent early bankruptcy, and they would only receive the show-up fee if the total point is negative.

4.7 Experimental Results

4.7.1 Aggregate-Level Analysis

The data includes two treatments (sequential and simultaneous) with 60 subjects in the sequential treatment and 58 subjects in the simultaneous treatment. Each subject participates in 12 games, resulting in 1024 observations for the sequential treatment and 1979 observations for the simultaneous treatment at the information set level.³⁵ Figure 4.6 provides a comprehensive overview of the data by plotting the distribution of stopping periods in both treatments, aggregating across all payoff configurations. This analysis considers scenarios where players encounter either a clean face or a dirty face.³⁶

A few key observations emerge from this figure. First, when players see a clean face, their behavior is consistent across both treatments. A majority of players seem to understand that their face is dirty and claim in period 1. Second, when players encounter a dirty face, it is evident that the distribution of stopping periods in the simultaneous treatment first-order stochastically dominates the distribution in the sequential treatment, implying that players are more inclined to claim earlier in the sequential treatment. Furthermore, a striking pattern in the right panel of Figure

³³To control for the amount of feedback in both treatments, in the simultaneous treatment, subjects would learn the other's exact plan if the other chose "I'm red" earlier or at the same period; otherwise, they would be told that the other subject was *later than you*.

³⁴That is, a correct claim in the first period would yield 100 points, while an incorrect claim in the first period would result in a penalty of $100/\alpha$ points.

³⁵For the simultaneous treatment, the choice data at the information set level are implied by the contingent strategies. For instance, choosing the contingent strategy "claim at period 4" implies that the subject will wait from period 1 to period 3 and claim in period 4.

³⁶In the sequential treatment, the cumulative density of stopping periods is derived from the choice probability at each information set. For example, the probability of stopping in period 1 corresponds to the empirical frequency of choosing *C*. Similarly, the probability of stopping in period 2 is the product of the empirical frequency of choosing *W* in period 1 and the empirical frequency of choosing *C* in period 2. The probabilities for other stopping periods are calculated in a similar manner.

4.6 is the prevalence of the “always wait” strategy in the simultaneous treatment, chosen by approximately 36.72% of participants. This is in stark contrast to the sequential treatment, where the proportion of participants employing the “always wait” strategy is only about 13.88%.




Figure 4.6: The CDFs of the stopping periods when players see either a clean face (left panel) or a dirty face (right panel). The blue solid and the red dashed are the distributions in the sequential and the simultaneous treatments, respectively.

Focusing on data from the first two periods, Figure 4.7 displays the empirical frequencies of choosing C at each information set during the first two periods. From the figure, it’s evident that at information set $(1, O)$, the behavior in the sequential treatment is not significantly different from the simultaneous treatment (Ranksum test p -value = 0.4423). In both treatments, the frequencies of C exceed 80%, indicating that the majority of subjects understand that choosing C in the first period is a strictly dominant strategy.

Despite the limited number of observations at information set $(2, O)$, it provides valuable insights into the behavioral strategies of level 0 players. This is because, from the perspective of DCH, information set $(2, O)$ is reached *only* when a player is level 0. As depicted in Figure 4.7, the frequencies of C in the sequential and simultaneous treatments are 43.8% and 25%, respectively.³⁷ These results align with DCH, which predicts that the frequencies of C in the sequential and simultaneous treatments should be 50% and 20%, respectively. This suggests that uniform randomization is a reasonable specification for level 0 players’ behavior.

³⁷A similar pattern is also found in Bayer and Chan (2007). In their dataset, the frequency of choosing C at information set $(2, O)$ is exactly 50%, which coincides with the prediction of DCH.

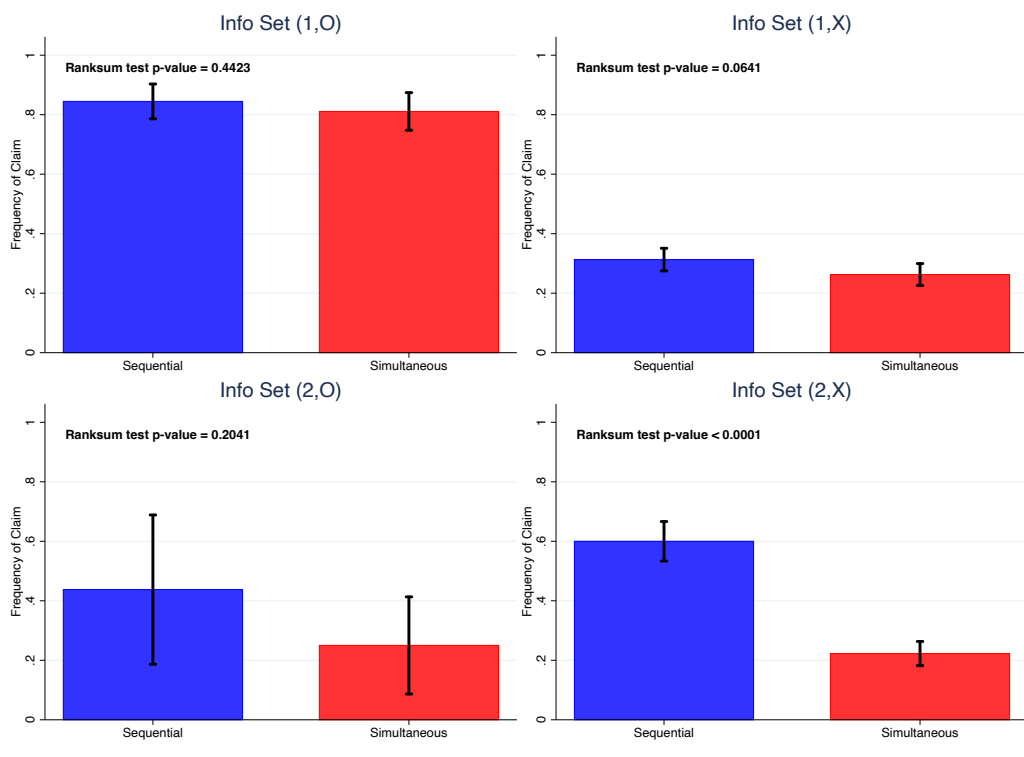


Figure 4.7: The empirical frequencies of C and 95% CI at each information set in period 1 and 2, aggregating across all configurations. Each panel represents an information set. The blue bars are the frequencies of the sequential treatment and the red bars are the frequencies of the simultaneous treatment.

On the other hand, when players see a dirty face, they need to make inferences about their own faces either from the opponent's actions or hypothetically. However, regardless of the treatment, claiming in period 1 is strictly dominated. Comparing the empirical frequencies of choosing C at information set $(1, X)$, we find that players in the simultaneous treatment are less likely to choose C (Ranksum test p-value = 0.0641). This observation is consistent with DCH, as level 0 players in the simultaneous treatment are less likely to claim at information set $(1, X)$. Furthermore, a strong treatment effect is detected at period 2. The frequency of choosing C at information set $(2, X)$ in the sequential treatment is 60.00%, while the frequency in the simultaneous treatment is 22.28% (Ranksum test p-value < 0.0001).

Result 1. (1) When observing a clean face in both treatments, over 80% of the subjects choose C in period 1, the strictly dominant strategy. Additionally, the behavior at information set $(2, O)$ aligns with the prediction of DCH about level 0

players' behavior. (2) When players observe a dirty face, a significant difference emerges: they are more likely to claim at information set $(2, X)$ in the sequential treatment. Furthermore, the most prevalent strategy in the simultaneous treatment when players see a dirty face is to "always wait."

The supplementary analysis can be found in Appendix D. In the following, I will focus on information set $(2, X)$, where a strong treatment effect is found, and I will test two hypotheses related to the sensitivity of behavior to the payoff structures and the interplay between the payoff structures and the magnitude of the effect.

4.7.2 The Payoff Effect

Focusing on information set $(2, X)$, DCH predicts that in both treatments, players' behavior is sensitive to the payoff configurations. Specifically, DCH predicts that the empirical frequencies of choosing C at information set $(2, X)$ will exhibit a monotonic relationship with δ and $\bar{\alpha}$. To test this prediction, I perform Kruskal-Wallis ranksum tests on the sequential and simultaneous treatments separately.




Figure 4.8: The empirical frequencies of C and 95% CI at information set $(2, X)$ in each payoff configuration $(\delta, \bar{\alpha})$. The data from the sequential and the simultaneous treatments are plotted in the left and the right panel, respectively.

In the sequential treatment, we find that the null hypothesis is marginally rejected ($\chi^2(5) = 9.856$, p -value = 0.0794), suggesting that behavior is influenced by variations in payoff structures. Furthermore, we can observe from the left panel of Figure 4.8 that the frequency of choosing C weakly increases with $\bar{\alpha}$ for any δ . This monotonic pattern aligns with the prediction of DCH.

Similarly, the null hypothesis is rejected for the simultaneous treatment ($\chi^2(5) = 11.831$, $p\text{-value} = 0.0372$), indicating that behavior in the simultaneous treatment significantly varies with the payoff parameters. Once again, we can observe from Figure 4.8 that for each δ , the frequency of choosing C weakly increases with $\bar{\alpha}$, aligning with DCH.

Result 2. *The behavior at information set $(2, X)$ in both treatments significantly varies with payoffs, aligning with the qualitative predictions of DCH.*

4.7.3 The Violation of Invariance under Strategic Equivalence

The behavior in both treatments significantly varies with the payoff structures. Additionally, the difference in behavior between the two treatments also varies with the payoff structures. This variability allows us to examine the predictions of DCH.

First, the left panel of Figure 4.9 displays the joint distribution of the empirical frequencies of choosing C at $(2, X)$ between the two treatments, where each point represents one payoff configuration. From the figure, we can observe that all six points are below the 45-degree line, implying that players are more likely to claim at $(2, X)$ in the sequential treatment than in the simultaneous treatment, regardless of the payoff configuration.




Figure 4.9: (Left) The empirical frequencies of C and 95% CI at information set $(2, X)$ in each payoff configuration $(\delta, \bar{\alpha})$. (Right) The difference in the frequencies of C at information set $(2, X)$ between the sequential and simultaneous treatments for each payoff configuration $(\delta, \bar{\alpha})$ with the 95% CIs. The standard errors are clustered at the session level.

Second, for each payoff configuration $(\delta, \bar{\alpha})$, I calculate $\Delta(\delta, \bar{\alpha})$, which represents the difference in the empirical frequencies of choosing C at $(2, X)$ between both treatments. The results are shown in the right panel of Figure 4.9. Focusing on the diagnostic game where $(\delta, \bar{\alpha}) = (0.6, 0.45)$, we observe a significant treatment effect with a magnitude of $\Delta(0.6, 0.45) = 30.77\%$ (95% CI = [15.90%, 45.64%], p-value = 0.001). This is highly consistent with the prediction of the calibrated DCH.

Furthermore, when focusing on the four control games, we find that the magnitudes of the treatment effects are similar in these games, which aligns with the qualitative predictions of DCH. However, these magnitudes are much stronger than the predictions of calibrated DCH, ranging from $\Delta(0.95, 0.45) = 30.26\%$ (95% CI = [13.85%, 46.68%], p-value = 0.002) to $\Delta(0.6, 0.8) = 42.88\%$ (95% CI = [6.73%, 79.03%], p-value = 0.025). Lastly, in the second diagnostic game with $(\delta, \bar{\alpha}) = (0.95, 0.8)$, DCH predicts a negative treatment effect based on the calibration results. However, the observed empirical difference for this game is $\Delta(0.95, 0.8) = 42.94\%$ (95% CI = [18.17%, 67.71%] and p-value = 0.004), which is inconsistent with the quantitative prediction of the calibrated DCH.

Result 3. *In the diagnostic game with $(\delta, \bar{\alpha}) = (0.6, 0.45)$, the frequency of C at information set $(2, X)$ is 30.77% higher in the sequential treatment compared to the simultaneous treatment. In the diagnostic game with $(\delta, \bar{\alpha}) = (0.95, 0.8)$, the difference is 42.94%. Furthermore, treatment effects are detected in all control games, with magnitudes exceeding the predictions of calibrated DCH.*

In summary, when analyzing the interplay between the violation of strategic equivalence and payoff structures, we observe that while calibrated DCH captures some qualitative patterns, the observed magnitudes are significantly larger. This suggests that the observed behavior might result from both the violation of mutual consistency and other behavioral biases. To delve deeper into this aspect, in the next subsection, I will compare DCH with other behavioral models that relax other requirements of the standard equilibrium theory.

4.7.4 Structural Estimation and Model Comparison

DCH relaxes the requirement of mutual consistency in sequential equilibrium while still adhering to the requirements of best response and Bayesian inference. Can the empirical pattern be better explained by relaxing other requirements? To assess

the relaxation of two other requirements, I estimate the “Quantal Cursed Sequential Equilibrium (QCSE),”³⁸ which is a hybrid model combining AQRE and CSE, thereby relaxing the requirements of best response and Bayesian inference.

QCSE assumes that players are unable to fully understand how other players’ actions depend on their private information.³⁹ In particular, for any strategy profile σ , any player i and any information set $\mathcal{I}_i = (h^{t-1}, \theta_i)$, the *average behavioral strategy of player $-i$* is

$$\bar{\sigma}_{-i}(a_{-i}^t | h^{t-1}, \theta_i) = \sum_{\theta'_{-i}} \mu_i(\theta'_{-i} | h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^{t+1} | h^{t-1}, \theta'_{-i}).$$

In QCSE, there is a parameter $\chi \in [0, 1]$. For any χ , χ -cursed player i believes the other players are playing the behavioral strategy:

$$\sigma_{-i}^\chi(a_{-i}^t | h^{t-1}, \theta) = \chi \bar{\sigma}_{-i}(a_{-i}^t | h^{t-1}, \theta_i) + (1 - \chi) \sigma_{-i}(a_{-i}^t | h^{t-1}, \theta_{-i}),$$

which is a linear combination between the average behavioral strategy (with χ weight) and the true behavioral strategy (with $1 - \chi$ weight). When $\chi = 0$, players have correct perceptions about others’ behavioral strategies. On the other extreme, when $\chi = 1$, players fail to understand the correlation between others’ actions and types. As the game progresses, players update their beliefs via Bayes’ rule, believing that other players are using σ_{-i}^χ instead of the true behavioral strategy σ_{-i} . As shown by Fong, Lin, and Palfrey (2023b), at any history $h^t = (h^{t-1}, a^t)$, player i ’s χ -cursed belief is

$$\mu_i^\chi(\theta_{-i} | h^t, \theta_i) = \chi \mu_i^\chi(\theta_{-i} | h^{t-1}, \theta_i) + (1 - \chi) \left[\frac{\mu_i^\chi(\theta_{-i} | h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^t | h^{t-1}, \theta_{-i})}{\sum_{\theta'_{-i}} \mu_i^\chi(\theta'_{-i} | h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^t | h^{t-1}, \theta'_{-i})} \right],$$

which is a linear combination between the belief from the previous period (with χ weight) and the Bayesian belief (with $1 - \chi$ weight).

Moreover, in QCSE, players make quantal responses rather than best responses. In particular, players make make logit quantal responses, and the precision is determined by a parameter $\lambda \in [0, \infty)$. Consider any information set \mathcal{I}_i . For any $a_i \in A_i(\mathcal{I}_i)$, let \bar{u}_{a_i} denote the continuation value of a_i in QCSE. The choice probability of a_i is given by a multinomial logit distribution:

$$\sigma_i(a_i | \mathcal{I}_i) = \frac{e^{\lambda \bar{u}_{a_i}}}{\sum_{a' \in A_i(\mathcal{I}_i)} e^{\lambda \bar{u}_{a'}}}.$$

³⁸See Appendix D for a detailed description of the model.

³⁹In the context of the dirty-faces game, QCSE assumes that players do not fully recognize how the other player’s actions depend on the observed face type.

When $\lambda = 0$, players become insensitive to the payoffs, behaving like level 0 players. As λ increases, players' behavior becomes more sensitive to the payoffs. In the limit as $\lambda \rightarrow \infty$, players become fully rational and make best responses. In summary, QCSE relaxes the requirements of best response and Bayesian inferences with two parameters, $\lambda \in [0, \infty)$ and $\chi \in [0, 1]$.

Remark 4.3. *When $\chi = 0$, QCSE reduces to AQRE, and as $\lambda \rightarrow \infty$, it reduces to CSE.*

To enable a fair comparison between DCH and QCSE, I estimate a Quantal DCH model (QDCH) where the prior distribution of levels follows Poisson(τ), and all levels ($k \geq 1$) of players make logit quantal responses instead of best responses. In essence, Quantal DCH relaxes the requirements of best response and mutual consistency with two parameters, $\lambda \in [0, \infty)$ and $\tau \in [0, \infty)$. A description of QDCH can be found in Appendix D.

Remark 4.4. *When $\lambda \rightarrow \infty$, QDCH reduces to DCH.*

In addition to QDCH and QCSE, I also estimate DCH and AQRE, which are nested within QDCH and QCSE, respectively.⁴⁰ These models are estimated using maximum likelihood estimation, and the construction of the likelihood functions can be found in Appendix D. Table 4.3 presents the estimation results for both the sequential and simultaneous treatments.⁴¹ The comparison between the models is summarized in Figure 4.10.

Comparing these four models, we first observe that in both the sequential and simultaneous treatments, QDCH fits the data significantly better than QCSE (Sequential: Vuong Test p-value = 0.0056; Simultaneous: Vuong Test p-value < 0.0001). Without relaxing the best response requirement, DCH's fitness is significantly better than QCSE in the simultaneous treatment (Vuong Test p-value < 0.0001). However, in the sequential treatment, QCSE fits the data significantly better than DCH (Vuong Test p-value = 0.0245).

QDCH outperforms other models in both treatments, indicating that the observed violation of strategic equivalence is primarily due to the relaxation of mutual consistency. However, there is evidence of the violation of other behavioral biases. In

⁴⁰CSE cannot be estimated independently as it lacks an error structure in the model.

⁴¹In the simultaneous treatment, due to the flatness of the log-likelihood functions for both QDCH and QCSE at the MLE estimates, the square roots of the inverse Hessian matrices are not well-defined.

Table 4.3: Estimation Results for the Sequential and Simultaneous Treatments

		Sequential Treatment				Simultaneous Treatment			
		QDCH	DCH	QCSE	AQRE	QDCH	DCH	QCSE	AQRE
Parameters	λ	12.371		5.672	5.484	1774.5		8.740	3.839
	S.E.	(2.062)		(0.821)	(0.426)	—		—	(0.452)
	τ	1.309	0.277			0.388	0.389		
	S.E.	(0.220)	(0.043)			—	(0.015)		
	χ			0.101				1.000	
	S.E.			(0.364)				—	
Fitness	LL	-634.72	-671.78	-648.36	-648.40	-1100.76	-1100.76	-1167.68	-1211.08
	AIC	1273.43	1345.57	1300.73	1298.80	2205.51	2203.52	2339.37	2424.16
	BIC	1283.29	1350.50	1310.59	1303.73	2214.61	2208.07	2348.46	2428.70

the sequential treatment, a significant quantal response effect is observed (QDCH vs. DCH: Likelihood Ratio Test p-value < 0.0001), but not in the simultaneous treatment (Likelihood Ratio Test p-value = 0.9340). Furthermore, in the simultaneous treatment, a significant cursed effect is detected ($\hat{\chi} = 1.000$, p-value < 0.0001), but not in the sequential treatment ($\hat{\chi} = 0.101$, p-value = 0.7846). This suggests that, players struggle to accurately understand how other players' actions depend on their private information and update their beliefs accordingly in the simultaneous treatment, but not in the sequential treatment.

Lastly, it's worth noting that DCH estimates a significantly lower $\hat{\tau}$ in the sequential treatment compared to the simultaneous treatment (Sequential: $\hat{\tau} = 0.277$; Simultaneous: $\hat{\tau} = 0.389$). In contrast, when introducing quantal responses into DCH, we observe a significantly higher $\hat{\tau}$ in the sequential treatment compared to the simultaneous treatment (Sequential: $\hat{\tau} = 1.309$; Simultaneous: $\hat{\tau} = 0.389$). This suggests that in the simultaneous treatment, all of the randomness can be attributed to level 0 behavior, whereas in the sequential treatment, some randomness is attributable to the mistakes of higher-level players.

Result 4. (1) In both the sequential and simultaneous treatments, QDCH outperforms QCSE in explaining the data. Additionally, in the sequential treatment, the fitness of DCH is not significantly different from QCSE and AQRE. In the simultaneous treatment, DCH significantly outperforms QCSE and AQRE. (2) In both treatments, there is evidence of the failure of Bayesian inferences. Additionally, in the sequential treatment, evidence of quantal responses is present, while it is not observed in the simultaneous treatment.

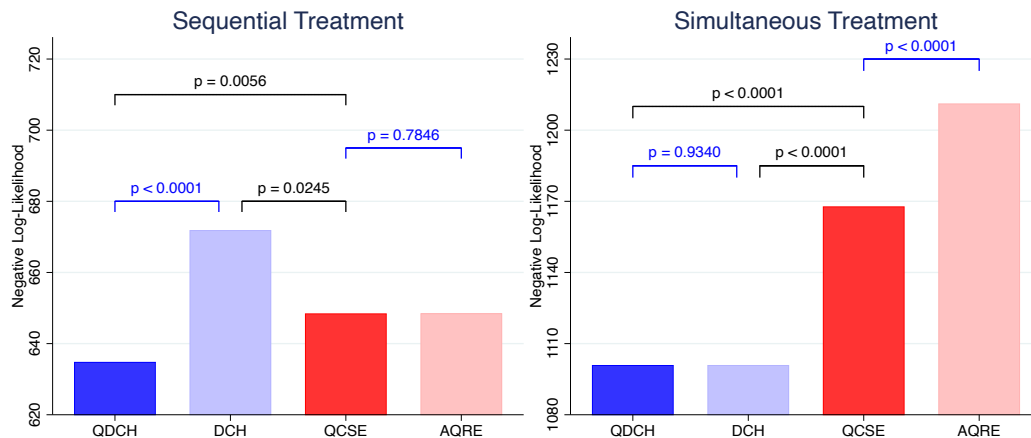


Figure 4.10: Negative log-likelihood of each model. The likelihood ratio test is performed when comparing two nested models, while the Vuong test is performed when comparing QDCH and DCH with QCSE.

4.7.5 The Analysis of Reaction Times

Besides the choice data, it is also interesting to see how long it takes individuals to make decisions. There is evidence suggesting that people tend to take an action faster if they adopt some simple decision-making heuristics or have strong preferences over the action (see, for example, Rubinstein, 2007; Chabris et al., 2009; Konovalov and Krajbich, 2019; Lin, Brown, et al., 2020 and Gill and Prowse, 2023).

Focusing on the case where players observe a dirty face, Figure 4.11 presents two panels. The left panel illustrates the distribution of reaction times at each period of the sequential games. The right panel displays the distribution of reaction times for each stopping strategy in the simultaneous games. In the sequential treatment, we can observe that the reaction times at each period are significantly different (Kruskal-Wallis ranksum test: $\chi^2(4) = 32.519$ and p-value = 0.0001). Moreover, the reaction time decreases as the game progresses to later periods, dropping from 11.29 seconds in period 1 to 7.66 seconds in period 5. Combined with the low frequencies of C in later periods (around 21.43%), we can conclude that players quickly decide to wait in later periods.

In the simultaneous treatment, we once again observe that the reaction times for each stopping strategy are significantly different (Kruskal-Wallis ranksum test: $\chi^2(5) = 54.291$ and p-value = 0.0001). The right panel of Figure 4.11 reveals a monotonic pattern: players take longer to decide to claim in later periods, with average reaction




Figure 4.11: (Left) The reaction time in the sequential treatment. The scatterplot of reaction time conditional on the current period is shown by the blue dots. The mean and the 95% CIs are overlaid. (Right) The reaction time in the simultaneous treatment when seeing a dirty face. The scatterplot of reaction time conditional on the choice of the stopping periods is shown by the red dots. The mean and the 95% CIs are overlaid.

time of 11.93 seconds for period 1 and 23.34 seconds for period 5. However, it only takes players approximately 14.42 seconds to decide to always wait.

The empirical patterns from both treatments provide suggestive evidence that the heuristic of choosing to “always wait” differs from the heuristic of claiming at a specific period. This finding aligns with the rationale of DCH—level 1 players will always wait upon seeing a dirty face, regardless of the payoff configurations. Conversely, higher-level strategic players will make inferences to determine their stopping strategies. Lastly, the observed monotonic increase in reaction times across stopping strategies in the simultaneous treatment aligns with the idea that choosing to claim at later periods requires more steps of reasoning.

Result 5. (1) *In the sequential treatment, when players see a dirty face, their reaction time is shorter in later periods.* (2) *In the simultaneous treatment, the reaction time of choosing to claim at some period when players see a dirty face is monotonically increasing in the stopping periods. However, players take much less time to decide to always wait.*

4.8 Conclusions

This paper theoretically and experimentally studies the DCH solution, an alternative model that relaxes the mutual consistency requirement of the standard equilibrium theory, in multi-stage games of incomplete information. Instead of mutual consistency, DCH posits that players are heterogeneous with respect to their levels of sophistication and incorrectly believe others are strictly less sophisticated than they are. As the dynamic game progresses, strategic players will update their beliefs about others' types and levels.

In this paper, I characterize some general properties of the DCH belief system in multi-stage games of incomplete information. Proposition 4.1 guarantees that the DCH belief system is a product measure across players when every player's payoff-relevant type is independently determined. On the other hand, when the prior distribution of types is correlated across players, Proposition 4.2 demonstrates the existence of a unique corresponding game, where the types are independently drawn, resulting in the DCH solution being invariant in both games. While solving the DCH solution does not require a fixed point argument, it could be computationally challenging in principle, especially when there are more players or information sets involved. To this end, Proposition 4.1 and 4.2 simplify the computation, preserving the tractability of DCH. In addition, Proposition 4.3 shows that strategic players always consider the possibility of others being non-strategic, causing the lack of common knowledge of rationality in DCH.

Furthermore, another feature of DCH is the violation of invariance under strategic equivalence, which arises because level 0 players' behavioral strategies are not always outcome-equivalent in different strategically equivalent games, leading to different behavior of higher-level players. To demonstrate the violation of invariance and contrast DCH with the standard equilibrium theory, I characterize the DCH solutions of the sequential and simultaneous two-person dirty-faces games. Despite the two versions of the game sharing the same reduced normal form, the DCH solutions of the two versions differ in a specific way, as characterized by Proposition 4.6. In summary, DCH predicts that higher-level (level $k \geq 2$) players tend to claim earlier in the sequential version when they are sufficiently impatient, and vice versa in the simultaneous version when they are patient enough.

To test the predictions of DCH, I design and run a laboratory experiment on two-person dirty-faces games where I manipulate both the timing (sequential vs. simultaneous) and payoff structures. The experimental design is guided by DCH, wherein

I first calibrate the model using an existing dirty-faces game experimental dataset and choose the payoff parameters to maximize the diagnosticity. Considering that the prior distributions of levels might significantly vary among different subject pools, this experimental design to some extent serves as a stress test for assessing the external validity of DCH.

Overall, a significant treatment effect is detected: players tend to claim earlier in the sequential treatment than in the simultaneous treatment. Some interesting patterns emerge from the data. First, players' behavior significantly varies with payoff structures in both treatments, aligning with the qualitative predictions of DCH. Second, players take longer to choose higher-level stopping strategies. Third, when comparing the fitness of different behavioral models, we find that QDCH outperforms other behavioral models in both treatments. Moreover, there is some evidence of the failure of best responses and Bayesian inferences.

Lastly, when comparing the observed treatment effects with the predictions of the calibrated DCH, we find in one of the diagnostic games where $(\delta, \bar{\alpha}) = (0.6, 0.45)$, the empirical frequency of choosing to claim at period 2 with the observation of a dirty face is 30.77% higher in the sequential treatment than in the simultaneous treatment, which is highly consistent with the calibrated DCH (approximately 31.15%). However, in all control games and the other diagnostic game, the treatment effect is significantly higher than the predictions of the calibrated DCH. Along with the estimation results, we can conclude that while the observed violation of invariance in the data is primarily attributed to the relaxation of mutual consistency, it is a joint consequence of the relaxation of all equilibrium requirements.

The key contribution of this paper is establishing the theoretical and empirical foundations of the DCH solution. However, there are considerable extensions and applications that might be fruitful for future research. The first extension worth pursuing is to endogenize the levels of sophistication, possibly using the cost-benefit analysis proposed by Alaoui and Penta (2016). This extension could be challenging in dynamic games because each player's level might vary in different information sets of the same game. Additionally, players not only form beliefs about others' current levels but also about their cognitive bounds, which might make the model less tractable.

Second, while the assumption of uniform randomization of level 0 players has some distinct advantages, exploring the actual behavior of level 0 players is another direction worth investigating. Inspired by X. Li and Camerer (2022), an alternative

assumption for level 0 players is that they will randomize across *visually salient* actions at each information set. In particular, due to the rapid development of machine learning algorithms, how visually salient an action is can be quantified even before any behavioral data is collected.

Finally, this last section lists several potential applications of dynamic games of incomplete information where the mutual consistency requirement is easily violated and the DCH solution might provide some new insights.

1. *Social learning*: In social learning games with repeated actions, players make inferences about the true state based on their private signals and publicly observed actions (Bala and Goyal, 1998; Harel et al., 2021). The DCH solution posits that players do not commonly believe others are able to make correct inferences. Specifically, level 0 players' actions do not convey any information about the true states, while level 1 players will always obey their private signals. For higher-level players, they will constantly update their beliefs about the true state and other players' levels of sophistication. An open question is whether higher-level players will eventually learn the true state.
2. *Sequential bargaining*: The equilibrium of a sequential bargaining game was first characterized by Rubinstein (1982). To reach the perfect equilibrium, players are required to choose the optimal proposal among a continuum of choices at every subgame, and believe the other player to optimally respond to each proposal. Later, McKelvey and Palfrey (1993) and McKelvey and Palfrey (1995) considered a two-person multi-stage bargaining game where each player has a private payoff-relevant type and makes a binary decision (whether to give in or hold out) in every period. The game continues until at least one of the players gives in. In this game, it is strictly dominant for the strong type of players to hold out forever, but not for the weak type—the weak type players need to trade-off between the reward of giving in earlier and the reputational benefit from mimicking the high type. This reasoning is behaviorally challenging. In contrast, DCH is not a solution of a fixed point problem but solved iteratively from lower to higher levels. Therefore, the DCH solution is expected to be sharply different from the standard equilibrium in the sequential bargaining game.
3. *Signaling*: In a multi-stage signaling game, an informed player will have

a persistent type and interact with an uninformed player repeatedly. Kaya (2009) analyzed such an environment, finding that the set of equilibrium signal sequences includes a large class of possibly complex signal sequences. In contrast, in the DCH solution, the uninformed player will learn about the informed player's true type and level when observing a new signal, and the informed player will also learn about the uninformed player's level at each stage. Given that the DCH solution is unique, it would be interesting to characterize the signal sequence of each level of informed players and test whether this is consistent with the behavioral data.

4. *Sequential voting*: There is a large class of voting rules that includes multiple rounds, such as sequential voting over agendas (Baron and Ferejohn, 1989) or elections based on repeated ballots and elimination of one candidate in each round (Bag, Sabourian, and Winter, 2009). To reach Condorcet consistent outcomes, players are required to behave strategically. However, in cases where voters are not strategic or believe others might not be strategic, the DCH solution becomes an ideal solution concept. In the DCH solution, voters will update their beliefs about others' preferences and levels of sophistication simultaneously, and vote according to their posterior beliefs in each round. Since the common knowledge of rationality is violated in DCH, it is natural to conjecture that higher-level players will vote more sincerely in DCH than in the equilibrium.

References

- Alaoui, Larbi, Katharina A Janezic, and Antonio Penta (2020). "Reasoning about others' reasoning". In: *Journal of Economic Theory* 189, p. 105091.
- Alaoui, Larbi and Antonio Penta (2016). "Endogenous depth of reasoning". In: *Review of Economic Studies* 83.4, pp. 1297–1333.
- (2022). "Cost-benefit analysis in reasoning". In: *Journal of Political Economy* 130.4, pp. 881–925.
- Aumann, Robert (1992). "Irrationality in game theory". In: *Economic Analysis of Markets and Games*, pp. 214–227.
- Bag, Parimal Kanti, Hamid Sabourian, and Eyal Winter (2009). "Multi-stage voting, sequential elimination and Condorcet consistency". In: *Journal of Economic Theory* 144.3, pp. 1278–1299.
- Bala, Venkatesh and Sanjeev Goyal (1998). "Learning from neighbours". In: *Review of Economic Studies* 65.3, pp. 595–621.

- Baron, David P and John A Ferejohn (1989). “Bargaining in legislatures”. In: *American Political Science Review* 83.4, pp. 1181–1206.
- Barwise, Jon (1981). “Scenes and other situations”. In: *Journal of Philosophy* 78.7, pp. 369–397.
- Bayer, Ralph C and Mickey Chan (2007). *The dirty faces game revisited*. Tech. rep. University of Adelaide, School of Economics.
- Bayer, Ralph C and Ludovic Renou (2016a). “Logical abilities and behavior in strategic-form games”. In: *Journal of Economic Psychology* 56, pp. 39–59.
- (2016b). “Logical omniscience at the laboratory”. In: *Journal of Behavioral and Experimental Economics* 64, pp. 41–49.
- Binmore, Ken and Adam Brandeburger (1988). “Common knowledge and game theory”. In: *MichU DeptE CenREST W89-06*.
- Bland, James R (2023). “Optimizing experiment design for estimating parametric models in economic experiments”. In: *Available at SSRN 4566855*.
- Bosch-Domenech, Antoni et al. (2002). “One, two,(three), infinity,...: Newspaper and lab beauty-contest experiments”. In: *American Economic Review* 92.5, pp. 1687–1701.
- Brandts, Jordi and Gary Charness (2011). “The strategy versus the direct-response method: A first survey of experimental comparisons”. In: *Experimental Economics* 14.3, pp. 375–398.
- Cai, Hongbin and Joseph Tao-Yi Wang (2006). “Overcommunication in strategic information transmission games”. In: *Games and Economic Behavior* 56.1, pp. 7–36.
- Camerer, Colin F (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.
- Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong (2004). “A cognitive hierarchy model of games”. In: *Quarterly Journal of Economics* 119.3, pp. 861–898.
- Camerer, Colin F and Dan Lovall (1999). “Overconfidence and excess entry: An experimental approach”. In: *American Economic Review* 89.1, pp. 306–318.
- Chabris, Christopher F et al. (2009). “The allocation of time in decision-making”. In: *Journal of the European Economic Association* 7.2-3, pp. 628–637.
- Chen, Daniel L and Martin Schonger (2023a). “Invariance of equilibrium to the strategy method I: Theory”. In: *Journal of the Economic Science Association*, pp. 1–16.
- (2023b). “Invariance of equilibrium to the strategy method II: Experimental evidence”. In: *Journal of the Economic Science Association*, pp. 1–13.

- Chen, Daniel L, Martin Schonger, and Chris Wickens (2016). “oTree—An open-source platform for laboratory, online, and field experiments”. In: *Journal of Behavioral and Experimental Finance* 9, pp. 88–97.
- Chen, Wei James, Meng-Jhang Fong, and Po-Hsuan Lin (2023). “Measuring higher-order rationality with belief control”. In: *arXiv preprint arXiv:2309.07427*.
- Choo, Lawrence and Xiaoyu Zhou (2022). “Can market selection reduce anomalous behaviour in games?” In: *European Economic Review* 141, p. 103958.
- Cohen, Shani and Shengwu Li (2023). “Sequential cursed equilibrium”. In: *arXiv preprint arXiv:2212.06025v3*.
- Costa-Gomes, Miguel and Vincent P Crawford (2006). “Cognition and behavior in two-person guessing games: An experimental study”. In: *American Economic Review* 96.5, pp. 1737–1768.
- Costa-Gomes, Miguel, Vincent P Crawford, and Bruno Broseta (2001). “Cognition and behavior in normal-form games: An experimental study”. In: *Econometrica* 69.5, pp. 1193–1235.
- Crawford, Vincent P and Nagore Iriberri (2007a). “Fatal attraction: Saliency, naivete, and sophistication in experimental “hide-and-seek” games”. In: *American Economic Review* 97.5, pp. 1731–1750.
- (2007b). “Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner’s curse and overbidding in private-value auctions?” In: *Econometrica* 75.6, pp. 1721–1770.
- Devetag, Giovanna and Massimo Warglien (2003). “Games and phone numbers: Do short-term memory bounds affect strategic behavior?” In: *Journal of Economic Psychology* 24.2, pp. 189–202.
- Elmes, Susan and Philip J Reny (1994). “On the strategic equivalence of extensive form games”. In: *Journal of Economic Theory* 62.1, pp. 1–23.
- Enke, Benjamin, Thomas Graeber, and Ryan Oprea (2023). “Confidence, self-selection, and bias in the aggregate”. In: *American Economic Review* 113.7, pp. 1933–1966.
- Fong, Meng-Jhang, Po-Hsuan Lin, and Thomas R Palfrey (2023a). “A note on cursed sequential equilibrium and sequential cursed equilibrium”. In: *arXiv preprint arXiv:2304.05515*. DOI: <https://doi.org/10.48550/arXiv.2304.05515>.
- (2023b). “Cursed sequential equilibrium”. In: *arXiv preprint arXiv:2301.11971*.
- Fudenberg, Drew and David Levine (1983). “Subgame-perfect equilibria of finite- and infinite-horizon games”. In: *Journal of Economic Theory* 31.2, pp. 251–268.
- Fudenberg, Drew and Jean Tirole (1991). “Perfect Bayesian equilibrium and sequential equilibrium”. In: *Journal of Economic Theory* 53.2, pp. 236–260.
- (1993). *Game theory*. MIT press.

- Gamow, G and M Stern (1958). “Forty unfaithful wives”. In: *Puzzle Math*, pp. 20–23.
- García-Pola, Bernardo, Nagore Iriberry, and Jaromír Kovářik (2020). “Hot versus cold behavior in centipede games”. In: *Journal of the Economic Science Association* 6.2, pp. 226–238.
- Geanakoplos, John (1994). “Common knowledge”. In: *Handbook of Game Theory with Economic Applications* 2, pp. 1437–1496.
- Gill, David and Victoria Prowse (2023). “Strategic complexity and the value of thinking”. In: *Economic Journal* 133.650, pp. 761–786.
- Grehl, Sascha and Andreas Tutić (2015). “Experimental evidence on iterated reasoning in games”. In: *PloS one* 10.8, e0136524.
- Halevy, Yoram, Johannes Hoelzemann, and Terri Kneeland (2021). “Magic mirror on the wall, who is the smartest one of all?” In: *Working Paper*.
- Halpern, Joseph Y and Yoram Moses (1990). “Knowledge and common knowledge in a distributed environment”. In: *Journal of the ACM (JACM)* 37.3, pp. 549–587.
- Hardin, Christopher S and Alan D Taylor (2008). “An introduction to infinite hat problems”. In: *The Mathematical Intelligencer* 30.4, pp. 20–25.
- Harel, Matan et al. (2021). “Rational groupthink”. In: *Quarterly Journal of Economics* 136.1, pp. 621–668.
- Ho, Teck-Hua, Colin F Camerer, and Keith Weigelt (1998). “Iterated dominance and iterated best response in experimental “p-beauty contests””. In: *American Economic Review* 88.4, pp. 947–969.
- Ho, Teck-Hua, So-Eun Park, and Xuanming Su (2021). “A Bayesian level-k model in n-person games”. In: *Management Science* 67.3, pp. 1622–1638.
- Ho, Teck-Hua and Xuanming Su (2013). “A dynamic level-k model in sequential games”. In: *Management Science* 59.2, pp. 452–469.
- Johanson, Michael (2013). “Measuring the size of large no-limit poker games”. In: *arXiv preprint arXiv:1302.7008*.
- Kaya, Ayça (2009). “Repeated signaling games”. In: *Games and Economic Behavior* 66.2, pp. 841–854.
- Kohlberg, Elon and Jean-Francois Mertens (1986). “On the strategic stability of equilibria”. In: *Econometrica*, pp. 1003–1037.
- Kononov, Arkady and Ian Krajbich (2019). “Revealed strength of preference: Inference from response times”. In: *Judgment and Decision making* 14.4, pp. 381–394.
- Levin, Dan and Luyao Zhang (2022). “Bridging level-k to Nash equilibrium”. In: *Review of Economics and Statistics* 104.6, pp. 1329–1340.

- Li, Shengwu (2017). “Obviously strategy-proof mechanisms”. In: *American Economic Review* 107.11, pp. 3257–3287.
- Li, Xiaomin and Colin F Camerer (2022). “Predictable effects of visual salience in experimental decisions and games”. In: *Quarterly Journal of Economics* 137.3, pp. 1849–1900.
- Lin, Po-Hsuan, Alexander L Brown, et al. (2020). “Evidence of general economic principles of bargaining and trade from 2,000 classroom experiments”. In: *Nature Human Behaviour* 4.9, pp. 917–927.
- Lin, Po-Hsuan and Thomas R Palfrey (2022). “Cognitive hierarchies in extensive form games”. In: *Social Science Working Paper 1460, California Institute of Technology*.
- Littlewood, John E (1953). *A mathematician’s miscellany*. London, England: Meuthen & Co. Ltd.
- Liu, Zhen (2008). “The dirty face problem with unawareness”. In: *The BE Journal of Theoretical Economics* 8.1, pp. 1935–1704.
- McKelvey, Richard D and Thomas R Palfrey (1993). “Engodeneity of alternating offers in a bargaining game”. In: *Caltech Social Science Working Paper*.
- (1995). “The holdout game: An experimental study of an infinitely repeated game with”. In: *Social Choice, Welfare, and Ethics: Proceedings of the Eighth International Symposium in Economic Theory and Econometrics*. Vol. 8. Cambridge University Press, p. 321.
- (1998). “Quantal response equilibria for extensive form games”. In: *Experimental Economics* 1.1, pp. 9–41.
- Moffatt, Peter (2020). *Experimetrics: Econometrics for experimental economics*. Bloomsbury Publishing.
- Moore, Don A and Paul J Healy (2008). “The trouble with overconfidence”. In: *Psychological Review* 115.2, p. 502.
- Moses, Yoram, Danny Dolev, and Joseph Y Halpern (1986). “Cheating husbands and other stories: A case study of knowledge, action, and communication”. In: *Distributed Computing* 1.3, pp. 167–176.
- Myerson, Roger B (1985). “Bayesian equilibrium and incentive-compatibility: An introduction”. In: *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*, pp. 229–260.
- Nagel, Rosemarie (1995). “Unraveling in guessing games: An experimental study”. In: *American Economic Review* 85.5, pp. 1313–1326.
- Rubinstein, Ariel (1982). “Perfect equilibrium in a bargaining model”. In: *Econometrica*, pp. 97–109.

- Rubinstein, Ariel (2007). "Instinctive and cognitive reasoning: A study of response times". In: *Economic Journal* 117.523, pp. 1243–1259.
- Selten, Reinhard (1967). "Die strategiemethode zur erforschung des eingeschr nkt rationale verhaltens im rahmen eines oligopolexperiments". In: *Beitr ge zur experimentellen Wirtschaftsforschung*, p. 136.
- (1991). "Anticipatory learning in two-person games". In: *Game Equilibrium Models I*. Springer, pp. 98–154.
 - (1998). "Features of experimentally observed bounded rationality". In: *European Economic Review* 42.3-5, pp. 413–436.
- Stahl, Dale O (1993). "Evolution of smartn players". In: *Games and Economic Behavior* 5.4, pp. 604–617.
- (1996). "Boundedly rational rule learning in a guessing game". In: *Games and Economic Behavior* 16.2, pp. 303–330.
- Stahl, Dale O and Paul W Wilson (1995). "On players' models of other players: Theory and experimental evidence". In: *Games and Economic Behavior* 10.1, pp. 218–254.
- Thompson, F B (1952). *Equivalence of games in extensive form*. Tech. rep. RAND Corporation.
- Wald, Abraham (1947). *Sequential analysis*. John Wiley and Sons, New York.
- Wang, Joseph Tao-Yi, Michael Spezio, and Colin F Camerer (2010). "Pinocchio's pupil: using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games". In: *American Economic Review* 100.3, pp. 984–1007.
- Weber, Roberto A (2001). "Behavior and learning in the "dirty faces" game". In: *Experimental Economics* 4.3, pp. 229–242.

Appendix A

PROOFS FOR CHAPTER 2

A.1 Proofs of General Properties

Proof of Lemma 2.1

By Definition 2.1, for any $(\mu, \sigma) \in \Psi^\chi$, any history h^{t-1} , any player i and any type profile $\theta = (\theta_i, \theta_{-i})$,

$$\begin{aligned}
 & \sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) [\chi \bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i) + (1 - \chi) \sigma_{-i}(a^t_{-i}|h^{t-1}, \theta'_{-i})] \\
 = & \chi \underbrace{\left[\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \right]}_{=1} \bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i) + (1 - \chi) \underbrace{\left[\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \sigma_{-i}(a^t_{-i}|h^{t-1}, \theta'_{-i}) \right]}_{=\bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i)} \\
 = & \bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i).
 \end{aligned}$$

Therefore, since $(\mu, \sigma) \in \Psi^\chi$, with some rearrangement, it follows that

$$\begin{aligned}
 \mu_i(\theta_{-i}|h^t, \theta_i) &= \frac{\mu_i(\theta_{-i}|h^{t-1}, \theta_i) \sigma_{-i}^\chi(a^t_{-i}|h^{t-1}, \theta_{-i}, \theta_i)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \sigma_{-i}^\chi(a^t_{-i}|h^{t-1}, \theta'_{-i}, \theta_i)} \\
 &= \frac{\mu_i(\theta_{-i}|h^{t-1}, \theta_i) [\chi \bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i) + (1 - \chi) \sigma_{-i}(a^t_{-i}|h^{t-1}, \theta_{-i})]}{\bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i)} \\
 &= \chi \mu_i(\theta_{-i}|h^{t-1}, \theta_i) + (1 - \chi) \left[\frac{\mu_i(\theta_{-i}|h^{t-1}, \theta_i) \sigma_{-i}(a^t_{-i}|h^{t-1}, \theta_{-i})}{\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \sigma_{-i}(a^t_{-i}|h^{t-1}, \theta'_{-i})} \right]. \blacksquare
 \end{aligned}$$

Proof of Proposition 2.1

The proof is similar to the proof for sequential equilibrium and proceeds in three steps. First, for any finite multi-stage games with observed actions, Γ , we construct an ϵ -perturbed game Γ^ϵ that is identical to Γ but every player in every information set has to play any available action with probability at least ϵ . Second, we defined a cursed best-response correspondence for Γ^ϵ and prove that the correspondence has a fixed point by Kakutani's fixed point theorem. Finally, in step 3, we use a sequence of fixed points in perturbed games, with ϵ converging to 0, where the limit of this sequence is a χ -CSE.

Step 1: Let Γ^ϵ be a game identical to Γ but for each player $i \in N$, player i must play

any available action in every information set $\mathcal{I}_i = (\theta_i, h^t)$ with probability at least ϵ where $\epsilon < \frac{1}{\sum_{j=1}^n |A_j|}$. Let $\Sigma^\epsilon = \times_{j=1}^n \Sigma_j^\epsilon$ be set of feasible behavioral strategy profiles for players in the perturbed game Γ^ϵ . For any behavioral strategy profile $\sigma \in \Sigma^\epsilon$, let $\mu^\chi(\cdot) \equiv (\mu_i^\chi(\cdot))_{i=1}^n$ be the belief system induced by σ via χ -cursed Bayes' rule. That is, for each player $i \in N$, information set $\mathcal{I}_i = (\theta_i, h^t)$ where $h^t = (h^{t-1}, a^t)$ and type profile $\theta_{-i} \in \Theta_{-i}$,

$$\mu_i^\chi(\theta_{-i}|h^t, \theta_i) = \chi \mu_i^\chi(\theta_{-i}|h^{t-1}, \theta_i) + (1 - \chi) \left[\frac{\mu_i^\chi(\theta_{-i}|h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^t|h^{t-1}, \theta_{-i})}{\sum_{\theta'_{-i} \in \Theta_{-i}} \mu_i^\chi(\theta'_{-i}|h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^t|h^{t-1}, \theta'_{-i})} \right].$$

Notice that the χ -cursed Bayes' rule is only defined on the family of multi-stage games with observed actions. As σ is fully mixed, the belief system is uniquely pinned down.

Finally, let $B^\epsilon : \Sigma^\epsilon \rightrightarrows \Sigma^\epsilon$ be the cursed best response correspondence which maps any behavioral strategy profile $\sigma \in \Sigma^\epsilon$ to the set of ϵ -constrained behavioral strategy profiles $\tilde{\sigma} \in \Sigma^\epsilon$ that are best replies given the belief system $\mu^\chi(\cdot)$.

Step 2: Next, fix any $0 < \epsilon < \frac{1}{\sum_{j=1}^n |A_j|}$ and show that B^ϵ has a fixed point by Kakutani's fixed point theorem. We check the conditions of the theorem:

1. It is straightforward that Σ^ϵ is compact and convex.
2. For any $\sigma \in \Sigma^\epsilon$, as $\mu^\chi(\cdot)$ is uniquely pinned down by χ -cursed Bayes' rule, it is straightforward that $B^\epsilon(\sigma)$ is non-empty and convex.
3. To verify that B^ϵ has a closed graph, take any sequence of ϵ -constrained behavioral strategy profiles $\{\sigma^k\}_{k=1}^\infty \subseteq \Sigma^\epsilon$ such that $\sigma^k \rightarrow \sigma \in \Sigma^\epsilon$ as $k \rightarrow \infty$, and any sequence $\{\tilde{\sigma}^k\}_{k=1}^\infty$ such that $\tilde{\sigma}^k \in B^\epsilon(\sigma^k)$ for any k and $\tilde{\sigma}^k \rightarrow \tilde{\sigma}$. We want to prove that $\tilde{\sigma} \in B^\epsilon(\sigma)$.

Fix any player $i \in N$ and information set $\mathcal{I}_i = (\theta_i, h^t)$. For any $\sigma \in \Sigma^\epsilon$, recall that $\sigma_{-i}^\chi(\cdot)$ is player i 's χ -cursed perceived behavioral strategies of other players induced by σ . Specifically, for any type profile $\theta \in \Theta$, non-terminal history h^{t-1} and action profile $a_{-i}^t \in A_{-i}(h^{t-1})$, $\sigma_{-i}^\chi(a_{-i}^t|h^{t-1}, \theta_{-i}, \theta_i) = \chi \tilde{\sigma}_{-i}(a_{-i}^t|h^{t-1}, \theta_i) + (1 - \chi) \sigma_{-i}(a_{-i}^t|h^{t-1}, \theta_{-i})$.

Additionally, recall that $\rho_i^\chi(\cdot)$ is player i 's belief about the terminal nodes (conditional on the history and type profile), which is also induced by σ . Since $\mu^\chi(\cdot)$ is continuous in σ we have that $\sigma_{-i}^\chi(\cdot)$ and $\rho_i^\chi(\cdot)$ are also continuous

in σ . We further define $\mathcal{S}_{\mathcal{I}_i}^k \equiv \{\sigma'_i \in \Sigma_i^\epsilon : \sigma'_i(\cdot | \mathcal{I}_i) = \tilde{\sigma}_i^k(\cdot | \mathcal{I}_i)\}$ and $\mathcal{S}_{\mathcal{I}_i} \equiv \{\sigma'_i \in \Sigma_i^\epsilon : \sigma'_i(\cdot | \mathcal{I}_i) = \tilde{\sigma}_i(\cdot | \mathcal{I}_i)\}$. Since $\tilde{\sigma}^k \in B^\epsilon(\sigma^k)$, for any $\sigma'_i \in \Sigma_i^\epsilon$, we can obtain that

$$\begin{aligned} & \max_{\sigma''_i \in \mathcal{S}_{\mathcal{I}_i}^k} \left\{ \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i^\chi[\sigma^k](\theta_{-i}|h^t, \theta_i) \rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi[\sigma^k], \sigma''_i) u_i(h^T, \theta_{-i}, \theta_i) \right\} \\ & \geq \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i^\chi[\sigma^k](\theta_{-i}|h^t, \theta_i) \rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi[\sigma^k], \sigma'_i) u_i(h^T, \theta_{-i}, \theta_i). \end{aligned}$$

By continuity, as we take limits on both sides, we can find $\tilde{\sigma} \in B^\epsilon(\sigma)$ because

$$\begin{aligned} & \max_{\sigma''_i \in \mathcal{S}_{\mathcal{I}_i}} \left\{ \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i^\chi[\sigma](\theta_{-i}|h^t, \theta_i) \rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi[\sigma], \sigma''_i) u_i(h^T, \theta_{-i}, \theta_i) \right\} \\ & \geq \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i^\chi[\sigma](\theta_{-i}|h^t, \theta_i) \rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi[\sigma], \sigma'_i) u_i(h^T, \theta_{-i}, \theta_i). \end{aligned}$$

By Kakutani's fixed point theorem, B^ϵ has a fixed point.

Step 3: For any ϵ , let σ^ϵ be a fixed point of B^ϵ and μ^ϵ be the belief system induced by σ^ϵ via χ -cursed Bayes' rule. We combine these two components and let $(\mu^\epsilon, \sigma^\epsilon)$ be the induced assessment. We now consider a sequence of $\epsilon \rightarrow 0$, where $\{(\mu^\epsilon, \sigma^\epsilon)\}$ is the corresponding sequence of assessments. By compactness and the finiteness of Γ , the Bolzano-Weierstrass theorem guarantees the existence of a convergent subsequence of the assessments. As $\epsilon \rightarrow 0$, let $(\mu^\epsilon, \sigma^\epsilon) \rightarrow (\mu^*, \sigma^*)$. By construction, the limit assessment (μ^*, σ^*) satisfies χ -consistency and sequential rationality. Hence, (μ^*, σ^*) is a χ -CSE. ■

Proof of Proposition 2.2

To prove $\Phi(\chi)$ is upper hemi-continuous in χ , consider any sequence of $\{\chi^k\}_{k=1}^\infty$ such that $\chi^k \rightarrow \chi^* \in [0, 1]$, and any sequence of CSE, $\{(\mu^k, \sigma^k)\}$, such that $(\mu^k, \sigma^k) \in \Phi(\chi^k)$ for all k . Let (μ^*, σ^*) be the limit assessment, i.e., $(\mu^k, \sigma^k) \rightarrow (\mu^*, \sigma^*)$. We need to show that $(\mu^*, \sigma^*) \in \Phi(\chi^*)$.

For simplicity, for any player $i \in N$, any information set $\mathcal{I}_i = (h^t, \theta_i)$, any $\sigma'_i \in \Sigma_i$, and any $\sigma \in \Sigma$, the expected payoff under the belief system $\mu^\chi(\cdot)$ induced by σ is denoted as

$$\mathbb{E}_{\mu^\chi[\sigma]} [u_i(\sigma'_i, \sigma_{-i}|h^t, \theta_i)] \equiv \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i^\chi(\theta_{-i}|h^t, \theta_i) \rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi, \sigma'_i) u_i(h^T, \theta_{-i}, \theta_i).$$

Suppose $(\mu^*, \sigma^*) \notin \Phi(\chi^*)$. Then there exists some player $i \in N$, some information set $\mathcal{I}_i = (h^t, \theta_i)$, some $\sigma'_i \in \Sigma_i$, and some $\epsilon > 0$ such that

$$\mathbb{E}_{\mu^{\chi^*}[\sigma^*]} [u_i(\sigma'_i, \sigma_{-i}^* | h^t, \theta_i)] - \mathbb{E}_{\mu^{\chi^*}[\sigma^*]} [u_i(\sigma_i^*, \sigma_{-i}^* | h^t, \theta_i)] > \epsilon. \quad (\text{A.1})$$

Since $\mu^\chi(\cdot)$ is continuous in χ , it follows that for any strategy profile σ , $\sigma_{-i}^\chi(\cdot)$ and $\rho_i^\chi(\cdot)$ are both continuous in χ . Thus, there exists a sufficiently large M_1 such that for every $k \geq M_1$,

$$\left| \mathbb{E}_{\mu^{\chi^k}[\sigma^k]} [u_i(\sigma_i^k, \sigma_{-i}^k | h^t, \theta_i)] - \mathbb{E}_{\mu^{\chi^*}[\sigma^*]} [u_i(\sigma_i^*, \sigma_{-i}^* | h^t, \theta_i)] \right| < \frac{\epsilon}{3}. \quad (\text{A.2})$$

Similarly, there exists a sufficiently large M_2 such that for every $k \geq M_2$,

$$\left| \mathbb{E}_{\mu^{\chi^k}[\sigma^k]} [u_i(\sigma'_i, \sigma_{-i}^k | h^t, \theta_i)] - \mathbb{E}_{\mu^{\chi^*}[\sigma^*]} [u_i(\sigma'_i, \sigma_{-i}^* | h^t, \theta_i)] \right| < \frac{\epsilon}{3}. \quad (\text{A.3})$$

Therefore, for any $k \geq \max\{M_1, M_2\}$, inequalities (A.1), (A.2) and (A.3) imply:

$$\mathbb{E}_{\mu^{\chi^k}[\sigma^k]} [u_i(\sigma'_i, \sigma_{-i}^k | h^t, \theta_i)] - \mathbb{E}_{\mu^{\chi^k}[\sigma^k]} [u_i(\sigma_i^k, \sigma_{-i}^k | h^t, \theta_i)] > \frac{\epsilon}{3},$$

implying that σ'_i is a profitable deviation for player i at information set $\mathcal{I}_i = (h^t, \theta_i)$, which contradicts $(\mu^k, \sigma^k) \in \Phi(\chi^k)$. Therefore, $(\mu^*, \sigma^*) \in \Phi(\chi^*)$, as desired. ■

Proof of Proposition 2.3

Fix any $\chi \in [0, 1]$ and let (μ, σ) be a χ -consistent assessment. We prove the result by contradiction. Suppose (μ, σ) does not satisfy χ -dampened updating property. Then there exists $i \in N$, $\tilde{\theta} \in \Theta$ and a non-terminal history h^t such that $\mu_i(\theta_{-i} | h^t, \tilde{\theta}_i) < \chi \mu_i(\theta_{-i} | h^{t-1}, \tilde{\theta}_i)$.

Since (μ, σ) is χ -consistent, there exists a sequence $\{(\mu^k, \sigma^k)\} \subseteq \Psi^\chi$ such that $(\mu^k, \sigma^k) \rightarrow (\mu, \sigma)$ as $k \rightarrow \infty$. By Lemma 2.1, we know for this $i, \tilde{\theta}$ and h^t , $\mu_i^k(\tilde{\theta}_{-i} | h^t, \tilde{\theta}_i)$ equals to

$$\chi \mu_i^k(\tilde{\theta}_{-i} | h^{t-1}, \tilde{\theta}_i) + (1 - \chi) \left[\frac{\mu_i^k(\tilde{\theta}_{-i} | h^{t-1}, \tilde{\theta}_i) \sigma_{-i}^k(a_{-i}^t | h^{t-1}, \tilde{\theta}_{-i})}{\sum_{\theta'_{-i}} \mu_i^k(\theta'_{-i} | h^{t-1}, \tilde{\theta}_i) \sigma_{-i}^k(a_{-i}^t | h^{t-1}, \theta'_{-i})} \right] \geq \chi \mu_i^k(\tilde{\theta}_{-i} | h^{t-1}, \tilde{\theta}_i).$$

There will be a contradiction as we take the limit $k \rightarrow \infty$ on both sides:

$$\mu_i(\tilde{\theta}_{-i} | h^t, \tilde{\theta}_i) = \lim_{k \rightarrow \infty} \mu_i^k(\tilde{\theta}_{-i} | h^t, \tilde{\theta}_i) \geq \lim_{k \rightarrow \infty} \chi \mu_i^k(\tilde{\theta}_{-i} | h^{t-1}, \tilde{\theta}_i) = \chi \mu_i(\tilde{\theta}_{-i} | h^{t-1}, \tilde{\theta}_i),$$

which yields a contradiction. ■

Proof of Corollary 2.2

We prove the statement by induction on t . For $t = 1$, by Proposition 2.3, $\mu_i(\theta_{-i}|h^1, \theta_i) \geq \chi\mu_i(\theta_{-i}|h_0, \theta_i) = \chi F(\theta_{-i}|\theta_i)$. Next, suppose there is t' such that the statement holds for all $1 \leq t \leq t' - 1$. At stage t' , by Proposition 2.3 and the induction hypothesis, we can find that

$$\mu_i(\theta_{-i}|h^{t'}, \theta_i) \geq \chi\mu_i(\theta_{-i}|h^{t'-1}, \theta_i) \geq \chi \left[\chi^{t'-1} F(\theta_{-i}|\theta_i) \right] = \chi^{t'} F(\theta_{-i}|\theta_i). \quad \blacksquare$$

Proof of Proposition 2.5

Let the assessment (μ, σ) be a pooling χ -CSE. We want to show that for any $\chi' \leq \chi$, the assessment (μ, σ) is also a χ' -CSE. Consider any non-terminal history h^{t-1} , any player i , any $a_i^t \in A_i(h^{t-1})$ and any $\theta \in \Theta$. We can first observe that

$$\begin{aligned} \bar{\sigma}_{-i}(a_{-i}^t|h^{t-1}, \theta_i) &= \sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^t|h^{t-1}, \theta'_{-i}) \\ &= \sigma_{-i}(a_{-i}^t|h^{t-1}, \theta_{-i}) \left[\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \right] = \sigma_{-i}(a_{-i}^t|h^{t-1}, \theta_{-i}), \end{aligned}$$

where the second equality holds because σ is a pooling behavioral strategy profile, so σ_{-i} is independent of other players' types. For this pooling χ -CSE, let G^σ be the set of on-path histories and \tilde{G}^σ be the set of off-path histories. We can first show that for every $h \in G^\sigma$, $i \in N$ and $\theta \in \Theta$, $\mu_i(\theta_{-i}|h, \theta_i) = F(\theta_{-i}|\theta_i)$.

This can be shown by induction on t . For $t = 1$, any $h^1 = (h_0, a^1)$ and any $\theta \in \Theta$, by Lemma 2.1, we can obtain that

$$\begin{aligned} \mu_i(\theta_{-i}|h^1, \theta_i) &= \chi\mu_i(\theta_{-i}|h_0, \theta_i) + (1 - \chi) \left[\frac{\mu_i(\theta_{-i}|h_0, \theta_i) \sigma_{-i}(a_{-i}^1|h_0, \theta_{-i})}{\bar{\sigma}_{-i}(a_{-i}^1|h_0, \theta_i)} \right] \\ &= \chi F(\theta_{-i}|\theta_i) + (1 - \chi) F(\theta_{-i}|\theta_i) \underbrace{\left[\frac{\sigma_{-i}(a_{-i}^1|h_0, \theta_{-i})}{\bar{\sigma}_{-i}(a_{-i}^1|h_0, \theta_i)} \right]}_{=1} = F(\theta_{-i}|\theta_i). \end{aligned}$$

Now, suppose there is t' such that the statement holds for $1 \leq t \leq t' - 1$. At stage t' and $h^{t'} = (h^{t'-1}, a^{t'}) \in G^\sigma$, by Lemma 2.1 and the induction hypothesis, we can

again obtain that the posterior belief is the prior belief

$$\begin{aligned}\mu_i(\theta_{-i}|h^t, \theta_i) &= \chi\mu_i(\theta_{-i}|h^{t-1}, \theta_i) + (1 - \chi) \left[\frac{\mu_i(\theta_{-i}|h^{t-1}, \theta_i)\sigma_{-i}(a'_{-i}|h^{t-1}, \theta_{-i})}{\bar{\sigma}_{-i}(a'_{-i}|h^{t-1}, \theta_i)} \right] \\ &= \chi F(\theta_{-i}|\theta_i) + (1 - \chi) \underbrace{F(\theta_{-i}|\theta_i) \left[\frac{\sigma_{-i}(a'_{-i}|h^{t-1}, \theta_{-i})}{\bar{\sigma}_{-i}(a'_{-i}|h^{t-1}, \theta_i)} \right]}_{=1} = F(\theta_{-i}|\theta_i).\end{aligned}$$

Therefore, we have shown that players will not update their beliefs at every on-path information set, so the belief system is independent of χ . Finally, for any off-path history $h^t \in \tilde{G}^\sigma$, by Proposition 2.3, we can find that the belief system satisfies for any $\theta \in \Theta$,

$$\mu_i(\theta_{-i}|h^t, \theta_i) \geq \chi\mu_i(\theta_{-i}|h^{t-1}, \theta_i) \geq \chi'\mu_i(\theta_{-i}|h^{t-1}, \theta_i),$$

implying that when $\chi' \leq \chi$, μ will still satisfy the dampened updating property. Therefore, (μ, σ) remains a χ' -CSE. ■

A.2 Proofs for Section 2.4

Proof of Claim 2.1

First, observe that after player 1 chooses B , it is strictly optimal for player 2 to choose R for all beliefs $\mu_2(\theta_1|B)$, and after player 1 chooses A , it is optimal for player 2 to choose L if and only if

$$2\mu_2(\theta_1|A) + [1 - \mu_2(\theta_1|A)] \geq 4\mu_2(\theta_1|A) \iff \mu_2(\theta_1|A) \leq 1/3.$$

Equilibrium 1. If both types of player 1 choose A , then $\mu_2(\theta_1|A) = 1/4$, so it is optimal for player 2 to choose L . Given $a(A) = L$ and $a(B) = R$, it is optimal for both types of player 1 to choose A as $2 > 1$. Hence $m(\theta_1) = m(\theta_2) = A$, $a(A) = L$ and $a(B) = R$ is a pooling χ -CSE for any $\chi \in [0, 1]$.

Equilibrium 2. In order to support $m(\theta_1) = m(\theta_2) = B$ to be an equilibrium, player 2 has to choose R at the off-path information set A , which is optimal if and only if $\mu_2(\theta_1|A) \geq 1/3$. In addition, by Proposition 3, we know in a χ -CSE, the belief system satisfies

$$\mu_2(\theta_2|A) \geq \frac{3}{4}\chi \iff \mu_2(\theta_1|A) \leq 1 - \frac{3}{4}\chi.$$

Thus, the belief system has to satisfy that $\mu_2(\theta_1|A) \in [\frac{1}{3}, 1 - \frac{3}{4}\chi]$, requiring $\chi \leq 8/9$. ■

Proof of Proposition 2.6

Here we provide a characterization of χ -CSE of Game *BH 3* and Game *BH 4*. For the analysis of both games, we denote $\mu_I \equiv \mu_2(\theta_1|m = I)$ and $\mu_S \equiv \mu_2(\theta_1|m = S)$.

Analysis of Game *BH 3*.

At information set S , given μ_S , the expected payoffs of C, D, E are $90\mu_S, 30 - 15\mu_S$ and 15 , respectively. Therefore, for any μ_S , E is never a best response. Moreover, C is the best response if and only if $90\mu_S \geq 30 - 15\mu_S$ or $\mu_S \geq 2/7$. Similarly, at information set I , given μ_I , the expected payoffs of C, D, E are $30, 45 - 45\mu_I$ and 15 , respectively. Therefore, E is strictly dominated, and C is the best response if and only if $30 \geq 45 - 45\mu_I$ or $\mu_I \geq 1/3$. Now we consider four cases.

Case 1 [$m(\theta_1) = I, m(\theta_2) = S$]: By Lemma 2.1, $\mu_I = 1 - \chi/2$ and $\mu_S = \chi/2$. Moreover, since $\mu_I = 1 - \chi/2 \geq 1/2$ for any χ , player 2 will choose C at information set I . To support this equilibrium, player 2 has to choose C at information set S . In other words, $[(I, S); (C, C)]$ is separating χ -CSE if and only if $\mu_S \geq 2/7$ or $\chi \geq 4/7$.

Case 2 [$m(\theta_1) = S, m(\theta_2) = I$]: By Lemma 2.1, $\mu_I = \chi/2$ and $\mu_S = 1 - \chi/2$. Because $\mu_S \geq 1 - \chi/2 \geq 1/2$, it is optimal for player 2 to choose C at information set S . To support this as an equilibrium, player 2 has to choose D at information set I . Yet in this case, type θ_2 player 1 will deviate to S . Therefore, this profile cannot be supported as an equilibrium.

Case 3 [$m(\theta_1) = I, m(\theta_2) = I$]: Since player 1 follows a pooling strategy, player 2 will not update his belief at information set I , i.e., $\mu_I = 1/2$. χ -dampened updating property implies $\chi/2 \leq \mu_S \leq 1 - \chi/2$. Since $\mu_I > 1/3$, player 2 will choose C at information set I . To support this profile to be an equilibrium, player 2 has to choose D at information set S , and hence, it must be the case that $\mu_S \leq 2/7$. Coupled with the requirement from χ -dampened updating, the off-path belief has to satisfy $\chi/2 \leq \mu_S \leq 2/7$. That is, $[(I, I); (C, D)]$ is pooling χ -CSE if and only if $\chi/2 \leq 2/7$ or $\chi \leq 4/7$.

Case 4 [$m(\theta_1) = S, m(\theta_2) = S$]: Similar to the previous case, since player 1 follows a pooling strategy, player 2 will not update his belief at information set S , i.e.,

$\mu_S = 1/2$. Also, the χ -dampened updating property suggests $\chi/2 \leq \mu_I \leq 1 - \chi/2$. Because $\mu_S > 2/7$, it is optimal for player 2 to choose C at information set S . To support this as an equilibrium, player 2 has to choose D at information set I . Therefore, it must be that $\mu_I \leq 1/3$. Combined with the requirement of χ -dampened updating, the off-path belief has to satisfy $\chi/2 \leq \mu_I \leq 1/3$. As a result, $[(S, S); (D, C)]$ is a pooling χ -CSE if and only if $\chi \leq 2/3$.

Analysis of Game BH 4.

At information set I , given μ_I , the expected payoffs of C , D , E are 30 , $45 - 45\mu_I$ and $35\mu_I$. Hence, D is the best response if and only if $\mu_I \leq 1/3$ while E is the best response if $\mu_I \geq 6/7$. For $1/3 \leq \mu_I \leq 6/7$, C is the best response. On the other hand, since player 2's payoffs at information set S are the same as in Game 1, player 2 will adopt the same decision rule—player 2 will choose C if and only if $\mu_S \geq 2/7$, and choose D if and only if $\mu_S \leq 2/7$. Now, we consider the following four cases.

Case 1 [$m(\theta_1) = I, m(\theta_2) = S$]: In this case, by Lemma 2.1, $\mu_I = 1 - \chi/2$ and $\mu_S = \chi/2$. To support this profile to be an equilibrium, player 2 has to choose E and C at information set I and S , respectively. To make it profitable for player 2 to choose E at information set I , it must be that

$$\mu_I = 1 - \chi/2 \geq 6/7 \iff \chi \leq 2/7.$$

On the other hand, player 2 will choose C at information set S if and only if $\chi/2 \geq 2/7$ or $\chi \geq 4/7$, which is not compatible with the previous inequality. Therefore, this profile cannot be supported as an equilibrium.

Case 2 [$m(\theta_1) = S, m(\theta_2) = I$]: In this case, by Lemma 2.1, $\mu_I = \chi/2$ and $\mu_S = 1 - \chi/2$. To support this as an equilibrium, player 2 has to choose D at both information sets. Yet $\mu_S = 1 - \chi/2 > 2/7$, implying that it is not a best reply for player 2 to choose D at information set S . Hence this profile also cannot be supported as an equilibrium.

Case 3 [$m(\theta_1) = I, m(\theta_2) = I$]: Since player 1 follows a pooling strategy, player 2 will not update his belief at information set I , i.e., $\mu_I = 1/2$. The χ -dampened updating property implies $\chi/2 \leq \mu_S \leq 1 - \chi/2$. Because $1/3 < \mu_I = 1/2 < 6/7$, player 2 will choose C at information set I . To support this profile as an equilibrium, player 2 has to choose D at information set S , and hence, it must be the case that

$\mu_S \leq 2/7$. Coupled with the requirement of χ -dampened updating, the off-path belief has to satisfy $\chi/2 \leq \mu_S \leq 2/7$. That is, $[(I, I); (C, D)]$ is pooling χ -CSE if and only if $\chi/2 \leq 2/7$ or $\chi \leq 4/7$.

Case 4 [$m(\theta_1) = S, m(\theta_2) = S$]: Similar to the previous case, since player 1 follows a pooling strategy, player 2 will not update his belief at information set S , i.e., $\mu_S = 1/2$. Also, the χ -dampened updating property implies $\chi/2 \leq \mu_I \leq 1 - \chi/2$. Because $\mu_S > 2/7$, it is optimal for player 2 to choose C at information set S . To support this as an equilibrium, player 2 can choose either C or D at I .

Case 4.1: To make it a best reply for player 2 to choose D at information set I , it must be that $\mu_I \leq 1/3$. Combined with the requirement from χ -dampened updating, the off-path belief has to satisfy $\chi/2 \leq \mu_I \leq 1/3$. As a result, $[(S, S); (D, C)]$ is a pooling χ -CSE if and only if $\chi \leq 2/3$.

Case 4.2: To make it a best reply for player 2 to choose C at information set I , it must be that $1/3 \leq \mu_I \leq 6/7$. Combined with the requirement from χ -dampened updating, the off-path belief has to satisfy

$$\max \left\{ \frac{1}{2}\chi, \frac{1}{3} \right\} \leq \mu_I \leq \min \left\{ \frac{6}{7}, 1 - \frac{1}{2}\chi \right\}.$$

For any $\chi \in [0, 1]$, one can find μ_I that satisfies both inequalities. Hence, $[(S, S); (C, C)]$ is a pooling χ -CSE for any χ .

This completes the analysis of Game *BH 3* and Game *BH 4*. ■

Proof of Proposition 2.7

To prove this set of cost cutoffs form a χ -CSE, we need to show that there is no profitable deviation for any type at any subgame. First, at the second stage where there are exactly $0 \leq k \leq N - 1$ players sending 1 in the first stage, since no players will contribute, setting $C_k^\chi = 0$ is indeed a best response. At the subgame where all N players send 1 in the first stage, we use $\mu_i^\chi(c_{-i}|N)$ to denote player i 's cursed belief density. By Lemma 2.1, the cursed belief about all other players having a cost lower than c is simply

$$\begin{aligned} F^\chi(c) &\equiv \int_{\{c_j \leq c, \forall j \neq i\}} \mu_i^\chi(c'_{-i}|N) dc'_{-i} \\ &= \begin{cases} \chi (c/K)^{N-1} + (1 - \chi) (c/C_c^\chi)^{N-1} & \text{if } c \leq C_c^\chi \\ 1 - \chi + \chi (c/C_c^\chi)^{N-1} & \text{if } c > C_c^\chi, \end{cases} \end{aligned}$$

and C_N^χ is the solution of the fixed point problem of $C_N^\chi = F^\chi(C_N^\chi)$.

Moreover, in equilibrium, C_c^χ type of players would be indifferent between sending 1 and 0 in the communication stage. Thus, given C_N^χ , C_c^χ is the solution of the following equation

$$0 = \left(\frac{C_c^\chi}{K} \right)^{N-1} [-C_c^\chi + F^\chi(C_N^\chi)].$$

As a result, we obtain that in equilibrium, $C_c^\chi = C_N^\chi = F^\chi(C_N^\chi) \leq 1$ and denote this cost cutoff by $C^*(N, K, \chi)$. Substituting it into $F^\chi(c)$, gives

$$C^*(N, K, \chi) - \chi \left[\frac{C^*(N, K, \chi)}{K} \right]^{N-1} = 1 - \chi.$$

Next, we show that for any $N \geq 2$ and χ , the cutoff $C^*(N, K, \chi)$ is unique.

Case 1: When $N = 2$, the cutoff $C^*(2, K, \chi)$ is the unique solution of the linear equation

$$C^*(2, K, \chi) - \chi \left[\frac{C^*(2, K, \chi)}{K} \right] = 1 - \chi \iff C^*(2, K, \chi) = \frac{K - K\chi}{K - \chi}.$$

Case 2: For $N \geq 3$, we define the function $h(y) : [0, 1] \rightarrow \mathbb{R}$ where

$$h(y) = y - \chi \left(\frac{y}{K} \right)^{N-1} - (1 - \chi).$$

It suffices to show that $h(y)$ has a unique root in $[0, 1]$. When $\chi = 0$, $h(y) = y - 1$ which has a unique root at $y = 1$. In the following, we will focus on the case where $\chi > 0$. Since $h(y)$ is continuous, $h(0) = -(1 - \chi) < 0$ and $h(1) = \chi [1 - (1/K)^{N-1}] > 0$, there exists a root $y^* \in (0, 1)$ by the intermediate value theorem. Moreover, as we take the second derivative, we can find that for any $y \in (0, 1)$,

$$h''(y) = - \left(\frac{\chi}{K^{N-1}} \right) (N-1)(N-2)y^{N-3} < 0,$$

implying that $h(y)$ is strictly concave in $[0, 1]$. Furthermore, $h(0) < 0$ and $h(1) > 0$, so the root is unique, as illustrated in the left panel of Figure 2.3. This completes the proof. ■

Proof of Corollary 2.3

By Proposition 2.7, we know the cutoff $C^*(N, K, \chi) \leq 1$ and it satisfies

$$C^*(N, K, \chi) - \chi \left[\frac{C^*(N, K, \chi)}{K} \right]^{N-1} = 1 - \chi.$$

Therefore, when $\chi = 0$, the condition becomes $C^*(N, K, 0) = 1$. In addition, when $\chi = 1$, the condition becomes

$$C^*(N, K, 1) - \left[\frac{C^*(N, K, 1)}{K} \right]^{N-1} = 0,$$

implying $C^*(N, K, 1) = 0$.

For $\chi \in (0, 1)$, to prove $C^*(N, K, \chi)$ is strictly decreasing in N, K and χ , we consider a function $g(y; N, K, \chi) : (0, 1) \rightarrow \mathbb{R}$ where $g(y; N, K, \chi) = y - \chi[y/K]^{N-1}$. For any $y \in (0, 1)$ and fix any K and χ , we can observe that when $N \geq 2$,

$$g(y; N+1, K) - g(y; N, K) = -\chi \left[\frac{y}{K} \right]^N + \chi \left[\frac{y}{K} \right]^{N-1} > 0,$$

so $g(\cdot; N, K, \chi)$ is strictly increasing in N . Thus, the cutoff $C^*(N, K, \chi)$ is strictly decreasing in N . Similarly, for any $y \in (0, 1)$ and fix any N and χ , observe that when $K > 1$,

$$\frac{\partial g}{\partial K} = \chi(N-1) \left(\frac{y^{N-1}}{K^N} \right) > 0,$$

which implies that cutoff $C^*(N, K, \chi)$ is also strictly decreasing in K . For the comparative statics of χ , we can rearrange the equilibrium condition where

$$\frac{1 - C^*(N, K, \chi)}{\chi} = 1 - \left[\frac{C^*(N, K, \chi)}{K} \right]^{N-1}.$$

Since LHS is strictly decreasing in χ , the equilibrium cutoff is also strictly decreasing in χ . Finally, taking the limit on both sides of the equilibrium condition, we obtain

$$\lim_{N \rightarrow \infty} C^*(N, K, \chi) = \lim_{K \rightarrow \infty} C^*(N, K, \chi) = 1 - \chi. \quad \blacksquare$$

Proof of Claim 2.2

By backward induction, we know selfish player two will choose $T4$ for sure. Given that player two will choose $T4$ at stage four, it is optimal for selfish player one to choose $T3$. Now, suppose selfish player one will choose $P1$ with probability q_1 and player two will choose $P2$ with probability q_2 . Given this behavioral strategy profile, player two's belief about the other player being altruistic at stage two is $\mu = \alpha / [\alpha + (1 - \alpha)q_1]$. In this case, it is optimal for selfish player two to pass if and only if $32\mu + 4(1 - \mu) \geq 8 \iff \mu \geq 1/7$.

At the equilibrium, selfish player two is indifferent between $T2$ and $P2$. If not, say $32\mu + 4(1 - \mu) > 8$, player two will choose $P2$. Given that player two will choose $P2$,

it is optimal for selfish player one to choose $P1$, which makes $\mu = \alpha$ and $\alpha > 1/7$. However, we know $\alpha \leq 1/7$ which yields a contradiction. On the other hand, if $32\mu + 4(1 - \mu) < 8$, then it is optimal for player two to choose $T2$ at stage two. As a result, selfish player one would choose $T1$ at stage one, causing $\mu = 1$. In this case, player two would deviate to choose $P2$, yielding a contradiction. To summarize, in equilibrium, player two must be indifferent between $T2$ and $P2$, i.e., $\mu = 1/7$. As rearranging the equality, we can obtain that

$$\frac{\alpha}{\alpha + (1 - \alpha)q_1^*} = \frac{1}{7} \iff q_1^* = \frac{6\alpha}{1 - \alpha}.$$

Finally, since the equilibrium requires selfish player one to mix at stage one, selfish player one has to be indifferent between $P1$ and $T1$. Therefore,

$$4 = 16q_2^* + 2(1 - q_2^*) \iff q_2^* = \frac{1}{7}. \blacksquare$$

Proof of Proposition 2.8

By backward induction, we know selfish player two will choose $T4$ for sure. Given this, it is optimal for selfish player one to choose $T3$. Now, suppose selfish player one will choose $P1$ with probability q_1 and player two will choose $P2$ with probability q_2 . Given this behavioral strategy profile, by Lemma 1, player two's *cursed* belief about the other player being altruistic at stage 2 is

$$\mu^\chi = \chi\alpha + (1 - \chi) \left[\frac{\alpha}{\alpha + (1 - \alpha)q_1} \right].$$

In this case, it is optimal for player two to pass if and only if $32\mu^\chi + 4(1 - \mu^\chi) \geq 8 \iff \mu^\chi \geq 1/7$. We can first show that in equilibrium, it must be that $\mu^\chi \leq 1/7$. If not, then it is strictly optimal for player two to choose $P2$. Therefore, it is optimal for selfish player one to choose $P1$ and hence $\mu^\chi = \alpha \leq 1/7$, which yields a contradiction. In the following, we separate the discussion into two cases.

Case 1 $\left[\chi \leq \frac{6}{7(1-\alpha)} \right]$: In this case, we argue that player two is indifferent between $P2$ and $T2$. If not, then $32\mu^\chi + 4(1 - \mu^\chi) < 8$ and it is strictly optimal for player two to choose $T2$. This would cause selfish player one to choose $T1$ and hence $\mu^\chi = 1 - (1 - \alpha)\chi$. This yields a contradiction because $\mu^\chi = 1 - (1 - \alpha)\chi < 1/7 \iff \chi > 6/[7(1 - \alpha)]$. Therefore, in this case, player two is indifferent

between $T2$ and $P2$ and thus,

$$\begin{aligned}\mu^\chi = \frac{1}{7} &\iff \chi\alpha + (1-\chi) \left[\frac{\alpha}{\alpha + (1-\alpha)q_1^\chi} \right] = \frac{1}{7} \\ &\iff q_1^\chi = \left[\frac{7\alpha - 7\alpha\chi}{1 - 7\alpha\chi} - \alpha \right] / (1-\alpha).\end{aligned}$$

Since the equilibrium requires selfish player one to mix at stage 1, selfish player one has to be indifferent between $P1$ and $T1$. Therefore,

$$4 = 16q_2^\chi + 2(1 - q_2^\chi) \iff q_2^\chi = \frac{1}{7}.$$

Case 2 $\left[\chi > \frac{6}{7(1-\alpha)} \right]$: In this case, we know for any $q_1^\chi \in [0, 1]$,

$$\mu^\chi = \chi\alpha + (1-\chi) \left[\frac{\alpha}{\alpha + (1-\alpha)q_1^\chi} \right] \leq 1 - (1-\alpha)\chi < \frac{1}{7},$$

implying that it is strictly optimal for player two to choose $T2$, and hence it is strictly optimal for selfish player one to choose $T1$ at stage 1. This completes the proof. ■

Proof of Proposition 2.9

If $a^1(\theta_1) = b$ and all other types of voters as well as type θ_1 at stage 2 vote sincerely, voter i 's χ -cursed belief in the second stage upon observing $a_{-i}^1 = (a, b)$ is

$$\mu_i^\chi(\theta_{-i} | a_{-i}^1 = (a, b)) = \begin{cases} p_1 p_3 \chi + \frac{p_1}{p_1 + p_2} (1 - \chi) & \text{if } \theta_{-i} = (\theta_3, \theta_1) \\ p_2 p_3 \chi + \frac{p_2}{p_1 + p_2} (1 - \chi) & \text{if } \theta_{-i} = (\theta_3, \theta_2) \\ p_k p_l \chi & \text{otherwise.} \end{cases}$$

As mentioned in Section 2.4.4, a voter would act as if he perceives the other voters' (behavioral) strategies correctly in the last stage. However, misunderstanding the link between the other voters' types and actions would distort a voter's belief updating process. In other words, a voter would perceive the strategies correctly but form beliefs incorrectly. As a result, the continuation value of the a vs c subgame to a type θ_1 voter is simply the voter's χ -cursed belief, conditional on being pivotal, about there being at least one type θ_1 voter among his opponents. Similarly, the continuation value of the b vs c subgame is equal to the voter's conditional χ -cursed belief about there being at least one type θ_1 or θ_2 voter among his opponents

multiplied by v . Therefore, the continuation values to a type θ_1 voter in the two possible subgames of the second stage are (let $\tilde{p}_2 \equiv \frac{p_1}{p_1+p_2}$):

$$\begin{aligned} a \text{ vs } c &: \chi \left(1 - (1 - p_1)^2\right) + (1 - \chi)\tilde{p}_2 \\ b \text{ vs } c &: \left(1 - p_3^2\chi\right)v \end{aligned}$$

It is thus optimal for a type θ_1 voter to vote for b in the first stage if

$$\begin{aligned} \chi \left(1 - (1 - p_1)^2\right) + (1 - \chi)\tilde{p}_2 &\leq \left(1 - p_3^2\chi\right)v \\ \iff [2p_1 - p_1^2 - \tilde{p}_2 + p_3^2v]\chi &\leq v - \tilde{p}_2. \end{aligned} \quad (\text{A.4})$$

Notice that the statement would automatically hold when $\chi = 0$. In the following, we want to show that given v and p , if condition (A.4) holds for some $\chi \in (0, 1]$, then it will hold for all $\chi' \leq \chi$. As $\chi > 0$, we can rewrite condition (A.4) as

$$2p_1 - p_1^2 - \tilde{p}_2 + p_3^2v \leq \frac{v - \tilde{p}_2}{\chi}. \quad (\text{A.4}')$$

Case 1 [$v - \tilde{p}_2 < 0$]: We want to show that voting b in the first stage is never optimal for type θ_1 voter. That is, we want to show condition (A.4') never holds for $v < \tilde{p}_2$. To see this, we can first observe that the RHS is strictly increasing in χ . Therefore, it suffices to show

$$2p_1 - p_1^2 - \tilde{p}_2 + p_3^2v > v - \tilde{p}_2.$$

This is true because

$$\begin{aligned} 2p_1 - p_1^2 - \tilde{p}_2 + p_3^2v - (v - \tilde{p}_2) &= 2p_1 - p_1^2 - (1 - p_3^2)v \\ &> 2p_1 - p_1^2 - (1 + p_3)p_1 = p_1p_2 \geq 0, \end{aligned}$$

where the second inequality holds as $v < \frac{p_1}{p_1+p_2}$.

Case 2 [$v - \tilde{p}_2 \geq 0$]: Since the RHS of condition (A.4') is greater or equal to 0, it will weakly increase as χ decreases. Thus, if condition (A.4') holds for some $\chi \in (0, 1]$, it will also hold for all $\chi' \leq \chi$. This completes the proof. ■

Proof of Proposition 2.10

Assuming that all voters vote sincerely in both stages, voter i 's χ -cursed belief in the second stage upon observing $a_{-i}^1 = (a, b)$ is

$$\mu_i^\chi(\theta_{-i} | a_{-i}^1 = (a, b)) = \begin{cases} p_1 p_2 \chi + \frac{p_1}{p_1 + p_3} (1 - \chi) & \text{if } \theta_{-i} = (\theta_1, \theta_2) \\ p_2 p_3 \chi + \frac{p_3}{p_1 + p_3} (1 - \chi) & \text{if } \theta_{-i} = (\theta_3, \theta_2) \\ p_k p_l \chi & \text{otherwise.} \end{cases}$$

Similar to the proof of Proposition 2.9, the continuation values to a type θ_1 voter in the two possible subgames of the second stage are (let $\tilde{p}_3 \equiv \frac{p_1}{p_1 + p_3}$):

$$\begin{aligned} a \text{ vs } c : & \quad \chi \left(1 - (1 - p_1)^2 \right) + (1 - \chi) \tilde{p}_3 \\ b \text{ vs } c : & \quad \left(1 - p_3^2 \chi \right) v \end{aligned}$$

Thus, it is optimal for a type θ_1 voter to vote for a in the first stage if

$$\begin{aligned} & \chi \left(1 - (1 - p_1)^2 \right) + (1 - \chi) \tilde{p}_3 \geq \left(1 - p_3^2 \chi \right) v \\ \iff & \chi \left(2p_1 - p_1^2 - \tilde{p}_3 + p_3^2 v \right) \geq v - \tilde{p}_3. \end{aligned} \tag{A.5}$$

Case 1 [$v - \tilde{p}_3 > 0$]: In this case, we want to show that given p and v , there exists $\tilde{\chi}$ such that condition (A.5) holds if and only if $\chi \geq \tilde{\chi}$. Let $\tau \equiv 2p_1 - p_1^2 - \tilde{p}_3 + p_3^2 v$. If $\tau > 0$, then condition (A.5) holds if and only if $\chi \geq \tilde{\chi} \equiv \frac{v - \tilde{p}_3}{\tau}$. On the other hand, if $\tau \leq 0$, condition (A.5) will not hold for all $\chi \in [0, 1]$ and hence we can set $\tilde{\chi} = 2$.

Case 2 [$v - \tilde{p}_3 \leq 0$]: In this case, we want to show that given p and v , there exists $\tilde{\chi}$ such that condition (A.5) holds if and only if $\chi \leq \tilde{\chi}$. If $\tau < 0$, then condition (A.5) holds if and only if $\chi \leq \frac{v - \tilde{p}_3}{\tau}$ where the RHS is greater or equal to 0. On the other hand, if $\tau \geq 0$, then condition (A.5) will hold for any $\chi \in [0, 1]$ and hence we can again set $\tilde{\chi} = 2$. ■

Proof of Proposition 2.11

When observing a clean face, a player will know that he has a dirty face immediately. Therefore, choosing 1 (i.e., choosing D at stage 1) when observing a clean face is a strictly dominant strategy. In other words, for any $\chi \in [0, 1]$, $\hat{\sigma}^\chi(O) = 1$.

The analysis of the case where one observes a dirty face is separated into two cases.

Case 1 [$\chi > \bar{\alpha}$]: In this case, we show that $\hat{\sigma}^\chi(X) = T + 1$ is the only χ -CE. If not, suppose $\hat{\sigma}^\chi(X) = t$ where $t \leq T$ can be supported as a χ -CE. We can first notice that $\hat{\sigma}^\chi(X) = 1$ cannot be supported as a χ -CE because it is strictly dominated to choose 1 when observing a dirty face. For $2 \leq t \leq T$, given the other player $-i$ chooses $\hat{\sigma}^\chi(X) = t$, we can find player $-i$'s *average strategy* is

$$\bar{\sigma}_{-i}(j) = \begin{cases} 1-p & \text{if } j = 1 \\ p & \text{if } j = t \\ 0 & \text{if } j \neq 1, t. \end{cases}$$

Therefore, the other player $-i$'s χ -cursed strategy is

$$\sigma_{-i}^\chi(j|x_i = O) = \begin{cases} \chi(1-p) + (1-\chi) & \text{if } j = 1 \\ \chi p & \text{if } j = t \\ 0 & \text{if } j \neq 1, t, \end{cases} \quad \text{and}$$

$$\sigma_{-i}^\chi(j|x_i = X) = \begin{cases} \chi(1-p) & \text{if } j = 1 \\ \chi p + (1-\chi) & \text{if } j = t \\ 0 & \text{if } j \neq 1, t. \end{cases}$$

In this case, given (player i perceives that) player $-i$ chooses the χ -cursed strategy, player i 's expected payoff to choose $2 \leq j \leq t$ when observing a dirty face is

$$(1-p) [-\delta^{j-1} \chi p] + p \{ \delta^{j-1} \alpha [\chi p + (1-\chi)] \} = p \delta^{j-1} \underbrace{[\alpha - \chi(1+\alpha)(1-p)]}_{<0 \iff \chi > \bar{\alpha}} < 0.$$

Hence, given the other player chooses t when observing a dirty face, it is strictly dominated to choose any $j \leq t$. Therefore, the only χ -CE is $\hat{\sigma}^\chi(X) = T + 1$.

Case 2 [$\chi < \bar{\alpha}$]: In this case, we want to show that $\hat{\sigma}^\chi(X) = 2$ is the only χ -CE. If not, suppose $\hat{\sigma}^\chi(X) = t$ for some $t \geq 3$ can be supported as a χ -CE. We can again notice that since when observing a dirty face, it is strictly dominated to choose 1, 1 is never a best response. Given player $-i$ chooses $\hat{\sigma}^\chi(X) = t$, by the same calculation as in **Case 1**, the expected payoff to choose $2 \leq j \leq t$ is

$$p \delta^{j-1} \underbrace{[\alpha - \chi(1+\alpha)(1-p)]}_{>0 \iff \chi < \bar{\alpha}} > 0,$$

which is decreasing in j . Therefore, the best response to $\hat{\sigma}^\chi(X) = t$ is to choose 2 when observing a dirty face. Hence, the only χ -CE in this case is $\hat{\sigma}^\chi(X) = 2$. ■

Proof of Proposition 2.12

When observing a clean face, the player would know that his face is dirty. Thus, choosing D at stage 1 is a strictly dominant strategy, and $\tilde{\sigma}^\chi(O) = 1$ for all $\chi \in [0, 1]$. On the other hand, the analysis for the case where the player observes a dirty face consists of several steps.

Step 1: Assume that both players choosing D at some stage \bar{t} . We claim that at stage $t \leq \bar{t}$, the cursed belief $\mu^\chi(X|t, X) = 1 - (1 - p)\chi^{t-1}$. We can prove this by induction on t . At stage $t = 1$, the belief about having a dirty face is simply the prior belief p . Hence this establishes the base case. Now suppose the statement holds for any stage $1 \leq t \leq t'$ (and $t' < \bar{t}$). At stage $t' + 1$, by Lemma 2.1,

$$\begin{aligned}\mu^\chi(X|t'+1, X) &= \chi\mu^\chi(X|t', X) + (1 - \chi) = \chi \left[1 - (1 - p)\chi^{t'-1} \right] + (1 - \chi) \\ &= 1 - (1 - p)\chi^{t'},\end{aligned}$$

where the second equality holds by the induction hypothesis. This proves the claim.

Step 2: Given the cursed belief computed in the previous step, the expected payoff to choose D at stage t is

$$\begin{aligned}\mu^\chi(X|t, X)\alpha - [1 - \mu^\chi(X|t, X)] &= [1 - (1 - p)\chi^{t-1}] \alpha - [(1 - p)\chi^{t-1}] \\ &= \alpha - (1 - p)(1 + \alpha)\chi^{t-1},\end{aligned}$$

which is increasing in t . Notice that at the first stage, the expected payoff is $\alpha - (1 - p)(1 + \alpha) < 0$ by Assumption (2.1), so choosing U at stage 1 is strictly dominated. Furthermore, the player would choose U at every stage when observing a dirty face if and only if

$$\begin{aligned}\mu^\chi(X|T, X)\alpha - [1 - \mu^\chi(X|T, X)] \leq 0 &\iff \alpha - (1 - p)(1 + \alpha)\chi^{T-1} \leq 0 \\ &\iff \chi \geq \bar{\alpha}^{\frac{1}{T+1}}.\end{aligned}$$

As a result, both players choosing $\tilde{\sigma}^\chi(X) = T + 1$ is a χ -CSE if and only if $\chi \geq \bar{\alpha}^{\frac{1}{T+1}}$.

Step 3: In this step, we show both players choosing $\tilde{\sigma}^\chi(X) = 2$ is a χ -CSE if and only if $\chi \leq \bar{\alpha}$. We can notice that given the other player chooses D at stage 2, the player would know stage 2 would be the last stage regardless of his face type. Therefore, it is optimal to choose D at stage 2 as long as the expected payoff of D

at stage 2 is positive. Consequently, both players choosing $\tilde{\sigma}^\chi(X) = 2$ is a χ -CSE if and only if

$$\mu^\chi(X|2, X)\alpha - [1 - \mu^\chi(X|2, X)] \geq 0 \iff \alpha - (1 - p)(1 + \alpha)\chi \iff \chi \leq \bar{\alpha}.$$

Step 4: Given the other player chooses $\tilde{\sigma}^\chi(X) > t$, as the game reaches stage t , the belief about the other player choosing U at stage t is

$$\begin{aligned} & \underbrace{\mu^\chi(X|t, X)}_{\text{prob. of dirty}} [\chi\mu^\chi(X|t, X) + (1 - \chi)] \\ & + \underbrace{[1 - \mu^\chi(X|t, X)]}_{\text{prob. of clean}} [\chi\mu^\chi(X|t, X)] = \mu^\chi(X|t, X). \end{aligned}$$

Furthermore, we denote the expected payoff of choosing D at stage t as

$$\mathbb{E}[u^\chi(D|t, X)] \equiv \mu^\chi(X|t, X)\alpha - (1 - \mu^\chi(X|t, X)).$$

In the following, we claim that for any stage $2 \leq t \leq T - 2$, given the other player will stop at some stage later than stage $t + 2$ or never stop, if it is optimal to choose U at stage $t + 1$, then it is also optimal for you to choose U at stage t . That is,

$$\begin{aligned} & \mathbb{E}[u^\chi(D|t + 1, X)] < \delta\mu^\chi(X|t + 1, X)\mathbb{E}[u^\chi(D|t + 2, X)] \\ \implies & \mathbb{E}[u^\chi(D|t, X)] < \delta\mu^\chi(X|t, X)\mathbb{E}[u^\chi(D|t + 1, X)]. \end{aligned}$$

To prove this claim, first observe that

$$\begin{aligned} & \mathbb{E}[u^\chi(D|t + 1, X)] < \delta\mu^\chi(X|t + 1, X)\mathbb{E}[u^\chi(D|t + 2, X)] \\ \iff & (1 + \alpha)\mu^\chi(X|t + 1, X) - 1 < \delta\mu^\chi(X|t + 1, X) [(1 + \alpha)\mu^\chi(X|t + 2, X) - 1]. \end{aligned}$$

After rearrangement, the inequality is equivalent to

$$\delta\chi [\mu^\chi(X|t + 1, X)]^2 + \left[\delta(1 - \chi) - \frac{\delta}{1 + \alpha} - 1 \right] \mu^\chi(X|t + 1, X) + \frac{1}{1 + \alpha} > 0.$$

Consider a function $F : [0, 1] \rightarrow \mathbb{R}$ where

$$F(y) = \delta\chi y^2 + \left[\delta(1 - \chi) - \frac{\delta}{1 + \alpha} - 1 \right] y + \frac{1}{1 + \alpha}.$$

Since $\mu^\chi(X|j, X) = 1 - (1 - p)\chi^{j-1}$ is increasing in j , it suffices to show that there exists a unique $y^* \in (0, 1)$ such that F is single-crossing on $[0, 1]$ where $F(y^*) = 0$, $F(y) < 0$ for all $y > y^*$, and $F(y) > 0$ for all $y < y^*$. Because F is continuous and

- $F(0) = \frac{1}{1+\alpha} > 0$,
- $F(1) = \delta\chi + \left[\delta(1-\chi) - \frac{\delta}{1+\alpha} - 1 \right] + \frac{1}{1+\alpha} = -\frac{\alpha(1-\delta)}{1+\alpha} < 0$.

By intermediate value theorem, there exists a $y^* \in (0, 1)$ such that $F(y^*) = 0$. Moreover, y^* is the unique root of F on $[0, 1]$ because F is a strictly convex parabola and $F(1) < 0$. This establishes the claim.

Step 5: For any $3 \leq t \leq T$, in this step, we find the conditions to support both players choosing $\tilde{\sigma}^\chi(X) = t$ as a χ -CSE. We can first notice that both players choosing $\tilde{\sigma}^\chi(X) = t$ is a χ -CSE if and only if

1. $\mathbb{E}[u^\chi(D|t, X)] \geq 0$
2. $\mathbb{E}[u^\chi(D|t-1, X)] \leq \delta\mu^\chi(X|t-1, X)\mathbb{E}[u^\chi(D|t, X)]$.

Condition 1 is necessary because if it fails, then it is better for the player to choose U at stage t and get at least 0. Condition 2 is also necessary because if the condition doesn't hold, it would be profitable for the player to choose D before stage t . Furthermore, these two conditions are jointly sufficient to support $\tilde{\sigma}^\chi(X) = t$ as a χ -CSE by the same argument as **step 3**. From condition 1, we can obtain that

$$\begin{aligned} \mathbb{E}[u^\chi(D|t, X)] \geq 0 &\iff (1+\alpha)\mu^\chi(X|t, X) - 1 \geq 0 \\ &\iff 1 - (1-p)\chi^{t-1} \geq \frac{1}{1+\alpha} \iff \chi \leq \bar{\alpha}^{\frac{1}{t-1}}. \end{aligned}$$

In addition, by the calculation of **step 4**, we know

$$\mathbb{E}[u^\chi(D|t-1, X)] \leq \delta\mu^\chi(X|t-1, X)\mathbb{E}[u^\chi(D|t, X)] \iff F(\mu^\chi(X|t-1, X)) \geq 0,$$

which is equivalent to

$$\begin{aligned} \mu^\chi(X|t-1, X) &\leq \frac{\left[1 + \frac{\delta}{1+\alpha} - \delta(1-\chi)\right] - \sqrt{\left[1 + \frac{\delta}{1+\alpha} - \delta(1-\chi)\right]^2 - 4\delta\chi\left(\frac{1}{1+\alpha}\right)}}{2\delta\chi} \\ &= \frac{[(1+\alpha)(1+\delta\chi) - \alpha\delta] - \sqrt{[(1+\alpha)(1+\delta\chi) - \alpha\delta]^2 - 4\delta\chi(1+\alpha)}}{2\delta\chi(1+\alpha)} \equiv \kappa(\chi). \end{aligned}$$

Therefore, condition 2 holds if and only if

$$1 - (1-p)\chi^{t-2} \leq \kappa(\chi) \iff \chi \geq \left(\frac{1 - \kappa(\chi)}{1-p}\right)^{\frac{1}{t-2}}.$$

In summary, both players choosing $\tilde{\sigma}^\chi(X) = t$ is a χ -CSE if and only if

$$\left(\frac{1 - \kappa(\chi)}{1 - p}\right)^{\frac{1}{i-2}} \leq \chi \leq \bar{\alpha}^{\frac{1}{i-1}}.$$

This completes the proof. ■

Appendix B

A NOTE ON CURSED SEQUENTIAL EQUILIBRIUM AND
SEQUENTIAL CURSED EQUILIBRIUM

B.1 Introduction

In this appendix, we compare two recently proposed theories: the cursed sequential equilibrium (CSE) of Fong, Lin, and Palfrey (2023); and the sequential cursed equilibrium (SCE) of Cohen and Li (2023). Both generalize the cursed equilibrium (CE) by Eyster and Rabin (2005) to dynamic games, yet in different ways.

Following the specification of CE, CSE extends the concept of cursedness—the failure to fully account for the correlation between the types and actions of other players—to multi-stage games with publicly observed actions. After each stage of the game, each player in the game updates their current beliefs about the profile of other players’ types in a cursed way, i.e., based on the *average* behavioral strategies of the other players rather than the type conditional behavioral strategies of the other players. As in CE, CSE is generalized to χ -CSE, where the parameter $\chi \in [0, 1]$ indicates the degree of cursedness, with $\chi = 0$ corresponding to standard fully rational equilibrium behavior (i.e., updating based on the type conditional behavior strategies of the other players), and the case of $\chi = 1$ is *fully cursed* (i.e., updating based on the *average* behavioral strategies of the other players). Values of $\chi \in (0, 1)$ correspond to a mixture of these two extremes, defined similarly to CE, but with respect to behavioral strategies rather than mixed strategies.

The SCE extension of CE to dynamic games differs from χ -CSE in a number of ways, and we will illustrate through a series of examples some of these differences. Two differences are immediate and are due to a major difference in basic approach. First, while χ -CSE is developed for a subclass of games in extensive form with perfect recall, SCE is developed for a larger class: extensive form games with perfect recall. In such games, the “stages” and “public histories” do not generally exist or are not well-defined. As a result, instead of applying the cursed updating stage-by-stage, a player’s cursedness is defined in terms of coarsening the partition of other players’ information sets. In this sense, the difference between SCE and CSE reminds one of the differences between the CE and Analogy Based Expectations (ABEE) approach by Jehiel (2005) and Jehiel and Koessler (2008), where the

latter approach is based on the bundling of nodes at which other players move into analogy classes, which formally is a coarsening of information sets. Second, SCE is generalized by introducing *two* free parameters, $(\chi_S, \psi_S) \in [0, 1]^2$.¹ With two parameters, (χ_S, ψ_S) -SCE can distinguish inferences (and degree of neglect) based on past observed actions by other players and the inferences (and degree of neglect) based on simultaneous or future (hypothetical) strategies of the other players.

In addition to the family of applicable games and the dimension of parameter space, this appendix identifies and illustrates six additional differences between CSE and SCE. To formally illustrate these technical differences, in this appendix, we will focus on the class of multi-stage games with observed actions, using the framework of Fudenberg and Tirole (1991). We next introduce the framework and the two solution concepts in Section B.2. After that, we will discuss and illustrate the rest of the six differences between CSE and SCE in Section B.3 which are organized as follows:

- (1) the belief updating process (Section B.3.1),
- (2) the way of treating public histories (Section B.3.2),
- (3) effects in games of complete information (Section B.3.3),
- (4) violations of subgame perfection and sequential rationality (Section B.3.4),
- (5) the effect of re-labeling actions (Section B.3.5), and
- (6) effects in one-stage simultaneous-move games (Section B.3.6).

B.2 Preliminary

B.2.1 Multi-Stage Games with Observed Actions

Let $N = \{1, \dots, n\}$ be a finite set of players. Each player $i \in N$ has a *type* θ_i drawn from a finite set Θ_i . Let $\theta \in \Theta \equiv \times_{i=1}^n \Theta_i$ be the type profile and θ_{-i} be the type profile without player i . All players have the common (full support) prior distribution $F : \Theta \rightarrow (0, 1)$. At the beginning of the game, each player is told his own type, but is not informed anything about the types of others. Therefore, each player i 's initial belief about the types of others when his type is θ_i is:

$$F(\theta_{-i}|\theta_i) = \frac{F(\theta_{-i}, \theta_i)}{\sum_{\theta'_{-i} \in \Theta_{-i}} F(\theta'_{-i}, \theta_i)}.$$

¹To avoid confusion, we will henceforth add the subscript ‘‘S’’ to the parameters of SCE.

If the types are independent across players, each player i 's initial belief about the types of others is $F_{-i}(\theta_{-i}) = \prod_{j \neq i} F_j(\theta_j)$ where $F_j(\theta_j)$ is the marginal distribution of player j 's type.

The game is played in stages $t = 1, 2, \dots, T$ where $T < \infty$. In each stage, players simultaneously choose their actions, which will be revealed at the end of the stage. The feasible set of actions can vary with histories, so games with alternating moves are also included. Let \mathcal{H}^{t-1} be the set of all available *public* histories² at stage t , where $\mathcal{H}^0 = \{h_0\}$ and \mathcal{H}^T is the set of terminal histories. Let $\mathcal{H} = \cup_{t=0}^T \mathcal{H}^t$ be the set of all available histories of the game, and let $\mathcal{H} \setminus \mathcal{H}^T$ be the set of non-terminal histories.

For every player i , the available information at stage t is in $\mathcal{H}^{t-1} \times \Theta_i$. Thus, each player i 's information sets can be specified as $\mathcal{I}_i \in \Pi_i = \{(h, \theta_i) : h \in \mathcal{H} \setminus \mathcal{H}^T, \theta_i \in \Theta_i\}$. That is, a type θ_i player i 's information set at the public history h^t can be defined as $\cup_{\theta_{-i} \in \Theta_{-i}} (h^t, \theta_{-i}, \theta_i)$. With a slight abuse of notation, it will be denoted as (h^t, θ_i) . For the sake of simplicity, the feasible set of actions for every player at every history is assumed to be type-independent. Let $A_i(h^{t-1})$ be the feasible set of actions for player i at history h^{t-1} and let $A_i = \times_{h \in \mathcal{H} \setminus \mathcal{H}^T} A_i(h)$ be the set of player i 's all feasible actions in the game. For each player i , A_i is assumed to be finite and $|A_i(h)| \geq 1$ for any $h \in \mathcal{H} \setminus \mathcal{H}^T$. Let $a_i^t \in A_i(h^{t-1})$ be player i 's action at history h^{t-1} , and let $a^t = (a_1^t, \dots, a_n^t) \in \times_{i=1}^n A_i(h^{t-1})$ denote the action profile at stage t . If a^t is the action profile chosen at stage t , then $h^t = (h^{t-1}, a^t)$.

A behavioral strategy for player i is a function $\sigma_i : \Pi_i \rightarrow \Delta(A_i)$ satisfying $\sigma_i(h^{t-1}, \theta_i) \in \Delta(A_i(h^{t-1}))$. Let $\sigma_i(a_i^t | h^{t-1}, \theta_i)$ denote the probability for player i to choose $a_i^t \in A_i(h^{t-1})$. A strategy profile $\sigma = (\sigma_i)_{i \in N}$ specifies a behavioral strategy for each player i . Lastly, each player i has a payoff function (in von Neumann-Morgenstern utilities) $u_i : \mathcal{H}^T \times \Theta \rightarrow \mathbb{R}$, and let $u = (u_1, \dots, u_n)$ be the profile of utility functions. A multi-stage game with observed actions, Γ , is defined by the tuple $\Gamma = \langle N, \mathcal{H}, \Theta, F, u \rangle$.

B.2.2 CSE in Multi-Stage Games

Consider an assessment (μ, σ) , where μ is a belief system and σ is a behavioral strategy profile. The belief system specifies, for each player, a conditional distribu-

²Technically, each public history is not a *history* in general extensive form games as a public history only specifies the past action profiles but not the realized type profile. In the following, we will use the term *public history* when referring to the past action profiles, and use the term *history* when referring to the past action profiles with a type profile.

tion over the set of type profiles conditional on each *public history*. Following the spirit of the cursed equilibrium, for player i at the public history h^{t-1} , we define the *average behavioral strategy profile of the other players* as

$$\bar{\sigma}_{-i}(a_{-i}^t|h^{t-1}, \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} \mu_i(\theta_{-i}|h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^t|h^{t-1}, \theta_{-i})$$

for any $i \in N$, $\theta_i \in \Theta_i$ and $h^{t-1} \in \mathcal{H}^{t-1}$.

A χ -CSE is parameterized by a single parameter $\chi \in [0, 1]$. Instead of thinking the other players are using σ_{-i} , a χ -cursed type θ_i player i would believe they are using a χ -weighted average of the average behavioral strategy and the true behavioral strategy:

$$\sigma_{-i}^\chi(a_{-i}^t|h^{t-1}, \theta_{-i}, \theta_i) = \chi \bar{\sigma}_{-i}(a_{-i}^t|h^{t-1}, \theta_i) + (1 - \chi) \sigma_{-i}(a_{-i}^t|h^{t-1}, \theta_{-i}).$$

The beliefs of player i about the type profile θ_{-i} are updated via Bayes' rule in χ -CSE, whenever possible, assuming others are using the χ -cursed behavioral strategy rather than the true behavioral strategy. Specifically, an assessment satisfies the *χ -cursed Bayes' rule* if the belief system is derived from Bayes' rule while perceiving others are using σ_{-i}^χ rather than σ_{-i} :

Definition B.1. (μ, σ) satisfies *χ -cursed Bayes' rule* if the following is applied to update the posterior beliefs as $\sum_{\theta'_{-i} \in \Theta_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \sigma_{-i}^\chi(a_{-i}^t|h^{t-1}, \theta'_{-i}, \theta_i) > 0$:

$$\mu_i(\theta_{-i}|h^t, \theta_i) = \frac{\mu_i(\theta_{-i}|h^{t-1}, \theta_i) \sigma_{-i}^\chi(a_{-i}^t|h^{t-1}, \theta_{-i}, \theta_i)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \sigma_{-i}^\chi(a_{-i}^t|h^{t-1}, \theta'_{-i}, \theta_i)}.$$

Finally, χ -CSE places a consistency restriction, analogous to consistent assessments in sequential equilibrium, on how χ -cursed beliefs are updated off the equilibrium path:

Definition B.2. Let Ψ^χ be the set of assessments (μ, σ) such that σ is a totally mixed behavioral strategy profile and μ is derived from σ using χ -cursed Bayes' rule. (μ, σ) satisfies *χ -consistency* if there is a sequence of assessments $\{(\mu^k, \sigma^k)\} \subseteq \Psi^\chi$ such that $\lim_{k \rightarrow \infty} (\mu^k, \sigma^k) = (\mu, \sigma)$.

For any $i \in N$, $\chi \in [0, 1]$, σ , and $\theta \in \Theta$, let $\rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi, \sigma_i)$ be player i 's perceived conditional realization probability of terminal history $h^T \in \mathcal{H}^T$ at the public history $h^t \in \mathcal{H} \setminus \mathcal{H}^T$ if the type profile is θ and player i uses the behavioral

strategy σ_i whereas perceives other players' using the cursed behavioral strategy σ_{-i}^χ . At every non-terminal public history h^t , a χ -cursed player in χ -CSE will use χ -cursed Bayes' rule (Definition B.1) to derive the posterior belief about the other players' types. Accordingly, a type θ_i player i 's conditional expected payoff at the public history h^t is given by

$$\mathbb{E}u_i(\sigma|h^t, \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i(\theta_{-i}|h^t, \theta_i) \rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi, \sigma_i) u_i(h^T, \theta_{-i}, \theta_i).$$

Definition B.3. An assessment (μ^*, σ^*) is a χ -cursed sequential equilibrium if it satisfies χ -consistency and $\sigma_i^*(h^t, \theta_i)$ maximizes $\mathbb{E}u_i(\sigma^*|h^t, \theta_i)$ for all i , θ_i , $h^t \in \mathcal{H} \setminus \mathcal{H}^T$.

B.2.3 SCE in Multi-Stage Games

Some additional notation is required to define sequential cursed equilibrium (SCE) in a multi-stage game. First, let λ denote nature. In the framework of multi-stage games with observed actions, λ only moves once in stage 0, i.e., at the initial history, denoted h_0 . The action set for λ is the set of states of nature, i.e., the set of profiles of player types, $\Theta = \times_{i=1}^n \Theta_i$. The totally mixed strategy of nature, denoted by $\sigma_\lambda(h_0) \in \Delta(\Theta)$, is common knowledge to all players $i \in N$ and exogenously given by F . Note that nature's information set is singleton. Also, note that after nature's action (i.e., type profile) is realized, θ_i will be observed only by player i but not by the other players. That is, each player's type is his own private information.

Second, let $\mu[\sigma]$ denote the probability measure on $\mathcal{H}^T \times \Theta$ that is induced by $\sigma \equiv (\sigma_1, \dots, \sigma_n, \sigma_\lambda)$. Following Cohen and Li (2023), we allow $\mu[\sigma]$ to denote the probability measure on $\mathcal{H}^T \times \Theta$ for some $t < T$ by viewing $((a^1, \dots, a^t), \theta') \in \mathcal{H}^t \times \Theta$ as $\{(h^T, \theta) \in \mathcal{H}^T \times \Theta : (a^1, \dots, a^t) < h^T, \theta = \theta'\}$.

Third, let \mathcal{P} denote the *coarsest (valid) partition* of the non-terminal histories $\mathcal{H} \setminus \mathcal{H}^T \times \Theta$,³ and $P(\mathcal{I}_i)$ denotes $P \in \mathcal{P}$ such that $\mathcal{I}_i \subseteq P$. Note that a partition \mathcal{P}' is coarser than \mathcal{P}'' if, given any $P'' \in \mathcal{P}''$, there exists a cell of \mathcal{P}' that contains P'' .

Finally, given $\mathcal{I}_j = (h^t, \theta_j)$, let $\sigma_i^{\mathcal{I}_j} : \Pi_i^{\mathcal{I}_j} \rightarrow \Delta(A_i)$ denote a *partial strategy* for player i at \mathcal{I}_j , where $\Pi_i^{\mathcal{I}_j} = \{(h', \theta'_i) : h' \in \mathcal{H} \setminus \mathcal{H}^T, h^t \leq h' \text{ or } h' \leq h^t, \theta'_i \in \Theta_i\}$. In other words, $\sigma_i^{\mathcal{I}_j}$ is a function from information sets containing public histories compatible with h^t to distributions over player i 's feasible actions.

³A partition of the non-terminal histories is invalid if it defines information sets that violate the assumption of perfect recall.

To define a SCE, Cohen and Li (2023) define three types of conjectures about opponents' and nature's strategies: the *sequentially cursed conjecture* ($\tilde{\sigma}$), the *typically cursed conjecture* ($\check{\sigma}$), and the *Bayesian conjecture* ($\hat{\sigma}$).

Consider a non-terminal information set $\mathcal{I}_i = (h^t, \theta_i)$ where $h^t = (a^1, \dots, a^t)$. Given a totally mixed behavioral strategy σ , Cohen and Li (2023) define the three types of conjectures of player i about player $j \neq i$ at information set \mathcal{I}_i as follows:⁴

1. *Sequentially cursed conjecture*: $\tilde{\sigma}_j^{\mathcal{I}_i}(\cdot)$ is said to be player i 's *sequentially cursed conjecture* about j 's behavioral strategy at \mathcal{I}_i if, for all $\mathcal{I}_j = (h', \theta_j)$ and all $a_j \in A_j(h')$,

$$\tilde{\sigma}_j^{\mathcal{I}_i}(a_j | \mathcal{I}_j) \equiv \mu[\sigma](\{(h, \theta), a_j : (h, \theta) \in P(\mathcal{I}_j)\} | \mathcal{I}_i \cap P(\mathcal{I}_j)).$$

2. *Typically cursed conjecture*: $\check{\sigma}_j^{\mathcal{I}_i}(\cdot)$ is said to be player i 's *typically cursed conjecture* about j 's behavioral strategy at \mathcal{I}_i if, for all $\mathcal{I}_j = (h', \theta_j)$ and all $a_j \in A(h')$,

$$\check{\sigma}_j^{\mathcal{I}_i}(a_j | \mathcal{I}_j) \equiv \mu[\sigma](\{(h, \theta), a_j : (h, \theta) \in P(\mathcal{I}_j)\} | (\cdot, \theta_i) \cap P(\mathcal{I}_j))$$

where $(\cdot, \theta_i) \equiv \{(h, \theta_{-i}, \theta_i) : h \in \mathcal{H} \setminus \mathcal{H}^T, \theta_{-i} \in \Theta_{-i}\}$.

3. *Bayesian conjecture*: $\hat{\sigma}_j^{\mathcal{I}_i}(\cdot)$ is said to be player i 's *Bayesian conjecture* about j 's behavioral strategy at \mathcal{I}_i if, for all $\mathcal{I}_j = (h', \theta_j)$ and all $a_j \in A(h')$,

$$\hat{\sigma}_j^{\mathcal{I}_i}(a_j | \mathcal{I}_j) \equiv \mu[\sigma](\{(h, \theta), a_j : (h, \theta) \in \mathcal{I}_j\} | \mathcal{I}_i \cap \mathcal{I}_j).$$

A (χ_S, ψ_S) -SCE has two parameters, $(\chi_S, \psi_S) \in [0, 1]^2$. At every information set \mathcal{I}_i , a (χ_S, ψ_S) -sequentially cursed player i would perceive that the game has been and will be played by the other players according to his sequentially cursed, typically cursed, and Bayesian conjectures with probability $\chi_S \psi_S$, $\chi_S(1 - \psi_S)$, and $1 - \chi_S$, respectively. Based on such perceptions, a (χ_S, ψ_S) -sequential cursed equilibrium is defined as follows:

Definition B.4. A strategy profile σ^* is said to be a (χ_S, ψ_S) -*sequential cursed equilibrium* if, given any $i \in N$ and $\mathcal{I}_i = (h^t, \theta_i)$ with $h^t = (a^1, \dots, a^t)$, there exists a partial strategy $\sigma_i^{\mathcal{I}_i}$ such that $\sigma_i^{\mathcal{I}_i}(\mathcal{I}_i) = \sigma_i^*(\mathcal{I}_i)$ and $\sigma_i^{\mathcal{I}_i}$ maximizes i 's expected utility under the belief that θ_{-i} is realized with probability $\mu[\tilde{\sigma}^{\mathcal{I}_i}](h^t, \theta_{-i}, \theta_i | \mathcal{I}_i)$ and that the game will proceed according to $\tilde{\sigma}_{-i}^{\mathcal{I}_i}$, where

⁴For general behavioral strategy profiles, SCE imposes a consistency requirement similar to Kreps and Wilson (1982) as it requires the limits of all three types of conjectures exist.

1. $\bar{\sigma}_i^{I_i}(a^{t'} | h^{t'-1}, \theta_i) = 1$ for all $t' \leq t$
2. $\bar{\sigma}_i^{I_i} = \check{\sigma}_i^{I_i}$ with probability $\chi_S \psi_S$
3. $\bar{\sigma}_i^{I_i} = \check{\sigma}_i^{I_i}$ with probability $\chi_S(1 - \psi_S)$
4. $\bar{\sigma}_i^{I_i} = \hat{\sigma}_i^{I_i}$ with probability $(1 - \chi_S)$

When $\chi_S = 0$, a (χ_S, ψ_S) -SCE reduces to a sequential equilibrium. A sequential cursed equilibrium refers to a boundary case when $\chi_S = \psi_S = 1$ in a (χ_S, ψ_S) -SCE.

B.3 Differences Between CSE and SCE

B.3.1 Difference in the Belief Updating Dynamics

One key difference between the two equilibrium concepts is in the belief updating process about nature's strategy. In the framework of CSE, players update their beliefs about the other players' type profile *at each public history* by χ -cursed Bayes' rule, which is stated in Claim B.1. In the case of totally mixed strategies, the updating process is very simple to characterize:

Claim B.1. *For any χ -cursed sequential equilibrium (μ, σ) where σ is a totally mixed behavioral strategy profile, any $i \in N$, any non-terminal history $h^t = (h^{t-1}, a^t) \in \mathcal{H} \setminus \mathcal{H}^T$ and any $\theta \in \Theta$,*

$$\mu_i(\theta_{-i} | h^t, \theta_i) = \chi \mu_i(\theta_{-i} | h^{t-1}, \theta_i) + (1 - \chi) \left[\frac{\mu_i(\theta_{-i} | h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^t | h^{t-1}, \theta_{-i})}{\sum_{\theta'_{-i}} \mu_i(\theta'_{-i} | h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^t | h^{t-1}, \theta'_{-i})} \right]. \quad (\text{B.1})$$

On the other hand, SCE is defined as a mixture of three different kinds of conjectures. The mathematical object that corresponds to the belief of SCE is the *conjecture about nature's strategy*. In the following, we will characterize the evolution of this conjecture. In general, the characterization of the belief updating process about nature's strategy can be quite complicated. To simplify the comparison between CSE and SCE, it is instructive to first examine the comparison in games where the coarsest valid partition, \mathcal{P} , accords with public histories, in the sense that \mathcal{P} is measurable with respect to public histories. Later examples in this appendix will explore some of the implications if this condition is not satisfied.

Definition B.5. *The coarsest valid partition, \mathcal{P} , satisfies **Public History Consistency (PHC)** if, for all $t < T, i \in N, \theta_i \in \Theta_i, h^t \in \mathcal{H}^t, \hat{h}^t \in \mathcal{H}^t$:*

$$h^t \neq \hat{h}^t \Rightarrow P(h^t, \theta_i) \neq P(\hat{h}^t, \theta_i).$$

Following Definition B.5, we can explicitly characterize $P(\mathcal{I}_i)$ under PHC, as demonstrated in Lemma B.1:

Lemma B.1. *The coarsest valid partition, \mathcal{P} , satisfies Public History Consistency if and only if, for all $t < T$, $i \in N$, $\theta_i \in \Theta_i$, and $h^t \in \mathcal{H}^t$, we can obtain that $P(h^t, \theta_i) = \bigcup_{\theta \in \Theta} (h^t, \theta)$.*

Proof: (\Rightarrow): We prove by contrapositive. Suppose that there exists some non-terminal information set (h^t, θ_i) such that $P(h^t, \theta_i) \neq \bigcup_{\theta \in \Theta} (h^t, \theta)$, then by the definition of the coarsest partition, we have $P(h^t, \theta_i) \supsetneq \bigcup_{\theta \in \Theta} (h^t, \theta)$. However, in games of perfect recall, this implies that there is some $P \in \mathcal{P}$ and $\hat{h}^t \in \mathcal{H}^t$ such that $h^t \neq \hat{h}^t$ but $(h^t, \theta_i) \subset P$ and $(\hat{h}^t, \theta_i) \subset P$, which contradicts Definition B.5.

(\Leftarrow): Definition B.5 follows immediately from the statement. ■

A key implication of PHC is that no matter how cursed the players are, they know that everyone's behavioral strategy depends on the realized public history. In other words, this condition shuts down one source of cursedness: players fully understand that other players' future actions are conditional on current and past actions. Other implications of the PHC condition will be discussed in later sections. Of course, it is still the case that at any history the players are cursed in the sense of neglecting the dependence of other players' current and future actions on the type profile (i.e., nature's initial move).⁵

Fix any totally mixed behavioral strategy profile σ and any information set $\mathcal{I}_i = (h^t, \theta_i)$ where $h^t = (h^{t-1}, a^t)$. Under SCE, player i 's belief about other players' type profile (Nature's initial move) is the weighted average of the conjectures about nature's strategy. Specifically, for any $(\chi_S, \psi_S) \in [0, 1]^2$, the belief is simply

$$\begin{aligned} \mu_i^{(\chi_S, \psi_S)}(\theta_{-i} | h^t, \theta_i) \equiv & \underbrace{\chi_S \psi_S \tilde{\sigma}_\lambda^{(h^t, \theta_i)}(\theta | h_\emptyset)}_{\text{sequentially cursed}} + \underbrace{\chi_S (1 - \psi_S) \check{\sigma}_\lambda^{(h^t, \theta_i)}(\theta | h_\emptyset)}_{\text{typically cursed}} \\ & + (1 - \chi_S) \underbrace{\hat{\sigma}_\lambda^{(h^t, \theta_i)}(\theta | h_\emptyset)}_{\text{Bayesian}}. \end{aligned}$$

⁵PHC may seem like a weak condition, but for most games the coarsest valid partition in SCE violates the condition. For example, it will be violated even in very simple games of perfect information, where there is no initial move by nature, and no simultaneous moves, as we illustrate later in this appendix. One alternative to avoid this would be to modify the definition of SCE so that the coarsest valid partition, \mathcal{P} , is *required* to satisfy PHC, leading to a different equilibrium concept, which could be called *Public Sequential Cursed Equilibrium (PSCE)*. For multi-stage games with *public* histories, PHC seems like a plausible requirement for valid coarsening.

If PHC is satisfied, one can relatively easily characterize the evolution of the (χ_S, ψ_S) -SCE conjecture about nature. This is done next in Claim B.2.

Claim B.2. *Under PHC, for any (χ_S, ψ_S) -sequential cursed equilibrium where σ is a totally mixed behavioral strategy profile, and any non-terminal information set (θ_i, h^t) where $h^t = (a^1, \dots, a^t)$, the conjecture about nature's strategy is:*

$$\mu_i^{(\chi_S, \psi_S)}(\theta_{-i}|h^t, \theta_i) = \chi_S(1 - \psi_S)F(\theta_{-i}|\theta_i) + [1 - \chi_S(1 - \psi_S)]\mu_i^*(\theta_{-i}|h^t, \theta_i) \quad (\text{B.2})$$

where

$$\mu_i^*(\theta_{-i}|h^t, \theta_i) = \frac{F(\theta_{-i}|\theta_i) \left[\prod_{j \neq i} \prod_{1 \leq l \leq t} \sigma_j(a_j^l|h^{l-1}, \theta_j) \right]}{\sum_{\theta'_i \in \Theta_{-i}} F(\theta'_{-i}|\theta_i) \left[\prod_{j \neq i} \prod_{1 \leq l \leq t} \sigma_j(a_j^l|h^{l-1}, \theta'_j) \right]}.$$

Proof: To obtain the belief of (χ_S, ψ_S) -SCE under PHC, we need to derive the sequentially cursed conjecture, typically cursed conjecture and the Bayesian conjecture of nature separately.

1. Player i 's **Bayesian conjecture** about nature's strategy is

$$\begin{aligned} \hat{\sigma}_\lambda^{(h^t, \theta_i)}(\theta|h_\emptyset) &= \mu[\sigma](h_\emptyset, \theta|(h^t, \theta_i) \cap h_\emptyset) \\ &= \frac{F(\theta_{-i}, \theta_i) \left[\prod_{j \neq i} \prod_{1 \leq l \leq t} \sigma_j(a_j^l|h^{l-1}, \theta_j) \right]}{\sum_{\theta'_i \in \Theta_{-i}} F(\theta'_{-i}, \theta_i) \left[\prod_{j \neq i} \prod_{1 \leq l \leq t} \sigma_j(a_j^l|h^{l-1}, \theta'_j) \right]} \\ &= \frac{F(\theta_{-i}|\theta_i) \left[\prod_{j \neq i} \prod_{1 \leq l \leq t} \sigma_j(a_j^l|h^{l-1}, \theta_j) \right]}{\sum_{\theta'_i \in \Theta_{-i}} F(\theta'_{-i}|\theta_i) \left[\prod_{j \neq i} \prod_{1 \leq l \leq t} \sigma_j(a_j^l|h^{l-1}, \theta'_j) \right]} \\ &\equiv \mu_i^*(\theta_{-i}|h^t, \theta_i). \end{aligned}$$

2. Player i 's **sequentially cursed conjecture** about nature's strategy is

$$\check{\sigma}_\lambda^{(h^t, \theta_i)}(\theta|h_\emptyset) = \mu[\sigma](h_\emptyset, \theta|(h^t, \theta_i) \cap P(h_\emptyset)) = \mu_i^*(\theta_{-i}|h^t, \theta_i),$$

which coincides with the Bayesian conjecture. Note that $P(h_\emptyset) = h_\emptyset$ since the nature's information set is singleton.

3. Player i 's **typically cursed conjecture** about nature's strategy is

$$\check{\sigma}_\lambda^{(h^t, \theta_i)}(\theta|h_\emptyset) = \mu[\sigma](h_\emptyset, \theta|(\cdot, \theta_i) \cap h_\emptyset) = \frac{F(\theta_{-i}, \theta_i)}{F(\theta_i)} = F(\theta_{-i}|\theta_i),$$

which is just the common prior conditional on player i 's type.

Combined these three conjectures, we can obtain that $\mu_i^{(\chi_S, \psi_S)}(\theta_{-i}|h^t, \theta_i) =$

$$\begin{aligned} & \chi_S \psi_S \tilde{\sigma}_\lambda^{(h^t, \theta_i)}(\theta|h_0) + \chi_S(1 - \psi_S) \check{\sigma}_\lambda^{(h^t, \theta_i)}(\theta|h_0) + (1 - \chi_S) \hat{\sigma}_\lambda^{(h^t, \theta_i)}(\theta|h_0) \\ & = \chi_S(1 - \psi_S) F(\theta_{-i}|\theta_i) + [1 - \chi_S(1 - \psi_S)] \mu_i^*(\theta_{-i}|h^t, \theta_i). \end{aligned}$$

This completes the characterization of the evolution of the conjecture. ■

From the proof of Claim B.2, we can find that under the typically cursed conjecture, players do not update their beliefs about others' types while under the sequentially cursed and the Bayesian conjectures, players will update their beliefs correctly—in the sense that their conjecture about nature's strategy coincides with the Bayesian posterior. By contrast, under the framework of CSE, players update their beliefs via Bayes' rule while having incorrect perceptions about other players' behavioral strategies. Therefore, as characterized by Claim B.1, in χ -CSE the posterior belief about others' types at stage t is a weighted average between the posterior belief at stage $t - 1$ and the Bayesian posterior.

The results from Claim B.1 and Claim B.2 sharply contrast the difference in the belief updating process of the two theories. Under the framework of (χ_S, ψ_S) -SCE, a player perceives the others' strategies by weighting over cursed and actual strategies formed under different partitions of *entire game trees*. As a result, the posterior belief induced by a (χ_S, ψ_S) -SCE at every stage is a weighted average over two extreme cases—the prior belief (i.e., no belief updating) and the Bayesian posterior belief implied by the actual strategy profile (i.e., correct Bayesian belief updating). Alternatively, under the framework of χ -CSE, a player perceives the others' strategies by weighting between cursed (average) and actual strategies *at each stage*. The player would update his belief by applying Bayes rule to his belief in the previous stage, but with an incorrect perception about the behavioral strategy used in that stage. As a result, the posterior belief induced by a χ -CSE at stage t is a weighted average over two cases—the belief at stage $t - 1$, and the Bayesian posterior belief implied jointly by that belief and actual behavioral strategy at $t - 1$.

Another implication of Claim B.1 and Claim B.2 concerns the anchoring of the belief evolution to the original prior (nature's move). Under the framework of (χ_S, ψ_S) -SCE, a player always puts $\chi_S(1 - \psi_S)$ weight on his prior belief about the joint distribution of types when forming the posterior belief. In this sense, the impact of prior on posterior is persistent and independent of which period the player is currently in. By contrast, under the framework of χ -CSE, a player at stage t puts a

χ weight on his belief in previous stage, which has typically evolved over $t - 1$ stages. The impact of prior on posterior belief thus diminishes over time compared to the posterior beliefs formed in the latest periods. The following example highlights this difference in the belief updating processes and sheds light on how they reflect distinct views about (cursed) learning behavior in dynamic games.

Before diving into the illustrative example, in Remark B.1, we derive a player's conjectures about other players' strategies in past and future events under PHC. This remark is useful for understanding the mechanism of the belief updating in SCE and useful for the calculations in the illustrative example.

Remark B.1. Fix any non-terminal information set $\mathcal{I}_i = (h^t, \theta_i)$ such that there are public histories $h^{t'}, h^{t''}$ with $h^{t'} < h^t \leq h^{t''}$. Let $h^{t''} \equiv (a^1, \dots, a^{t'}, \dots, a^t, \dots, a^{t''})$. Following the definition of SCE, under PHC, we can find that player i 's conjectures about type θ_j player j 's partial strategy at the past public history $h^{t'}$ are:

1. The **sequentially cursed conjecture** coincides with the **Bayesian conjecture**:

$$\tilde{\sigma}_j^{\mathcal{I}_i}(\tilde{a}_j^{t'+1}|h^{t'}, \theta_j) = \hat{\sigma}_j^{\mathcal{I}_i}(\tilde{a}_j^{t'+1}|h^{t'}, \theta_j) = \begin{cases} 1 & \text{if } \tilde{a}_j^{t'+1} = a_j^{t'+1} \\ 0 & \text{if } \tilde{a}_j^{t'+1} \neq a_j^{t'+1} \end{cases}$$

2. The **typically cursed conjecture** is the average behavioral strategy:

$$\tilde{\sigma}_j^{\mathcal{I}_i}(a_j^{t'+1}|h^{t'}, \theta_j) = \sum_{\theta'_j \in \Theta_j} \mu_i^*(\theta'_j|h^{t'}, \theta_i) \sigma_j(a_j^{t'+1}|h^{t'}, \theta'_j).$$

On the other hand, player i 's conjectures about type θ_j player j 's partial strategy at the future public history $h^{t''}$ are:

1. The **sequentially cursed conjecture** coincides with the **typically cursed conjecture**:

$$\tilde{\sigma}_j^{\mathcal{I}_i}(a_j^{t''+1}|h^{t''}, \theta_j) = \check{\sigma}_j^{\mathcal{I}_i}(a_j^{t''+1}|h^{t''}, \theta_j) = \sum_{\theta'_j \in \Theta_j} \mu_i^*(\theta'_j|h^{t''}, \theta_i) \sigma_j(a_j^{t''+1}|h^{t''}, \theta'_j).$$

2. The **Bayesian conjecture** coincides with the true behavioral strategy:

$$\hat{\sigma}_j^{\mathcal{I}_i}(\tilde{a}_j^{t''+1}|h^{t''}, \theta_j) = \sigma_j(\tilde{a}_j^{t''+1}|h^{t''}, \theta_j).$$

In summary, a player who is sequentially cursed accurately perceives another player's past behavior by assigning a probability of 1 to the actions that have been played, but assumes that the other player will use an average strategy (averaged based on μ^*) going forward. On the other hand, a player who is typically cursed perceives both the past and future strategies of another player in the same way, assuming that the other player uses the average strategy.

Illustrative Example

Consider the following game where there is one broadcaster (B) and two other players. We denote the set of players as $N = \{B, 1, 2\}$. There are two possible states—good (θ_g) or bad (θ_b)—with equal probability. At the beginning of the game, nature will draw the true state. Only the broadcaster observes the state, not the other two players. Therefore, the true state is the broadcaster's private information.

In this game, all players except for the broadcaster will take turns to choose either a safe option (s) or a risky option (r). Without loss of generality, we let player i be the i -th mover for any $i \in \{1, 2\}$. Each of these two players only takes one action and all actions of all players are public. Moreover, before each player $i \in \{1, 2\}$ makes a move, the broadcaster makes a public announcement about the state being good (g) or bad (b).

Technically speaking, this is a three-player four-stage multi-stage game with observed actions where player B moves at odd stages ($t = 1, 3$) with the action set $A_B(h^t) = \{g^{h^t}, b^{h^t}\}$ and each player $i \in \{1, 2\}$ moves at stage $2i$ with the action set $A_i(h^{2i}) = \{s^{h^{2i}}, r^{h^{2i}}\}$. Formally, this game satisfies PHC since the available actions at different histories are labeled differently.⁶ We drop the superscripts h^t and h^{2i} when there is no risk of confusion.

The broadcaster will get one unit payoff for each truthful announcement and 0 otherwise. Therefore, it is strictly optimal for the broadcaster to always report the true state. For other players, they will get 0 for sure if they choose the safe option s . If they choose the risky option r , they will get $\alpha \in (0, 1)$ if the state is θ_g and -1 if the stage is θ_b .

In summary, the broadcaster has private information about the true state, and is incentivized to always truthfully report the state. All other players do not know the

⁶See Section B.3.5 for more discussion on the labeling issue.

true state, but they are incentivized to choose r in their turn if they are sufficiently confident about the state being θ_g . The standard equilibrium theory predicts the broadcaster will always announce the true state and every player $i \in \{1, 2\}$ will choose s if the announcement is θ_b and choose r otherwise.

To highlight how the belief updating processes differ under SCE and CSE, we will focus on the case in which the state is good (θ_g) and characterize player i 's best response to the broadcaster's announcement(s) for $i \in \{1, 2\}$. Because the solutions have the structure that if it is optimal for some player i to choose r , then it is also optimal for $i+1$ to choose r , in the following, we characterize the solutions by finding the first player to choose r . Notice that there are three possible cases, $\{1, 2, \bar{s}\}$, which corresponds to the cutoff player or the case where both players choose s . We use i^* and i^{**} to denote the cutoff players predicted by SCE and CSE, respectively. Claim B.3 shows that, in every SCE, the two players other than the broadcaster would take identical actions—either $i^* = 1$ (both choose r) or $i^* = \bar{s}$ (both choose s).

Claim B.3. *In a (χ_S, ψ_S) -SCE, when the broadcaster announces g , the cutoff player i^* is characterized as the following:*

1. $i^* = 1$ if and only if $\chi_S(1 - \psi_S) \leq \frac{2\alpha}{(1+\alpha)}$, and
2. $i^* = \bar{s}$ if and only if $\chi_S(1 - \psi_S) \geq \frac{2\alpha}{(1+\alpha)}$.

Proof: Consider the case where the broadcaster announces g and fix any player $i \in \{1, 2\}$. By Claim B.2, we can obtain that player i 's belief about θ_g is

$$\begin{aligned} \mu_i^{(\chi_S, \psi_S)}(\theta_g \mid \text{observing } g \text{ for } i \text{ times}) &= 0.5\chi_S(1 - \psi_S) + [1 - \chi_S(1 - \psi_S)] \\ &= 1 - 0.5\chi_S(1 - \psi_S). \end{aligned}$$

In other words, every player i will hold the same belief. Therefore, it is optimal for player i to choose r if and only if

$$\alpha [1 - 0.5\chi_S(1 - \psi_S)] - 0.5\chi_S(1 - \psi_S) \geq 0 \iff \chi_S(1 - \psi_S) \leq \frac{2\alpha}{(1 + \alpha)}.$$

Because every player has the same belief, they will take the same action. Coupled with the calculation above, we can conclude that

1. $i^* = 1$ if and only if $\chi_S(1 - \psi_S) \leq \frac{2\alpha}{(1+\alpha)}$, and
2. $i^* = \bar{s}$ if and only if $\chi_S(1 - \psi_S) \geq \frac{2\alpha}{(1+\alpha)}$,

which completes the proof. ■

On the contrary, in the framework of χ -CSE, each player i will gradually update their beliefs as they gather more observations of g announcements, so player 2 will be better informed than player 1. Claim B.4 shows that, for intermediate values of χ , the first mover would choose the safe option but the second mover would be confident enough about state being good and choose the risky option in a χ -CSE.

Claim B.4. *In a χ -CSE, when the broadcaster announces g , the cutoff player i^{**} is characterized as the following:*

1. $i^{**} = 1$ if and only if $\chi \leq \frac{2\alpha}{1+\alpha}$,
2. $i^{**} = 2$ if and only if $\chi \in \left[\left(\frac{2\alpha}{1+\alpha} \right), \left(\frac{2\alpha}{1+\alpha} \right)^{\frac{1}{2}} \right]$, and
3. $i^{**} = \bar{s}$ if and only if $\chi \geq \left(\frac{2\alpha}{1+\alpha} \right)^{\frac{1}{2}}$.

Proof: Consider the case where the broadcaster announces g and fix any player $i \in \{1, 2\}$. By Claim B.1, we can obtain that player i 's belief about θ_g is

$$\mu_i^\chi(\theta_g \mid \text{observing } g \text{ for } i \text{ times}) = \chi \mu_i^\chi(\theta_g \mid \text{observing } g \text{ for } i-1 \text{ times}) + (1-\chi)$$

As we iteratively apply Claim B.1, we can obtain that for any player i , the belief about θ_g is

$$\mu_i^\chi(\theta_g \mid \text{observing } g \text{ for } i \text{ times}) = 1 - 0.5\chi^i,$$

suggesting that players are more certain about the state being good as they observe more g announcements. Thus, it is optimal for player i to choose r if and only if

$$\alpha [1 - 0.5\chi^i] - 0.5\chi^i \geq 0 \iff \chi \leq \left(\frac{2\alpha}{1+\alpha} \right)^{\frac{1}{i}}.$$

This completes the proof since the RHS is strictly increasing in i . ■

It is noteworthy that the above example as well as Claim B.3 and B.4 can be easily extended to the case with n players (other than the broadcaster) for any $n > 2$. In a (χ_S, ψ_S) -SCE, the n players would take the same action, while in a χ -CSE, the early movers would choose s but the late movers would choose r for some intermediate

value of χ . Besides, for any value of $\chi < 1$, there is a critical value n_χ such that for all $n > n_\chi$, there will be a switch point at which later movers will all choose r .

The (χ_S, ψ_S) -SCE predicts that in this example, even though later players ($i \geq 2$) repeatedly observe the broadcaster's behavior, they only believe the first announcement contains useful information. Under the SCE framework, a cursed player i acts as if she fully understands the correlation between the broadcaster's announcements, so collecting more of the broadcaster's announcements does not provide new information for updating her belief. In contrast, the χ -CSE predicts that a (partially) cursed player i will eventually learn the true state if she can observe the broadcaster's announcements infinitely many times. Under the CSE framework, although a cursed player i may be unsure about the broadcaster's private information at the beginning, she will gain confidence from observing that the broadcaster makes the same announcement over time. In this sense, player i gradually learns how another player's strategy depends on his type as she observes the other's actions over time, which is fundamentally different from the learning process characterized in SCE, where beliefs can get stuck after the first stage of the game.

B.3.2 Difference in Treating Public Histories

In addition to the difference in the belief updating process, another key difference between the CSE and SCE is about the *publicness* of public histories. In the framework of CSE, players have correct understanding about how choosing different actions would result in different histories and they know all other players know this. In other words, from the perspective of CSE, *public histories are essentially public*.

However, public histories are not necessarily public in the framework of SCE. Players in SCE are allowed to be cursed about endogenous information—namely, players may incorrectly believe other players will not respond to changes of actions. Technically speaking, when the coarsest valid partition is not consistent with public histories, i.e., PHC is violated, there is some information set where the player would believe regardless of what he chooses, others will have the same perception about the course of game play.

To demonstrate how this difference leads to different predictions in a specific game, we next consider a simple signaling game and compare χ -CSE and (χ_S, ψ_S) -SCE. This comparison illustrates when PHC is violated, the (χ_S, ψ_S) -SCE can look dramatically different from χ -CSE.

Illustrative Example




Figure B.1: Game Tree for the Illustrative Example.

Consider the signaling game depicted in Figure B.1 (Example 1 Fong, Lin, and Palfrey, 2023): The sender has two possible types drawn from the set $\Theta = \{\theta_1, \theta_2\}$ with $\Pr(\theta_1) = 1/4$. The receiver does not have any private information. After the sender's type is drawn, the sender observes his type and decides to send a message $m \in \{A, B\}$ (or any mixture between the two). After that, the receiver decides between action $a \in \{L, R\}$ (or any mixture between the two), and the game ends. In the following, we will focus on equilibrium in *pure* behavioral strategies and use a four-tuple $[(m(\theta_1), m(\theta_2)); (a(A), a(B))]$ to denote a behavioral strategy profile.

In Claim B.5 and Claim B.6, we summarize the solutions of χ -CSE and (χ_S, ψ_S) -SCE, respectively.

Claim B.5. *There are two pure pooling χ -CSE, which are:*

1. $[(A, A); (L, R)]$ is a pooling χ -CSE for any $\chi \in [0, 1]$.
2. $[(B, B); (R, R)]$ is a pooling χ -CSE if and only if $\chi \leq 8/9$.

Proof: See Appendix A. ■

Claim B.6. *There are four pure (χ_S, ψ_S) -SCE, which are:*

1. $[(A, A); (L, R)]$ is a (χ_S, ψ_S) -SCE if and only if $\chi_S \leq 1/3$.
2. $[(B, B); (L, R)]$ is a (χ_S, ψ_S) -SCE if and only if $\chi_S \geq 1/3$.

3. $[(B, A); (L, R)]$ is a (χ_S, ψ_S) -SCE if and only if $\chi_S = 1/3$.
4. $[(B, B); (R, R)]$ is a (χ_S, ψ_S) -SCE if and only if $\chi_S(1 - \psi_S) \leq 8/9$.

Proof: First, we can observe that the presence or absence of a PHC restriction on the coarsest partition does not affect player 2's best response to player 1's strategy. Since player 2 moves at the last stage after receiving player 1's message, player 2's conjecture about the other player's (and nature's) strategy is independent of the condition on the publicness of public history. Therefore, we can still apply Claim B.2 to pin down player 2's beliefs and the corresponding best responses under player 1's different strategies. For player 2, it is strictly optimal to choose R after seeing $m = B$ for any belief, and it is optimal to choose L if and only if $\mu_2^{(\chi_S, \psi_S)}(\theta_1|A) \leq 1/3$. By Claim B.2, we can find that:

1. $m(\theta_1) = m(\theta_2) = A$: $\mu_2^{(\chi_S, \psi_S)}(\theta_1|A) = 1/4$, so it is optimal for player 2 to choose L when $m = A$, implying that $[(A, A); (R, R)]$ is not a (χ_S, ψ_S) -SCE.
2. $m(\theta_1) = m(\theta_2) = B$: $\mu_2^{(\chi_S, \psi_S)}(\theta_1|A) \in [\frac{1}{4}\chi_S(1 - \psi_S), 1 - \frac{3}{4}\chi_S(1 - \psi_S)]$, so $m(A) = R$ can be supported as a best response if and only if $\chi_S(1 - \psi_S) \leq 8/9$; alternatively, $m(A) = L$ can be supported as a best response for all $(\chi_S, \psi_S) \in [0, 1]^2$.
3. $m(\theta_1) = A$ and $m(\theta_2) = B$: $\mu_2^{(\chi_S, \psi_S)}(\theta_1|A) = 1 - \frac{3}{4}\chi_S(1 - \psi_S)$, so it is optimal for player 2 to choose L when receiving A if and only if $\chi_S(1 - \psi_S) \geq 8/9$.
4. $m(\theta_1) = B$ and $m(\theta_2) = A$: $\mu_2^{(\chi_S, \psi_S)}(\theta_1|A) = \frac{1}{4}\chi_S(1 - \psi_S) \leq 1/3$, so it is optimal for player 2 to choose L when receiving A .

Player 1's conjecture about player 2's strategy, however, depends on whether player 1 realizes that the message m will be public. In the framework of SCE, the element of the coarsest valid partition at the message m is

$$P(m) = \{(m', \theta) : m' \in \{A, B\}, \theta \in \{\theta_1, \theta_2\}\},$$

which is the same for any $m \in \{A, B\}$. This implies player 1 would (incorrectly) believe player 2 will behave the same regardless of which message is sent. Given a strategy profile, $\sigma = [(m(\theta_1), m(\theta_2)); (a(A), a(B))]$, a type θ' player 1's (sequentially and typically) cursed conjecture about player 2's strategy when seeing the

message $\tilde{m} \in \{A, B\}$ becomes

$$\begin{aligned}\bar{\sigma}_2^{(\theta')}(a'|\tilde{m}) &= \mu[\sigma](\{(m, \theta), a'\} : (m, \theta) \in P(\tilde{m}) \mid \theta' \cap P(\tilde{m})) \\ &= \frac{\sigma_\lambda(\theta') \sum_{m \in \{A, B\}} \sigma_1(m|\theta') \sigma_2(a'|m)}{\sigma_\lambda(\theta')} \\ &= 1 \cdot \sigma_2(a'|m = m(\theta')) + 0 \cdot \sigma_2(a'|m \neq m(\theta')) \\ &= \sigma_2(a'|m(\theta')).\end{aligned}$$

In other words, with probability χ_S , player 1 conjectures that player 2's off-path strategy would be the same as her on-path strategy. Notice that both types of player 1 have the same payoff function. Hence:

1. $a(A) = a(B) = R$: It is optimal for player 1 to send the message B as $1 > -1$, implying $[(B, B); (R, R)]$ is a (χ_S, ψ_S) -SCE when $\chi_S(1 - \psi_S) \leq 8/9$.
2. $a(A) = L$ and $a(B) = R$: It is optimal for player 1 to send A if and only if

$$2 \geq 4\chi_S + (1 - \chi_S) \iff \chi_S \leq 1/3.$$

Alternatively, it is optimal for player 1 to send B if and only if

$$1 \geq -\chi_S + 2(1 - \chi_S) \iff \chi_S \geq 1/3.$$

Therefore, $[(A, A); (L, R)]$ is a (χ_S, ψ_S) -SCE when $\chi_S \leq 1/3$. In addition, $[(B, B); (L, R)]$ is a (χ_S, ψ_S) -SCE when $\chi_S \geq 1/3$. Moreover, $[(B, A); (L, R)]$ is also a (χ_S, ψ_S) -SCE when $\chi_S = 1/3$. It is worth noticing that $[(A, B); (L, R)]$ cannot be supported as a (χ_S, ψ_S) -SCE since it requires two contradicting conditions to hold ($\chi_S = 1/3$ and $\chi_S(1 - \psi_S) \geq 8/9$).

This completes the characterization of (χ_S, ψ_S) -SCE. ■

To visualize the effect of PHC, we plot the solutions characterized in Claim B.6 in Figure B.2. In a (χ_S, ψ_S) -SCE, a cursed player 1 would incorrectly believe player 2 would behave the same regardless of which m is chosen. As a result, (χ_S, ψ_S) -SCE differs dramatically from χ -CSE when PHC is violated. The pooling equilibrium $[(A, A); (L, R)]$ is a (χ_S, ψ_S) -SCE if and only if players are not too cursed about endogenous information, i.e., $\chi_S \leq 1/3$, while this is a χ -CSE for any $\chi \in [0, 1]$. Moreover, when players are sufficiently cursed about endogenous information ($\chi_S \geq 1/3$), there is an additional (χ_S, ψ_S) -SCE $[(B, B); (L, R)]$ which

is neither a PBE, a CE nor a χ -CSE. Lastly, it is worth mentioning that there exists an knife-edge case that $[(B, A); (L, R)]$ is a separating (χ_S, ψ_S) -SCE if and only if $\chi_S = 1/3$. Since the incentives of both types of player 1 are perfectly aligned, the existence of such a separating (χ_S, ψ_S) -SCE is unexpected.




Figure B.2: The Solution Set of (χ_S, ψ_S) -SCE.

In summary, a key difference between CSE and SCE is about the way of treating public histories. From the illustrative example, we can find that when the coarsest valid partition is not consistent with the public histories, (χ_S, ψ_S) -SCE and χ -CSE make extremely different predictions. In the following sections, we will explore the implications of the difference in the publicness of public histories. Surprisingly, the impact of the publicness of public histories not only arises in multistage games with incomplete information but also appears in games of perfect information, which will be formally discussed in the next section.

B.3.3 Difference in Games of Complete Information

A significant difference between CSE and SCE is about the predictions of the games with complete information, i.e., $|\Theta| = 1$. As shown in Claim B.7, in games with complete information, CSE coincides with sequential equilibrium. As the type space is singleton, there is no possibility for players to make mistaken inferences

about the types of other players, so neglect of the correlation between types and actions is a moot issue, as is also the case in CE.

Claim B.7. *If $|\Theta| = 1$, χ -CSE is equivalent to the sequential equilibrium for any $\chi \in [0, 1]$.*

Proof: Let $\Theta = \{(\bar{\theta}_1, \dots, \bar{\theta}_n)\}$ and consider any behavioral strategy profile. Because $|\Theta| = 1$, for any player i and public history h^{t-1} , we can obtain that $\mu_i(\bar{\theta}_{-i}|h^t, \bar{\theta}_i) = 1$. Consequently, the average behavioral strategy profile of $-i$ is simply:

$$\bar{\sigma}_{-i}(a_{-i}^t|h^{t-1}, \bar{\theta}_i) = \mu_i(\bar{\theta}_{-i}|h^t, \bar{\theta}_i)\sigma_{-i}(a_{-i}^t|h^{t-1}, \bar{\theta}_i) = \sigma_{-i}(a_{-i}^t|h^{t-1}, \bar{\theta}_i)$$

for any $a_{-i}^t \in A_{-i}(h^{t-1})$, suggesting $\sigma_{-i}^\chi(a_{-i}^t|h^{t-1}, \bar{\theta}) = \sigma_{-i}(a_{-i}^t|h^{t-1}, \bar{\theta}_{-i})$ for any $\chi \in [0, 1]$. Since the perception about others' strategy profile always aligns with the true strategy profile, χ -CSE is equivalent to the sequential equilibrium for any $\chi \in [0, 1]$. ■

In contrast, SCE does not coincide with the sequential equilibrium in games with complete information. In fact, this is even true for games of *perfect information*, i.e., nature is not a player and every information set is a singleton. This is illustrated with the following simple game of perfect information. The intuition of this phenomenon ties in with the discussion earlier in this appendix, since in this example PHC is violated. The coarsest partition is not consistent with the public history—even though the type space and information sets are singleton sets. As a result, when the coarsest partition bundles multiple public histories, players neglect how their current action affects another player's future action.

One of the motivations for SCE is to extend the notion of cursedness to *endogenous* information, i.e., observed actions, which are endogenous to the game. However, it is useful to distinguish between private endogenous information and public endogenous information. In CSE, cursedness arises because players neglect the jointness of other players' actions and their *private information*, but understand the link between other players' actions and *public information*. In multistage games with public histories, there is no “private endogenous” information. We conjecture that if SCE were modified so that the coarsest valid partition is required to be measurable with respect to public histories, then cursedness with respect to *private endogenous* information could still arise, but the phenomenon of cursedness with respect to *public endogenous* information would not arise.

Illustrative Example

Consider the two-player game of perfect information in Figure B.3. In the first stage, player 1 makes a choice from $A_1 = \{B, R\}$. After observing player 1's decision, player 2 then makes a choice from $A_2 \in \{b, r\}$. The payoffs are shown in the game tree where $x \in \mathbb{R}$ and $y < 1$. In the following, we will denote the players' (pure) behavioral strategy profile by $[a_1; (a_2(B), a_2(R))]$.




Figure B.3: A Two-Player Game with Complete Information.

In this game, the only subgame perfect Nash equilibrium is $[B; (b, r)]$. However, Claim B.8 shows that $[R; (b, r)]$ can also be supported as a (χ_S, ψ_S) -SCE if χ_S is sufficiently large.

Claim B.8. *For any $x \in \mathbb{R}$ and $y < 1$, $[R; (b, r)]$ is a (χ_S, ψ_S) -SCE if and only if $\chi_S \geq 1/2$.*

Proof: First, player 2's best responses to B and R are b and r , respectively. Note that player 1's information set is singleton, so coarsening it has no effect on player 2's conjecture about, and thus best response to, player 1's action. Second, under the coarsest partition, $P(h^1 = R)$ is $\{a_1 : a_1 \in \{B, R\}\}$. Given $[R; (b, r)]$, player 1 conjectures that player 2 when observing B will choose r with probability χ_S and choose b with probability $1 - \chi_S$. Therefore, player 1 will not have an incentive to deviate to B if and only if $1 \geq 2(1 - \chi_S) \iff \chi_S \geq 1/2$. ■

It is noteworthy that, given any x and $y < 1$, choosing R is a (χ_S, ψ_S) -SCE strategy for player 1 if $\chi_S > 1/2$ and this threshold is independent of x and y . In other words, under the SCE framework, a cursed player 1 could choose R in equilibrium even

when such choice is extremely risky for player 1 due to a huge potential loss (e.g., $x = -1000000$) or when player 2 is extremely unlikely to go for b conditional on observing R due to a large negative payoff (e.g., $y = -1000000$), and such possibility does not change in the values of x and y .

B.3.4 Difference in Consistency with Subgame Perfection and Sequential Rationality

The example above shows that SCE can yield predictions that violate subgame perfection, which cannot happen in CSE. A similar issue arises in the context of the signaling game analyzed in section B.3.2. In the CSE analysis of the signaling game, players correctly understand that choosing different actions would bring the game to different public histories although they update their beliefs about the types via the χ -cursed Bayes' rule. They further believe that future players will choose optimally (perhaps with non-Bayesian beliefs) at future public histories. Thus, CSE has built into it the property of sequential rationality, analogous to subgame perfection in games of perfect information.⁷

As we observed in the signaling game, when $\chi_S \geq 1/3$, there exists a (χ_S, ψ_S) -SCE where both types of player 1 pool at B and player 2 will choose L and R in public histories A and B , respectively. This solution violates sequential rationality in the sense that, in this SCE, player 1 incorrectly believes player 2 would behave the same at both public histories. That is, player 1 believes player 2 will irrationally choose R in response to A . But if player 1 correctly believes player 2 will rationally choose L at public history A , it would be profitable for both types of player 1 to choose A .

In summary, the violation of sequential rationality of SCE in these examples is indeed a consequence of the coarsest valid partition being incompatible with public histories. In the next section, we show how the coarseness of the partition can also be affected by how the actions are labelled, which in CSE is an inessential technical detail of the formal game representation. For some additional discussion about the labeling issue see Cohen and Li (2023).

B.3.5 Difference in the Effect of Re-Labeling Actions

From the discussion in the previous section, we can find that the violation of subgame perfection of SCE is a consequence of coarsening public histories into an information set, which in fact, is sensitive to the labels of actions available at these

⁷Here we use the term sequential rationality instead of subgame perfection, because in the signaling game there are no proper subgames.

public histories. The requirement of an information set is that the set of actions for every history in this information set is exactly the same. Therefore, if the actions at every public histories are all labelled differently, then the coarsest valid partition is consistent with the public histories. This observation is stated in Claim B.9.

Definition B.6. A multi-stage game with observed actions is *scrambled* if for any $i \in N$, any $t < T$, and any $h, h' \in \mathcal{H}^t$ such that $h \neq h'$, then

$$\forall s \in A_i(h) \text{ and } s' \in A_i(h') \implies s \neq s'.$$

Claim B.9. A scrambled multi-stage game with observed actions satisfies PHC.

Proof: If not, there is some $P(h^t, \theta_i)$ which contains two public histories $h, h' \in \mathcal{H}^t$ where $h \neq h'$. However, because the game is scrambled, $A_i(h) \neq A_i(h')$. Therefore, h and h' cannot belong to the same cell of a partition under any partition, which yields a contradiction. ■

This observation provides an alternative interpretation of PHC—we can view this as an additional requirement of SCE such that the solution concept is immune to the effect of re-labeling actions.⁸ On the other hand, because the average behavioral strategy of CSE is defined at every public history, the immunity of CSE is built in the model setup.

The illustrative example in Section B.3.2 can also demonstrate the effect of re-labeling. Let $A_2(A)$ and $A_2(B)$ be the action sets at the public history A and B , respectively. If $A_2(A) = \{L, R\}$ and $A_2(B) = \{L', R'\}$, then this is a scrambled game and hence satisfies PHC. As characterized in Claim B.10, we can find that CSE and SCE are equivalent if $\chi = \chi_S(1 - \psi_S)$. However, if $A_2(A) = A_2(B) = \{L, R\}$, then two public histories A and B belong to the same information set under the coarsest valid partition. As shown in Claim B.6, the (χ_S, ψ_S) -SCE solution looks dramatically different.

Claim B.10. If the signaling game depicted in Figure B.1 is scrambled, then there are two pure pooling (χ_S, ψ_S) -SCE, which are:

1. $[(A, A); (L, R')]$ is a pooling (χ_S, ψ_S) -SCE for any $(\chi_S, \psi_S) \in [0, 1]^2$.

⁸To deal with this issue, Cohen and Li propose the concept of *casual SCE*. See Appendix D of Cohen and Li (2023) for details.

2. $[(B, B); (R, R')]$ is a pooling (χ_S, ψ_S) -SCE if and only if $\chi_S(1 - \psi_S) \leq 8/9$.

Proof: If the signaling game is scrambled, then by PHC, we can find that the partition illustrated in Figure B.1 is the coarsest valid partition. We can observe that at the public history $m = B$, it is strictly optimal for player 2 to choose R' for any belief. On the other hand, at the public history $m = A$, it is optimal for player 2 to choose L if and only if $\mu_2^{(\chi_S, \psi_S)}(\theta_1|A) \leq 1/3$, as shown in the proof of Claim B.5.

Equilibrium 1. If both types of player 1 choose A , then by Claim B.2, we can find that $\mu_2^{(\chi_S, \psi_S)}(\theta_1|A) = 1/4$, so it's optimal for player 2 to choose L at the public history $m = A$. Because of PHC, player 1 knows player 2 is choosing a contingent strategy for different public histories. As a result, given $a(A) = L$ and $a(B) = R'$, it is optimal for both types of player 1 to choose A . This shows under PHC, $[(A, A); (L, R')]$ is a pooling (χ_S, ψ_S) -SCE for any $(\chi_S, \psi_S) \in [0, 1]^2$.

Equilibrium 2. To support $m(\theta_1) = m(\theta_2) = B$ to be an equilibrium, because of PHC, player 2 has to choose R at the public history $m = A$. By Claim B.2 and the consistency requirement, the belief system has to satisfy

$$\mu_2^{(\chi_S, \psi_S)}(\theta_1|A) \in \left[\frac{1}{3}, 1 - \frac{3}{4}\chi_S(1 - \psi_S) \right],$$

which is valid if and only if $\chi_S(1 - \psi_S) \leq 8/9$. Finally, because of PHC, both players know that player 2's strategy is a contingent strategy which is responsive to different public histories. Consequently, we can use the standard equilibrium argument to show that there is no separating pure strategy equilibrium. This completes the characterization of SCE under PHC. ■

Figure B.4 depicts the SCE of the signaling game when it is unscrambled (i.e., PHC is violated) and scrambled (i.e., PHC is satisfied). With PHC, all players understand that every player's strategy is a contingent strategy conditional on the public histories. Therefore, if this two-stage signaling game is scrambled, (χ_S, ψ_S) -SCE coincides with χ -CSE when $\chi = \chi_S(1 - \psi_S)$. Moreover, the strategy profiles that could only be supported as a (χ_S, ψ_S) -SCE in the unscrambled signaling game by a relatively large χ_S (i.e., $[(B, B); (L, R')]$ and $[(B, A); (L, R')]$) would no longer be a (χ_S, ψ_S) -SCE when PHC is required.

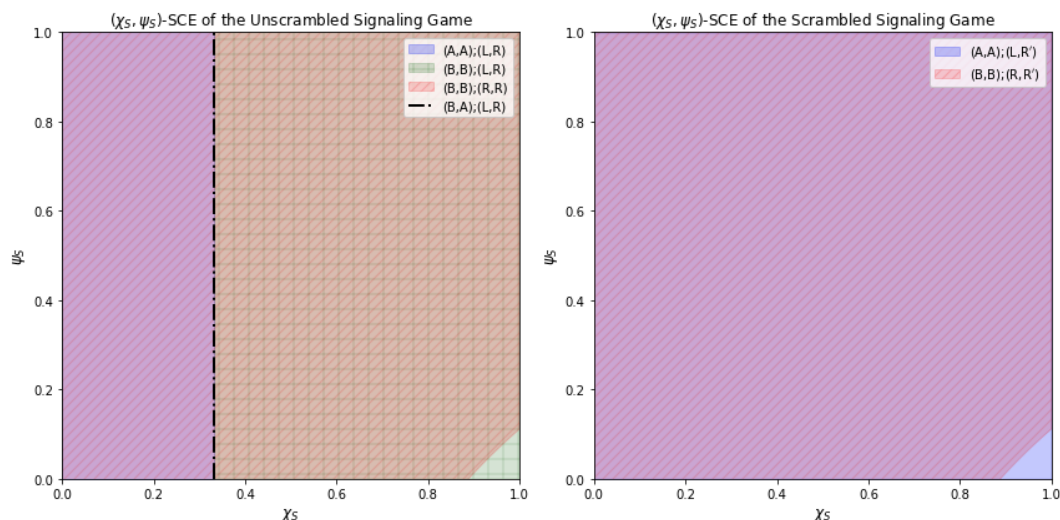


Figure B.4: (χ_S, ψ_S) -SCE of an unscrambled and scrambled signaling game.

B.3.6 Difference in One-Stage Simultaneous-Move Games

The last important difference between CSE and SCE is in their relations with the standard cursed equilibrium in one-stage games. As shown by Fong, Lin, and Palfrey (2023), CSE coincides with the standard CE for any one-stage game. Yet Cohen and Li (2023) show that SCE in one-stage games is equivalent to *independently cursed equilibrium* (ICE), under which players are cursed about not only the dependence of opponents' actions on private information but also the correlation between opponents' actions. In the following, we will first summarize the definitions of CE and ICE, and then illustrate how substantial their predictions can differ in an example of a three-player game.

Definition B.7 describes the standard CE in an one-stage game. Under CE, a player fails to account for how the other players' action profile may depend on their types, and best responds to the *average strategy profile of the other players*.

Definition B.7. A strategy profile σ is a cursed equilibrium (CE) if for each player i , type $\theta_i \in \Theta_i$ and each $a_i^1 \in A_i(h_\emptyset)$ such that $\sigma_i(a_i^1 | h_\emptyset, \theta_i) > 0$,

$$a_i^1 \in \operatorname{argmax}_{a_i^1 \in A_i(h_\emptyset)} \sum_{\theta_{-i} \in \Theta_{-i}} F(\theta_{-i} | \theta_i) \times \sum_{a_{-i}^1 \in A_{-i}(h_\emptyset)} \left[\sum_{\theta_{-i} \in \Theta_{-i}} F(\theta_{-i} | \theta_i) \sigma_{-i}(a_{-i}^1 | h_\emptyset, \theta_{-i}) \right] u_i(a_{-i}^1, a_i^1, \theta_{-i}, \theta_i),$$

Definition B.8 provides the definition of ICE in an one-stage game. Under ICE, a player fails to account for how each player's action may depend on her own type, and how it may correlate with another player's action (via the correlation in type distribution). Therefore, a player would best respond as if the *average strategies across the other players are independent*.

Definition B.8. A strategy profile σ is an *independently cursed equilibrium (ICE)* if for each player i , type $\theta_i \in \Theta_i$ and each $a_i^1 \in A_i(h_0)$ such that $\sigma_i(a_i^1|h_0, \theta_i) > 0$,

$$a_i^1 \in \operatorname{argmax}_{a_i^1 \in A_i(h_0)} \sum_{\theta_{-i} \in \Theta_{-i}} F(\theta_{-i}|\theta_i) \times \sum_{a_{-i}^1 \in A_{-i}(h_0)} \prod_{j \in N \setminus \{i\}} \left[\sum_{\theta_j \in \Theta_j} F(\theta_j|\theta_i) \sigma_j(a_j^1|h_0, \theta_j) \right] u_i(a_{-i}^1, a_i^1, \theta_{-i}, \theta_i).$$

Claim B.11 and B.12 summarize the relations between CSE, SCE, and CE in one-stage games. The proofs can be found in Fong, Lin, and Palfrey (2023) and Cohen and Li (2023), respectively.

Claim B.11. For any one-stage game and for any $\chi \in [0, 1]$, χ -CSE and χ -CE are equivalent.

Claim B.12. For any finite one-stage game, ICE and SCE are equivalent.

Although CE and ICE are equivalent in an one-stage two-person game, we show in the following example that the sets of CE and ICE may be non-overlapping in a three-person game. This finding suggests that CSE and SCE are generally different when there are more than two players.

Illustrative Example

Consider the following three-player one-stage game. Player 1 and 2 have two possible types drawn from the set $\Theta = \{b, r\}$ with the joint distribution $F(\theta_1 = \theta_2 = b) = F(\theta_1 = \theta_2 = r) = 0.5 - \epsilon$ and $F(\theta_1 = b, \theta_2 = r) = F(\theta_1 = r, \theta_2 = b) = \epsilon$ where $\epsilon \in (0, 0.5)$. Player 3 has no private information. Each player makes a choice from the set $A = \{b, r, m\}$. Player 1 and 2 will get one unit of payoff if his choice matches his type (and 0 otherwise). Player 3 will get one unit of payoff if his choice matches player 1's and 2's choices when $a_1 = a_2$, or if he chooses m when $a_1 \neq a_2$ (and 0 otherwise).

To summarize, player 1 and 2 have private information, and their payoffs will be maximized if their actions match their types. Player 1's (and 2's) type is b or r with equal probabilities. However, their types can be the same with probability $1 - 2\epsilon$. Player 3 has no private information, and his goal is to guess his opponents' actions by following them when they act the same and choosing m when they act differently.

Claim B.13. *When $\epsilon < 1/6$, player 3 will choose b or r in CE but m in ICE. That is, CE and ICE do not overlap.*

Proof: In both CE and ICE, $a_i^*(\theta_i) = \theta_i$ for $i \in \{1, 2\}$, which is a strictly dominant strategy for Player 1 and 2. Player 3's expected payoff of choosing a_3 under CE is thus

$$\mathbb{E}u_3(a_3, a_{-3} | \sigma_{-3}^*) = \begin{cases} F(\theta_1 = \theta_2 = a_3) & \text{if } a_3 = b \text{ or } r \\ F(\theta_1 \neq \theta_2) & \text{if } a_3 = m. \end{cases}$$

Therefore, it is optimal for player 3 to choose b or r in CE if

$$0.5 - \epsilon > 2\epsilon \iff \epsilon < 1/6.$$

Alternatively, player 3's expected payoff of choosing a_3 under ICE is

$$\begin{aligned} \mathbb{E}u_3(a_3, a_{-3} | \sigma_{-3}^*) &= \begin{cases} F(\theta_1 = a_3)F(\theta_2 = a_3) & \text{if } a_3 = b \text{ or } r \\ F(\theta_1 = b)F(\theta_2 = r) + F(\theta_1 = r)F(\theta_2 = b) & \text{if } a_3 = m \end{cases} \\ \Rightarrow \mathbb{E}u_3(a_3, a_{-3} | \sigma_{-3}^*) &= \begin{cases} 0.25 & \text{if } a_3 = b \text{ or } r \\ 0.5 & \text{if } a_3 = m. \end{cases} \end{aligned}$$

Therefore, it is optimal for player 3 to choose m for all $\epsilon \in [0, 0.5]$. ■

Players' private information may serve as a coordinating device for the players' actions if their types are correlated. As a result, when a player neglects the possible correlation in the other players' strategies under ICE, he would respond as if he neglects the dependence of types across players in the prior distribution.

References

- Cohen, Shani and Shengwu Li (2023). "Sequential cursed equilibrium". In: *arXiv preprint arXiv:2212.06025v3*.
- Eyster, Erik and Matthew Rabin (2005). "Cursed equilibrium". In: *Econometrica* 73.5, pp. 1623–1672.

- Fong, Meng-Jhang, Po-Hsuan Lin, and Thomas R Palfrey (2023). “Cursed sequential equilibrium”. In: *arXiv preprint arXiv:2301.11971*.
- Fudenberg, Drew and Jean Tirole (1991). “Perfect Bayesian equilibrium and sequential equilibrium”. In: *Journal of Economic Theory* 53.2, pp. 236–260.
- Jehiel, Philippe (2005). “Analogy-based expectation equilibrium”. In: *Journal of Economic Theory* 123.2, pp. 81–104.
- Jehiel, Philippe and Frédéric Koessler (2008). “Revisiting games of incomplete information with analogy-based expectations”. In: *Games and Economic Behavior* 62.2, pp. 533–557.
- Kreps, David M and Robert Wilson (1982). “Sequential equilibria”. In: *Econometrica* 50.4, pp. 863–894.

Appendix C

PROOFS FOR CHAPTER 3

C.1 Proofs for Section 3.4

Let τ_i be player i 's level. Following previous notations, we use $\sigma_i(h)$ to denote player i 's (pure) action at h . In addition, $\mu_i^k(\tau_{-i})$ is level- k player i 's prior belief about the opponent's level, and $\nu_i^k(\tau_{-i} | h)$ is level- k player i 's posterior belief about the opponent's level at history h . Finally, level-0 players would uniformly randomize at every node. The analysis of the examples is summarized in the following claims.

Example 3.4.2.1

Claim C.1. *In Example 3.4.2.1, each level of players' strategies are:*

1. for any $k \in \mathbb{N}$, $\sigma_1^k(1b) = r_{1b}$, $\sigma_1^k(1c) = l_{1c}$, $\sigma_2^k(2b) = l_{2b}$, and $\sigma_2^k(2c) = r_{2c}$;
2. $\sigma_1^1(1a) = r_{1a}$ and $\sigma_1^k(1a) = l_{2a}$ for $k \geq 2$; $\sigma_2^1(2a) = \sigma_2^2(2a) = l_{2a}$ and $\sigma_2^k(2a) = r_{2a}$ for $k \geq 3$.

Proof: The calculation consists of two parts.

1. First, all strategic levels of players would choose the action with a higher payoff at the last node. Hence, $\sigma_1^k(1b) = r_{1b}$ and $\sigma_2^k(2c) = r_{2c}$ for all $k \geq 1$. Player 2 has a dominant action at history $h = 2b$, so $\sigma_2^k(2b) = l_{2b}$ for all $k \geq 1$. Notice that whenever a dominant action is not chosen, players would believe the opponent is level-0 with certainty. At history $h = 1c$, every level of player 1 thinks player 2 is level-0 and hence for all $k \geq 1$, $\sigma_1^k(1c) = l_{1c}$ since the expected payoff is $13/2 > 6$.

2. Level-1 players believe the other player would randomize at every node. On the one hand, $\sigma_1^1(1a) = r_{1a}$ and $\sigma_2^1(2a) = l_{2a}$ so that they can maximize the expected payoff. On the other hand, level-2 players' initial beliefs are $\mu_i^2(0) = e^{-1.5}/(e^{-1.5} + 1.5e^{-1.5}) = 2/5$ and $\mu_i^2(1) = 3/5$. Thus, $\sigma_1^2(1a) = l_{1a}$ since the expected payoff for l_{1a} is $19/5 > 29/10$. On the other hand, when history $h = 2a$ is realized, level-2 player 2 would believe the opponent is definitely level-0 and hence $\sigma_2^2(2a) = \sigma_2^1(2a) = l_{2a}$.

The behavior of higher-level players can be solved by induction. Level-3 players' prior beliefs are $\mu_i^3(0) = 8/29$, $\mu_i^3(1) = 12/29$, and $\mu_i^3(2) = 9/29$. In this case, $\sigma_1^3(1a) = l_{1a}$ since the expected payoff for l_{1a} is $112/29 > 76/29$. In addition, when history $h = 2a$ is realized, level-3 player 2's posterior belief becomes $v_2^3(0 | 2a) = 0.5e^{-1.5}/(0.5e^{-1.5} + 1.125e^{-1.5}) = 4/13$ and $v_2^3(2 | 2a) = 9/13$, and hence $\sigma_2^3(2a) = r_{2a}$ since $4 > 45/13$. Suppose for some $k > 3$, $\sigma_1^k(1a) = l_{1a}$ for all $2 \leq \kappa \leq k$ and $\sigma_2^k(2a) = r_{2a}$ for all $3 \leq \kappa \leq k$. We want to show that $\sigma_1^{k+1}(1a) = l_{1a}$ and $\sigma_2^{k+1}(2a) = r_{2a}$. Level- $(k+1)$ players' prior beliefs are $\mu_i^{k+1}(\kappa) = p_\kappa/(\sum_{\kappa=0}^k p_\kappa)$ for $0 \leq \kappa \leq k$. By the induction hypothesis, $\sigma_1^{k+1}(1a) = l_{1a}$ if and only if

$$\frac{7}{2} \left(\frac{p_0}{\sum_{\kappa=0}^k p_\kappa} \right) + 4 \left(\frac{p_1 + p_2}{\sum_{\kappa=0}^k p_\kappa} \right) + 3 \left(\frac{\sum_{\kappa=3}^k p_\kappa}{\sum_{\kappa=0}^k p_\kappa} \right) > \frac{17}{4} \left(\frac{p_0}{\sum_{\kappa=0}^k p_\kappa} \right) + 2 \left(1 - \frac{p_0}{\sum_{\kappa=0}^k p_\kappa} \right),$$

which is equivalent to $(7/4)p_0 - p_1 - p_2 < \sum_{\kappa=0}^k p_\kappa$. This holds because $(7/4)p_0 - p_1 - p_2 = -(7/8)e^{-1.5} < 0$. Finally, by the induction hypothesis, level- $(k+1)$ player 2's posterior belief at $h = 2a$ is $v_2^{k+1}(0 | 2a) = 0.5p_0/(0.5p_0 + \sum_{\kappa=2}^k p_\kappa)$ and $v_2^{k+1}(j | 2a) = p_j/(0.5p_0 + \sum_{\kappa=2}^k p_\kappa)$ where $2 \leq j \leq k$. Thus, $\sigma_2^{k+1}(2a) = r_{2a}$ if and only if

$$\frac{9}{2}v_2^{k+1}(0 | 2a) + 3 \left(1 - v_2^{k+1}(0 | 2a) \right) < 4 \iff v_2^{k+1}(0 | 2a) < \frac{2}{3}.$$

Moreover, the induction hypothesis suggests that

$$v_2^{k+1}(0 | 2a) = \frac{\frac{1}{2}p_0}{\frac{1}{2}p_0 + \sum_{\kappa=2}^k p_\kappa} < \frac{\frac{1}{2}p_0}{\frac{1}{2}p_0 + \sum_{\kappa=2}^{k-1} p_\kappa} = v_2^k(0 | 2a) < \frac{2}{3},$$

implying the optimal choice for level- $(k+1)$ player 2 is r_{2a} . ■

Example 3.4.2.3

Claim C.2. Suppose τ_i 's are independently drawn from $p = (p_k)_{k=0}^\infty$, then in Example 3.4.2.3,

1. for any $k \in \mathbb{N}$, $\sigma_1^k(1a) = r_{1a}$, $\sigma_1^k(1b) = r_{1b}$, $\sigma_1^k(1c) = l_{1c}$, $\sigma_2^k(2a) = l_{2a}$, $\sigma_2^k(2b) = l_{2b}$, and $\sigma_2^k(2c) = r_{2c}$;
2. the ex ante probability of the subgame perfect equilibrium path being realized converges to 0 as $p_0 \rightarrow 0^+$.

Proof: The calculation consists of two parts.

1. By the analysis of Example 3.4.2.1, we only need to check player 1's action at the initial node and player 2's action at history $h = 2a$. We can prove the statement by induction on k . For $k = 1$, players would think the opponent is level-0. In this case, $\sigma_1^1(1a) = r_{1a}$ since the expected payoff is $17/4 > 9/4$ and $\sigma_2^1(2a) = l_{2a}$ with the expected payoff being $9/2 > 4$. Suppose there is some K such that $\sigma_1^k(1a) = r_{1a}$ and $\sigma_2^k(2a) = l_{2a}$ for all $1 \leq k \leq K$. For level- $(K+1)$ player 1, the prior belief is $\mu_1^{K+1}(0) = p_0 / (\sum_{\kappa=0}^K p_\kappa)$ and $\sigma_1^{K+1}(1a) = r_{1a}$ if and only if

$$\frac{17}{4}\mu_1^{K+1}(0) + 2\left(1 - \mu_1^{K+1}(0)\right) > \frac{9}{4}\mu_1^{K+1}(0) + \frac{3}{2}\left(1 - \mu_1^{K+1}(0)\right),$$

which holds as $\mu_1^{K+1}(0) > 0$. On the other hand, by the induction hypothesis, player 2 would believe player 1 is level-0 with certainty when history $h = 2a$ is realized, so $\sigma_2^{K+1}(2a) = \sigma_2^1(2a) = l_{2a}$.

2. Statement **1** implies the probability of the subgame perfect equilibrium path r_{2a} being realized is

$$\Pr(r_{2a}) = \Pr((1a, 2a) \mid 1a) \Pr(r_{2a} \mid 2a) = [\sigma_1^0(1a, 2a)p_0] [\sigma_{2,2a}^0(r_{2a})p_0] = \frac{1}{4}p_0^2.$$

Therefore, we can find the limit of the probability is

$$\lim_{p_0 \rightarrow 0^+} \Pr(r_{2a}) = \lim_{p_0 \rightarrow 0^+} \frac{1}{4}p_0^2 = 0.$$

This completes the proof. ■

Example 3.4.3

Claim C.3. *In Game Γ , each level of players' strategies are:*

1. for any $k \in \mathbb{N}$, $\sigma_1^k(1b) = r_{1b}$, $\sigma_1^k(1c) = l_{1c}$, $\sigma_2^k(2b) = l_{2b}$, and $\sigma_2^k(2c) = r_{2c}$;
2. $\sigma_1^k(1a) = l_{1a}$ for all $k \neq 2$, and $\sigma_1^2(1a) = r_{1a}$; $\sigma_2^1(2a) = l_{2a}$, and $\sigma_2^k(2a) = r_{2a}$ for all $k \geq 2$.

Proof: The proof consists of two parts.

1. The proof is the same as the proof of Claim C.1.

2. First, level-1 players believe the other player randomizes everywhere, so $\sigma_1^1(1a) = l_{1a}$ and $\sigma_2^1(2a) = l_{2a}$ in order to maximize their expected payoffs. Level-2 players'

prior beliefs are $\mu_i^2(0) = 2/5$ and $\mu_i^2(1) = 3/5$. Therefore, $\sigma_1^2(1a) = r_{1a}$ since the expected payoff is $29/10 > 28/10$. Level-2 player 2's posterior belief at history $h = 2a$ is $v_2^2(0 | 2a) = 0.5e^{-1.5}/(0.5e^{-1.5} + 1.5e^{-1.5}) = 1/4$ and $v_2^2(1 | 2a) = 3/4$. In this case, $\sigma_2^2(2a) = r_{2a}$ because $4 > 27/8$.

Finally, we can solve higher-level players' behavior by induction. Level-3 players' prior beliefs are $\mu_i^3(0) = 8/29$, $\mu_i^3(1) = 12/29$, and $\mu_i^3(2) = 9/29$, and hence $\sigma_1^3(1a) = l_{1a}$ since the expected payoff is $128/29 > 76/29$. At history $h = 2a$, level-3 player 2's posterior belief is the same as level-2, and so $\sigma_2^3(2a) = \sigma_2^2(2a) = r_{2a}$. Suppose there is some $K > 3$ such that $\sigma_1^k(1a) = l_{1a}$ for all $3 \leq k \leq K$ and $\sigma_2^k(2a) = r_{2a}$ for all $2 \leq k \leq K$. Level- $(K+1)$ players' prior beliefs are $\mu_i^{K+1}(j) = p_j / \sum_{i=0}^K p_i$ for $0 \leq j \leq K$. By the induction hypothesis, $\sigma_1^{K+1}(1a) = l_{1a}$ if and only if

$$\frac{19}{4} \left(\frac{p_0}{\sum_{i=0}^K p_i} \right) + \frac{3}{2} \left(\frac{p_1}{\sum_{i=0}^K p_i} \right) + 8 \left(\frac{\sum_{i=2}^K p_i}{\sum_{i=0}^K p_i} \right) > \frac{17}{4} \left(\frac{p_0}{\sum_{i=0}^K p_i} \right) + 2 \left(1 - \frac{p_0}{\sum_{i=0}^K p_i} \right),$$

which is equivalent to $5.5p_0 + 6.5p_1 < 6 \sum_{i=0}^K p_i$. This holds when the distribution of levels follows Poisson(1.5). On the other hand, by the induction hypothesis, level- $(K+1)$ player 2's posterior belief at history $h = 2a$ is $v_2^{K+1}(0 | 2a) = 0.5p_0 / (0.5p_0 + p_1 + \sum_{i=3}^K p_i)$ and $v_2^{K+1}(j | 2a) = p_j / (0.5p_0 + p_1 + \sum_{i=3}^K p_i)$ where $j \neq 0$ or 2 , and hence $\sigma_2^{K+1}(2a) = r_{2a}$ if and only if

$$\frac{9}{2} v_2^{K+1}(0 | 2a) + 3 \left(1 - v_2^{K+1}(0 | 2a) \right) < 4 \iff v_2^{K+1}(0 | 2a) < \frac{2}{3}.$$

Moreover, the induction hypothesis implies:

$$v_2^{K+1}(0 | 2a) = \frac{\frac{1}{2}p_0}{\frac{1}{2}p_0 + p_1 + \sum_{i=3}^K p_i} < \frac{\frac{1}{2}p_0}{\frac{1}{2}p_0 + p_1 + \sum_{i=3}^{K-1} p_i} = v_2^K(0 | 2a) < \frac{2}{3},$$

as desired. ■

C.2 Proofs for Section 3.5

Proof of Lemma 3.1

1. Since stage $2S$ is the last stage of the game, for any $k \geq 1$, player 2 would take at this stage if and only if

$$1 + (2S - 1)c > 2Sc \iff 1 > c,$$

which holds by assumption. Therefore, $\sigma_{2,S}^k = 1$ for all $k \geq 1$.

2. Consider a level-1 type of player 1 and any of player 1's decision nodes $j \in \{1, \dots, S\}$. The payoff from Take is $1 + (2j - 2)c$ and the expected payoff from Pass is greater than or equal to $\frac{1}{2}(2j - 1)c + \frac{1}{2}(1 + (2j)c)$. Thus, $\sigma_{1,j}^k = 0$ is strictly optimal if and only if:

$$\begin{aligned} 1 + (2j - 2)c &< \frac{1}{2}(2j - 1)c + \frac{1}{2}(1 + (2j)c) \\ &\iff \frac{1}{3} < c. \end{aligned}$$

Hence, $\sigma_1^1 = (0, \dots, 0)$. A similar argument shows that $\sigma_2^1 = (0, \dots, 0, 1)$.

3. The argument is similar to the proof of the first statement. Consider a level- k type of player 1 and any of player 1's decision nodes $j \in \{1, \dots, S - 1\}$, and suppose $\sigma_{2,j}^m = 0$ for every $1 \leq m \leq k - 1$. Then the payoff from Take is $1 + (2j - 2)c$ and the expected payoff from Pass is greater than or equal to $\frac{1}{2}v_{2j-1}^k(0)(2j - 1)c + (1 - \frac{1}{2}v_{2j-1}^k(0))(1 + (2j)c)$, which in turn is greater than or equal to $\frac{1}{2}(2j - 1)c + \frac{1}{2}(1 + (2j)c)$ because $v_{2j-1}^k(0) \leq 1$. Thus $\sigma_{1,j}^k = 0$ is strictly optimal if and only if:

$$\begin{aligned} 1 + (2j - 2)c &< \frac{1}{2}(2j - 1)c + \frac{1}{2}(1 + (2j)c) \\ &\iff \frac{1}{3} < c. \end{aligned}$$

Hence, $\sigma_{1,j}^k = 0$. A similar argument shows that $\sigma_{2,j}^k = 0$ if $\sigma_{1,j+1}^m = 0$ for every $1 \leq m \leq k - 1$. This completes the proof. ■

Proof of Proposition 3.4

1. The statement can be proved by induction. Consider stage $2S - 1$. By Lemma 3.1, we know $K_{2S}^* = 1$ and $\sigma_{1,S}^1 = 0$, suggesting $K_{2S-1}^* \geq 2 = K_{2S}^* + 1$. Now, fix any $2 \leq m \leq 2S - 1$ and suppose the statement holds for all stages $m \leq j \leq 2S - 1$. Without loss of generality, we consider an even m . We want to show that if $K_m^* < \infty$, then $K_{m-1}^* \geq K_m^* + 1$. By construction, we know $\sigma_{2,\frac{m}{2}}^k = 0$ for all $1 \leq k \leq K_m^* - 1$. Therefore, Lemma 3.1 implies $\sigma_{1,\frac{m}{2}} = 0$ for all $1 \leq k \leq K_m^*$, and $K_{m-1}^* \geq K_m^* + 1$.

2. Consider any j such that $K_{j+1}^* = \infty$. Without loss of generality, we consider an odd j . Hence, $\sigma_{2,\frac{j+1}{2}}^k = 0$ for all $k \geq 1$ and we want to show $\sigma_{1,\frac{j+1}{2}}^k = 0$ for all $k \geq 1$ by induction. Lemma 3.1 implies $\sigma_{1,\frac{j+1}{2}}^1 = 0$. Suppose there is $\bar{K} \geq 2$ such that $\sigma_{1,\frac{j+1}{2}}^k = 0$ for all $1 \leq k \leq \bar{K}$. Since $\sigma_{2,\frac{j+1}{2}}^k = 0$ for all $k \geq 1$, Lemma 3.1 implies $\sigma_{1,\frac{j+1}{2}}^{\bar{K}+1} = 0$, as desired. ■

Proof of Proposition 3.5

We prove this by induction. Consider stage $2S - 1$. Level K_{2S-1}^* player 1 believes that only level-0 player 2 will pass at stage $2S$, so

$$\begin{aligned} 1 + (2S - 2)c &> \left(1 - \frac{1}{2}v_{2S-1}^{K_{2S-1}^*}(0)\right) [(2S - 1)c] + \frac{1}{2}v_{2S-1}^{K_{2S-1}^*}(0)[1 + 2Sc] \\ &> \left(1 - \frac{1}{2}v_{2S-1}^k(0)\right) [(2S - 1)c] + \frac{1}{2}v_{2S-1}^k(0)[1 + 2Sc] \text{ for all } k > K_{2S-1}^* \end{aligned}$$

since $\frac{1}{2}v_{2S-1}^k(0) < \frac{1}{2}v_{2S-1}^{K_{2S-1}^*}(0)$ and therefore $\sigma_{1,S}^k = 1$ for all $k \geq K_{2S-1}^*$.

Next, suppose for any m where $2 \leq m \leq 2S - 1$, the statement holds for all j such that $m \leq j \leq 2S - 1$. Suppose m is odd and $K_{m-1}^* < \infty$. (A similar argument applies if m is even.) By construction, $\sigma_{2, \frac{m-1}{2}}^k = 0$ for all $1 \leq k \leq K_{m-1}^* - 1$. By Lemma 3.1, we have $\sigma_{1,s}^k = 0$ for all $1 \leq s \leq \frac{m-1}{2}$ and for all $1 \leq k \leq K_{m-1}^*$. Level K_{m-1}^* player 2's belief at stage $m - 1$ that the other player would pass at stage m is

$$\begin{aligned} \frac{1}{2}v_{m-1}^{K_{m-1}^*}(0) + \sum_{\kappa=1}^{K_{m-1}^*} v_{m-1}^{K_{m-1}^*}(\kappa) &= \frac{1}{2} \frac{p_0 \left(\frac{1}{2}\right)^{\frac{m-1}{2}}}{p_0 \left(\frac{1}{2}\right)^{\frac{m-1}{2}} + \sum_{\kappa=1}^{K_{m-1}^*-1} p_\kappa} + \frac{\sum_{\kappa=1}^{K_{m-1}^*-1} p_\kappa}{p_0 \left(\frac{1}{2}\right)^{\frac{m-1}{2}} + \sum_{\kappa=1}^{K_{m-1}^*-1} p_\kappa} \\ &= \frac{p_0 \left(\frac{1}{2}\right)^{\frac{m+1}{2}} + \sum_{\kappa=1}^{K_{m-1}^*-1} p_\kappa}{p_0 \left(\frac{1}{2}\right)^{\frac{m-1}{2}} + \sum_{\kappa=1}^{K_{m-1}^*-1} p_\kappa}. \end{aligned}$$

Since $\sigma_{1,s}^{K_{m-1}^*} = 0$ for all $1 \leq s \leq \frac{m-1}{2}$, then for any $k > K_{m-1}^*$, at stage $m - 1$ level- k player 2's belief about the probability that the other player would pass at stage m is

$$\begin{aligned} \frac{1}{2}v_{m-1}^k(0) + \sum_{\kappa=1}^{K_{m-1}^*} v_{m-1}^k(\kappa) &\leq \frac{1}{2}v_{m-1}^{K_{m-1}^*+1}(0) + \sum_{\kappa=1}^{K_{m-1}^*} v_{m-1}^{K_{m-1}^*+1}(\kappa) \\ &= \frac{1}{2} \frac{p_0 \left(\frac{1}{2}\right)^{\frac{m-1}{2}}}{p_0 \left(\frac{1}{2}\right)^{\frac{m-1}{2}} + \sum_{\kappa=1}^{K_{m-1}^*} p_\kappa} + \frac{\sum_{\kappa=1}^{K_{m-1}^*} p_\kappa}{p_0 \left(\frac{1}{2}\right)^{\frac{m-1}{2}} + \sum_{\kappa=1}^{K_{m-1}^*} p_\kappa} \\ &< \frac{p_0 \left(\frac{1}{2}\right)^{\frac{m+1}{2}} + \sum_{\kappa=1}^{K_{m-1}^*} p_\kappa}{p_0 \left(\frac{1}{2}\right)^{\frac{m-1}{2}} + \sum_{\kappa=1}^{K_{m-1}^*} p_\kappa}. \end{aligned}$$

since $\sum_{\kappa=1}^{K_{m-1}^*} p_\kappa > \sum_{\kappa=1}^{K_{m-1}^*-1} p_\kappa$. This implies that for any level $k > K_{m-1}^*$, higher level of player 2 at stage $m - 1$ would think the other player is less likely to pass at stage m . Since it is already profitable for level K_{m-1}^* player 2 to take at stage $m - 1$, we can conclude that $\sigma_{2, \frac{m-1}{2}}^k = 1$ for all $k \geq K_{m-1}^*$. ■

Proof of Proposition 3.7

Without loss of generality, we can consider an even j , so $\lfloor \frac{j}{2} \rfloor = \frac{j}{2}$ and it is player 2's turn at stage j .

Only if: Suppose $K_j^* < \infty$. Then from Proposition 3.4, $K_{j'}^* \geq K_j^* + 1$ for all $j' < j$. Hence, the belief of level K_j^* of player 2 that player 1 is level-0 at stage j equals to

$$v_j^{K_j^*}(0) = \frac{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}}}{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}} + \sum_{\kappa=1}^{K_j^*-1} p_\kappa} > \frac{p_0 \left(\frac{1}{2}\right)^{S-1}}{p_0 \left(\frac{1}{2}\right)^{S-1} + (1-p_0)}.$$

Level K_j^* player 2's belief at stage j that the player 1 would pass at stage $j+1$ is

$$\begin{aligned} \frac{1}{2} v_j^{K_j^*}(0) + \sum_{\kappa=1}^{K_{j+1}^*-1} v_j^{K_j^*}(\kappa) &= \frac{1}{2} \frac{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}}}{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}} + \sum_{\kappa=1}^{K_j^*-1} p_\kappa} + \frac{\sum_{\kappa=1}^{K_{j+1}^*-1} p_\kappa}{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}} + \sum_{\kappa=1}^{K_j^*-1} p_\kappa} \\ &= \frac{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}+1} + \sum_{\kappa=1}^{K_{j+1}^*-1} p_\kappa}{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}} + \sum_{\kappa=1}^{K_j^*-1} p_\kappa}, \end{aligned}$$

where we know that $K_{j+1}^* \leq K_j^* - 1$. Since it is optimal for level $K_j^* < \infty$ to take at j

this implies $\frac{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}+1} + \sum_{\kappa=1}^{K_{j+1}^*-1} p_\kappa}{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}} + \sum_{\kappa=1}^{K_j^*-1} p_\kappa} < \frac{1-c}{1+c}$ and therefore

$$\frac{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}+1} + \sum_{\kappa=1}^{K_{j+1}^*-1} p_\kappa}{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}} + (1-p_0)} < \frac{1-c}{1+c}.$$

If: Suppose $K_j^* = \infty$. Then from Proposition 3.4, $K_{j'}^* = \infty$ for all $j' < j$. That is, all levels of both players pass at every stage up to and including j . Hence the belief of level $k \geq 1$ of player 2 that player 1 is level-0 at stage j equals to

$$v_j^k(0) = \frac{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}}}{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}} + \sum_{\kappa=1}^{k-1} p_\kappa} > \frac{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}}}{p_0 \left(\frac{1}{2}\right)^{\frac{j}{2}} + (1-p_0)}.$$

Since $K_j^* = \infty$ it is optimal to pass at j for all levels $k \geq 1$ of player 2, which implies

$$\frac{\frac{1}{2}p_0\left(\frac{1}{2}\right)^{\frac{j}{2} + \sum_{\kappa=1}^{K_{j+1}^*-1} p_\kappa}}{p_0\left(\frac{1}{2}\right)^{\frac{j}{2} + \sum_{\kappa=1}^{k-1} p_\kappa}} \geq \frac{1-c}{1+c}, \text{ for all } k, \text{ where possibly } K_{j+1}^* = \infty, \text{ so}$$

$$\frac{\frac{1}{2}p_0\left(\frac{1}{2}\right)^{\frac{j}{2} + \sum_{\kappa=1}^{K_{j+1}^*-1} p_\kappa}}{p_0\left(\frac{1}{2}\right)^{\frac{j}{2} + (1-p_0)}} \geq \frac{1-c}{1+c},$$

as desired. ■

Proof of Lemma 3.2

1. To prove the statement, we can discuss player 1 and 2 separately.

Player 2:

(i) $a_2^1 = S$ strictly dominates $a_2^1 = S + 1$: $\mathbb{E}[u_2(a_1^0, S)] - \mathbb{E}[u_2(a_1^0, S + 1)] = \frac{1-c}{S+1} > 0$ since $c < 1$.

(ii) $a_2^1 = j + 1$ strictly dominates $a_2^1 = j$ for all $1 \leq j \leq S - 1$: For $1 \leq j \leq S - 1$, since $c > \frac{1}{3}$,

$$\mathbb{E}[u_2(a_1^0, j + 1)] - \mathbb{E}[u_2(a_1^0, j)] = \frac{1}{S+1} [-1 + (2S - 2j + 1)c] \geq \frac{1}{S+1} (-1 + 3c) > 0.$$

Hence, we can obtain that $a_2^1 = S$.

Player 1: By the same logic as (ii) above, $a_1^1 = j + 1$ strictly dominates $a_1^1 = j$ for all $1 \leq j \leq S$: For $1 \leq j \leq S$, since $c > \frac{1}{3}$,

$$\mathbb{E}[u_1(j + 1, a_2^0)] - \mathbb{E}[u_1(j, a_2^0)] = \frac{1}{S+1} [-1 + (2S - 2j + 3)c] \geq \frac{1}{S+1} (-1 + 3c) > 0.$$

Hence, we can obtain that $a_1^1 = S + 1$.

2. (i) Notice that for any a_2 , $u_1(a_1, a_2)$ is maximized at $a_1 = a_2$. Fix level $k \geq 2$. If level- k player 1 chooses s , then the expected payoff is

$$\begin{aligned} V_1^k(s) &\equiv \sum_{\kappa=0}^{k-1} \tilde{p}_\kappa^k \mathbb{E}[u_1(s, a_2^\kappa)] \\ &= \tilde{p}_0^k \mathbb{E}[u_1(s, a_2^0)] + \sum_{\kappa=1}^{k-1} \tilde{p}_\kappa^k u_1(s, a_2^\kappa). \end{aligned}$$

Suppose $\min\{a_2^m : 1 \leq m \leq k-1\} = 1$, then (i) holds trivially. If $\min\{a_2^m : 1 \leq m \leq k-1\} \geq 2$, then we can prove the statement by contradiction. Suppose $a_1^k < \min\{a_2^m : 1 \leq m \leq k-1\}$, then

$$\begin{aligned} V_1^k(a_1^k) &= \tilde{p}_0^k \underbrace{\mathbb{E}[u_1(a_1^k, a_2^0)]}_{< \mathbb{E}[u_1(a_1^{k+1}, a_2^0)]} + \sum_{\kappa=1}^{k-1} \tilde{p}_\kappa^k \underbrace{u_1(a_1^k, a_2^\kappa)}_{\leq u_1(a_1^{k+1}, a_2^\kappa)} \\ &< V_1^k(a_1^{k+1}). \end{aligned}$$

$\mathbb{E}[u_1(a_1^k, a_2^0)] < \mathbb{E}[u_1(a_1^{k+1}, a_2^0)]$ follows from the first statement. Furthermore, $a_1^k < \min\{a_2^m : 1 \leq m \leq k-1\}$ implies $u_1(a_1^k, a_2^\kappa) \leq u_1(a_1^{k+1}, a_2^\kappa)$ for all $1 \leq \kappa \leq k-1$. Hence, $a_1^k < \min\{a_2^m : 1 \leq m \leq k-1\}$ is not optimal for level- k player 1, a contradiction.

(ii) The logic is similar for player 2.

3. We prove this statement by induction on k . First, it holds for $k = 1$, by the first statement. Next, we suppose it holds for any k where $1 \leq k \leq K-1$ and prove it holds for $k = K$. For level $K+1$ player 1, the expected payoff for choosing s is

$$\begin{aligned} V_1^{K+1}(s) &= \tilde{p}_0^{K+1} \mathbb{E}[u_1(s, a_2^0)] + \sum_{\kappa=1}^K \tilde{p}_\kappa^{K+1} u_1(s, a_2^\kappa) \\ &= \left(\frac{\sum_{\kappa=0}^{K-1} p_\kappa}{\sum_{\kappa=0}^K p_\kappa} \right) V_1^K(s) + \tilde{p}_K^{K+1} u_1(s, a_2^K). \end{aligned}$$

Suppose, by way of contradiction, that $a_1^{K+1} > a_1^K$. Then $V_1^K(a_1^{K+1}) < V_1^K(a_1^K)$. From the induction hypothesis, $a_2^K \leq a_2^{K-1}$, and from the second statement, $a_1^K \geq a_2^{K-1}$ and hence $a_1^{K+1} > a_1^K \geq a_2^{K-1} \geq a_1^K$. This implies $u_1(a_1^{K+1}, a_2^K) \leq u_1(a_1^K, a_2^K)$, so $V_1^{K+1}(a_1^{K+1}) < V_1^{K+1}(a_1^K)$, which contradicts that a_1^{K+1} is the optimal strategy for level $K+1$ player 1. Hence $a_1^{K+1} \leq a_1^K$, so the result is proved for $i = 1$. A similar argument proves the result for $i = 2$. ■

Proof of Proposition 3.8

With slight abuse of notation, denote a level- k player's prior belief that the opponent is level- κ by $\mu_\kappa^k \equiv \frac{p_\kappa}{\sum_{j=0}^{k-1} p_j}$, $\kappa = 1, \dots, k-1$.

Only if: Suppose $p_0 \geq \frac{S+1}{(S+1) + (\frac{3c-1}{1-c})}$, then we want to show that $a_1^k = S+1$ for all $k \geq 1$. We can prove this statement by induction on k . By Lemma 3.2, we know $a_1^1 = S+1$. Now, suppose this statement holds for all $1 \leq k \leq K$ for some $K \in \mathbb{N}$,

then we want to show this holds for level $K + 1$ player 1. First, by Lemma 3.2, we have $a_2^k = S$ for all $1 \leq k \leq K$. Level $K + 1$ player 1 would choose S if and only if

$$\begin{aligned} & \mu_0^{K+1} \left[\frac{1}{S+1} \left[1 + 2Sc + \sum_{i=2}^{S+1} (2i-3)c \right] \right] + (1 - \mu_0^{K+1}) (2S-1)c \\ & < \mu_0^{K+1} \left[\frac{1}{S+1} \left[2(1 + (2S-2)c) + \sum_{i=2}^S (2i-3)c \right] \right] + (1 - \mu_0^{K+1}) (1 + (2S-2)c) \\ & \iff \mu_0^{K+1} \left[\frac{1}{S+1} (1-3c) \right] + (1 - \mu_0^{K+1}) (1-c) > 0 \\ & \iff \mu_0^{K+1} < \frac{S+1}{(S+1) + \left(\frac{3c-1}{1-c}\right)}. \end{aligned}$$

However, we know $\mu_0^K > p_0$ and we have assumed $p_0 \geq \frac{S+1}{(S+1) + \left(\frac{-1+3c}{1-c}\right)}$, so $\mu_0^{K+1} > p_0 \geq \frac{S+1}{(S+1) + \left(\frac{3c-1}{1-c}\right)}$, implying that $a_1^{K+1} = S + 1$.

If: Suppose $p_0 < \frac{S+1}{(S+1) + \left(\frac{3c-1}{1-c}\right)}$, then there exists $N^* < \infty$ such that $\mu_0^{N^*} < \frac{S+1}{(S+1) + \left(\frac{3c-1}{1-c}\right)}$. Therefore, by a previous calculation we have that

$$\tilde{K}_{2S-1}^* = \arg \min_{N^*} \left\{ \mu_0^{N^*} < \frac{S+1}{(S+1) + \left(\frac{3c-1}{1-c}\right)} \right\} < \infty,$$

which is the lowest level of player 1 who would take at no later than stage $2S - 1$. ■

Proof of Proposition 3.9

First, an immediate implication of Lemma 3.2 is that for all level $k \geq 1$, the optimal choice for level- $(k+1)$ is either the same as level- k or to take at one stage earlier. Given this observation, the logic of the proof is similar to Proposition 3.7.

Only if: For any $1 \leq j \leq 2S - 2$, suppose $p_0 \left(\frac{S}{S+1} - \frac{2\lfloor \frac{j}{2} \rfloor c}{(S+1)(1+c)} \right) + \sum_{\kappa=1}^{\tilde{K}_{j+1}^*-1} p_\kappa \geq \frac{1-c}{1+c}$, then we want to show $\tilde{K}_j^* = \infty$. Without loss of generality, we consider an odd j . If $\tilde{K}_{j+1}^* = \infty$, then the statement holds immediately. Otherwise, we can prove $a_1^k > \frac{j+1}{2}$ for all $k \geq 1$ by induction. By construction, we know $a_2^m > \frac{j+1}{2}$ for all $1 \leq m \leq \tilde{K}_{j+1}^* - 1$ and $a_1^k > \frac{j+1}{2}$ for all $1 \leq k \leq \tilde{K}_{j+1}^*$ by Lemma 3.2. Suppose there is some $K \geq \tilde{K}_{j+1}^* + 1$ such that $a_1^k > \frac{j+1}{2}$ for all $1 \leq k \leq K$. We want to show this holds for level $K + 1$ player 1. Level $K + 1$ player 1 would choose $\frac{j+1}{2} + 1$ if and only if

$$p_0 \left[\frac{1}{S+1} (1 - (2S - j + 2)c) \right] + \left(\sum_{\kappa=1}^{\tilde{K}_{j+1}^*-1} p_\kappa \right) (-2c) + \left(\sum_{\kappa=\tilde{K}_{j+1}^*}^K p_\kappa \right) (1-c) \leq 0.$$

Moreover, we can observe that this condition is implied by

$$p_0 \left[\frac{1}{S+1} (1 - (2S - j + 2)c) \right] + \left(\sum_{\kappa=1}^{\tilde{K}_{j+1}^*-1} p_{\kappa} \right) (-2c) + \left(1 - p_0 - \sum_{\kappa=1}^{\tilde{K}_{j+1}^*-1} p_{\kappa} \right) (1 - c) \leq 0$$

$$\iff p_0 \left[\frac{S}{S+1} - \frac{(j-1)c}{(S+1)(1+c)} \right] + \sum_{\kappa=1}^{\tilde{K}_{j+1}^*-1} p_{\kappa} \geq \frac{1-c}{1+c}.$$

By our assumption, we can conclude that the optimal choice for level $(K+1)$ player 1 is $\frac{j+1}{2} + 1$,¹ which completes the only if part of the proof.

If: For any $1 \leq j \leq 2S - 2$, suppose

$$p_0 \left(\frac{S}{S+1} - \frac{2\lfloor \frac{j}{2} \rfloor c}{(S+1)(1+c)} \right) + \sum_{\kappa=1}^{\tilde{K}_{j+1}^*-1} p_{\kappa} < \frac{1-c}{1+c},$$

then there exists N^* where $\tilde{K}_{j+1}^* + 1 \leq N^* < \infty$ such that

$$\mu_0^{N^*} \left(\frac{S}{S+1} - \frac{2\lfloor \frac{j}{2} \rfloor c}{(S+1)(1+c)} \right) + \frac{\sum_{\kappa=1}^{\tilde{K}_{j+1}^*-1} p_{\kappa}}{\sum_{\kappa=0}^{N^*-1} p_{\kappa}} < \frac{1-c}{1+c}.$$

Therefore, by previous calculation and the existence of such $N^* < \infty$, we can obtain that

$$\tilde{K}_j^* = \arg \min_{N^*} \left\{ \mu_0^{N^*} \left(\frac{S}{S+1} - \frac{2\lfloor \frac{j}{2} \rfloor c}{(S+1)(1+c)} \right) + \frac{\sum_{\kappa=1}^{\tilde{K}_{j+1}^*-1} p_{\kappa}}{\sum_{\kappa=0}^{N^*-1} p_{\kappa}} < \frac{1-c}{1+c} \right\} < \infty,$$

which is the lowest level of player who would take at no later than stage j . ■

Proof of Theorem 3.1

Step 1: By Lemma 3.1 and Lemma 3.2, we can obtain that $1 = K_{2S}^* \leq \tilde{K}_{2S}^* = 1$, suggesting that the inequality holds at stage $2S$.

Step 2: By Proposition 3.6 and 3.8, we know K_{2S-1}^* and \tilde{K}_{2S-1}^* are the lowest levels such that

¹If j is even, by the same argument, we can obtain level $(K+1)$ player 2 would choose $\frac{j}{2} + 1$ as

$$p_0 \left(\frac{S}{S+1} - \frac{jc}{(S+1)(1+c)} \right) + \sum_{l=1}^{\tilde{K}_{j+1}^*-1} p_l \geq \frac{1-c}{1+c}.$$

$$\frac{p_0}{\sum_{\kappa=0}^{K_{2S-1}^*} p_\kappa} < \frac{2^S}{2^S + \left(\frac{-1+3c}{1-c}\right)}, \text{ and}$$

$$\frac{p_0}{\sum_{\kappa=0}^{\tilde{K}_{2S-1}^*} p_\kappa} < \frac{S+1}{(S+1) + \left(\frac{-1+3c}{1-c}\right)}, \text{ respectively.}$$

We can observe that $\frac{S+1}{(S+1) + \left(\frac{-1+3c}{1-c}\right)} < \frac{2^S}{2^S + \left(\frac{-1+3c}{1-c}\right)}$, suggesting the inequality for the dynamic model is less stringent. Hence, we can obtain that $K_{2S-1}^* \leq \tilde{K}_{2S-1}^*$.

Step 3: We can finish the proof by induction on the stages. At stage $2S - 2$, as we rearrange the condition from Proposition 3.7, we can obtain K_{2S-2}^* is the lowest level such that

$$\sum_{\kappa=1}^{K_{2S-2}^*} p_\kappa > p_0 \left(\frac{1}{2}\right)^S \left(\frac{-1+3c}{1-c}\right) + \left(\sum_{\kappa=1}^{K_{2S-1}^*} p_\kappa\right) \left(\frac{1+c}{1-c}\right). \quad (\text{C.1})$$

Similarly, as we rearrange the necessary and sufficient condition from Proposition 3.9, we can find that \tilde{K}_{2S-2}^* is the lowest level such that

$$\sum_{\kappa=1}^{\tilde{K}_{2S-2}^*} p_\kappa > p_0 \left(\frac{1}{S+1}\right) \left(\frac{-1+3c}{1-c}\right) + \left(\sum_{\kappa=1}^{\tilde{K}_{2S-1}^*} p_\kappa\right) \left(\frac{1+c}{1-c}\right). \quad (\text{C.2})$$

It suffices to prove $K_{2S-2}^* \leq \tilde{K}_{2S-2}^*$ by showing the right-hand side of Condition (C.1) is smaller than the right-hand side of (C.2). This holds because $\left(\frac{1}{2}\right)^S < \frac{1}{S+1}$ for all $S \geq 2$ and $K_{2S-1}^* \leq \tilde{K}_{2S-1}^*$ as we have shown in *step 2*.

Step 4: Consider any j where $3 \leq j \leq 2S - 1$ and suppose $K_{2S-i}^* \leq \tilde{K}_{2S-i}^*$ for all $0 \leq i \leq j - 1$. We want to show $K_{2S-j}^* \leq \tilde{K}_{2S-j}^*$. Without loss of generality, we consider an odd j . That is, player 1 owns stage $2S - j$. By Proposition 3.7, we know K_{2S-j}^* is the lowest level such that

$$\sum_{\kappa=1}^{K_{2S-j}^*} p_\kappa > p_0 \left(\frac{1}{2}\right)^{S-\frac{j+1}{2}+1} \left(\frac{-1+3c}{1-c}\right) + \left(\sum_{\kappa=1}^{K_{2S-j+1}^*} p_\kappa\right) \left(\frac{1+c}{1-c}\right). \quad (\text{C.3})$$

Similarly, as we rearrange the necessary and sufficient condition from Proposition 3.9, we can obtain that \tilde{K}_{2S-j}^* is the lowest level such that

$$\sum_{\kappa=1}^{\tilde{K}_{2S-j}^*} p_\kappa > p_0 \left(\frac{1}{S+1}\right) \left[\frac{-1+(j+2)c}{1-c}\right] + \left(\sum_{\kappa=1}^{\tilde{K}_{2S-j+1}^*} p_\kappa\right) \left(\frac{1+c}{1-c}\right). \quad (\text{C.4})$$

Similar to the previous step, we can finish the proof by showing the right-hand side of Condition (C.3) is smaller than the right-hand side of (C.4). The induction hypothesis implies the second term of (C.4) is larger than the second term of (C.3). Hence, the only thing left to show is

$$\left(\frac{1}{2}\right)^{S-\frac{j+1}{2}+1} \left(\frac{-1+3c}{1-c}\right) < \left(\frac{1}{S+1}\right) \left[\frac{-1+(j+2)c}{1-c}\right].$$

Or equivalently,

$$(S+1)(-1+3c) < 2^{S-\frac{j+1}{2}+1}(-1+(j+2)c). \quad (\text{C.5})$$

Since $3 \leq j \leq 2S-1$, there is nothing to show if $S < \frac{j+1}{2}$. When $S \geq \frac{j+1}{2}$, we know (C.5) would hold in the following three different cases.

– *Case 1:* If $S+1 = 2^{S-\frac{j+1}{2}+1}$, then (C.5) becomes $-1+3c < -1+(j+2)c \iff j > 1$.

– *Case 2:* If $S+1 < 2^{S-\frac{j+1}{2}+1}$, then (C.5) is equivalent to

$$2^{S-\frac{j+1}{2}+1} - (S+1) < \left[(j+2)2^{S-\frac{j+1}{2}+1} - 3(S+1) \right] c \iff 1 < \left[3 + \underbrace{\frac{(j-1)2^{S-\frac{j+1}{2}+1}}{2^{S-\frac{j+1}{2}+1} - (S+1)}}_{>0 \text{ as } j \geq 3} \right] c,$$

which holds under our assumption $c > \frac{1}{3}$.

– *Case 3:* If $S+1 > 2^{S-\frac{j+1}{2}+1}$, then (C.5) can be rearranged as

$$(S+1) - 2^{S-\frac{j+1}{2}+1} > \left[3(S+1) - (j+2)2^{S-\frac{j+1}{2}+1} \right] c \iff 1 > \left[3 - \frac{j-1}{\left(\frac{S+1}{2^{S-\frac{j+1}{2}+1}}\right) - 1} \right] c.$$

The right-hand side of the inequality is negative since

$$3 - \frac{j-1}{\left(\frac{S+1}{2^{S-\frac{j+1}{2}+1}}\right) - 1} \leq 3 - \frac{j-1}{\left(\frac{\frac{j+1}{2}+1}{2}\right) - 1} = -1.$$

This completes the proof. ■

Proof of Proposition 3.10

By Proposition 3.6, we know

$$K_{2S-1}^* < \infty \iff p_0 < \frac{2^S}{2^S \left(\frac{-1+3c}{1-c} \right)}.$$

As the prior distribution follows $\text{Poisson}(\lambda)$, the condition becomes

$$\begin{aligned} K_{2S-1}^*(\lambda) < \infty &\iff e^{-\lambda} < \frac{2^S}{2^S \left(\frac{-1+3c}{1-c} \right)} \\ &\iff \lambda > \ln \left[1 + \left(\frac{1}{2} \right)^S \left(\frac{-1+3c}{1-c} \right) \right]. \end{aligned}$$

Similarly, by Proposition 3.8, we know

$$\tilde{K}_{2S-1}^* < \infty \iff p_0 < \frac{S+1}{(S+1) \left(\frac{-1+3c}{1-c} \right)},$$

which can be rearranged to the following expression when the prior distribution follows $\text{Poisson}(\lambda)$:

$$\begin{aligned} \tilde{K}_{2S-1}^*(\lambda) < \infty &\iff e^{-\lambda} < \frac{S+1}{(S+1) \left(\frac{-1+3c}{1-c} \right)} \\ &\iff \lambda > \ln \left[1 + \left(\frac{1}{S+1} \right) \left(\frac{-1+3c}{1-c} \right) \right]. \end{aligned}$$

This completes the proof. ■

Proof of Proposition 3.11

Here we show the existence of λ^* . The existence of $\tilde{\lambda}^*$ can be proven by the same argument.

Step 1: By Proposition 3.1, we know for all $\lambda > 0$, $K_{2S}^*(\lambda) = 1$.

Step 2: By Proposition 3.10, we know

$$K_{2S-1}^*(\lambda) < \infty \iff \lambda > \ln \left[1 + \left(\frac{1}{2} \right)^S \left(\frac{-1+3c}{1-c} \right) \right] \equiv \lambda_{2S-1}^*.$$

Since $\lambda_{2S-1}^* < \infty$, we know $K_{2S-1}^*(\lambda) < \infty \iff \lambda > \lambda_{2S-1}^*$.

Step 3: By Proposition 3.7, we know

$$K_{2S-2}^*(\lambda) < \infty \iff \frac{e^{-\lambda} \left(\frac{1}{2}\right)^S + e^{-\lambda} \sum_{l=1}^{K_{2S-1}^*(\lambda)-1} \frac{\lambda^l}{l!}}{e^{-\lambda} \left(\frac{1}{2}\right)^{S-1} + (1 - e^{-\lambda})} < \frac{1-c}{1+c}$$

$$\iff 1 - e^{-\lambda} \left[1 + \left(\frac{1}{2}\right)^S \left(\frac{-1+3c}{1-c}\right) \right] - e^{-\lambda} \sum_{\kappa=1}^{K_{2S-1}^*(\lambda)-1} \frac{\lambda^\kappa}{\kappa!} \left(\frac{1+c}{1-c}\right) > 0.$$

Notice that by **step 2**, we know there exists some $M < \infty$, such that for all $\lambda > \lambda_{2S-1}^*, K_{2S-1}^*(\lambda) < M$. Moreover, by Proposition 3.4, we know $K_{2S-1}^*(\lambda) \geq 2$. Hence,

$$0 = \lim_{\lambda \rightarrow \infty} \frac{\lambda}{e^\lambda} \leq \lim_{\lambda \rightarrow \infty} e^{-\lambda} \sum_{\kappa=1}^{K_{2S-1}^*(\lambda)-1} \frac{\lambda^\kappa}{\kappa!} \leq \lim_{\lambda \rightarrow \infty} e^{-\lambda} \sum_{\kappa=1}^{M-1} \frac{\lambda^\kappa}{\kappa!} = 0.$$

Coupled with the fact that $\lim_{\lambda \rightarrow \infty} e^{-\lambda} = 0$, we can conclude that there exists λ_{2S-2}^* such that $\lambda_{2S-1}^* < \lambda_{2S-2}^* < \infty$ and $K_{2S-2}^*(\lambda) < \infty \iff \lambda > \lambda_{2S-2}^*$.

Step 4: Now we can prove this statement by induction on each stage. Consider any j where $3 \leq j \leq 2S-1$ and suppose there exists $\lambda_{2S-j+1}^* < \infty$ such that $K_{2S-j+1}^*(\lambda) < \infty$ for all $\lambda > \lambda_{2S-j+1}^*$. By Proposition 3.7, we know

$$K_{2S-j}^*(\lambda) < \infty \iff \frac{e^{-\lambda} \left(\frac{1}{2}\right)^{\lfloor \frac{2S-j}{2} \rfloor + 1} + e^{-\lambda} \sum_{\kappa=1}^{K_{2S-j+1}^*(\lambda)-1} \frac{\lambda^\kappa}{\kappa!}}{e^{-\lambda} \left(\frac{1}{2}\right)^{\lfloor \frac{2S-j}{2} \rfloor} + (1 - e^{-\lambda})} < \frac{1-c}{1+c}$$

$$\iff 1 - e^{-\lambda} \left[1 + \left(\frac{1}{2}\right)^{\lfloor \frac{2S-j}{2} \rfloor + 1} \left(\frac{-1+3c}{1-c}\right) \right] - e^{-\lambda} \sum_{\kappa=1}^{K_{2S-j+1}^*(\lambda)-1} \frac{\lambda^\kappa}{\kappa!} \left(\frac{1+c}{1-c}\right) > 0.$$

By the induction hypothesis, we know there exists some $L < \infty$ such that for all $\lambda > \lambda_{2S-j+1}^*, K_{2S-j+1}^*(\lambda) < L$. Proposition 3.4 gives us $K_{2S-j+1}^*(\lambda) \geq j$, and hence,

$$0 = \lim_{\lambda \rightarrow \infty} e^{-\lambda} \left(\sum_{\kappa=1}^{j-1} \frac{\lambda^\kappa}{\kappa!} \right) \leq \lim_{\lambda \rightarrow \infty} e^{-\lambda} \sum_{\kappa=1}^{K_{2S-j+1}^*(\lambda)-1} \frac{\lambda^\kappa}{\kappa!} \leq \lim_{\lambda \rightarrow \infty} e^{-\lambda} \sum_{\kappa=1}^{L-1} \frac{\lambda^\kappa}{\kappa!} = 0.$$

Combined with the fact $\lim_{\lambda \rightarrow \infty} e^{-\lambda} = 0$, we have proved that there exists λ_{2S-j}^* such that $\lambda_{2S-j+1}^* < \lambda_{2S-j}^* < \infty$ and $K_{2S-j}^*(\lambda) < \infty \iff \lambda > \lambda_{2S-j}^*$. Thus, λ_1^* is the desired λ^* . ■

Proof of Proposition 3.12

The proofs of Propositions 3.6 through 3.9 provide a recipe to derive the necessary and sufficient conditions for complete unraveling at each stage. That is, when the prior distribution follows Poisson distribution, we can compute the minimum λ for both models such that the predictions coincide with the standard level- k model. In the dynamic model, we can obtain from Proposition 3.6 and 3.7 that for any stage $2S - j$ where $1 \leq j \leq 2S - 1$,

$$\begin{aligned}
 K_{2S-1}^*(\lambda) = 2 &\iff \frac{e^{-\lambda}}{e^{-\lambda} + \lambda e^{-\lambda}} < \frac{2^S}{2^S + \left(\frac{-1+3c}{1-c}\right)} \iff \lambda > \left(\frac{1}{2}\right)^S \left(\frac{-1+3c}{1-c}\right) \equiv \lambda_{2S-1}^{**}, \text{ and} \\
 K_{2S-j}^*(\lambda) = j + 1 &\iff \sum_{\kappa=1}^j \frac{\lambda^\kappa e^{-\lambda}}{\kappa!} > e^{-\lambda} \left(\frac{1}{2}\right)^{S-\lfloor \frac{j+1}{2} \rfloor + 1} \left(\frac{-1+3c}{1-c}\right) + \left(\sum_{\kappa=1}^{j-1} \frac{\lambda^\kappa e^{-\lambda}}{\kappa!}\right) \left(\frac{1+c}{1-c}\right) \\
 &\iff \frac{1}{j!} \lambda^j - \left(\frac{2c}{1-c}\right) \left(\sum_{\kappa=1}^{j-1} \frac{\lambda^\kappa}{\kappa!}\right) > \left(\frac{1}{2}\right)^{S-\lfloor \frac{j+1}{2} \rfloor + 1} \left(\frac{-1+3c}{1-c}\right) \equiv M_{2S-j}^{**}.
 \end{aligned}$$

Similarly, we know from Proposition 3.8 and Proposition 3.9 that for any stage $2S - j$ where $1 \leq j \leq 2S - 1$,

$$\begin{aligned}
 \tilde{K}_{2S-1}^*(\lambda) = 2 &\iff \frac{e^{-\lambda}}{e^{-\lambda} + \lambda e^{-\lambda}} < \frac{S+1}{(S+1) + \left(\frac{-1+3c}{1-c}\right)} \iff \lambda > \frac{-1+3c}{(S+1)(1-c)} \equiv \tilde{\lambda}_{2S-1}^{**}, \text{ and} \\
 \tilde{K}_{2S-j}^*(\lambda) = j + 1 &\iff \sum_{\kappa=1}^j \frac{\lambda^\kappa e^{-\lambda}}{\kappa!} > e^{-\lambda} \left(\frac{-1 + (2\lfloor \frac{j+1}{2} \rfloor + 1)c}{(S+1)(1-c)}\right) + \left(\sum_{\kappa=1}^{j-1} \frac{\lambda^\kappa e^{-\lambda}}{\kappa!}\right) \left(\frac{1+c}{1-c}\right) \\
 &\iff \frac{1}{j!} \lambda^j - \left(\frac{2c}{1-c}\right) \left(\sum_{\kappa=1}^{j-1} \frac{\lambda^\kappa}{\kappa!}\right) > \left(\frac{-1 + (2\lfloor \frac{j+1}{2} \rfloor + 1)c}{(S+1)(1-c)}\right) \equiv \tilde{M}_{2S-j}^{**}.
 \end{aligned}$$

First, we can find $\lambda_{2S-1}^{**} < \tilde{\lambda}_{2S-1}^{**}$ since $\left(\frac{1}{2}\right)^S < \frac{1}{S+1}$. Moreover, because the LHS of each inequality is a degree of j polynomial of λ , it has only one positive root by Descartes' rule of signs. Hence, it suffices to prove $M_{2S-j}^{**} < \tilde{M}_{2S-j}^{**}$, or equivalently, $\frac{\tilde{M}_{2S-j}^{**}}{M_{2S-j}^{**}} > 1$, for all $2 \leq j \leq 2S - 1$. Due to the property of floor functions, we can focus on odd j without loss of generality. Also, we can observe that this ratio is

decreasing in j since for any odd j where $3 \leq j \leq 2S - 3$,

$$\begin{aligned} \frac{\tilde{M}_{2S-(j+2)}^{**}}{M_{2S-(j+2)}^{**}} &= \left(\frac{2^{S-\frac{j+1}{2}}}{S+1} \right) \left(\frac{-1+(j+4)c}{-1+3c} \right) \\ &= \frac{1}{2} \left(\frac{\tilde{M}_{2S-j}^{**}}{M_{2S-j}^{**}} \right) + \frac{1}{2} \left[\left(\frac{2^{S-\frac{j-1}{2}}}{S+1} \right) \left(\frac{2c}{-1+3c} \right) \right] < \frac{\tilde{M}_{2S-j}^{**}}{M_{2S-j}^{**}} \\ &\iff 2c < -1+(j+2)c \iff 1 < jc, \end{aligned}$$

which holds because of the assumption $c > \frac{1}{3}$. The monotonicity implies that the ratio is minimized when $j = 2S - 1$, and we can obtain the conclusion by showing $\frac{\tilde{M}_1^{**}}{M_1^{**}} > 1$:

$$\frac{\tilde{M}_1^{**}}{M_1^{**}} = \left(\frac{2}{S+1} \right) \left(\frac{-1+(2S+1)c}{-1+3c} \right) > 1 \iff (S-1)(1+c) > 0,$$

as desired. ■

Proof of Proposition 3.13

Here we only provide the proof for the case where the game is played in extensive form. A very similar argument can be applied to the case where the game is played in reduced normal form. First of all, by Proposition 3.1, we know $K_{2S}^*(\lambda) = 1$ for all $\lambda > 0$. Therefore, it is weakly decreasing in λ .

To show the monotonicity of $K_{2S-1}^*(\lambda)$, we need to introduce the function $F_k(\lambda) : \mathbb{R}_{++} \rightarrow \mathbb{R}$ where $k \in \mathbb{N}$ and $F_k(\lambda) = \sum_{\kappa=1}^k \frac{\lambda^\kappa}{\kappa!}$. Notice that $F_{k+1}(\lambda) > F_k(\lambda)$ for all $\lambda > 0$, and $F_k(\lambda)$ is strictly increasing since $F'_k(\lambda) = \sum_{\kappa=0}^{k-1} \frac{\lambda^\kappa}{\kappa!} > 0$ for all $\lambda > 0$. We prove the monotonicity toward contradiction. By Proposition 3.6, we know $K_{2S-1}^*(\lambda)$ is the lowest level such that

$$F_{K_{2S-1}^*(\lambda)-1}(\lambda) > \left(\frac{1}{2} \right)^S \left(\frac{-1+3c}{1-c} \right).$$

If $K_{2S-1}^*(\lambda)$ is not weakly decreasing in λ , then there exists $\lambda' > \lambda$ such that $K_{2S-1}^*(\lambda') > K_{2S-1}^*(\lambda)$. By the construction and the monotonicity of $F_k(\lambda)$, we can find that

$$F_{K_{2S-1}^*(\lambda)-1}(\lambda') > F_{K_{2S-1}^*(\lambda)-1}(\lambda) > \left(\frac{1}{2} \right)^S \left(\frac{-1+3c}{1-c} \right).$$

Also, $K_{2S-1}^*(\lambda')$ is the lowest level such that

$$F_{K_{2S-1}^*(\lambda')-1}(\lambda') > \left(\frac{1}{2} \right)^S \left(\frac{-1+3c}{1-c} \right),$$

implying that

$$F_{K_{2S-1}^*(\lambda)-1}(\lambda') > F_{K_{2S-1}^*(\lambda')-1}(\lambda') > \left(\frac{1}{2}\right)^S \left(\frac{-1+3c}{1-c}\right).$$

This contradicts the assumption that $K_{2S-1}^*(\lambda') > K_{2S-1}^*(\lambda)$. ■

Appendix D

SUPPLEMENTARY INFORMATION FOR CHAPTER 4

D.1 Proofs for General Properties

Proof of Lemma 4.1

The uniqueness of DCH can be proven by induction on levels. That is, it suffices to prove that for every level $k \geq 1$, the optimal behavior strategy profile is unique. Consider any player $i \in N$ any type $\theta_i \in \Theta_i$. Level 1 type θ_i player i believes all other players are level 0, and will uniformly randomize at every history. Therefore, at every history h^t , level 1 type θ_i player i 's DCH belief is $\mu_i^1((0, \dots, 0), \theta_{-i} | h^t, \theta_i) = F(\theta_{-i} | \theta_i)$. By one-deviation principle, level 1 type θ_i player i 's best response satisfies that for any history h^t and action $a \in A_i(h^t)$, $\sigma_i^1(a|h^t, \theta_i) > 0$ if and only if

$$a \in \operatorname{argmax}_{a' \in A_i(h^t)} \mathbb{E}u_i^1((\sigma_{-i}^0, \bar{\sigma}_i^1(a')) | h^t, \theta_i),$$

where $\bar{\sigma}_i^1$ agrees with σ_i^1 except at (θ_i, h^t) where $\bar{\sigma}_i^1(a')$ chooses a' with probability 1. Since players are assumed to uniformly randomize over optimal actions when they are indifferent, σ_i^1 is uniquely pinned down.

Suppose there is $K > 2$ such that the optimal strategy profiles for level 1 to $K - 1$ are unique. In this case, level K player i 's conjecture about other players' behavior strategy profile $\tilde{\sigma}_{-i}^{-K}$ is also unique and totally mixed. By one-deviation principle, level K type θ_i player i 's best response satisfies that for any history h^t and $a \in A_i(h^t)$, $\sigma_i^K(a|h^t, \theta_i) > 0$ if and only if

$$a \in \operatorname{argmax}_{a' \in A_i(h^t)} \mathbb{E}u_i^K((\tilde{\sigma}_{-i}^{-K}, \bar{\sigma}_i^K(a')) | h^t, \theta_i).$$

Since players are assumed to uniformly randomize over optimal actions when they are indifferent, σ_i^K is again uniquely pinned down. This completes the proof. ■

Proof of Proposition 4.1

To prove this proposition, I first characterize the posterior beliefs in Lemma D.1 then prove that the beliefs are independent across players if the types are independently drawn.

Lemma D.1. Consider any multi-stage game with observed actions Γ , any $i \in N$, $\theta_i \in \Theta_i$, $h \in \mathcal{H} \setminus \mathcal{H}^T$, and every level $k \in \mathbb{N}$. For every information set $\mathcal{I}_i = (h, \theta_i)$, level k player i 's belief at \mathcal{I}_i can be characterized as follows.

1. Level k player i 's prior belief about other players' types and levels are independent. That is, $\mu_i^k(\tau_{-i}, \theta_{-i} | h_\emptyset, \theta_i) = F(\theta_{-i} | \theta_i) \prod_{j \neq i} \hat{P}_{ij}^k(\tau_j)$.
2. For any $1 \leq t < T$, and $h^t \in \mathcal{H}^t$, level k player i 's belief at information set (h^t, θ_i) where $h^t = (a^1, \dots, a^t)$ is

$$\mu_i^k(\tau_{-i}, \theta_{-i} | h^t, \theta_i) = \frac{F(\theta_{-i} | \theta_i) \prod_{j \neq i} \left\{ \hat{P}_{ij}^k(\tau_j) \prod_{l=1}^t \sigma_j^{\tau_j}(a_j^l | h^{l-1}, \theta_j) \right\}}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < k \forall j \neq i\}} F(\theta'_{-i} | \theta_i) \prod_{j \neq i} \left\{ \hat{P}_{ij}^k(\tau'_j) \prod_{l=1}^t \sigma_j^{\tau'_j}(a_j^l | h^{l-1}, \theta'_j) \right\}}.$$

Proof of Lemma D.1:

1. At the beginning of the game, the only information available to player i is his own type θ_i and his level of sophistication $\tau_i = k$. Therefore, the prior belief is the probability of the opponents' types and levels conditional on θ_i and τ_i , which is

$$\begin{aligned} \mu_i^k(\tau_{-i}, \theta_{-i} | h_\emptyset, \theta_i) &= \Pr(\tau_{-i}, \theta_{-i} | \tau_i = k, \theta_i) \\ &= \Pr(\theta_{-i} | \theta_i) \Pr(\tau_{-i} | \tau_i = k) \\ &= F(\theta_{-i} | \theta_i) \prod_{j \neq i} \hat{P}_{ij}^k(\tau_j). \end{aligned}$$

The second equality holds because the types and levels are independently drawn.

2. This can be shown by induction on t . Consider any available history at period 2, $h^1 \in \mathcal{H}^1$. Level k player i 's belief at information set (h^1, θ_i) is

$$\mu_i^k(\tau_{-i}, \theta_{-i} | h^1, \theta_i) = \frac{\mu_i^k(\tau_{-i}, \theta_{-i} | h_\emptyset, \theta_i) \prod_{j \neq i} \sigma_j^{\tau_j}(a_j^1 | h_\emptyset, \theta_j)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < k \forall j \neq i\}} \mu_i^k(\tau'_{-i}, \theta'_{-i} | h_\emptyset, \theta_i) \prod_{j \neq i} \sigma_j^{\tau'_j}(a_j^1 | h_\emptyset, \theta'_j)}. \quad (\text{D.1})$$

By step 1, we can obtain that $\mu_i^k(\tau_{-i}, \theta_{-i} | h_\emptyset, \theta_i) = F(\theta_{-i} | \theta_i) \prod_{j \neq i} \hat{P}_{ij}^k(\tau_j)$. Plugging in Equation (D.1), we can further obtain that

$$\begin{aligned} \mu_i^k(\tau_{-i}, \theta_{-i} | h^1, \theta_i) &= \frac{\mu_i^k(\tau_{-i}, \theta_{-i} | h_\emptyset, \theta_i) \prod_{j \neq i} \sigma_j^{\tau_j}(a_j^1 | h_\emptyset, \theta_j)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < k \forall j \neq i\}} \mu_i^k(\tau'_{-i}, \theta'_{-i} | h_\emptyset, \theta_i) \prod_{j \neq i} \sigma_j^{\tau'_j}(a_j^1 | h_\emptyset, \theta'_j)} \\ &= \frac{F(\theta_{-i} | \theta_i) \prod_{j \neq i} \left\{ \hat{P}_{ij}^k(\tau_j) \sigma_j^{\tau_j}(a_j^1 | h_\emptyset, \theta_j) \right\}}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < k \forall j \neq i\}} F(\theta'_{-i} | \theta_i) \prod_{j \neq i} \left\{ \hat{P}_{ij}^k(\tau'_j) \sigma_j^{\tau'_j}(a_j^1 | h_\emptyset, \theta'_j) \right\}}. \end{aligned}$$

Next, suppose there is t' such that the statement holds for every period $t = 2, \dots, t'$. Consider period $t' + 1$ and any history available at period $t' + 1$, $h^{t'} \in \mathcal{H}^{t'}$ where $h^{t'} = (a^1, \dots, a^{t'})$. Then level k player i 's belief at information set $(h^{t'}, \theta_i)$ is

$$\begin{aligned} \mu_i^k(\tau_{-i}, \theta_{-i} | h^{t'}, \theta_i) &= \frac{\mu_i^k(\tau_{-i}, \theta_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau_j}(a_j^{t'} | h^{t'-1}, \theta_j)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < k \forall j \neq i\}} \mu_i^k(\tau'_{-i}, \theta'_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau'_j}(a_j^{t'} | h^{t'-1}, \theta'_j)} \\ &= \frac{F(\theta_{-i} | \theta_i) \prod_{j \neq i} \left\{ \hat{P}_{ij}^k(\tau_j) \prod_{l=1}^{t'-1} \sigma_j^{\tau_j}(a_j^l | h^{l-1}, \theta_j) \right\} \prod_{j \neq i} \sigma_j^{\tau_j}(a_j^{t'} | h^{t'-1}, \theta_j)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < k \forall j \neq i\}} F(\theta'_{-i} | \theta_i) \prod_{j \neq i} \left\{ \hat{P}_{ij}^k(\tau'_j) \prod_{l=1}^{t'-1} \sigma_j^{\tau'_j}(a_j^l | h^{l-1}, \theta'_j) \right\} \prod_{j \neq i} \sigma_j^{\tau'_j}(a_j^{t'} | h^{t'-1}, \theta'_j)} \\ &= \frac{F(\theta_{-i} | \theta_i) \prod_{j \neq i} \left\{ \hat{P}_{ij}^k(\tau_j) \prod_{l=1}^{t'} \sigma_j^{\tau_j}(a_j^l | h^{l-1}, \theta_j) \right\}}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < k \forall j \neq i\}} F(\theta'_{-i} | \theta_i) \prod_{j \neq i} \left\{ \hat{P}_{ij}^k(\tau'_j) \prod_{l=1}^{t'} \sigma_j^{\tau'_j}(a_j^l | h^{l-1}, \theta'_j) \right\}}. \end{aligned}$$

The second equality holds because of the induction hypothesis, as desired. \square

Proof of Proposition 4.1:

We prove this by induction on t . Let σ be any level-dependent strategy profile and F and P be any distributions of types and levels. First, consider $t = 1$. By Lemma D.1, we know $\mu_i^k(\tau_{-i}, \theta_{-i} | h_0, \theta_i) = F(\theta_{-i} | \theta_i) \prod_{j \neq i} \hat{P}_{ij}^k(\tau_j)$. As the prior distribution of types is independent across players, we can obtain that

$$\begin{aligned} \mu_i^k(\tau_{-i}, \theta_{-i} | h_0, \theta_i) &= F(\theta_{-i} | \theta_i) \prod_{j \neq i} \hat{P}_{ij}^k(\tau_j) \\ &= \prod_{j \neq i} F_j(\theta_j) \prod_{j \neq i} \hat{P}_{ij}^k(\tau_j) \\ &= \prod_{j \neq i} \left[F_j(\theta_j) \hat{P}_{ij}^k(\tau_j) \right] = \prod_{j \neq i} \mu_{ij}^k(\tau_j, \theta_j | h_0, \theta_i). \end{aligned}$$

Therefore, we know the result is true at $t = 1$. Next, suppose there is $t' > 1$ such that the result holds for all $t = 1, \dots, t'$. We want to show that the result holds at period $t' + 1$. Let $h^{t'} \in \mathcal{H}^{t'}$ be any available history in period $t' + 1$ where $h^{t'} = (h^{t'-1}, a^{t'})$. Therefore, player i 's posterior belief at history $h^{t'}$ is

$$\mu_i^k(\tau_{-i}, \theta_{-i} | h^{t'}, \theta_i) = \frac{\mu_i^k(\tau_{-i}, \theta_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau_j}(a_j^{t'} | h^{t'-1}, \theta_j)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < k \forall j \neq i\}} \mu_i^k(\tau'_{-i}, \theta'_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau'_j}(a_j^{t'} | h^{t'-1}, \theta'_j)}.$$

By induction hypothesis, we know

$$\mu_i^k(\tau_{-i}, \theta_{-i} | h^{t'-1}, \theta_i) = \prod_{j \neq i} \mu_{ij}^k(\tau_j, \theta_j | h^{t'-1}, \theta_i).$$

Therefore, as we rearrange the posterior belief $\mu_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i)$, we can obtain that

$$\begin{aligned}
\mu_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i) &= \frac{\mu_i^k(\tau_{-i}, \theta_{-i}|h^{t-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau_j}(a_j^t|h^{t-1}, \theta_j)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}; \tau'_j < k \forall j \neq i\}} \mu_i^k(\tau'_{-i}, \theta'_{-i}|h^{t-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau'_j}(a_j^t|h^{t-1}, \theta'_j)} \\
&= \frac{\prod_{j \neq i} \left[\mu_{ij}^k(\tau_j, \theta_j|h^{t-1}, \theta_i) \sigma_j^{\tau_j}(a_j^t|h^{t-1}, \theta_j) \right]}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}; \tau'_j < k \forall j \neq i\}} \prod_{j \neq i} \left[\mu_{ij}^k(\tau'_j, \theta'_j|h^{t-1}, \theta_i) \sigma_j^{\tau'_j}(a_j^t|h^{t-1}, \theta'_j) \right]} \\
&= \frac{\prod_{j \neq i} \left[\mu_{ij}^k(\tau_j, \theta_j|h^{t-1}, \theta_i) \sigma_j^{\tau_j}(a_j^t|h^{t-1}, \theta_j) \right]}{\sum_{\theta'_{-i} \in \Theta_{-i}} \prod_{j \neq i} \left[\sum_{\tau'_j < k} \mu_{ij}^k(\tau'_j, \theta'_j|h^{t-1}, \theta_i) \sigma_j^{\tau'_j}(a_j^t|h^{t-1}, \theta'_j) \right]} \\
&= \frac{\prod_{j \neq i} \left[\mu_{ij}^k(\tau_j, \theta_j|h^{t-1}, \theta_i) \sigma_j^{\tau_j}(a_j^t|h^{t-1}, \theta_j) \right]}{\prod_{j \neq i} \left[\sum_{\theta'_j \in \Theta_j} \sum_{\tau'_j < k} \mu_{ij}^k(\tau'_j, \theta'_j|h^{t-1}, \theta_i) \sigma_j^{\tau'_j}(a_j^t|h^{t-1}, \theta'_j) \right]}.
\end{aligned}$$

As a result, we can conclude that

$$\begin{aligned}
\mu_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i) &= \prod_{j \neq i} \left[\frac{\mu_{ij}^k(\tau_j, \theta_j|h^{t-1}, \theta_i) \sigma_j^{\tau_j}(a_j^t|h^{t-1}, \theta_j)}{\sum_{\theta'_j \in \Theta_j} \sum_{\tau'_j < k} \mu_{ij}^k(\tau'_j, \theta'_j|h^{t-1}, \theta_i) \sigma_j^{\tau'_j}(a_j^t|h^{t-1}, \theta'_j)} \right] \\
&= \prod_{j \neq i} \mu_{ij}^k(\tau_j, \theta_j|h^{t-1}, \theta_i).
\end{aligned}$$

This completes the proof of the proposition. ■

Proof of Proposition 4.2

By Lemma D.1, we know that in the transformed game (with independent types) $\hat{\Gamma}$, level k player i 's belief at $h^t \in \mathcal{H}^t$ is

$$\begin{aligned}
\hat{\mu}_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i) &= \frac{\hat{F}(\theta_{-i}|\theta_i) \prod_{j \neq i} \left\{ \hat{\rho}_{ij}^k(\tau_j) \prod_{l=1}^t \sigma_j^{\tau_j}(a_j^l|h^{l-1}, \theta_j) \right\}}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}; \tau'_j < k \forall j \neq i\}} \hat{F}(\theta'_{-i}|\theta_i) \prod_{j \neq i} \left\{ \hat{\rho}_{ij}^k(\tau'_j) \prod_{l=1}^t \sigma_j^{\tau'_j}(a_j^l|h^{l-1}, \theta'_j) \right\}} \\
&= \frac{\prod_{j \neq i} \left\{ \hat{\rho}_{ij}^k(\tau_j) \prod_{l=1}^t \sigma_j^{\tau_j}(a_j^l|h^{l-1}, \theta_j) \right\}}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}; \tau'_j < k \forall j \neq i\}} \prod_{j \neq i} \left\{ \hat{\rho}_{ij}^k(\tau'_j) \prod_{l=1}^t \sigma_j^{\tau'_j}(a_j^l|h^{l-1}, \theta'_j) \right\}}.
\end{aligned}$$

Therefore, we can obtain that

$$\begin{aligned}
\mu_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i) &= \frac{F(\theta_{-i}|\theta_i) \prod_{j \neq i} \left\{ \hat{\rho}_{ij}^k(\tau_j) \prod_{l=1}^t \sigma_j^{\tau_j}(a_j^l|h^{l-1}, \theta_j) \right\}}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}; \tau'_j < k \forall j \neq i\}} F(\theta'_{-i}|\theta_i) \prod_{j \neq i} \left\{ \hat{\rho}_{ij}^k(\tau'_j) \prod_{l=1}^t \sigma_j^{\tau'_j}(a_j^l|h^{l-1}, \theta'_j) \right\}} \\
&= \frac{F(\theta_{-i}|\theta_i) \hat{\mu}_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}; \tau'_j < k \forall j \neq i\}} F(\theta'_{-i}|\theta_i) \hat{\mu}_i^k(\tau'_{-i}, \theta'_{-i}|h^t, \theta_i)}.
\end{aligned}$$

To complete the proof, it suffices to show that for each level k player i and every $h^t \in \mathcal{H} \setminus \mathcal{H}^T$, maximizing $\mathbb{E}u_i^k$ given belief μ_i^k and $\sigma_{-i}^{<k}$ is equivalent to maximizing $\mathbb{E}\hat{u}_i^k$ given belief $\hat{\mu}_i^k$ and $\hat{\sigma}_{-i}^{<k} = \sigma_{-i}^{<k}$. This is true because the expected payoff in the original game (with correlated types) is:

$$\mathbb{E}u_i^k(\sigma|h^t, \theta_i) = \sum_{h^T \in \mathcal{H}^T} \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{\{\tau_{-i}: \tau_j < k \forall j \neq i\}} \mu_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i) P_i^k(h^T|h^t, \theta, \tau_{-i}, \sigma_{-i}^{<k}, \sigma_i^k) u_i(h^T, \theta_{-i}, \theta_i),$$

which is proportional to

$$\begin{aligned} \mathbb{E}\hat{u}_i^k(\sigma|h^t, \theta_i) &= \sum_{h^T \in \mathcal{H}^T} \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{\{\tau_{-i}: \tau_j < k \forall j \neq i\}} F(\theta_{-i}|\theta_i) \hat{\mu}_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i) P_i^k(h^T|h^t, \theta, \tau_{-i}, \sigma_{-i}^{<k}, \sigma_i^k) u_i(h^T, \theta_{-i}, \theta_i) \\ &= \sum_{h^T \in \mathcal{H}^T} \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{\{\tau_{-i}: \tau_j < k \forall j \neq i\}} \hat{\mu}_i^k(\tau_{-i}, \theta_{-i}|h^t, \theta_i) P_i^k(h^T|h^t, \theta, \tau_{-i}, \sigma_{-i}^{<k}, \sigma_i^k) \hat{u}_i(h^T, \theta_{-i}, \theta_i). \end{aligned}$$

This completes the proof of the proposition. ■

Proof of Proposition 4.3

Proof of statement 1:

Consider any player $i \in N$, any level τ_i , any type θ_i and any non-terminal history $h^t = (h^{t-1}, a^t) \in \mathcal{H}^t \setminus \mathcal{H}^T$. To prove the statement, it suffices to show that for any $\tilde{\tau}_{-i}$, if $\tilde{\tau}_{-i} \notin \text{supp}_i(\tau_{-i}|h^{t-1}, \theta_i, \tau_i)$, then $\tilde{\tau}_{-i} \notin \text{supp}_i(\tau_{-i}|h^t, \theta_i, \tau_i)$.

If $\tilde{\tau}_{-i} \notin \text{supp}_i(\tau_{-i}|h^{t-1}, \theta_i, \tau_i)$, then $\mu_i^{\tau_i}(\tilde{\tau}_{-i}, \theta_{-i}|h^{t-1}, \theta_i) = 0$ for any θ_{-i} . By Lemma D.1, we can find that for any θ_{-i} ,

$$\mu_i^{\tau_i}(\tilde{\tau}_{-i}, \theta_{-i}|h^t, \theta_i) = \frac{\mu_i^{\tau_i}(\tilde{\tau}_{-i}, \theta_{-i}|h^{t-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tilde{\tau}_j}(a_j^t|h^{t-1}, \theta_j)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < k \forall j \neq i\}} \mu_i^{\tau_i}(\tau'_{-i}, \theta'_{-i}|h^{t-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau'_j}(a_j^t|h^{t-1}, \theta'_j)} = 0,$$

implying that $\mu_i^{\tau_i}(\tilde{\tau}_{-i}|\theta_i, h^t) = 0$ and hence $\tilde{\tau}_{-i} \notin \text{supp}_i(\tau_{-i}|\tau_i, \theta_i, h^t)$. □

Proof of statement 2:

The second statement can be proven by induction on t . First, consider $t = 1$. For any $i \in N$, $\tau_i \in \mathbb{N}$ and $\theta_i \in \Theta_i$, by Lemma D.1, we know the belief about other players' types and levels is $\mu_i^{\tau_i}(\tau_{-i}, \theta_{-i}|h_0, \theta_i) = F(\theta_{-i}|\theta_i) \prod_{j \neq i} \hat{P}_{ij}^{\tau_i}(\tau_j)$. Since F has full support, for any $\theta_{-i} \in \Theta_{-i}$,

$$\sum_{\{\tau_{-i}: \tau_j < \tau_i \forall j \neq i\}} \mu_i^{\tau_i}(\tau_{-i}, \theta_{-i}|h_0, \theta_i) = \sum_{\{\tau_{-i}: \tau_j < \tau_i \forall j \neq i\}} F(\theta_{-i}|\theta_i) \prod_{j \neq i} \hat{P}_{ij}^{\tau_i}(\tau_j) = F(\theta_{-i}|\theta_i) > 0.$$

Hence, the statement is true at period 1.

Next, suppose there is $t' > 1$ such that the result holds for all $t = 1, \dots, t'$. We want to show the statement holds at period $t' + 1$. Let $h^{t'}$ be any available history at period $t' + 1$ where $h^{t'} = (h^{t'-1}, a^{t'})$. Therefore, player i 's posterior belief at $h^{t'}$ is

$$\mu_i^{\tau_i}(\tau_{-i}, \theta_{-i} | h^{t'}, \theta_i) = \frac{\mu_i^{\tau_i}(\tau_{-i}, \theta_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau_j}(a_j^{t'} | h^{t'-1}, \theta_j)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < \tau_i \forall j \neq i\}} \mu_i^{\tau_i}(\tau'_{-i}, \theta'_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau'_j}(a_j^{t'} | h^{t'-1}, \theta'_j)},$$

which is well-defined as level 0s are always in the support and $\sigma_j^0(a_j^{t'} | h^{t'-1}, \theta'_j) = \frac{1}{|A_j(h^{t'-1})|} > 0$ for all j . By induction hypothesis, we know $\text{supp}_i(\theta_{-i} | h^{t'-1}, \theta_i, \tau_i) = \Theta_{-i}$. Therefore, as we fix any $\theta_{-i} \in \Theta_{-i}$, we know $\mu_i^{\tau_i}((0, \dots, 0), \theta_{-i} | h^{t'-1}, \theta_i) > 0$, suggesting that $\theta_{-i} \in \text{supp}_i(\theta_{-i} | h^{t'}, \theta_i, \tau_i)$ because

$$\begin{aligned} \mu_i^{\tau_i}(\theta_{-i} | h^{t'}, \theta_i) &= \frac{\sum_{\{\tau_{-i}: \tau_j < \tau_i \forall j \neq i\}} \mu_i^{\tau_i}(\tau_{-i}, \theta_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau_j}(a_j^{t'} | h^{t'-1}, \theta_j)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < \tau_i \forall j \neq i\}} \mu_i^{\tau_i}(\tau'_{-i}, \theta'_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau'_j}(a_j^{t'} | h^{t'-1}, \theta'_j)} \\ &\geq \frac{\mu_i^{\tau_i}((0, \dots, 0), \theta_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \sigma_j^0(a_j^{t'} | h^{t'-1}, \theta_j)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < \tau_i \forall j \neq i\}} \mu_i^{\tau_i}(\tau'_{-i}, \theta'_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau'_j}(a_j^{t'} | h^{t'-1}, \theta'_j)} \\ &= \frac{\mu_i^{\tau_i}((0, \dots, 0), \theta_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \frac{1}{|A_j(h^{t'-1})|}}{\sum_{\theta'_{-i} \in \Theta_{-i}} \sum_{\{\tau'_{-i}: \tau'_j < \tau_i \forall j \neq i\}} \mu_i^{\tau_i}(\tau'_{-i}, \theta'_{-i} | h^{t'-1}, \theta_i) \prod_{j \neq i} \sigma_j^{\tau'_j}(a_j^{t'} | h^{t'-1}, \theta'_j)} > 0. \end{aligned}$$

This completes the proof of the proposition. ■

D.2 Proofs for Two-Person Dirty-Faces Games

Proof of Proposition 4.4

Step 1: Consider any $i \in N$. If $x_{-i} = O$, then player i knows his face is dirty immediately. Therefore, C is a dominant strategy, suggesting $\sigma_i^k(t, O) = 1$ for all $k \geq 1$ and $1 \leq t \leq T$. If $x_{-i} = X$, player i 's belief of having a dirty face at period 1 is p . Hence, the expected payoff of choosing C at period 1 is $p\alpha - (1-p) < 0$, implying $\sigma_i^k(1, X) = 0$ for all $k \geq 1$. Finally, since level 1 players believe the other player's actions don't convey any information about their own face types, the expected payoff of C at each period is $p\alpha - (1-p) < 0$, implying $\sigma_i^1(t, X) = 0$ for any $1 \leq t \leq T$.

Step 2: Consider any level $k \geq 2$, and period $2 \leq t \leq T$. This step characterizes the DCH posterior belief when $x_{-i} = X$. When the game proceeds to period t , the

posterior belief of $(x_i, \tau_{-i}) = (f, l)$ for any $f \in \{O, X\}$ and $0 \leq l \leq k-1$ is:

$$\mu_i^k(f, l|t, X) = \frac{\left[\prod_{t'=1}^{t-1} (1 - \sigma_{-i}^l(t', f)) \right] p_l \Pr(f)}{\sum_{x \in \{O, X\}} \sum_{j=0}^{k-1} \left[\prod_{t'=1}^{t-1} (1 - \sigma_{-i}^j(t', x)) \right] p_j \Pr(x)}. \quad (\text{D.2})$$

By step 1, since strategic players will claim immediately when seeing a clean face, $\sigma_{-i}^l(t', O) = 1$ for all $1 \leq t' \leq t-1$. Therefore, as the game proceeds to period $t \geq 2$, level k player i would know that *it is impossible for the other player to observe a dirty face and have a positive level of sophistication at the same time*. Furthermore, let $\mathcal{M}_i^k(t)$ be the support of level k player i 's marginal belief of τ_{-i} at period t . For any $0 \leq l \leq k-1$,

$$l \in \mathcal{M}_i^k(t) \iff \sum_{x_i \in \{O, X\}} \prod_{t'=1}^{t-1} (1 - \sigma_{-i}^l(t', x_i)) > 0,$$

and we let $\mathcal{M}_{i+}^k(t) \equiv \mathcal{M}_i^k(t) \setminus \{0\}$. Therefore, equation (D.2) implies that for any $t \geq 2$,

$$\mu_i^k(X, 0|t, X) = \frac{\left(\frac{1}{2}\right)^{t-1} p p_0}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j \in \mathcal{M}_{i+}^k(t)} p_j}, \quad \mu_i^k(O, 0|t, X) = \frac{\left(\frac{1}{2}\right)^{t-1} (1-p)p_0}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j \in \mathcal{M}_{i+}^k(t)} p_j}.$$

Moreover, for any $1 \leq k' \leq k-1$, $\mu_i^k(O, k'|t, X) = 0$, and for any $l \in \mathcal{M}_{i+}^k(t)$,

$$\mu_i^k(X, l|t, X) = \frac{p p_l}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j \in \mathcal{M}_{i+}^k(t)} p_j}.$$

Consequently, the marginal belief of having a dirty face at period $2 \leq t \leq T$ is

$$\mu_i^k(X|t, X) = \sum_{j=0}^{k-1} \mu_i^k(X, j|t, X) = \frac{p \left[\left(\frac{1}{2}\right)^{t-1} p_0 + \sum_{j \in \mathcal{M}_{i+}^k(t)} p_j \right]}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j \in \mathcal{M}_{i+}^k(t)} p_j}.$$

Thus, the expected payoff of choosing C at period t is $\delta^{t-1} [(1 + \alpha)\mu_i^k(X|t, X) - 1]$, which equals to $\mathbb{E}u_i^k(C|t, X) =$

$$\frac{\delta^{t-1}}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j \in \mathcal{M}_{i+}^k(t)} p_j} \left\{ p\alpha \left[\left(\frac{1}{2}\right)^{t-1} p_0 + \sum_{j \in \mathcal{M}_{i+}^k(t)} p_j \right] - (1-p) \left[\left(\frac{1}{2}\right)^{t-1} p_0 \right] \right\}. \quad (\text{D.3})$$

Finally, at period t , level k player i believes the other player will wait with probability

$$\frac{1}{2}\mu_i^k(0|t, X) + \sum_{j \in \mathcal{M}_{i+}^k(t+1)} \mu_i^k(j|t, X) = \frac{\left(\frac{1}{2}\right)^t p_0 + p \sum_{j \in \mathcal{M}_{i+}^k(t+1)} p_j}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j \in \mathcal{M}_{i+}^k(t)} p_j}. \quad (\text{D.4})$$

Step 3: This step proves a monotonicity result: if $\sigma_i^k(t, X) = 1$, then $\sigma_i^{k+1}(t, X) = 1$ for any $k \geq 2$ and $2 \leq t \leq T$. The proof consists of two cases. First consider period T . Equation (D.3) implies $\sigma_i^k(T, X) = 1$ if and only if

$$\frac{\delta^{T-1}}{\left(\frac{1}{2}\right)^{T-1} p_0 + p \sum_{j \in \mathcal{M}_{i+}^k(T)} p_j} \left\{ p\alpha \left[\left(\frac{1}{2}\right)^{T-1} p_0 + \sum_{j \in \mathcal{M}_{i+}^k(T)} p_j \right] - (1-p) \left[\left(\frac{1}{2}\right)^{T-1} p_0 \right] \right\} \geq 0$$

$$\iff \bar{\alpha} \geq \frac{\left(\frac{1}{2}\right)^{T-1} p_0}{\left(\frac{1}{2}\right)^{T-1} p_0 + \sum_{j \in \mathcal{M}_{i+}^k(T)} p_j}.$$

Because $\mathcal{M}_i^k(T) \subseteq \mathcal{M}_i^{k+1}(T)$, it can be proven that

$$\bar{\alpha} \geq \frac{\left(\frac{1}{2}\right)^{T-1} p_0}{\left(\frac{1}{2}\right)^{T-1} p_0 + \sum_{j \in \mathcal{M}_{i+}^k(T)} p_j} \geq \frac{\left(\frac{1}{2}\right)^{T-1} p_0}{\left(\frac{1}{2}\right)^{T-1} p_0 + \sum_{j \in \mathcal{M}_{i+}^{k+1}(T)} p_j},$$

implying that if it is optimal for level k player i to claim at period T , it is also optimal for level $k+1$ player i to claim.

Second, consider any period $2 \leq t \leq T-1$. Note that if level k players would choose C at period t , $k \notin \mathcal{M}_{i+}^{k+1}(t+1)$ and hence $\mathcal{M}_{i+}^k(t') = \mathcal{M}_{i+}^{k+1}(t')$ for any $t+1 \leq t' \leq T$. Therefore, as the game proceeds beyond period t , level k and level $k+1$ players will have the same continuation value. Let $V_{\tilde{t}}^{\tilde{k}}$ be level \tilde{k} players' continuation value at period \tilde{t} . The observation implies $V_{t+1}^k = V_{t+1}^{k+1}$. Coupled with that $\mathcal{M}_{i+}^k(t) \subseteq \mathcal{M}_{i+}^{k+1}(t)$, level $k+1$ player i 's expected payoff of choosing W at period t satisfies

$$\frac{\left(\frac{1}{2}\right)^t p_0 + p \sum_{j \in \mathcal{M}_{i+}^{k+1}(t+1)} p_j}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j \in \mathcal{M}_{i+}^{k+1}(t)} p_j} V_{t+1}^{k+1} \leq \frac{\left(\frac{1}{2}\right)^t p_0 + p \sum_{j \in \mathcal{M}_{i+}^k(t+1)} p_j}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j \in \mathcal{M}_{i+}^k(t)} p_j} V_{t+1}^k,$$

where the RHS is level k player's expected payoff of choosing W at period t . The inequality shows level k player's expected payoff of choosing W is weakly higher

than level $k + 1$ player's expected payoff of choosing W . To complete the proof, it suffices to argue that level $k + 1$ player's expected payoff of C at period t is higher than level k player's expected payoff of C . This is true because $\mathcal{M}_{i+}^k(t) \subseteq \mathcal{M}_{i+}^{k+1}(t)$ implies $\mu_i^{k+1}(X|t, X) \geq \mu_i^k(X|t, X)$, and hence,

$$\delta^{t-1} [(1 + \alpha)\mu_i^{k+1}(X|t, X) - 1] \geq \delta^{t-1} [(1 + \alpha)\mu_i^k(X|t, X) - 1].$$

Step 4: The proposition is proven by induction on k . This step establishes the base case for level 2 players. By step 1, $\sigma_i^1(t, X) = 0$ for all $1 \leq t \leq T$, and hence $\mathcal{M}_{i+}^2(t) = \{1\}$ for all $1 \leq t \leq T$. Therefore, equation (D.3) suggests the expected payoff of C at period T is

$$\mathbb{E}u_i^2(C|T, X) = \frac{\delta^{T-1}}{\left(\frac{1}{2}\right)^{T-1} p_0 + pp_1} \left\{ p\alpha \left[\left(\frac{1}{2}\right)^{T-1} p_0 + p_1 \right] - (1-p) \left[\left(\frac{1}{2}\right)^{T-1} p_0 \right] \right\}.$$

Therefore, C is optimal at period T if and only if

$$\mathbb{E}u_i^2(C|T, X) \geq 0 \iff \bar{\alpha} \geq \frac{\left(\frac{1}{2}\right)^{T-1} p_0}{\left(\frac{1}{2}\right)^{T-1} p_0 + pp_1}.$$

For any period $2 \leq t \leq T - 1$, I first prove the direction of necessity. Equation (D.4) implies level 2 player i 's belief about the other player choosing W at period t is

$$\frac{1}{2}\mu_i^2(0|t, X) + \mu_i^2(1|t, X) = \frac{\left(\frac{1}{2}\right)^t p_0 + pp_1}{\left(\frac{1}{2}\right)^{t-1} p_0 + pp_1}.$$

Therefore, the expected payoff of W at period t is at least $\left[\frac{\left(\frac{1}{2}\right)^t p_0 + pp_1}{\left(\frac{1}{2}\right)^{t-1} p_0 + pp_1} \right] \mathbb{E}u_i^2(C|t + 1, X) =$

$$\frac{\delta^t}{\left(\frac{1}{2}\right)^{t-1} p_0 + pp_1} \left\{ p\alpha \left[\left(\frac{1}{2}\right)^t p_0 + p_1 \right] - (1-p) \left[\left(\frac{1}{2}\right)^t p_0 \right] \right\}.$$

Since W is always available, C is strictly dominated at period t for level 2 player i if

$$\begin{aligned} \mathbb{E}u_i^2(C|t, X) &< \left[\frac{\left(\frac{1}{2}\right)^t p_0 + pp_1}{\left(\frac{1}{2}\right)^{t-1} p_0 + pp_1} \right] \mathbb{E}u_i^2(C|t + 1, X) \\ &\iff \bar{\alpha} < \frac{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta \right] p_0 + (1 - \delta)p_1}. \end{aligned}$$

This proves the direction of necessity.

Second, the sufficiency is proven by induction on the periods. Namely, I show the sufficiency holds for any period $T - t'$ where $1 \leq t' \leq T - 2$. Consider the base case for period $T - 1$. Because

$$\bar{\alpha} \geq \frac{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0 + (1 - \delta)p_1} > \frac{\left(\frac{1}{2}\right)^{T-1} p_0}{\left(\frac{1}{2}\right)^{T-1} p_0 + p_1},$$

level 2 players will choose C at period T , it is optimal to choose C at period $T - 1$ if

$$\begin{aligned} \mathbb{E}u_i^2(C|T - 1, X) &\geq \left[\frac{\left(\frac{1}{2}\right)^{T-1} p_0 + p p_1}{\left(\frac{1}{2}\right)^{T-2} p_0 + p p_1} \right] \mathbb{E}u_i^2(C|T, X) \\ &\iff \bar{\alpha} \geq \frac{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0 + (1 - \delta)p_1}. \end{aligned}$$

Now, suppose there is $t' \leq T - 2$ such that the statement holds at any period $T - t$ where $1 \leq t \leq t' - 1$. It can be proven that the sufficiency also holds at period $T - t'$. Because

$$\bar{\alpha} \geq \frac{\left[\left(\frac{1}{2}\right)^{T-t'-1} - \left(\frac{1}{2}\right)^{T-t'} \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{T-t'-1} - \left(\frac{1}{2}\right)^{T-t'} \delta \right] p_0 + (1 - \delta)p_1} > \frac{\left[\left(\frac{1}{2}\right)^{T-t'} - \left(\frac{1}{2}\right)^{T-t'+1} \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{T-t'} - \left(\frac{1}{2}\right)^{T-t'+1} \delta \right] p_0 + (1 - \delta)p_1},$$

level 2 players will choose C at period $T - t' + 1$ by induction hypothesis. Therefore, it is optimal to choose C at period $T - t'$ if

$$\begin{aligned} \mathbb{E}u_i^2(C|T - t', X) &\geq \left[\frac{\left(\frac{1}{2}\right)^{T-t'} p_0 + p p_1}{\left(\frac{1}{2}\right)^{T-t'-1} p_0 + p p_1} \right] \mathbb{E}u_i^2(C|T - t' + 1, X) \\ &\iff \bar{\alpha} \geq \frac{\left[\left(\frac{1}{2}\right)^{T-t'-1} - \left(\frac{1}{2}\right)^{T-t'} \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{T-t'-1} - \left(\frac{1}{2}\right)^{T-t'} \delta \right] p_0 + (1 - \delta)p_1}. \end{aligned}$$

This completes the proof of sufficiency.

Step 5: Step 4 establishes the base case for $k = 2$. Now, suppose there is $K > 2$ such that the statement holds for all $2 \leq k \leq K$. It suffices to prove the statement holds for level $K + 1$ players. The proof for period T is straightforward. From step 3, we know if $\sigma_i^K(T, X) = 1$, then $\sigma_i^{K+1}(T, X) = 1$. Hence, the only case that needs to be considered is when

$$\bar{\alpha} < \frac{\left(\frac{1}{2}\right)^{T-1} p_0}{\left(\frac{1}{2}\right)^{T-1} p_0 + \sum_{j=1}^{K-1} p_j}.$$

By induction hypothesis, $\sigma_{-i}^l(t, X) = 0$ for all $1 \leq l \leq K$ and for all $1 \leq t \leq T$. Therefore, $\sigma_i^{K+1}(T, X) = 1$ if and only if $\mathbb{E}u_i^{K+1}(C|T, X) \geq 0$, which is equivalent to

$$\bar{\alpha} \geq \frac{\left(\frac{1}{2}\right)^{T-1} p_0}{\left(\frac{1}{2}\right)^{T-1} p_0 + \sum_{j=1}^K p_j}.$$

For any period $2 \leq t \leq T - 1$, I first prove the direction of necessity. If

$$\bar{\alpha} < \frac{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta\right] p_0}{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta\right] p_0 + (1 - \delta) \sum_{j=1}^K p_j},$$

then by induction hypothesis, $\sigma_{-i}^l(t', X) = 0$ for all $1 \leq l \leq K$ and $1 \leq t' \leq t$, implying that $\mathcal{M}_{i+}^{K+1}(t) = \{1, \dots, K\}$. Then equation (D.3) suggests that the expected payoff of C at period t is

$$\frac{\delta^{t-1}}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j=1}^K p_j} \left\{ p \alpha \left[\left(\frac{1}{2}\right)^{t-1} p_0 + \sum_{j=1}^K p_j \right] - (1 - p) \left[\left(\frac{1}{2}\right)^{t-1} p_0 \right] \right\}.$$

Furthermore, equation (D.4) suggests level $K + 1$ players believe the other player will wait at period t with probability

$$\frac{1}{2} \mu_i^{K+1}(0|t, X) + \sum_{j=1}^K \mu_i^{K+1}(j|t, X) = \frac{\left(\frac{1}{2}\right)^t p_0 + p \sum_{j=1}^K p_j}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j=1}^K p_j}.$$

Therefore, by similar calculation as in step 4, choosing C is strictly dominated if

$$\begin{aligned} & \frac{\delta^{t-1}}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j=1}^K p_j} \left\{ p\alpha \left[\left(\frac{1}{2}\right)^{t-1} p_0 + \sum_{j=1}^K p_j \right] - (1-p) \left[\left(\frac{1}{2}\right)^{t-1} p_0 \right] \right\} \\ & < \frac{\delta^t}{\left(\frac{1}{2}\right)^{t-1} p_0 + p \sum_{j=1}^K p_j} \left\{ p\alpha \left[\left(\frac{1}{2}\right)^t p_0 + \sum_{j=1}^K p_j \right] - (1-p) \left[\left(\frac{1}{2}\right)^t p_0 \right] \right\}, \end{aligned}$$

which is implied by

$$\bar{\alpha} < \frac{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta \right] p_0 + (1-\delta) \sum_{j=1}^K p_j}.$$

This proves the direction of necessity.

Second, the sufficiency is proven by induction on the periods. Namely, I show the sufficiency holds for any period $T - t'$ where $1 \leq t' \leq T - 2$. Consider the base case for period $T - 1$. By step 3, if $\sigma_i^K(T - 1, X) = 1$, then $\sigma_i^{K+1}(T - 1, X) = 1$.

Therefore, it suffices to consider the case where

$$\frac{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0 + (1-\delta) \sum_{j=1}^K p_j} \leq \bar{\alpha} < \frac{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0 + (1-\delta) \sum_{j=1}^{K-1} p_j}.$$

By induction hypothesis, $\sigma_{-i}^l(t, X) = 0$ for all $1 \leq t \leq T - 1$ and $1 \leq l \leq K$.

Moreover, $\sigma_i^{K+1}(T, X) = 1$ because

$$\frac{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0 + (1-\delta) \sum_{j=1}^K p_j} > \frac{\left(\frac{1}{2}\right)^{T-1} p_0}{\left(\frac{1}{2}\right)^{T-1} p_0 + \sum_{j=1}^K p_j}.$$

Therefore, by a similar calculation as in step 4, it can be proven that it is optimal for level $K + 1$ players to choose C at period $T - 1$ if

$$\begin{aligned} \mathbb{E}u_i^{K+1}(C|T - 1, X) & \geq \frac{\left(\frac{1}{2}\right)^{T-1} p_0 + p \sum_{j=1}^K p_j}{\left(\frac{1}{2}\right)^{T-2} p_0 + p \sum_{j=1}^K p_j} \mathbb{E}u_i^{K+1}(C|T, X) \\ & \iff \bar{\alpha} \geq \frac{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{T-2} - \left(\frac{1}{2}\right)^{T-1} \delta \right] p_0 + (1-\delta) \sum_{j=1}^K p_j}. \end{aligned}$$

Now, suppose there is $t' \leq T - 2$ such that the statement holds for any period $T - t$ where $1 \leq t \leq t' - 1$. It can be proven that the statement also holds at period $T - t'$. By step 3, if $\sigma_i^K(T - t', X) = 1$, then $\sigma_i^{K+1}(T - t', X) = 1$ and it suffices to consider the case:

$$\frac{\left[\left(\frac{1}{2} \right)^{T-t'-1} - \left(\frac{1}{2} \right)^{T-t'} \delta \right] p_0}{\left[\left(\frac{1}{2} \right)^{T-t'-1} - \left(\frac{1}{2} \right)^{T-t'} \delta \right] p_0 + (1 - \delta) \sum_{j=1}^K p_j} \leq \bar{\alpha}$$

$$< \frac{\left[\left(\frac{1}{2} \right)^{T-t'-1} - \left(\frac{1}{2} \right)^{T-t'} \delta \right] p_0}{\left[\left(\frac{1}{2} \right)^{T-t'-1} - \left(\frac{1}{2} \right)^{T-t'} \delta \right] p_0 + (1 - \delta) \sum_{j=1}^{K-1} p_j}.$$

By induction hypothesis, $\sigma_{-i}^l(t, X) = 0$ for all $1 \leq t \leq T - t'$ and $1 \leq l \leq K$, and $\sigma_i^{K+1}(T - t' + 1, X) = 1$. Therefore, by a similar calculation as in step 4, it can be proven that it is optimal for level $K + 1$ players to choose C at period $T - t'$ if

$$\bar{\alpha} \geq \frac{\left[\left(\frac{1}{2} \right)^{T-t'-1} - \left(\frac{1}{2} \right)^{T-t'} \delta \right] p_0}{\left[\left(\frac{1}{2} \right)^{T-t'-1} - \left(\frac{1}{2} \right)^{T-t'} \delta \right] p_0 + (1 - \delta) \sum_{j=1}^K p_j}.$$

This completes the proof of the proposition. ■

Proof of Corollary 4.1

By Proposition 4.4, we know $\sigma_i^k(t, O) = 1$ for all t and $k \geq 1$, and $\sigma_i^1(t, X) = 0$ for all t . Then by Definition 4.2, we can obtain that $\hat{\sigma}_i^k(O) = 1$ for all $k \geq 1$, and $\hat{\sigma}_i^1(X) = T + 1$. In addition, since $\sigma_i^k(1, X) = 0$ for all $k \geq 2$, $\hat{\sigma}_i^k(X) \neq 1$. Moreover, the DCH solution can be equivalently characterized by optimal stopping periods because for any $t \geq 2$ and $k \geq 2$,

$$\hat{\sigma}_i^k(X) = t \iff \sigma_i^k(t - 1, X) = 0 \quad \text{and} \quad \sigma_i^k(t, X) = 1,$$

$$\hat{\sigma}_i^k(X) = T + 1 \iff \sigma_i^k(t', X) = 0 \quad \text{for any} \quad 1 \leq t' \leq T.$$

Lastly, to show the monotonicity, it suffices to show that for any $k' > k \geq 2$ and any $2 \leq t \leq T$, if $\sigma_i^k(t, X) = 1$, then $\sigma_i^{k'}(t, X) = 1$. The discussion is separated into two cases. First, if $t = T$, then by Proposition 4.4, $\sigma_i^k(T, X) = 1$ suggests

$$\bar{\alpha} \geq \frac{\left(\frac{1}{2} \right)^{T-1} p_0}{\left(\frac{1}{2} \right)^{T-1} p_0 + \sum_{j=1}^{k-1} p_j} > \frac{\left(\frac{1}{2} \right)^{T-1} p_0}{\left(\frac{1}{2} \right)^{T-1} p_0 + \sum_{j=1}^{k'-1} p_j},$$

implying $\sigma_i^{k'}(T, X) = 1$. Second, if $2 \leq t \leq T - 1$, by Proposition 4.4, $\sigma_i^k(t, X) = 1$ suggests

$$\bar{\alpha} \geq \frac{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta \right] p_0 + (1 - \delta) \sum_{j=1}^{k-1} p_j} > \frac{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta \right] p_0}{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta \right] p_0 + (1 - \delta) \sum_{j=1}^{k'-1} p_j},$$

implying $\sigma_i^{k'}(t, X) = 1$. This completes the proof. ■

Proof of Proposition 4.5

Step 1: Consider any $i \in N$. If $x_{-i} = O$, player i knows his face is dirty immediately, suggesting 1 is a dominant strategy and $\tilde{\sigma}_i^k(O) = 1$ for any $k \geq 1$. If $x_{-i} = X$, the expected payoff of 1 is $p\alpha - (1 - p) < 0$, implying $\tilde{\sigma}_i^k(X) \geq 2$ for any $k \geq 1$. Moreover, level 1 players believe the other player is level 0, so when observing X , the expected payoff of $2 \leq j \leq T$ is

$$p \left[\frac{T+2-j}{T+1} \delta^{j-1} \alpha \right] - (1-p) \left[\frac{T+2-j}{T+1} \delta^{j-1} \right] = \delta^{j-1} \left(\frac{T+2-j}{T+1} \right) [p\alpha - (1-p)] < 0.$$

implying $\tilde{\sigma}_i^1(X) = T + 1$.

Step 2: This step proves for any $K > 1$, if $\tilde{\sigma}_i^{l+1}(X) \leq \tilde{\sigma}_i^l(X)$ for all $1 \leq l \leq K - 1$, then $\tilde{\sigma}_i^{K+1}(X) \leq \tilde{\sigma}_i^K(X)$. Note that if $\tilde{\sigma}_i^K(X) = T + 1$, then there is nothing to prove. Let $s^* \equiv \tilde{\sigma}_i^K(X)$ and focus on the case where $2 \leq s^* \leq T$. If $s^* = T$, then level $K + 1$ player's expected payoff of choosing T is

$$\begin{aligned} & \delta^{T-1} \left[p\alpha \left(\frac{2}{T+1} \frac{p_0}{\sum_{j=0}^K p_j} + \frac{\sum_{j=1}^K p_j}{\sum_{j=0}^K p_j} \right) - (1-p) \left(\frac{2}{T+1} \frac{p_0}{\sum_{j=0}^K p_j} \right) \right] \\ & > \delta^{T-1} \left[p\alpha \left(\frac{2}{T+1} \frac{p_0}{\sum_{j=0}^{K-1} p_j} + \frac{\sum_{j=1}^{K-1} p_j}{\sum_{j=0}^{K-1} p_j} \right) - (1-p) \left(\frac{2}{T+1} \frac{p_0}{\sum_{j=0}^{K-1} p_j} \right) \right] \geq 0. \end{aligned}$$

The last inequality holds because it is optimal for level K players to choose T . This suggests that $T + 1$ is dominated by T and hence $\tilde{\sigma}_i^{K+1}(X) \leq T = \tilde{\sigma}_i^K(X)$.

On the other hand, consider $2 \leq s^* \leq T - 1$. If level $K + 1$ player i chooses some s where $s^* < s < T + 1$ that yields a non-negative expected payoff, then choosing s is strictly suboptimal because

$$\begin{aligned}
& \delta^{s-1} \left[p\alpha \left(\frac{T+2-s}{T+1} \frac{p_0}{\sum_{j=0}^K p_j} + \frac{\sum_{j=1}^{K-1} p_j}{\sum_{j=0}^K p_j} \right) - (1-p) \left(\frac{T+2-s}{T+1} \frac{p_0}{\sum_{j=0}^K p_j} \right) \right] \\
& < \delta^{s-1} \left[p\alpha \left(\frac{T+2-s}{T+1} \frac{p_0}{\sum_{j=0}^{K-1} p_j} + \frac{\sum_{j=1}^{K-1} p_j}{\sum_{j=0}^{K-1} p_j} \right) - (1-p) \left(\frac{T+2-s}{T+1} \frac{p_0}{\sum_{j=0}^{K-1} p_j} \right) \right] \\
& \leq \delta^{s^*-1} \left[p\alpha \left(\frac{T+2-s^*}{T+1} \frac{p_0}{\sum_{j=0}^{K-1} p_j} + \frac{\sum_{j=1}^{K-1} p_j}{\sum_{j=0}^{K-1} p_j} \right) - (1-p) \left(\frac{T+2-s^*}{T+1} \frac{p_0}{\sum_{j=0}^{K-1} p_j} \right) \right] \\
& < \delta^{s^*-1} \left[p\alpha \left(\frac{T+2-s^*}{T+1} \frac{p_0}{\sum_{j=0}^K p_j} + \frac{\sum_{j=1}^K p_j}{\sum_{j=0}^K p_j} \right) - (1-p) \left(\frac{T+2-s^*}{T+1} \frac{p_0}{\sum_{j=0}^K p_j} \right) \right].
\end{aligned}$$

Note that the second inequality holds because s^* is level K player's optimal choice, and the RHS of the last inequality is level $K+1$ player's expected payoff of choosing s^* . These inequalities show that it is not optimal for level $K+1$ players to choose any $s > s^*$, suggesting that $\tilde{\sigma}_i^{K+1}(X) \leq \tilde{\sigma}_i^K(X)$.

Step 3: The proposition is proven by induction on k . This step establishes the base case for level 2 players. For any $2 \leq j \leq T$, the expected payoff of choosing j is $\mathbb{E}u_i^2(j|X) =$

$$p \left[\left(\frac{T+2-j}{T+1} \delta^{j-1} \alpha \right) \frac{p_0}{p_0+p_1} + \left(\delta^{j-1} \alpha \right) \frac{p_1}{p_0+p_1} \right] - (1-p) \left[\left(\frac{T+2-j}{T+1} \delta^{j-1} \right) \frac{p_0}{p_0+p_1} \right].$$

For level 2 players and any $2 \leq j \leq T-1$, let $\Delta_j^2 \equiv \mathbb{E}u_i^2(j|X) - \mathbb{E}u_i^2(j+1|X)$ be the difference of expected payoffs between j and $j+1$. That is,

$$\begin{aligned}
\Delta_j^2 = & \delta^{j-1} p\alpha \left[\left(\frac{T+2-j}{T+1} - \frac{T+1-j}{T+1} \delta \right) \frac{p_0}{p_0+p_1} + (1-\delta) \frac{p_1}{p_0+p_1} \right] \\
& - \delta^{j-1} (1-p) \left[\left(\frac{T+2-j}{T+1} - \frac{T+1-j}{T+1} \delta \right) \frac{p_0}{p_0+p_1} \right],
\end{aligned}$$

suggesting that j dominates $j+1$ if and only if

$$\Delta_j^2 \geq 0 \iff \bar{\alpha} \geq \frac{\left[\frac{T+2-j}{T+1} - \frac{T+1-j}{T+1} \delta \right] p_0}{\left[\frac{T+2-j}{T+1} - \frac{T+1-j}{T+1} \delta \right] p_0 + (1-\delta)p_1}.$$

Because the RHS is decreasing function in j , $\Delta_j^2 \geq 0$ implies $\Delta_{j+1}^2 \geq 0$. Moreover, since

$$\mathbb{E}u_i^2(j|X) \geq 0 \iff \bar{\alpha} \geq \frac{\frac{T+2-j}{T+1} p_0}{\frac{T+2-j}{T+1} p_0 + p_1},$$

$\Delta_j^2 \geq 0$ implies $\mathbb{E}u_i^2(j|X) \geq 0$ because

$$\bar{\alpha} \geq \frac{\left[\frac{T+2-j}{T+1} - \frac{T+1-j}{T+1} \delta \right] p_0}{\left[\frac{T+2-j}{T+1} - \frac{T+1-j}{T+1} \delta \right] p_0 + (1-\delta)p_1} > \frac{\frac{T+2-j}{T+1} p_0}{\frac{T+2-j}{T+1} p_0 + p_1}.$$

As a result, $\tilde{\sigma}_i^2(X) \leq T$ if and only if $\mathbb{E}u_i^2(T|X) \geq 0$, which is equivalent to

$$\bar{\alpha} \geq \frac{\frac{2}{T+1} p_0}{\frac{2}{T+1} p_0 + p_1},$$

and for any other $2 \leq t \leq T-1$, $\tilde{\sigma}_i^2(X) \leq t$ if and only if

$$\Delta_t^2 \geq 0 \iff \bar{\alpha} \geq \frac{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1} \delta \right] p_0}{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1} \delta \right] p_0 + (1-\delta)p_1}.$$

Step 4: Step 3 establishes the base case where $k = 2$. Now suppose there is $K > 2$ such that the statement holds for any $2 \leq k \leq K$. It suffices to prove that the statement also holds for level $K+1$ players. By step 1, $\tilde{\sigma}_i^{K+1}(X) \geq 2$. Besides, note that for any $1 \leq t \leq T$ and $1 \leq l \leq K$, if $\tilde{\sigma}_{-i}^l(X) > t$, then level $K+1$ player i 's expected payoff of choosing $2 \leq j \leq t+1$ is $\mathbb{E}u_i^{K+1}(j|X) =$

$$\delta^{j-1} \left[p\alpha \left(\frac{T+2-j}{T+1} \frac{p_0}{\sum_{j=0}^K p_j} + \frac{\sum_{j=1}^K p_j}{\sum_{j=0}^K p_j} \right) - (1-p) \left(\frac{T+2-j}{T+1} \frac{p_0}{\sum_{j=0}^K p_j} \right) \right].$$

Similar to step 3, we define $\Delta_{t'}^{K+1}$ for any $2 \leq t' \leq t$ where $\Delta_{t'}^{K+1}$ is the difference of expected payoff between choosing t' and $t'+1$. That is,

$$\begin{aligned} \Delta_{t'}^{K+1} \equiv & \delta^{t'-1} p\alpha \left[\left(\frac{T+2-t'}{T+1} - \frac{T+1-t'}{T+1} \delta \right) \frac{p_0}{\sum_{j=0}^K p_j} + (1-\delta) \frac{\sum_{j=1}^K p_j}{\sum_{j=0}^K p_j} \right] \\ & - \delta^{t'-1} (1-p) \left[\left(\frac{T+2-t'}{T+1} - \frac{T+1-t'}{T+1} \delta \right) \frac{p_0}{\sum_{j=0}^K p_j} \right]. \end{aligned}$$

By the same argument as in step 3, $\Delta_{t'}^{K+1} < 0$ implies $\Delta_{t'-1}^{K+1} < 0$. Therefore, if $\tilde{\sigma}_{-i}^l(X) > t$ for any $1 \leq l \leq K$, it is strictly dominated for level $K+1$ players to choose t' (and all strategies $s < t'$) where $2 \leq t' \leq t$ if

$$\bar{\alpha} < \frac{\left[\frac{T+2-t'}{T+1} - \frac{T+1-t'}{T+1} \delta \right] p_0}{\left[\frac{T+2-t'}{T+1} - \frac{T+1-t'}{T+1} \delta \right] p_0 + (1-\delta) \sum_{j=1}^K p_j}, \quad (\text{D.5})$$

and by a similar argument as in step 3, $\Delta_{t'}^{K+1} \geq 0$ implies $\mathbb{E}u_i^{K+1}(t'|X) \geq 0$.

The proof for period T is straightforward. The implication of the induction hypothesis is that $\tilde{\sigma}_i^{l+1}(X) \leq \tilde{\sigma}_i^l(X)$ for all $1 \leq l \leq K-1$. By step 2, $\tilde{\sigma}_i^{K+1}(X) \leq T$ if $\tilde{\sigma}_i^K(X) \leq T$. Thus, it suffices to consider the case where

$$\bar{\alpha} < \frac{\frac{2}{T+1}p_0}{\frac{2}{T+1}p_0 + \sum_{j=1}^{K-1} p_j}.$$

By induction hypothesis, $\tilde{\sigma}_i^l(X) = T+1$ for all $1 \leq l \leq K$, so $\tilde{\sigma}_i^{K+1}(X) \leq T$ if and only if

$$\mathbb{E}u_i^{K+1}(T|X) \geq 0 \iff \bar{\alpha} \geq \frac{\frac{2}{T+1}p_0}{\frac{2}{T+1}p_0 + \sum_{j=1}^K p_j}.$$

Next, consider any $2 \leq t \leq T-1$. By induction hypothesis and step 2, if $\tilde{\sigma}_i^K(X) \leq t$, then $\tilde{\sigma}_i^{K+1}(X) \leq t$. Hence, it suffices to complete the proof by considering

$$\bar{\alpha} < \frac{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1}\delta\right] p_0}{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1}\delta\right] p_0 + (1-\delta) \sum_{j=1}^{K-1} p_j}.$$

In this case, $t < \tilde{\sigma}_i^{l+1}(X) \leq \tilde{\sigma}_i^l(X)$ for all $1 \leq l \leq K-1$. Therefore, inequality (D.5) implies that $\tilde{\sigma}_i^{K+1}(X) \leq t$ if and only if

$$\bar{\alpha} \geq \frac{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1}\delta\right] p_0}{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1}\delta\right] p_0 + (1-\delta) \sum_{j=1}^K p_j}.$$

This completes the proof of this proposition. ■

Proof of Corollary 4.2

It suffices to prove the monotonicity by showing for all $k' > k \geq 2$, if $\tilde{\sigma}_i^k(X) \leq t$, then $\tilde{\sigma}_i^{k'}(X) \leq t$ for any $2 \leq t \leq T$. We can separate the analysis into two cases. First, if $t = T$, then by Proposition 4.5, $\tilde{\sigma}_i^k(X) \leq T$ suggests

$$\bar{\alpha} \geq \frac{\frac{2}{T+1}p_0}{\frac{2}{T+1}p_0 + \sum_{j=1}^{k-1} p_j} > \frac{\frac{2}{T+1}p_0}{\frac{2}{T+1}p_0 + \sum_{j=1}^{k'-1} p_j},$$

implying $\tilde{\sigma}_i^{k'}(X) \leq T$. Second, for any $2 \leq t \leq T-1$, by Proposition 4.5, $\tilde{\sigma}_i^k(X) \leq t$ suggests

$$\bar{\alpha} \geq \frac{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1}\delta\right] p_0}{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1}\delta\right] p_0 + (1-\delta) \sum_{j=1}^{k-1} p_j} > \frac{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1}\delta\right] p_0}{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1}\delta\right] p_0 + (1-\delta) \sum_{j=1}^{k'-1} p_j},$$

implying $\tilde{\sigma}_i^{k'}(X) \leq t$. This completes the proof. ■

Proof of Proposition 4.6

First, for any $k \geq 2$, it suffices to prove $\mathcal{S}_T^k \subset \mathcal{E}_T^k$ by showing if $\tilde{\sigma}_i^k(X) \leq T$, then $\hat{\sigma}_i^k(X) \leq T$. This is true because

$$\frac{\frac{2}{T+1}p_0}{\frac{2}{T+1}p_0 + \sum_{j=1}^{k-1} p_j} > \frac{\left(\frac{1}{2}\right)^{T-1} p_0}{\left(\frac{1}{2}\right)^{T-1} p_0 + \sum_{j=1}^{k-1} p_j}.$$

Similarly, for $2 \leq t \leq T - 1$, we can first observe that $\mathcal{S}_t^k \subset \mathcal{E}_t^k$ if and only if

$$\begin{aligned} \frac{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1}\delta\right] p_0}{\left[\frac{T+2-t}{T+1} - \frac{T+1-t}{T+1}\delta\right] p_0 + (1-\delta) \sum_{j=1}^{k-1} p_j} &\geq \frac{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta\right] p_0}{\left[\left(\frac{1}{2}\right)^{t-1} - \left(\frac{1}{2}\right)^t \delta\right] p_0 + (1-\delta) \sum_{j=1}^{k-1} p_j} \\ \iff \delta &\leq \frac{(2^t - 2)(T + 1) - (t - 1)2^t}{(2^t - 1)(T + 1) - t2^t} \equiv \bar{\delta}(T, t), \end{aligned} \tag{D.6}$$

where $\bar{\delta}(T, t) > 0$ because $(2^t - 2)(T + 1) - (t - 1)2^t \geq 2(T + 1) - 4 > 0$ and $(2^t - 1)(T + 1) - t2^t \geq 3(T + 1) - 8 > 0$. If $\bar{\delta}(T, t) > 1$, then the inequality holds for any $\delta \in (0, 1)$, and hence $\mathcal{S}_t^k \subset \mathcal{E}_t^k$. Otherwise, if $\bar{\delta}(T, t) < 1$, the inequality does not hold for all δ , implying there is no set inclusion relationship between \mathcal{S}_t^k and \mathcal{E}_t^k . In addition, inequality (D.6) suggests $\hat{\sigma}_i^k(X) < \tilde{\sigma}_i^k(X)$ if $\delta < \bar{\delta}(T, t)$ and $\hat{\sigma}_i^k(X) > \tilde{\sigma}_i^k(X)$ if $\delta > \bar{\delta}(T, t)$. Lastly, as we rearrange the inequality, we can obtain

$$\bar{\delta}(T, t) < 1 \iff \frac{(2^t - 2)(T + 1) - (t - 1)2^t}{(2^t - 1)(T + 1) - t2^t} < 1 \iff t < \frac{\ln(T + 1)}{\ln(2)}.$$

This completes the proof of this proposition. ■

Proof of Corollary 4.3

By Proposition 4.6, we know for any $k \geq 2$, there is no set inclusion relationship between \mathcal{S}_t^k and \mathcal{E}_t^k if $2 \leq t < [\ln(T + 1)/\ln(2)]$. When $T \rightarrow \infty$, this condition holds for any $t \geq 2$. Moreover, from Proposition 4.6, we can obtain that

$$\bar{\delta}^*(t) = \lim_{T \rightarrow \infty} \bar{\delta}(T, t) = \lim_{T \rightarrow \infty} \frac{(2^t - 2)(T + 1) - (t - 1)2^t}{(2^t - 1)(T + 1) - t2^t} = \frac{2^t - 2}{2^t - 1}.$$

This completes the proof. ■

Additional Result for Poisson-DCH

One feature of the Poisson-DCH model is that as $\tau \rightarrow \infty$, the aggregate choice frequencies converge to the equilibrium prediction. This provides a second interpretation for the parameter τ : the higher the value of τ , the closer the predictions

are to the equilibrium. It is worth noting that, as highlighted by Camerer, Ho, and Chong (2004), this convergence property does not hold for general classes of games.

For the sake of simplicity, I will prove the result for sequential two-person games. A similar argument holds for the simultaneous version. For any two-person dirty faces game, conditional on there is an announcement, there are two possible states: one dirty face or two dirty faces, which are denoted as $\Omega = \{OX, XX\}$. For each $\omega \in \Omega$, equilibrium predicts a deterministic terminal period. We use $F_\omega^*(t)$ to express the (degenerated) distribution of terminal periods at the equilibrium. The equilibrium predicts that players will choose C at period 1 when seeing O , and choose W at period 1 and C at period 2 when seeing X . Therefore, when $\omega = OX$, the game will end at period 1, and when $\omega = XX$, the game will end at period 2. In other words,

$$F_{OX}^*(t) = \begin{cases} 0 & \text{if } t < 1 \\ 1 & \text{if } t \geq 1, \end{cases} \quad \text{and} \quad F_{XX}^*(t) = \begin{cases} 0 & \text{if } t < 2 \\ 1 & \text{if } t \geq 2. \end{cases}$$

In contrast, given any $\tau > 0$ and $\omega \in \Omega$, the Poisson-DCH model predicts a non-degenerated distribution over all possible terminal periods. We use $F_\omega^D(t|\tau)$ to denote the distribution predicted by the Poisson-DCH. Proposition D.1 states that when $\tau \rightarrow \infty$, the max norm between $F_\omega^D(t|\tau)$ and $F_\omega^*(t)$ will converge to 0 for any $\omega \in \Omega$.

Proposition D.1. *Consider any sequential two-person dirty faces game. When the prior distribution of levels follows $\text{Poisson}(\tau)$, for any $\omega \in \Omega$,*

$$\lim_{\tau \rightarrow \infty} \|F_\omega^*(t) - F_\omega^D(t|\tau)\|_\infty = 0.$$

Proof: When $\omega = OX$, a strategic player that sees a clean face will choose C in period 1. Therefore,

$$F_{OX}^D(1|\tau) = 1 - \left(\frac{1}{2}e^{-\tau}\right) \left(1 - \frac{1}{2}e^{-\tau}\right).$$

To show $\|F_{OX}^*(t) - F_{OX}^D(t|\tau)\|_\infty \rightarrow 0$, it suffices to show $F_{OX}^D(1|\tau) \rightarrow 1$ as $\tau \rightarrow \infty$. This is true because

$$\lim_{\tau \rightarrow \infty} F_{OX}^D(1|\tau) = \lim_{\tau \rightarrow \infty} 1 - \left(\frac{1}{2}e^{-\tau}\right) \left(1 - \frac{1}{2}e^{-\tau}\right) = 1.$$

When $\omega = XX$, it suffices to prove the convergence by showing $F_{XX}^D(1|\tau) \rightarrow 0$ and $F_{XX}^D(2|\tau) \rightarrow 1$ as $\tau \rightarrow \infty$. Since every level $k \geq 1$ will choose W in period 1 when

seeing a dirty face, $F_{XX}^D(1|\tau) = 1 - [1 - (1/2)e^{-\tau}]^2$, implying that

$$\lim_{\tau \rightarrow \infty} F_{XX}^D(1|\tau) = \lim_{\tau \rightarrow \infty} 1 - \left[1 - \frac{1}{2}e^{-\tau}\right]^2 = 0.$$

Lastly, let $K^*(\tau)$ be the lowest level of players to choose C at period 2 when seeing a dirty face with the prior distribution of levels being $\text{Poisson}(\tau)$. By Proposition 4.4, $K^*(\tau)$ is weakly decreasing in τ , and $K^*(\tau) \rightarrow 2$ as $\tau \rightarrow \infty$. Hence,

$$F_{XX}^D(2|\tau) = 1 - \left[(1/4)e^{-\tau} + \sum_{j=1}^{K^*(\tau)-1} e^{-\tau} \tau^j / j! \right]^2,$$

suggesting the limit is

$$\lim_{\tau \rightarrow \infty} F_{XX}^D(2|\tau) = \lim_{\tau \rightarrow \infty} 1 - \left[\frac{1}{4}e^{-\tau} + \sum_{j=1}^{K^*(\tau)-1} \frac{e^{-\tau} \tau^j}{j!} \right]^2 = \lim_{\tau \rightarrow \infty} 1 - \left[\frac{1}{4}e^{-\tau} + \tau e^{-\tau} \right]^2 = 1. \blacksquare$$

D.3 Three-Person Three-Period Dirty-Faces Games

In this section, I characterize the DCH solutions for the sequential and simultaneous three-person three-period dirty-faces games. In a three-person dirty-faces game, each player i will be random assigned a face type type, either clean (O) or dirty (X). The face types are i.i.d. drawn from the distribution $p = \Pr(x_i = X) = 1 - \Pr(x_i = O)$ where $p > 0$. After the face types are drawn, each player i can observe the other two players' faces x_{-i} but not their own face. If there is at least one player having a dirty face, a public announcement is made, informing all players of this fact.

There are up to 3 periods. In each period, all three players simultaneously choose either to claim to have a dirty face (C) or wait (W) and the actions are revealed at the end of each period. Similar to the two-person games, the game will end after any period where some player chooses C or after period 3. A player's payoff depends on their own face types and their actions in the terminal period. If player i waits in the terminal period, his payoff is 0 regardless of his face type. If player i chooses C in the terminal period, say period t , then he will receive $\delta^{t-1}\alpha$ if his face is dirty but $-\delta^{t-1}$ if his face is clean. Following the analysis of two-person games, I will focus on the case where there is a public announcement. Moreover, the assumption that $0 < \bar{\alpha} \equiv \alpha p / (1 - p) < 1$ is maintained so it is strictly dominated to choose C in period 1 when seeing one or two dirty faces.

D.3.1 DCH Solution for the Sequential Games

In the sequential three-person three-period dirty-faces game, a behavioral strategy for player i is a mapping from the period and the observed face types ($x_{-i} \in \{OO, OX, XX\}$) to the probability of claiming to have a dirty face. The behavioral strategy is denoted by

$$\sigma_i : \{1, 2, 3\} \times \{OO, OX, XX\} \rightarrow [0, 1].$$

For the sake of simplicity, I assume that each player i 's level is i.i.d. drawn from the distribution $p = (p_k)_{k=0}^{\infty}$ where $p_k > 0$ for all k . Proposition D.2 characterizes the DCH solution for the sequential three-person three-period dirty-faces games.

Proposition D.2. *For any sequential three-person three-period dirty-faces game, the level-dependent strategy profile of the DCH solution satisfies that for any $i \in N$,*

1. $\sigma_i^k(t, OO) = 1$ for all $k \geq 1$ and $1 \leq t \leq 3$.
2. $\sigma_i^1(t, OX) = 0$ for any $1 \leq t \leq 3$. Moreover, for any $k \geq 2$,

- (1) $\sigma_i^k(1, OX) = 0$,

- (2) $\sigma_i^k(2, OX) = 1$ if and only if

$$\bar{\alpha} \geq \frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_k\delta\right)p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_k\delta\right)p_0 + (1 - \gamma_k\delta)\sum_{j=1}^{k-1}p_j}$$

where $\gamma_k \equiv \left[\frac{1}{4}p_0 + \sum_{j=1}^{k-1}p_j\right] / \left[\frac{1}{2}p_0 + \sum_{j=1}^{k-1}p_j\right]$,

- (3) $\sigma_i^k(3, OX) = 1$ if and only if

$$\bar{\alpha} \geq \frac{\frac{1}{4}p_0}{\frac{1}{4}p_0 + \sum_{j=1}^{k-1}p_j},$$

3. $\sigma_i^1(t, XX) = \sigma_i^2(t, XX) = 0$ for any $1 \leq t \leq 3$. Moreover, for any $k \geq 3$,

- (1) $\sigma_i^k(1, XX) = \sigma_i^k(2, XX) = 0$,

- (2) $\sigma_i^k(3, XX) = 1$ if and only if there exists $2 \leq l \leq k - 1$ such that $\sigma_i^l(2, OX) = 1$ with $L_k^* \equiv \min_j \left\{j < k : \sigma_i^j(2, OX) = 1\right\}$, and

$$\bar{\alpha} \geq \max \left\{ \frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_{L_k^*}\delta\right)p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_{L_k^*}\delta\right)p_0 + (1 - \gamma_{L_k^*}\delta)\sum_{j=1}^{L_k^*-1}p_j}, \left(\frac{\frac{1}{4}p_0 + \sum_{j=1}^{L_k^*-1}p_j}{\frac{1}{4}p_0 + \sum_{j=1}^{k-1}p_j}\right)^2 \right\}.$$

Proof: Step 1: Consider any $i \in N$. If $x_{-i} = OO$, then player i knows his face is dirty immediately. Therefore, C is a dominant strategy, suggesting $\sigma_i^k(t, OO) = 1$ for all $k \geq 1$ and $1 \leq t \leq 3$. If $x_{-i} = OX$, player i 's belief of having a dirty face at period 1 is p . Hence, the expected payoff of choosing C at period 1 is $p\alpha - (1-p) < 0$, implying $\sigma_i^k(1, OX) = 0$ for all $k \geq 1$. Similarly, if $x_{-i} = XX$, the beliefs of having a dirty face at period 1 and 2 are p , which suggests $\sigma_i^k(1, XX) = \sigma_i^k(2, XX) = 0$ for all $k \geq 1$.

In addition, level 1 players believe other players' actions don't convey any information about their own face types, so $\sigma_i^1(t, OX) = \sigma_i^1(t, XX) = 0$ for any $1 \leq t \leq 3$. Since level 1 players behave exactly the same when observing OX and XX , level 2 player i 's belief about having a dirty face at period 3 is still p when $x_{-i} = XX$, implying $\sigma_i^2(3, XX) = 0$.

Step 2: In this step, I prove that

$$\frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_k\delta\right) p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_k\delta\right) p_0 + (1 - \gamma_k\delta) \sum_{j=1}^{k-1} p_j}$$

is decreasing in k for all $k \geq 2$ where $\gamma_k \equiv \left[\frac{1}{4}p_0 + \sum_{j=1}^{k-1} p_j\right] / \left[\frac{1}{2}p_0 + \sum_{j=1}^{k-1} p_j\right]$. To prove this, it suffices to prove that for any $l \geq 2$,

$$\begin{aligned} & \frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_l\delta\right) p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_l\delta\right) p_0 + (1 - \gamma_l\delta) \sum_{j=1}^{l-1} p_j} \geq \frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_{l+1}\delta\right) p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_{l+1}\delta\right) p_0 + (1 - \gamma_{l+1}\delta) \sum_{j=1}^l p_j} \\ \iff & \left(-\frac{1}{4}\gamma_{l+1}\delta + \frac{1}{4}\gamma_l\delta\right) \sum_{j=1}^{l-1} p_j + (1 - \gamma_{l+1}\delta) \left(\frac{1}{2} - \frac{1}{4}\gamma_l\delta\right) p_l \geq 0 \end{aligned}$$

Notice that the LHS of the inequality is decreasing in δ since

$$\begin{aligned} & \frac{d}{d\delta} \left[\left(-\frac{1}{4}\gamma_{l+1}\delta + \frac{1}{4}\gamma_l\delta\right) \sum_{j=1}^{l-1} p_j + (1 - \gamma_{l+1}\delta) \left(\frac{1}{2} - \frac{1}{4}\gamma_l\delta\right) p_l \right] \\ & = \underbrace{\left(-\frac{1}{4}\gamma_{l+1} + \frac{1}{4}\gamma_l\right)}_{<0} \sum_{j=1}^{l-1} p_j + \underbrace{\left(-\frac{1}{2}\gamma_{l+1} - \frac{1}{4}\gamma_l + \frac{1}{2}\gamma_l\gamma_{l+1}\delta\right)}_{<-\frac{1}{2}\gamma_{l+1} - \frac{1}{4}\gamma_l + \frac{1}{2}\gamma_l\gamma_{l+1}} p_l < 0. \\ & \leq -\sqrt{\frac{1}{2}\gamma_l\gamma_{l+1} + \frac{1}{2}\gamma_l\gamma_{l+1}} < 0 \end{aligned}$$

Therefore, it suffices to prove that the inequality holds when $\delta = 1$. That is,

$$\left(-\frac{1}{4}\gamma_{l+1} + \frac{1}{4}\gamma_l\right) \sum_{j=1}^{l-1} p_j + (1 - \gamma_{l+1}) \left(\frac{1}{2} - \frac{1}{4}\gamma_l\right) p_l \geq 0,$$

which holds because the inequality is equivalent to

$$\frac{p_l}{\sum_{j=1}^{l-1} p_j} \geq \frac{\frac{1}{4}(\gamma_{l+1} - \gamma_l)}{(1 - \gamma_{l+1})\left(\frac{1}{2} - \frac{1}{4}\gamma_l\right)} = \frac{p_l}{\frac{3}{4}p_0 + \sum_{j=1}^{l-1} p_j}.$$

Step 3: In this step, I characterize level k player i 's behavior when $x_{-i} = OX$ for all $k \geq 2$ by induction on k . I first prove the base case where $k = 2$. At period 3, level 2 player i 's belief about having a dirty face is

$$\mu_i^2(X|3, OX) = \sum_{\tau_{-i}} \mu_i^2(X, \tau_{-i}|3, OX) = \frac{p\left(\frac{1}{4}p_0 + p_1\right)}{\frac{1}{4}p_0 + pp_1}.$$

Therefore, it is optimal to choose C at period 3 if and only if

$$\mu_i^2(X|3, OX)\alpha - (1 - \mu_i^2(X|3, OX)) \geq 0 \iff \bar{\alpha} \geq \frac{\frac{1}{4}p_0}{\frac{1}{4}p_0 + p_1}.$$

At period 2, level 2 player i 's belief about having a dirty face is

$$\mu_i^2(X|2, OX) = \sum_{\tau_{-i}} \mu_i^2(X, \tau_{-i}|2, OX) = \frac{p\left(\frac{1}{2}p_0 + p_1\right)}{\frac{1}{2}p_0 + pp_1},$$

and the belief about that the two other players wait at period 2 is

$$\frac{\left(\frac{1}{4}p_0 + p_1\right)\left(\frac{1}{4}p_0 + pp_1\right)}{\left(\frac{1}{2}p_0 + p_1\right)\left(\frac{1}{2}p_0 + pp_1\right)} \equiv \gamma_2 \left(\frac{\frac{1}{4}p_0 + pp_1}{\frac{1}{2}p_0 + pp_1}\right).$$

Conditional on reaching period 3, the payoff of waiting is 0, and the expected payoff of C is

$$\frac{\delta^2}{\frac{1}{4}p_0 + pp_1} \left[p\alpha \left(\frac{1}{4}p_0 + p_1\right) - (1 - p) \left(\frac{1}{4}p_0\right) \right].$$

Therefore, it is optimal to choose C at period 2 if and only if

$$\begin{aligned} & \frac{\delta}{\frac{1}{2}p_0 + pp_1} \left[p\alpha \left(\frac{1}{2}p_0 + p_1\right) - (1 - p) \left(\frac{1}{2}p_0\right) \right] \\ & \geq \max \left\{ \gamma_2 \left(\frac{\frac{1}{4}p_0 + pp_1}{\frac{1}{2}p_0 + pp_1}\right) \frac{\delta^2}{\frac{1}{4}p_0 + pp_1} \left[p\alpha \left(\frac{1}{4}p_0 + p_1\right) - (1 - p) \left(\frac{1}{4}p_0\right) \right], 0 \right\} \\ \iff \bar{\alpha} & \geq \max \left\{ \frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_2\delta\right)p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_2\delta\right)p_0 + (1 - \gamma_2\delta)p_1}, \frac{\frac{1}{2}p_0}{\frac{1}{2}p_0 + p_1} \right\}. \end{aligned}$$

Furthermore, because for any $\delta \in (0, 1)$,

$$\frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_2\delta\right)p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_2\delta\right)p_0 + (1 - \gamma_2\delta)p_1} > \frac{\frac{1}{2}p_0}{\frac{1}{2}p_0 + p_1},$$

it is optimal for level 2 players to claim at period 2 if and only if

$$\bar{\alpha} \geq \frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_2\delta\right)p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_2\delta\right)p_0 + (1 - \gamma_2\delta)p_1}.$$

This completes the proof for level 2 players.

Suppose there is $K > 2$ such that the statement holds for any level $2 \leq k \leq K$. I now prove the statement holds for level $K + 1$ players. By the same argument as in the proof of Proposition 4.4, level $K + 1$ players would choose C when it is already optimal for level K players to choose C . Therefore, for period 3, it suffices to consider the case where

$$\bar{\alpha} < \frac{\frac{1}{4}p_0}{\frac{1}{4}p_0 + \sum_{j=1}^{K-1} p_j}.$$

By induction hypothesis, we know for every level $1 \leq k \leq K$ players, they will wait for three periods when observing one dirty face. Therefore, level $K + 1$ player i 's belief about having a dirty face at period 3 when $x_{-i} = OX$ is

$$\mu_i^{K+1}(X|3, OX) = \sum_{\tau_i} \mu_i^{K+1}(X, \tau_i|3, OX) = \frac{p \left(\frac{1}{4}p_0 + \sum_{j=1}^K p_j\right)}{\frac{1}{4}p_0 + p \sum_{j=1}^K p_j}.$$

Consequently, level $K + 1$ players would choose C at period 3 if and only if

$$\mu_i^{K+1}(X|3, OX)\alpha - (1 - \mu_i^{K+1}(X|3, OX)) \geq 0 \iff \bar{\alpha} \geq \frac{\frac{1}{4}p_0}{\frac{1}{4}p_0 + \sum_{j=1}^K p_j}.$$

For period 2, by step 2 and the induction hypothesis, it suffices to consider

$$\bar{\alpha} < \frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_{K+1}\delta\right)p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_{K+1}\delta\right)p_0 + (1 - \gamma_{K+1}\delta)\sum_{j=1}^K p_j};$$

otherwise, level K players would choose C at period 2 and so do level $K + 1$ players.

By similar argument, level $K + 1$ player i will choose C at period 2 if and only if

$$\frac{\delta}{\frac{1}{2}p_0 + p \sum_{j=1}^K p_j} \left[p\alpha \left(\frac{1}{2}p_0 + \sum_{j=1}^K p_j \right) - (1 - p) \left(\frac{1}{2}p_0 \right) \right] \\ \geq \max \left\{ \gamma_{K+1} \left(\frac{\delta^2}{\frac{1}{2}p_0 + p \sum_{j=1}^K p_j} \right) \left[p\alpha \left(\frac{1}{4}p_0 + \sum_{j=1}^K p_j \right) - (1 - p) \left(\frac{1}{4}p_0 \right) \right], 0 \right\},$$

which is equivalent to

$$\bar{\alpha} \geq \frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_{K+1}\delta\right) p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_{K+1}\delta\right) p_0 + (1 - \gamma_{K+1}\delta) \sum_{j=1}^K p_j}.$$

Step 4: This step characterizes level k player i 's behavior when $x_{-i} = XX$ for all $k \geq 3$. Consider any level $k \geq 3$. For level k players, they update their beliefs about having a dirty face at period 3 only if there is some lower level of players that chooses C at period 2 when observing one dirty face. That is, $\sigma_i^k(3, XX) = 1$ only if there is $2 \leq l \leq k - 1$ such that

$$\bar{\alpha} \geq \frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_l\delta\right) p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_l\delta\right) p_0 + (1 - \gamma_l\delta) \sum_{j=1}^{l-1} p_j}.$$

If there exists such level of players, let L_k^* denote the lowest level below k that would choose C at period 2 when observing one dirty face. In this case, level k player i 's belief about having a dirty face at period 3 is

$$\mu_i^k(X|3, XX) = \frac{p \left(\frac{1}{4}p_0 + \sum_{j=1}^{k-1} p_j\right)^2}{p \left(\frac{1}{4}p_0 + \sum_{j=1}^{k-1} p_j\right)^2 + (1 - p) \left(\frac{1}{4}p_0 + \sum_{j=1}^{L_k^*-1} p_j\right)^2},$$

and expected payoff of C is greater than 0 if and only if

$$\mu_i^k(X|3, XX)\alpha - (1 - \mu_i^k(X|3, XX)) \geq 0 \iff \bar{\alpha} \geq \left(\frac{\frac{1}{4}p_0 + \sum_{j=1}^{L_k^*-1} p_j}{\frac{1}{4}p_0 + \sum_{j=1}^{k-1} p_j}\right)^2.$$

Therefore, we can conclude that $\sigma_i^k(3, XX) = 1$ if and only if

$$\bar{\alpha} \geq \max \left\{ \frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_{L_k^*}\delta\right) p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_{L_k^*}\delta\right) p_0 + (1 - \gamma_{L_k^*}\delta) \sum_{j=1}^{L_k^*-1} p_j}, \left(\frac{\frac{1}{4}p_0 + \sum_{j=1}^{L_k^*-1} p_j}{\frac{1}{4}p_0 + \sum_{j=1}^{k-1} p_j}\right)^2 \right\}.$$

This completes the proof of step 4 and this proposition. ■

D.3.2 DCH Solution for Simultaneous Games

The strategically equivalent simultaneous three-person three-period dirty-faces game is a one-period game, in which all three players simultaneously choose an action from the set $S = \{1, 2, 3, 4\}$. Action $t \leq 3$ represents the plan to wait from period 1

to $t-1$ and claim in period t . Action 4 is the plan to always wait. In the simultaneous three-person three-period dirty-faces game, a mixed strategy is a mapping from the observed face type ($x_{-i} \in \{OO, OX, XX\}$) to a probability distribution over the action set. That is,

$$\tilde{\sigma}_i : \{OO, OX, XX\} \rightarrow \Delta(S).$$

Suppose (s_i, s_{-i}) is the realized action profile. If s_i is the smallest number, then the payoff for player i is computed as the case where player i claims to have a dirty face at period s_i ; otherwise, player i 's payoff is 0. The equilibrium analysis for the simultaneous game is the same as the sequential game. However, as characterized by Proposition D.3, the DCH solution for the simultaneous games differs from the DCH solution for the sequential games.

Proposition D.3. *For any simultaneous three-person three-period dirty-faces game, the level-dependent strategy profile of the DCH solution satisfies that for $i \in N$,*

1. $\tilde{\sigma}_i^k(OO) = 1$ for all $k \geq 1$.
2. $\tilde{\sigma}_i^1(OX) = 4$. Moreover, for any $k \geq 2$, $\tilde{\sigma}_i^k(OX) > 1$ and

(1) $\tilde{\sigma}_i^k(OX) = 2$ if and only if

$$\bar{\alpha} \geq \frac{\frac{3}{4}p_0 \left(\frac{3}{4}p_0 + \sum_{j=1}^{k-1} p_j \right) - \delta \left(\frac{1}{2}p_0 \right) \left(\frac{1}{2}p_0 + \sum_{j=1}^{k-1} p_j \right)}{\left(\frac{3}{4}p_0 + \sum_{j=1}^{k-1} p_j \right)^2 - \delta \left(\frac{1}{2}p_0 + \sum_{j=1}^{k-1} p_j \right)^2},$$

(2) $\tilde{\sigma}_i^k(OX) \leq 3$ if and only if

$$\bar{\alpha} \geq \frac{\frac{1}{2}p_0}{\frac{1}{2}p_0 + \sum_{j=1}^{k-1} p_j},$$

- (3) $\tilde{\sigma}_i^1(XX) = \tilde{\sigma}_i^2(XX) = 4$. Furthermore, for any $k \geq 3$, $\tilde{\sigma}_i^k(XX) > 2$, and $\tilde{\sigma}_i^k(XX) = 3$ if and only if there exists $2 \leq l \leq k-1$ such that $\tilde{\sigma}_i^l(OX) = 2$ with $\tilde{L}_k^* = \min_j \left\{ j < k : \tilde{\sigma}_i^j(OX) = 2 \right\}$, and

$$\bar{\alpha} \geq \max \left\{ \frac{\frac{3}{4}p_0 \left(\frac{3}{4}p_0 + \sum_{j=1}^{\tilde{L}_k^*-1} p_j \right) - \delta \left(\frac{1}{2}p_0 \right) \left(\frac{1}{2}p_0 + \sum_{j=1}^{\tilde{L}_k^*-1} p_j \right)}{\left(\frac{3}{4}p_0 + \sum_{j=1}^{\tilde{L}_k^*-1} p_j \right)^2 - \delta \left(\frac{1}{2}p_0 + \sum_{j=1}^{\tilde{L}_k^*-1} p_j \right)^2}, \frac{\frac{1}{2}p_0 + \sum_{j=1}^{\tilde{L}_k^*-1} p_j}{\frac{1}{2}p_0 + \sum_{j=1}^{k-1} p_j} \right\}.$$

Proof: Step 1: Consider any $i \in N$. If $x_{-i} = OO$, player i knows his face is dirty immediately, suggesting 1 is a dominant strategy and $\tilde{\sigma}_i^k(OO) = 1$ for any $k \geq 1$. If $x_{-i} = OX$ or XX , the expected payoff of 1 is $p\alpha - (1-p) < 0$, implying $\tilde{\sigma}_i^k(OX) \geq 2$ and $\tilde{\sigma}_i^k(XX) \geq 2$ for any $k \geq 1$. Moreover, level 1 players believe all other players are level 0, so when observing OX or XX , the expected payoff of $t \in \{2, 3\}$ is

$$p \left[\delta^{t-1} \alpha \left(\frac{5-t}{4} \right)^2 \right] + (1-p) \left[-\delta^2 \left(\frac{5-t}{4} \right)^2 \right] = \delta^{t-1} \left(\frac{5-t}{4} \right)^2 [p\alpha - (1-p)] < 0,$$

implying $\tilde{\sigma}_i^1(OX) = \tilde{\sigma}_i^1(XX) = 4$.

In addition, $\tilde{\sigma}_i^k(XX) \geq 3$ for all $k \geq 1$ can be proven by induction on k . From the previous calculation, we know $\tilde{\sigma}_i^1(XX) = 4$, which establishes the base case. Suppose $\tilde{\sigma}_i^k(XX) \geq 3$ for all $1 \leq i \leq K$ for some $K > 1$. It suffices to prove $\tilde{\sigma}_i^{K+1}(XX) \geq 3$ by showing 2 is a strictly dominated strategy for level $K+1$ players. This is strictly dominated because

$$\begin{aligned} p \left[\delta \alpha \left(\frac{3}{4} \frac{p_0}{\sum_{j=0}^K p_j} + \frac{\sum_{j=1}^K p_j}{\sum_{j=0}^K p_j} \right)^2 \right] + (1-p) \left[-\delta \left(\frac{3}{4} \frac{p_0}{\sum_{j=0}^K p_j} + \frac{\sum_{j=1}^K p_j}{\sum_{j=0}^K p_j} \right)^2 \right] \\ = \delta \left(\frac{3}{4} \frac{p_0}{\sum_{j=0}^K p_j} + \frac{\sum_{j=1}^K p_j}{\sum_{j=0}^K p_j} \right)^2 [p\alpha - (1-p)] < 0. \end{aligned}$$

Step 2: This step establishes a monotonic result: for any $K > 1$, if $\tilde{\sigma}_i^{l+1}(OX) \leq \tilde{\sigma}_i^l(OX)$ for all $1 \leq l \leq K-1$, then $\tilde{\sigma}_i^{K+1}(OX) \leq \tilde{\sigma}_i^K(OX)$. If $\tilde{\sigma}_i^K(OX) = 4$, then there is nothing to prove. Suppose $\tilde{\sigma}_i^{l+1}(OX) \leq \tilde{\sigma}_i^l(OX)$ for all $1 \leq l \leq K-1$. If $\tilde{\sigma}_i^K(OX) = 3$, then it is necessary that level K player's expected payoff of choosing 3 is non-negative. Namely,

$$\delta^2 \left(\frac{1}{2} \frac{p_0}{\sum_{j=0}^{K-1} p_j} + \frac{\sum_{j=1}^{K-1} p_j}{\sum_{j=0}^{K-1} p_j} \right) \left[p\alpha \left(\frac{1}{2} \frac{p_0}{\sum_{j=0}^{K-1} p_j} + \frac{\sum_{j=1}^{K-1} p_j}{\sum_{j=0}^{K-1} p_j} \right) - (1-p) \left(\frac{1}{2} \frac{p_0}{\sum_{j=0}^{K-1} p_j} \right) \right] \geq 0,$$

which implies

$$\delta^2 \left(\frac{1}{2} \frac{p_0}{\sum_{j=0}^K p_j} + \frac{\sum_{j=1}^K p_j}{\sum_{j=0}^K p_j} \right) \left[p\alpha \left(\frac{1}{2} \frac{p_0}{\sum_{j=0}^K p_j} + \frac{\sum_{j=1}^K p_j}{\sum_{j=0}^K p_j} \right) - (1-p) \left(\frac{1}{2} \frac{p_0}{\sum_{j=0}^K p_j} \right) \right] > 0,$$

suggesting $\tilde{\sigma}_i^{K+1}(OX) \leq 3$. If $\tilde{\sigma}_i^K(OX) = 2$, it suffices to prove $\tilde{\sigma}_i^{K+1}(OX) = 2$ as well. Notice that if $\tilde{\sigma}_i^K(OX) = 2$, then it is necessary for level K players that 2

dominates 3 and 4. Let M be the lowest level of players that would choose 2 when observing OX . Then level K player's expected payoff of choosing 2 would satisfy

$$\delta \left(\frac{3}{4}p_0 + \sum_{j=1}^{K-1} p_j \right) \left[p\alpha \left(\frac{3}{4}p_0 + \sum_{j=1}^{K-1} p_j \right) - (1-p) \left(\frac{3}{4}p_0 \right) \right] \\ \geq \max \left\{ \delta^2 \left(\frac{1}{2}p_0 + \sum_{j=1}^{M-1} p_j \right) \left[p\alpha \left(\frac{1}{2}p_0 + \sum_{j=1}^{M-1} p_j \right) - (1-p) \left(\frac{1}{2}p_0 \right) \right], 0 \right\},$$

which implies

$$\delta \left(\frac{3}{4}p_0 + \sum_{j=1}^K p_j \right) \left[p\alpha \left(\frac{3}{4}p_0 + \sum_{j=1}^K p_j \right) - (1-p) \left(\frac{3}{4}p_0 \right) \right] \\ \geq \max \left\{ \delta^2 \left(\frac{1}{2}p_0 + \sum_{j=1}^{M-1} p_j \right) \left[p\alpha \left(\frac{1}{2}p_0 + \sum_{j=1}^{M-1} p_j \right) - (1-p) \left(\frac{1}{2}p_0 \right) \right], 0 \right\},$$

suggesting that $\tilde{\sigma}_i^{K+1}(OX) = 2$.

Step 3: In this step, I characterize level k player i 's behavior as $x_{-i} = OX$ for all $k \geq 2$ by induction on k . Level 2 player i 's expected payoff of choosing $t \in \{2, 3\}$ is

$$\delta^{t-1} \left(\frac{5-t}{4} \frac{p_0}{p_0+p_1} + \frac{p_1}{p_0+p_1} \right) \underbrace{\left[p\alpha \left(\frac{5-t}{4} \frac{p_0}{p_0+p_1} + \frac{p_1}{p_0+p_1} \right) - (1-p) \left(\frac{5-t}{4} \frac{p_0}{p_0+p_1} \right) \right]}_{\text{increasing in } t}.$$

Therefore, $\tilde{\sigma}_i^2(OX) \leq 3$ if and only if

$$p\alpha \left(\frac{1}{2}p_0 + p_1 \right) - (1-p) \left(\frac{1}{2}p_0 \right) \geq 0 \iff \bar{\alpha} \geq \frac{\frac{1}{2}p_0}{\frac{1}{2}p_0 + p_1},$$

and $\tilde{\sigma}_i^2(OX) = 2$ if and only if

$$\delta \left(\frac{3}{4}p_0 + p_1 \right) \left[p\alpha \left(\frac{3}{4}p_0 + p_1 \right) - (1-p) \left(\frac{3}{4}p_0 \right) \right] \\ \geq \max \left\{ \delta^2 \left(\frac{1}{2}p_0 + p_1 \right) \left[p\alpha \left(\frac{1}{2}p_0 + p_1 \right) - (1-p) \left(\frac{1}{2}p_0 \right) \right], 0 \right\} \\ \iff \bar{\alpha} \geq \max \left\{ \frac{\frac{3}{4}p_0 \left(\frac{3}{4}p_0 + p_1 \right) - \delta \left(\frac{1}{2}p_0 \right) \left(\frac{1}{2}p_0 + p_1 \right)}{\left(\frac{3}{4}p_0 + p_1 \right)^2 - \delta \left(\frac{1}{2}p_0 + p_1 \right)^2}, \frac{\frac{3}{4}p_0}{\frac{3}{4}p_0 + p_1} \right\}.$$

Since for any $\delta \in (0, 1)$,

$$\frac{\frac{3}{4}p_0 \left(\frac{3}{4}p_0 + p_1 \right) - \delta \left(\frac{1}{2}p_0 \right) \left(\frac{1}{2}p_0 + p_1 \right)}{\left(\frac{3}{4}p_0 + p_1 \right)^2 - \delta \left(\frac{1}{2}p_0 + p_1 \right)^2} > \frac{\frac{3}{4}p_0}{\frac{3}{4}p_0 + p_1},$$

2 is optimal for level 2 players if and only if

$$\bar{\alpha} \geq \frac{\frac{3}{4}p_0 \left(\frac{3}{4}p_0 + p_1 \right) - \delta \left(\frac{1}{2}p_0 \right) \left(\frac{1}{2}p_0 + p_1 \right)}{\left(\frac{3}{4}p_0 + p_1 \right)^2 - \delta \left(\frac{1}{2}p_0 + p_1 \right)^2}.$$

Now suppose there is $K > 1$ such that the statement holds for any $1 \leq k \leq K$. We want to show it also holds for level $K+1$ players. Notice that by induction hypothesis, $\tilde{\sigma}_i^{l+1}(OX) \leq \tilde{\sigma}_i^l(OX)$ for all $1 \leq l \leq K-1$, implying $\tilde{\sigma}_i^{K+1}(OX) \leq \tilde{\sigma}_i^K(OX)$ by step 2. If $\tilde{\sigma}_i^K(OX) \leq 3$, then $\tilde{\sigma}_i^{K+1}(OX) \leq 3$ by step 2. Therefore, it suffices to focus on the case where $\tilde{\sigma}_i^l(OX) = 4$ for all $1 \leq l \leq K$. In this case, level $K+1$ player's expected payoff of choosing $t \in \{2, 3\}$ is

$$\delta^{t-1} \left(\frac{5-t}{4} \frac{p_0}{\sum_{j=0}^K p_j} + \frac{\sum_{j=1}^K p_j}{\sum_{j=0}^K p_j} \right) \left[p\alpha \left(\frac{5-t}{4} \frac{p_0}{\sum_{j=0}^K p_j} + \frac{\sum_{j=1}^K p_j}{\sum_{j=0}^K p_j} \right) - (1-p) \left(\frac{5-t}{4} \frac{p_0}{\sum_{j=0}^K p_j} \right) \right],$$

suggesting 4 is a dominated strategy if and only if

$$\bar{\alpha} \geq \frac{\frac{1}{2}p_0}{\frac{1}{2}p_0 + \sum_{j=1}^K p_j}.$$

If $\tilde{\sigma}_i^K(OX) = 2$, then $\tilde{\sigma}_i^{K+1}(OX) = 2$ by step 2. Thus, it suffices to consider the case where $\tilde{\sigma}_i^l(OX) \geq 3$ for all $1 \leq l \leq K$. In this case, $\tilde{\sigma}_i^{K+1}(OX) = 2$ if and only if

$$\begin{aligned} & \delta \left(\frac{3}{4}p_0 + \sum_{j=1}^K p_j \right) \left[p\alpha \left(\frac{3}{4}p_0 + \sum_{j=1}^K p_j \right) - (1-p) \left(\frac{3}{4}p_0 \right) \right] \\ & \geq \max \left\{ \delta^2 \left(\frac{1}{2}p_0 + \sum_{j=1}^K p_j \right) \left[p\alpha \left(\frac{1}{2}p_0 + \sum_{j=1}^K p_j \right) - (1-p) \left(\frac{1}{2}p_0 \right) \right], 0 \right\} \\ \iff & \bar{\alpha} \geq \frac{\frac{3}{4}p_0 \left(\frac{3}{4}p_0 + \sum_{j=1}^K p_j \right) - \delta \left(\frac{1}{2}p_0 \right) \left(\frac{1}{2}p_0 + \sum_{j=1}^K p_j \right)}{\left(\frac{3}{4}p_0 + \sum_{j=1}^K p_j \right)^2 - \delta \left(\frac{1}{2}p_0 + \sum_{j=1}^K p_j \right)^2}. \end{aligned}$$

Step 4: Lastly, this step characterizes level k player i 's behavior when $x_{-i} = XX$ for level $k \geq 3$. Consider any level $K \geq 3$. For level k players, they would choose

3 only if there is some level $2 \leq l \leq k - 1$ such that $\tilde{\sigma}_i^l(OX) = 2$. Let \tilde{L}_k^* be the lowest level below k that would choose 2 when seeing one dirty face. Then level k player i 's expected payoff of 3 is

$$\delta^2 \left(\frac{1}{2} \frac{p_0}{\sum_{j=0}^{k-1} p_j} + \frac{\sum_{j=1}^{k-1} p_j}{\sum_{j=0}^{k-1} p_j} \right) \left[p\alpha \left(\frac{1}{2} \frac{p_0}{\sum_{j=0}^{k-1} p_j} + \frac{\sum_{j=1}^{k-1} p_j}{\sum_{j=0}^{k-1} p_j} \right) - (1-p) \left(\frac{1}{2} \frac{p_0}{\sum_{j=0}^{k-1} p_j} + \frac{\sum_{j=1}^{\tilde{L}_k^*-1} p_j}{\sum_{j=0}^{k-1} p_j} \right) \right],$$

which dominates 4 if and only if

$$\bar{\alpha} \geq \frac{\frac{1}{2}p_0 + \sum_{j=0}^{\tilde{L}_k^*-1} p_j}{\frac{1}{2}p_0 + \sum_{j=0}^k p_j}.$$

Coupled with the existence of \tilde{L}_k^* , $\tilde{\sigma}_i^k(XX) = 3$ if and only if

$$\bar{\alpha} \geq \max \left\{ \frac{\frac{3}{4}p_0 \left(\frac{3}{4}p_0 + \sum_{j=1}^{\tilde{L}_k^*-1} p_j \right) - \delta \left(\frac{1}{2}p_0 \right) \left(\frac{1}{2}p_0 + \sum_{j=1}^{\tilde{L}_k^*-1} p_j \right)}{\left(\frac{3}{4}p_0 + \sum_{j=1}^{\tilde{L}_k^*-1} p_j \right)^2 - \delta \left(\frac{1}{2}p_0 + \sum_{j=1}^{\tilde{L}_k^*-1} p_j \right)^2}, \frac{\frac{1}{2}p_0 + \sum_{j=1}^{\tilde{L}_k^*-1} p_j}{\frac{1}{2}p_0 + \sum_{j=1}^{k-1} p_j} \right\}.$$

This completes the proof of this proposition. ■

D.3.3 Illustrative Example

To illustrate the representation effect in three-person games, I characterize level 3 players' behavior in both the sequential and simultaneous games when the distribution of levels follows Poisson(1.5). Similar to the analysis of two-person games, the set of dirty-faces games is the unit square on the $(\delta, \bar{\alpha})$ -plane.

When observing one dirty face and one clean face, level 3 players cannot tell their faces for sure in period 1, so they will wait, no matter in sequential games or in simultaneous games. In period 2, level 3 players will claim to have a dirty face if and only if the expected payoff of C is higher than the continuation value of choosing W . By Proposition D.2 and D.3, level 3 players will claim in period 2 in the sequential game if and only if

$$\bar{\alpha} \geq \frac{\left(\frac{1}{2} - \frac{1}{4}\gamma_3\delta \right) p_0}{\left(\frac{1}{2} - \frac{1}{4}\gamma_3\delta \right) p_0 + (1 - \gamma_3\delta)(p_1 + p_2)} = \frac{100 - 46\delta}{625 - 529\delta},$$

and choose 2 in the simultaneous game if and only if

$$\bar{\alpha} \geq \frac{\frac{3}{4}p_0 \left(\frac{3}{4}p_0 + p_1 + p_2 \right) - \delta \left(\frac{1}{2}p_0 \right) \left(\frac{1}{2}p_0 + p_1 + p_2 \right)}{\left(\frac{3}{4}p_0 + p_1 + p_2 \right)^2 - \delta \left(\frac{1}{2}p_0 + p_1 + p_2 \right)^2} = \frac{162 - 100\delta}{729 - 625\delta}.$$

At period 3, level 3 players will claim to have a dirty face if and only if the expected payoff of C is positive. Therefore, in the sequential game, level 3 players will claim in period 3 if and only if

$$\bar{\alpha} \geq \frac{\frac{1}{4}p_0}{\frac{1}{4}p_0 + p_1 + p_2} = \frac{2}{23},$$

while in the simultaneous game, they will not choose always wait if and only if

$$\bar{\alpha} \geq \frac{\frac{1}{4}p_0}{\frac{1}{4}p_0 + p_1 + p_2} = \frac{4}{25}.$$

When observing two dirty faces, level 3 players cannot tell their face types in the first two periods, so they will wait in the first two periods. At period 3, level 3 players will claim if and only if (1) level 2 players will claim at period 2 when seeing only one dirty face,¹ and (2) the expected payoff of C is positive. Therefore, in the sequential game, it is optimal to claim at period 3 if and only if

$$\bar{\alpha} \geq \max \left\{ \frac{16 - 7\delta}{64 - 49\delta}, \frac{196}{529} \right\}.$$

In the simultaneous game, it is optimal to claim at period 3 when observing two dirty faces if and only if

$$\bar{\alpha} \geq \max \left\{ \frac{27 - 16\delta}{81 - 64\delta}, \frac{16}{25} \right\}.$$

Level 3 players' DCH optimal stopping periods in both sequential and simultaneous games are plotted in Figure D.1. The definition of optimal stopping periods is naturally extended to three-person games. From this figure, we can observe two features that are different from the two-person games. First, when observing one dirty face and $\delta \rightarrow 1$, level 3 players will claim at period 2 if $\bar{\alpha} \geq 9/16$. However, in two-person games, when $\delta \rightarrow 1$, players will always wait till the last period. This is because when there are more players, the game is more likely to be randomly terminated, causing the players to claim earlier even if the payoff is not discounted. Second, when observing two dirty faces, level 3 players' behavior at period 3 depends on δ *even if this is the last period*. The reason is that level 3 players' belief at period 3 depends on level 2 players' behavior at period 2 which depends on δ .

¹Otherwise, if both level 1 and 2 players wait at period 2 when seeing only one dirty face, level 3 players cannot make inferences about their face types when the game proceeds to period 3.

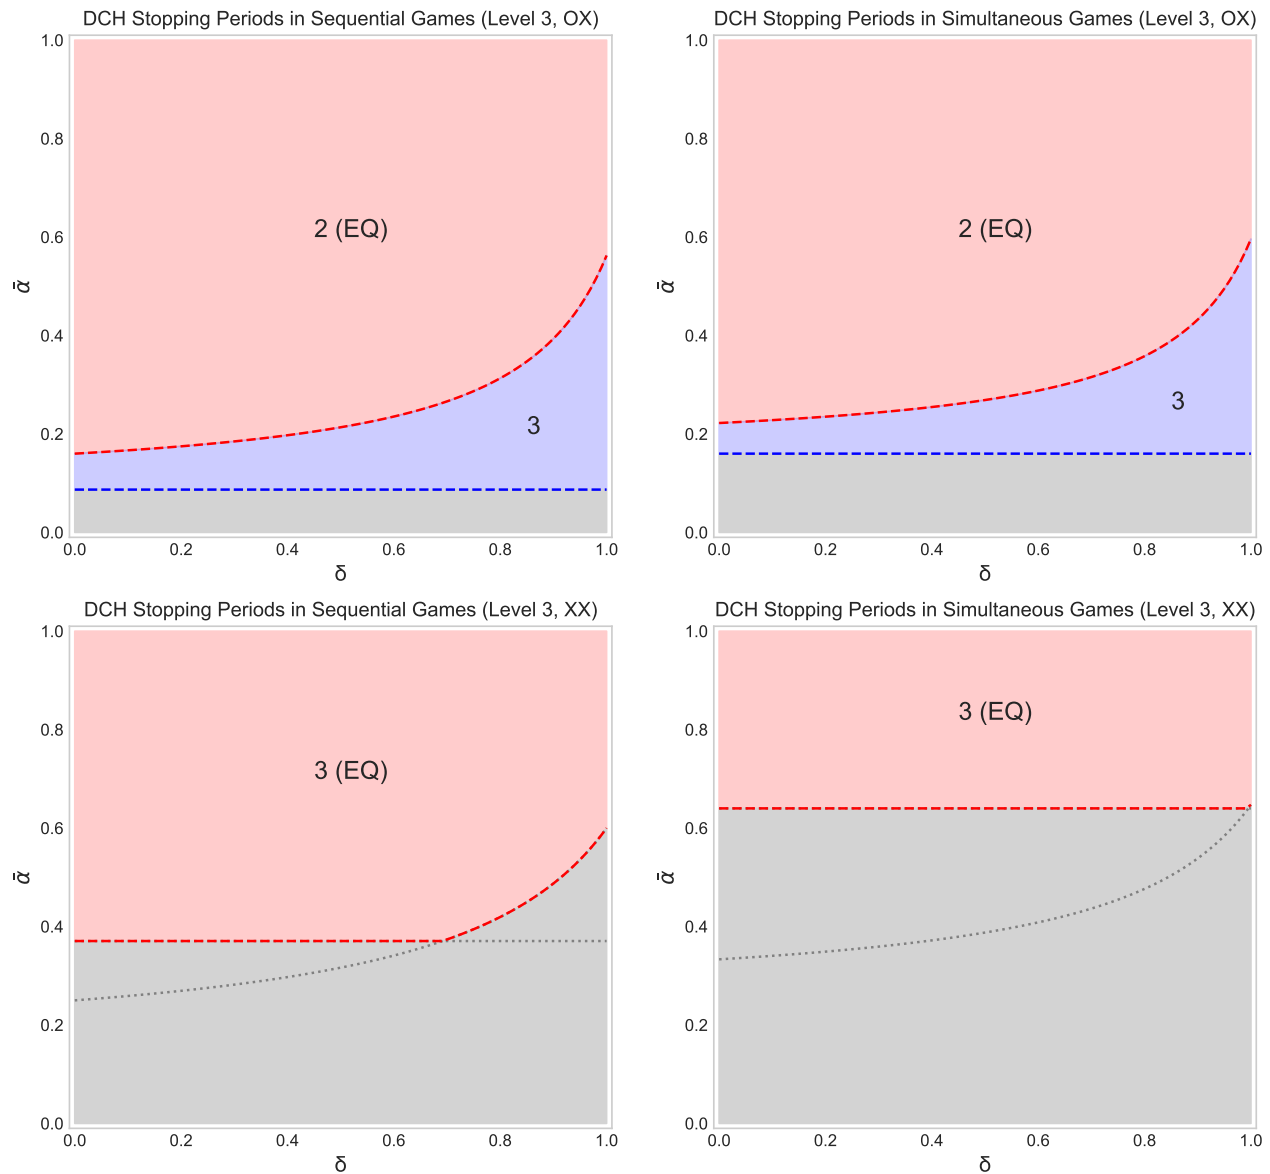


Figure D.1: Level 3 players' DCH stopping periods in sequential (left column) and simultaneous (right column) three-person three-period dirty-faces games where the distribution of levels follows Poisson(1.5).

Remark D.1. *In this illustrative example, DCH predicts level 3 players tend to claim earlier in sequential games than in simultaneous games because (1) at information set (2, OX), $(162 - 100\delta)/(729 - 625\delta) > (100 - 46\delta)/(625 - 529\delta)$, (2) at information set (3, OX), $4/25 > 2/23$, and (3) at information set (3, XX),*

$$\max \left\{ \frac{27 - 16\delta}{81 - 64\delta}, \frac{16}{25} \right\} > \max \left\{ \frac{16 - 7\delta}{64 - 49\delta}, \frac{196}{529} \right\}.$$

To summarize, the analysis of three-person three-period games demonstrates how the DCH solution varies with the representations and the number of players in a game. The prediction of the equilibrium theory only depends on the number of dirty faces, not the number of players. This sharply contrasts with DCH. The intuition is that when there are more players, the game is more likely to be randomly terminated by level 0 players, and hence strategic players' behavior is affected.

D.4 Detailed Analysis of Bayer and Chan (2007) Data

D.4.1 Data Description

This section revisits the dirty-faces experimental data by Bayer and Chan (2007). The description of the experimental setting can be found in the main text section 4.6.1, and the instructions and screenshots can be found in Bayer and Chan (2007) Appendix A.

Following previous notations, I use (t, x_{-i}) to denote the situation where subject i sees type x_{-i} at period t . After excluding the data from the case where there is no public announcement, the raw data at each information set is reported in Table D.1. Each entry in the table states the number of observations and the percentage of the choices that follow the equilibrium predictions. For instance, at information set $(t, x_{-i}) = (2, X)$, there are 170 choices and 62 percent of the choices are C , which is the action predicted by the equilibrium.

From Table D.1, we can observe that the behavior aligns with the equilibrium prediction when players do not see any dirty face. In this situation ($x_{-i} = O$ or OO), players are aware that their face type is X and choose C in period 1. However, the behavior becomes less consistent with the equilibrium as the reasoning complexity increases. When players see only one dirty face ($x_{-i} = X$ or OX), they should realize that their face type is X as the game progresses to period 2. However, the empirical data show that only 62% and 58% of players in Treatment 1 and 2, respectively, are able to do so. Furthermore, when players see two dirty faces ($x_{-i} = XX$), only 30% of the players claim to have a dirty face in period 3.

Table D.1: Experimental Data from Bayer and Chan (2007).

x_{-i}	Number of Players				
	2		3		
	<i>O</i>	<i>X</i>	<i>OO</i>	<i>OX</i>	<i>XX</i>
EQ	<i>C</i>	<i>WC</i>	<i>C</i>	<i>WC</i>	<i>WWC</i>
Period	Number of Obs (EQ %)				
1	123 (0.94)	391 (0.79)	48 (0.92)	280 (0.61)	320 (0.76)
2	6 (0.50)	170 (0.62)	2 (0.50)	60 (0.58)	145 (0.79)
3	—	—	—	10 (0.20)	56 (0.36)

Note: In Treatment 1, there are 21 groups of subjects (42 subjects in total), and in Treatment 2, there are 16 groups of subjects (48 subjects in total). Because each group plays 14 rounds, the data set consists of $(21 + 16) \times 14 = 518$ games.

These observations suggest that the equilibrium fails to explain a significant portion of the data. In the following analysis, I compare the fitness of the DCH model with that of the standard CH model and the agent quantal response equilibrium (AQRE) proposed by McKelvey and Palfrey (1998). By comparing the DCH and the standard CH models, I can quantify the improvement achieved by incorporating learning from past actions into the CH framework. On the other hand, AQRE is an equilibrium model designed for extensive games, where players make stochastic choices and assume that other players do the same. The comparison between the DCH and AQRE demonstrates how hierarchical thinking models can generate statistically comparable predictions as equilibrium-based models.

D.4.2 Likelihood Functions

This section derives the likelihood functions. For the cognitive hierarchy theories, I follow Camerer, Ho, and Chong (2004) to assume the prior distribution of levels follows Poisson distribution. Therefore, for both of the Poisson-DCH and the standard Poisson-CH, there is one parameter to be estimated—the average number of levels τ . For AQRE, I follow McKelvey and Palfrey (1998) to estimate the logit-AQRE which has a single parameter λ .

D.4.2.1 Poisson-CH Models

The Poisson-CH models assume each player's level is i.i.d. drawn from $(p_k)_{k=0}^{\infty}$ where

$$p_k \equiv \frac{e^{-\tau} \tau^k}{k!}, \quad \text{for all } k = 0, 1, 2, \dots$$

and $\tau > 0$. Because τ is the mean and variance of the Poisson distribution, the economic meaning of τ is the average level of sophistication among the population.

I first construct the likelihood function for the Poisson-DCH model. For each subject i , let Π_i denote the set of information sets that subject i has encountered in the game, and let $\mathcal{I}_i = (t, x_{-i})$ denote a generic information set. At any information set \mathcal{I}_i , subject i can choose $c_i \in \{C, W\}$. Let $P_k(c_i|\mathcal{I}_i, \tau)$ be the probability of level k players choosing c_i at information set \mathcal{I}_i . Moreover, let $f(k|\mathcal{I}_i, \tau)$ be the posterior distribution of levels at information set \mathcal{I}_i . At period 1, $f(k|\mathcal{I}_i, \tau) = e^{-\tau} \tau^k / k!$. For later periods, $f(k|\mathcal{I}_i, \tau)$ given any τ can be analytically solved by Proposition 4.4 (two-person games) and Proposition D.2 (three-person games). Finally, the predicted choice probability for c_i at information set \mathcal{I}_i is simply the aggregation of best responses from all levels weighted by the proportion $f(k|\mathcal{I}_i, \tau)$:

$$\mathcal{D}(c_i|\mathcal{I}_i, \tau) \equiv \sum_{k=0}^{\infty} f(k|\mathcal{I}_i, \tau) P_k(c_i|\mathcal{I}_i, \tau).$$

Consequently, the log-likelihood function for the DCH model can be formed by aggregating over every subject i , actions c_i and information set \mathcal{I}_i :

$$\ln L^D(\tau) = \sum_i \sum_{\mathcal{I}_i \in \Pi_i} \sum_{c_i \in \{W, C\}} \mathbb{1}\{c_i, \mathcal{I}_i\} \ln [\mathcal{D}(c_i|\mathcal{I}_i, \tau)],$$

where $\mathbb{1}\{c_i, \mathcal{I}_i\}$ is the indicator function which is 1 when subject i chooses c_i at \mathcal{I}_i .

Second, the log-likelihood function for the standard Poisson-CH model can be constructed in the similar way. Given any τ , the standard Poisson-CH model predicts a probability distribution over $\{1, \dots, T, T+1\}$ (earliest period to choose C or always W) for each level of players conditional on the announcement and other players' faces. Following previous notations, the probability of level k subject i choosing t conditional on x_{-i} is denoted by $\tilde{\sigma}_i^k(t|x_{-i})$, which can be analytically solved by Proposition 4.5 (two-person games) and Proposition D.3 (three-person games). Therefore, subject i 's predicted choice probability for $t \in \{1, \dots, T, T+1\}$ conditional on x_{-i} is the aggregation of choice frequencies of all levels weighted by

Poisson(τ):

$$\tilde{S}(t|x_{-i}, \tau) = \sum_{k=0}^{\infty} \frac{e^{-\tau} \tau^k}{k!} \tilde{\sigma}_i^k(t|x_{-i}).$$

Since $\tilde{\sigma}_i^0(t|x_{-i}) = \frac{1}{T+1}$ for all t , $\tilde{S}(t|x_{-i}, \tau) > 0$ for all t . Moreover, the conditional probability to choose C or W at information set $\mathcal{I}_i = (t, x_{-i})$ can be computed by

$$\mathcal{S}(C|\mathcal{I}_i, \tau) = \frac{\tilde{S}(t|x_{-i}, \tau)}{\sum_{t' \geq t} \tilde{S}(t'|x_{-i}, \tau)} \quad \text{and} \quad \mathcal{S}(W|\mathcal{I}_i, \tau) = 1 - \mathcal{S}(C|\mathcal{I}_i, \tau).$$

Finally, the log-likelihood function for the standard CH model can be constructed by aggregating over every subjects i , actions c_i , and information set \mathcal{I}_i :

$$\ln L^S(\tau) = \sum_i \sum_{\mathcal{I}_i \in \Pi_i} \sum_{c_i \in \{W, C\}} \mathbb{1}\{c_i, \mathcal{I}_i\} \ln [\mathcal{S}(c_i|\mathcal{I}_i, \tau)].$$

D.4.2.2 Logit-AQRE Model

For the purpose of illustrate, I only derive the likelihood function for two-person games. The likelihood function for three-person games can be derived by a similar calculation.

Let $Q(c_i|\mathcal{I}_i, \lambda)$ be the probability of subject i choosing c_i at information set \mathcal{I}_i predicted by the logit-AQRE. In the two-person two-period dirty-faces game, each player's strategy is defined by a four-tuple (q_1, q_2, r_1, r_2) which corresponds to $Q(C|1, O, \lambda)$, $Q(C|2, O, \lambda)$, $Q(C|1, X, \lambda)$, and $Q(C|2, X, \lambda)$, respectively. At information set $(t, x_{-i}) = (1, O)$, players would estimate the payoff of C and W by

$$\begin{aligned} U_{1,O}(C) &= \alpha + \epsilon_{1,O,C} \\ U_{1,O}(W) &= \delta\alpha(1 - r_1)q_2 + \epsilon_{1,O,W}, \end{aligned}$$

where $\epsilon_{1,O,C}$ and $\epsilon_{1,O,W}$ are independent random variables with a Weibull distribution with the precision parameter λ . Then the logit formula suggests

$$q_1 = \frac{1}{1 + \exp\{\lambda[\delta\alpha(1 - r_1)q_2 - \alpha]\}}.$$

Similarly, q_2 can be expressed by

$$q_2 = \frac{1}{1 + \exp\{-\delta\alpha\lambda\}}.$$

On the other hand, when observing a dirty face and the game proceeds to period 2, players' posterior beliefs become

$$\mu \equiv \Pr(X|2, X) = \frac{p(1 - r_1)}{p(1 - r_1) + (1 - p)(1 - q_1)} = \frac{1}{1 + \left(\frac{1-p}{p}\right)\left(\frac{1-q_1}{1-r_1}\right)},$$

and hence the expected payoff to choose C at information set $(2, X)$ is

$$\delta [\alpha\mu - (1 - \mu)] = \delta [(1 + \alpha)\mu - 1].$$

As a result, r_2 satisfies that

$$r_2 = \frac{1}{1 + \exp \{ \lambda \delta [1 - (1 + \alpha)\mu] \}}.$$

Finally, the expected payoff of choosing C at information set $(1, X)$ is $\alpha p - (1 - p)$, while the expected payoff of W is

$$\underbrace{[p(1 - r_1) + (1 - p)(1 - q_1)]}_{\text{prob. to reach period 2}} r_2 \delta [(1 + \alpha)\mu - 1] \equiv A,$$

and therefore, r_1 can be expressed by

$$r_1 = \frac{1}{1 + \exp \{ \lambda [A + (1 - p) - \alpha p] \}}.$$

As plugging $p = 2/3$, $\delta = 4/5$ and $\alpha = 2/3$ into the choice probabilities, we can obtain that

$$\begin{aligned} r_1 &= \frac{1}{1 + \exp \left\{ \lambda \left[\frac{2}{15}(1 - r_1)r_2 - \frac{4}{15}(1 - q_1)r_2 + \frac{1}{6} \right] \right\}} \\ r_2 &= \frac{1}{1 + \exp \left\{ \lambda \left[\frac{4}{5} - \frac{2 - 2r_1}{3 - 2r_1 - q_1} \right] \right\}} \\ q_1 &= \frac{1}{1 + \exp \left\{ \lambda \left[\frac{1}{5}(1 - r_1)q_2 - \frac{1}{4} \right] \right\}} \\ q_2 &= \frac{1}{1 + \exp \left\{ -\frac{1}{5}\lambda \right\}}. \end{aligned}$$

Given each λ , the system of four equations with four unknowns can be solved uniquely. Besides, for each \mathcal{I}_i , $Q(W|\mathcal{I}_i, \lambda) = 1 - Q(C|\mathcal{I}_i, \lambda)$. Thus, the log-likelihood function can be formed by aggregating over every subject i , action c_i , and information set \mathcal{I}_i :

$$\ln L^Q(\lambda) = \sum_i \sum_{\mathcal{I}_i \in \Pi_i} \sum_{c_i \in \{W, C\}} \mathbf{1}\{c_i, \mathcal{I}_i\} \ln [Q(c_i|\mathcal{I}_i, \lambda)].$$

D.4.3 Estimation Results

The Poisson-DCH, standard Poisson-CH and the AQRE models are estimated by maximum likelihood estimation. Table D.2 reports the estimation results on Treatment 1 and Treatment 2 data, showing the estimated parameters and the fitness of

each model. Comparing the fitness of these models, I find that the log-likelihood of DCH is significantly higher than standard CH (Vuong Test p-value < 0.001 for both treatments), while it is not significantly different from AQRE (Treatment 1: p-value = 0.144; Treatment 2: p-value = 0.184). This result suggests in both Treatment 1 and 2, DCH outperforms the standard CH in capturing the empirical patterns and generates predictions that are statistically comparable to other equilibrium-based behavioral solution concepts.

Table D.2: Estimation Results for Treatment 1 and Treatment 2 Data.

		Two-Person Games			Three-Person Games		
		DCH	Standard CH	AQRE	DCH	Standard CH	AQRE
Parameters	τ	1.269	1.161	—	0.370	0.140	—
	S.E.	(0.090)	(0.095)	—	(0.043)	(0.039)	—
	λ	—	—	7.663	—	—	5.278
	S.E.	—	—	(0.493)	—	—	(0.404)
Fitness	LL	-360.75	-381.46	-368.38	-575.30	-608.45	-565.05
	AIC	723.50	764.91	738.76	1152.61	1218.89	1132.11
	BIC	728.04	769.45	743.29	1157.43	1223.72	1136.93
Vuong Test			6.517	1.463		3.535	-1.330
p-value			< 0.001	0.144		< 0.001	0.184

Note: There are 294 games (rounds \times groups) in Treatment 1 and 224 games in Treatment 2.

Comparing the estimation results of Treatment 1 and 2, I observe that there is more randomness in three-person games compared to two-person games. In two-person games, the DCH estimates indicate that players can think 1.269 steps (95% C.I. = [1.093, 1.445]) on average, while in three-person games, players can only think an average of 0.370 steps (95% C.I. = [0.286, 0.454]). Additionally, the estimation result of AQRE suggests that as the game changes from two-person games to three-person games, the precision of decision-making decreases significantly (from 7.663 to 5.278). This implies that players are less likely to make best responses in three-person games.

To analyze the differences between the models in detail, I compare the choice probabilities predicted by each model. Figure D.2 illustrates the choice probabilities in two-person games, while Figure D.3 displays the choice probabilities in three-person games. Comparing the DCH and the standard CH models, I observe that the standard CH model generally underestimates the probability of choosing *C* in period

1. In two-person games, the empirical frequencies of choosing C at information sets $(1, O)$ and $(1, X)$ are 0.943 and 0.210, respectively. Yet the predictions of the standard CH model are 0.791 and 0.104 for the same information sets. A similar pattern of underestimation is also evident in three-person games.

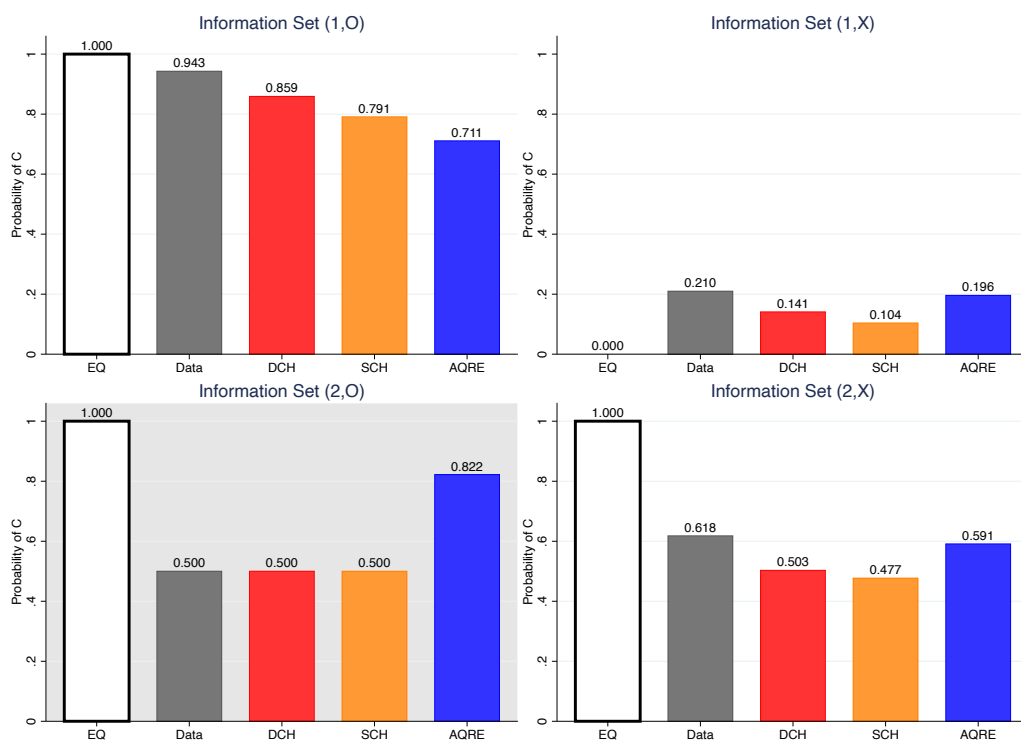


Figure D.2: The choice probabilities in two-person games at different information sets. Each panel plots the empirical choice frequencies and the predictions of different models at one information set. The gray panel represents the off-equilibrium path information set.

The underestimation is primarily caused by the difference in the specifications of level 0 players' behavior. In two-person games, the standard CH model assumes that level 0 players uniformly randomize across the set $\{1, 2, 3\}$. Consequently, the probability of level 0 players choosing C in period 1 according to the standard CH model is $1/3$. In contrast, in the DCH model, level 0 players uniformly randomize at every information set, resulting in a probability of $1/2$ for them to choose C . Similarly, in three-person games, the standard CH model assumes that level 0 players uniformly randomize across the set $\{1, 2, 3, 4\}$, leading to a probability of $1/4$ for them to choose C in period 1. In contrast, in the DCH model, level 0 players' behavior remains the same across both two-person and three-person games.

These differences in level 0 players' behavior contribute to the underestimation of the probability of choosing C in the standard CH model compared to DCH.

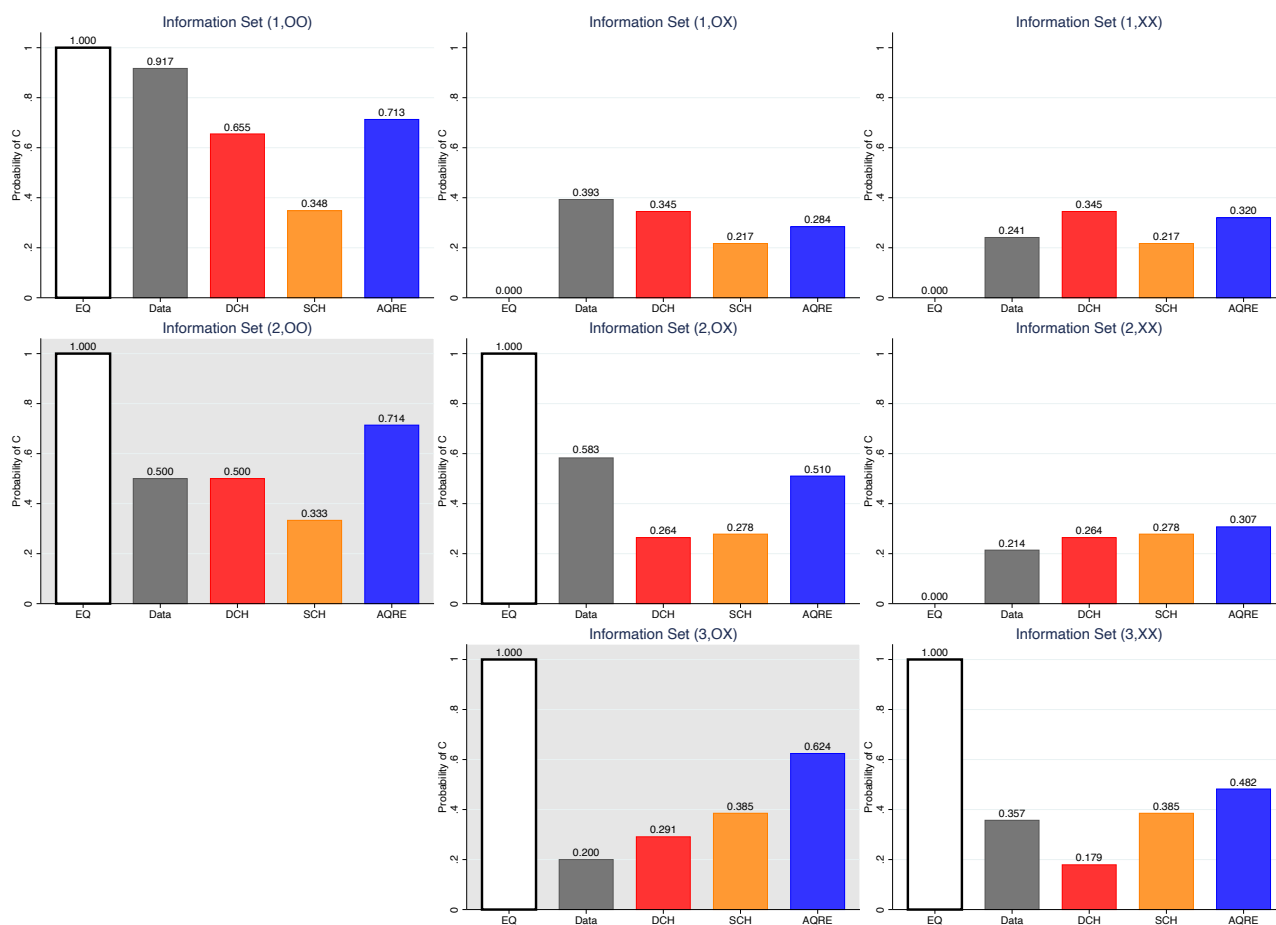


Figure D.3: The choice probabilities in three-person games at different information sets. Each panel plots the empirical choice frequencies and the predictions of different models at one information set. The gray panels represent the off-equilibrium-path information sets.

Moreover, the key difference between the CH approach and AQRE is highlighted in the off-equilibrium-path information sets.² Conceptually, the reason why the game could proceed to the off-equilibrium-path information sets differs between the CH approach and AQRE. From the perspective of AQRE, the off-equilibrium-path information sets are reached due to mistakes. As a result, AQRE predicts a high

²When $x_{-i} = OO$, the equilibrium predicts that players will choose C in period 1, resulting in the game not proceeding beyond period 2. Similarly, when $x_{-i} = OX$, the equilibrium suggests that players should choose C in period 2, preventing the game from progressing to period 3.

probability of choosing C at these off-equilibrium-path information sets because the expected payoff of choosing C is much higher than W at these information sets. By contrary, in the CH approach, the off-equilibrium-path information sets are reached because the players are not sophisticated enough. For instance, when observing no dirty face, players should immediately choose C since it is a dominant strategy. If someone doesn't choose C , they are definitely a level 0 player.

From the choice probabilities, it can be observed that DCH provides the most accurate predictions at off-path information sets, regardless of whether it is in two-person or three-person games. At information sets $(2, O)$ and $(2, OO)$, the empirical choice probabilities of C are 0.5, which are correctly predicted by DCH. Furthermore, at the information set $(3, OX)$, the empirical choice probability of C is 0.2, while the predictions of DCH, standard CH, and AQRE are 0.291, 0.385, and 0.624, respectively.

Table D.3: Estimation Results for Pooled Data.

		DCH	Standard CH	AQRE
Parameters	τ	1.030	0.241	—
	S.E.	(0.060)	(0.033)	—
	λ	—	—	6.235
	S.E.	—	—	(0.302)
Fitness	LL	-956.92	-1047.12	-940.65
	AIC	1915.84	2096.23	1883.30
	BIC	1921.22	2101.62	1888.69
Vuong Test			7.513	-1.363
p-value			< 0.001	0.173
LR Test $\chi^2_{(1)}$		41.74	114.42	14.44
p-value		< 0.001	< 0.001	< 0.001

Note: The likelihood ratio test is testing if the log-likelihood of two-parameter models (Treatment 1 and 2) is significantly higher than the log-likelihood of one-parameter models.

In addition, I estimate the three models using the pooled data, and the results are reported in Table D.3. Consistent with the results from the two-person games and three-person games, it can be observed that DCH provides a significantly better fit to the data compared to the standard CH model (Vuong test: p-value < 0.001). However, there is no statistically significant difference between DCH and AQRE (Vuong

test p-value = 0.173). Furthermore, I conduct a likelihood ratio test on all three models to assess whether allowing different parameters for two-person and three-person games can significantly improve the model fit. The results indicate that the heterogeneous models are significantly better than the homogeneous models. Taken together, these findings lead to the conclusion that both the level of sophistication and the precision vary with the complexity of the games.

To summarize, it is not surprising that DCH can provide a better explanation for the data compared to the misspecified standard CH model in dynamic games. However, what is surprising is that when the CH model is correctly specified, the estimated average level of sophistication is 1.03, which falls within the expected range of a “regular” τ value between 1 and 2, as predicted by Camerer, Ho, and Chong (2004).

D.5 Supplementary Analysis for Experimental Data

D.5.1 Supplementary Tables

This section includes all the supplementary tables from the experiment. Table D.4 lists the empirical frequencies of choosing C at each information set for both treatments. Table D.5 presents the empirical frequencies of choosing C at information set $(2, X)$ for different payoff structures and treatments.

Table D.4: The Empirical Frequencies of C at Each Information Set.

x_{-i}	Sequential Treatment						Simultaneous Treatment					
	O			X			O			X		
	Obs	Claim %	s.d.	Obs	Claim %	s.d.	Obs	Claim %	s.d.	Obs	Claim %	s.d.
Periods												
1	148	0.845	0.364	572	0.313	0.464	148	0.811	0.393	548	0.263	0.441
2	16	0.438	0.512	210	0.600	0.491	28	0.250	0.441	404	0.223	0.417
3	4	0.000	0.000	34	0.206	0.410	21	0.190	0.402	314	0.172	0.378
4	3	0.000	0.000	21	0.190	0.402	16	0.250	0.447	259	0.131	0.338
5	2	0.500	0.707	14	0.214	0.426	14	0.143	0.363	227	0.172	0.378

Note: For the simultaneous treatment, the choice data at the information set level are implied by the contingent strategies. For instance, choosing the contingent strategy “claim at period 4” implies that the subject will wait from period 1 to period 3 and claim in period 4.

To compute the measure of violation of invariance under strategic equivalence of each payoff structure $(\delta, \bar{\alpha})$, I run the following regression on the data of information set $(2, X)$:

$$\mathbb{1}\{\text{claim}\}_i = \alpha_0 + \alpha_1 \mathbb{1}\{\text{sequential}\}_i + \epsilon_i, \quad (\text{D.7})$$

where $\mathbb{1}\{\text{claim}\}_i$ is the dummy variable for player i choosing C and $\mathbb{1}\{\text{sequential}\}_i$ is the dummy variable for the sequential treatment. Table D.6 reports the results for

Table D.5: The Empirical Frequencies of C at $(2, X)$ for Different Games.

$(\delta, \bar{\alpha})$	Sequential Treatment			Simultaneous Treatment			
	Obs	Claim %	s.d.	Obs	Claim %	s.d.	s.d.
Diagnostic Games							
(0.60, 0.45)	39	0.564	0.502	78	0.256	0.439	
(0.95, 0.80)	36	0.667	0.479	59	0.237	0.429	
Control Games							
(0.60, 0.80)	38	0.789	0.413	61	0.361	0.484	
(0.80, 0.45)	35	0.543	0.505	67	0.134	0.344	
(0.80, 0.80)	24	0.542	0.509	63	0.190	0.396	
(0.95, 0.45)	38	0.474	0.506	76	0.171	0.379	

all payoff structures. The standard errors are clustered at the session level.

Table D.6: The Magnitude of the Treatment Effect for Different Games.

Payoff Structure						
$(\delta, \bar{\alpha})$	(0.60, 0.45)	(0.60, 0.80)	(0.80, 0.45)	(0.80, 0.80)	(0.95, 0.45)	(0.95, 0.80)
Sequential Treatment	0.308*** (0.066)	0.429* (0.160)	0.409* (0.160)	0.351** (0.101)	0.303** (0.073)	0.429** (0.110)
Constant	0.256*** (0.057)	0.361*** (0.070)	0.134*** (0.018)	0.190*** (0.025)	0.171** (0.049)	0.237** (0.062)
N	117	99	102	87	114	95
R-squared	0.0914	0.1744	0.1889	0.1203	0.1028	0.1808

Note: The standard errors are clustered at the session level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Finally, Table D.7 and D.8 report the distributions of reaction times when players see a dirty face in the sequential and simultaneous treatments, respectively.

Table D.7: Reaction Times (seconds) when Seeing X in the Sequential Treatment.

Periods	Obs	Mean	s.d.	Q1	Median	Q3
1	572	11.29	9.845	5.359	7.574	13.72
2	210	8.172	6.531	4.495	6.106	9.745
3	34	7.530	3.816	4.785	6.277	11.44
4	21	6.901	4.135	3.150	6.299	9.767
5	14	7.663	5.597	3.231	6.677	8.856
All	851	10.20	8.917	5.000	7.152	12.08

Table D.8: Reaction Times (seconds) when Seeing X in the Simultaneous Treatment.

Stopping Strategies	Obs	Mean	s.d.	Q1	Median	Q3
Claim at 1	144	11.93	8.804	5.811	8.958	15.91
Claim at 2	90	13.85	10.18	7.191	11.32	16.62
Claim at 3	54	17.64	12.81	8.095	14.30	24.61
Claim at 4	34	21.15	15.16	12.32	16.46	22.96
Claim at 5	39	23.34	14.43	13.84	19.76	27.83
Always Wait	187	14.24	11.09	7.342	10.48	20.24
All	548	15.04	11.58	7.413	11.32	19.48

D.5.2 Likelihood Functions

Quantal Cursed Sequential Equilibrium

The “Quantal Cursed Sequential Equilibrium (QCSE)” is a model applicable to multi-stage games with observed actions. This model relaxes both the requirements of best responses and Bayesian inference. Specifically, QCSE is a hybrid model, combining the Agent Quantal Response Equilibrium (AQRE) proposed by McKelvey and Palfrey (1998) and the Cursed Sequential Equilibrium introduced by Fong, Lin, and Palfrey (2023).

Consider an assessment (μ, σ) . For any player i and any history h^{t-1} , the *average behavioral strategy profile of $-i$* is defined as

$$\bar{\sigma}_{-i}(a_{-i}^t | h^{t-1}, \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} \mu_i(\theta_{-i} | h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^t | h^{t-1}, \theta_{-i}).$$

In QCSE, players have incorrect perceptions about the behavioral strategies of other players. Instead of thinking they are using σ_{-i} , a χ -cursed type θ_i player i would believe the other players are using a χ -weighted average of the average behavioral strategy and the true behavioral strategy:

$$\sigma_{-i}^\chi(a_{-i}^t | h^{t-1}, \theta_{-i}, \theta_i) = \chi \bar{\sigma}_{-i}(a_{-i}^t | h^{t-1}, \theta_i) + (1 - \chi) \sigma_{-i}(a_{-i}^t | h^{t-1}, \theta_{-i}).$$

The beliefs of player i about θ_{-i} in QCSE are updated via Bayes’ rule, whenever possible, assuming other players are using the χ -cursed behavioral strategy rather than the true behavioral strategy. This updating rule is called the *χ -cursed Bayes’ rule*. Specifically, an assessment satisfies the χ -cursed Bayes’ rule if the belief system is derived from the Bayes’ rule while perceiving others are using σ_{-i}^χ rather than σ_{-i} .

Consider any totally mixed strategy profile $\sigma \in \Sigma^0$. As shown by Fong, Lin, and Palfrey (2023), if the belief system μ is derived from the χ -cursed Bayes' rule, then player i 's cursed belief is simply a linear combination of player i 's cursed belief at the beginning of that stage (with χ weight) and the Bayesian posterior belief (with $1 - \chi$ weight). That is, for any $h^t = (h^{t-1}, a^t)$,

$$\mu_i(\theta_{-i}|h^t, \theta_i) = \chi \mu_i(\theta_{-i}|h^{t-1}, \theta_i) + (1 - \chi) \left[\frac{\mu_i(\theta_{-i}|h^{t-1}, \theta_i) \sigma_{-i}(a^t|h^{t-1}, \theta_{-i})}{\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \sigma_{-i}(a^t|h^{t-1}, \theta'_{-i})} \right].$$

For any player i , any $\chi \in [0, 1]$, $\sigma \in \Sigma^0$, and any $\theta \in \Theta$, let $\rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi, \sigma_i)$ be i 's perceived conditional realization probability of terminal history $h^T \in \mathcal{H}^T$ at history $h^t \in \mathcal{H} \setminus \mathcal{H}^T$ if the type profile is θ and i uses the behavioral strategy σ_i whereas perceives other players' using the cursed behavioral strategy σ_{-i}^χ . At every non-terminal history h^t , a χ -cursed player in QCSE will use χ -cursed Bayes' rule to derive the posterior belief about the other players' types. Accordingly, a type θ_i player i 's conditional expected payoff at history h^t is

$$\bar{u}_i(\sigma|h^t, \theta_i) \equiv \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i(\theta_{-i}|h^t, \theta_i) \rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi, \sigma_i) u_i(h^T, \theta_{-i}, \theta_i).$$

Moreover, let $\bar{u}_i(a, \sigma|h^t, \theta_i)$ be the conditional expected payoff of player i of using $a \in A_i(h^t)$ with probability one, and using σ_i elsewhere.

In QCSE, there is a parameter $\lambda \in [0, \infty)$ that governs the precision of choices. Given an assessment (μ, σ) where $\sigma \in \Sigma^0$ and μ is derived from the χ -cursed Bayes' rule, if (μ, σ) is a QCSE for any player i , history h^t , and type θ_i , then type θ_i player i will have choice probabilities at h^t that follow a multinomial logit distribution. In particular, the probability of player i choosing $a \in A_i(h^t)$ is

$$\frac{e^{\lambda \bar{u}_i(a, \sigma|h^t, \theta_i)}}{\sum_{a' \in A_i(h^t)} e^{\lambda \bar{u}_i(a', \sigma|h^t, \theta_i)}}.$$

In summary, for each $\lambda \in [0, \infty)$ and $\chi \in [0, 1]$, an assessment (μ, σ) is a QCSE if

1. The belief system is derived from the χ -cursed Bayes' rule, and
2. For any player i , type θ_i , history h^t and $a \in A_i(h^t)$,

$$\sigma_i(a|h^t, \theta_i) = \frac{e^{\lambda \bar{u}_i(a, \sigma|h^t, \theta_i)}}{\sum_{a' \in A_i(h^t)} e^{\lambda \bar{u}_i(a', \sigma|h^t, \theta_i)}}.$$

When estimating QCSE, constructing the likelihood function follows a similar process as described in Appendix D.4.2. For each information set \mathcal{I}_i , QCSE uniquely predicts the choice probability of each a_i , denoted as $\bar{Q}(a_i|\mathcal{I}_i, \lambda, \chi)$, given λ and χ . The log-likelihood function can be formed by aggregating over every subject i , action a_i , and information set \mathcal{I}_i :

$$\ln L^{\bar{Q}}(\lambda, \chi) = \sum_i \sum_{\mathcal{I}_i \in \Pi_i} \sum_{a_i \in A_i(\mathcal{I}_i)} \mathbb{1}\{a_i, \mathcal{I}_i\} \ln [\bar{Q}(a_i|\mathcal{I}_i, \lambda, \chi)].$$

Quantal Dynamic Cognitive Hierarchy Solution

The ‘‘Quantal Dynamic Cognitive Hierarchy Solution (QDCH)’’ is a natural extension of DCH, where all strategic levels of players make quantal responses instead of best responses. In particular, following previous notations, for any $i \in N$, $\tau_i \geq 1$, $\theta \in \Theta$, σ , and τ_{-i} such that $\tau_j < \tau_i$ for any $j \neq i$, let $P_i^{\tau_i}(h^T|h^{t-1}, \theta, \tau_{-i}, \sigma_{-i}^{<\tau_i}, \sigma_i^{\tau_i})$ be level τ_i player i 's belief about the conditional realization probability of $h^T \in \mathcal{H}^T$ at history $h^{t-1} \in \mathcal{H} \setminus \mathcal{H}^T$ if the type profile is θ , the level profile is τ , and player i uses $\sigma_i^{\tau_i}$. In this case, level τ_i player i 's expected payoff at any $h^t \in \mathcal{H} \setminus \mathcal{H}^T$ is

$$\bar{u}_i^{\tau_i}(\sigma|h^t, \theta_i) \equiv \sum_{h^T \in \mathcal{H}^T} \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{\{\tau_{-i}: \tau_j < \tau_i \forall j \neq i\}} \mu_i^{\tau_i}(\tau_{-i}, \theta_{-i} | h^t, \theta_i) P_i^{\tau_i}(h^T|h^t, \theta, \tau_{-i}, \sigma_{-i}^{<\tau_i}, \sigma_i^{\tau_i}) u_i(h^T, \theta_{-i}, \theta_i).$$

Similar to QCSE, in QDCH, there is a parameter $\lambda \in [0, \infty)$ that governs the precision of choices. Let $\bar{u}_i^{\tau_i}(a, \sigma|h^t, \theta_i)$ be the conditional expected payoff of level τ_i player i of using $a \in A_i(h^t)$ with probability one, and using $\sigma_i^{\tau_i}$ elsewhere. In QDCH, players' choice probabilities follow multinomial logit distributions. That is, in QDCH, the probability of level τ_i player i choosing $a \in A_i(h^t)$ at history h^t is

$$\sigma_i^{\tau_i}(a|h^t, \theta_i) = \frac{e^{\lambda \bar{u}_i^{\tau_i}(a, \sigma|h^t, \theta_i)}}{\sum_{a' \in A_i(h^t)} e^{\lambda \bar{u}_i^{\tau_i}(a', \sigma|h^t, \theta_i)}}.$$

When estimating QDCH, I assume the prior distribution of levels follows Poisson(τ). At any information set \mathcal{I}_i , let $f(k|\mathcal{I}_i, \lambda, \tau)$ be the posterior distribution of levels at information set \mathcal{I}_i given λ and τ . In this case, the predicted choice probability for a_i at \mathcal{I}_i is the aggregation of quantal responses from all levels weighted by the proportion $f(k|\mathcal{I}_i, \lambda, \tau)$:

$$\bar{D}(a_i|\mathcal{I}_i, \lambda, \tau) \equiv \sum_{k=0}^{\infty} f(k|\mathcal{I}_i, \lambda, \tau) P_k(a_i|\mathcal{I}_i, \lambda, \tau),$$

where $P_k(a_i|\mathcal{I}_i, \lambda, \tau)$ is the probability of level k players choosing a_i at \mathcal{I}_i . Consequently, the log-likelihood function for QDCH can be formed by aggregating over every subject i , actions a_i and information set \mathcal{I}_i :

$$\ln L^{\bar{D}}(\lambda, \tau) = \sum_i \sum_{\mathcal{I}_i \in \Pi_i} \sum_{a_i \in A_i(\mathcal{I}_i)} \mathbb{1}\{a_i, \mathcal{I}_i\} \ln [\bar{\mathcal{D}}(a_i|\mathcal{I}_i, \lambda, \tau)],$$

where $\mathbb{1}\{a_i, \mathcal{I}_i\}$ is the indicator function which is 1 when subject i chooses a_i at \mathcal{I}_i .

D.6 Experimental Instructions

D.6.1 Sequential Treatment

General Instructions

Thank you for participating in the experiment. You are about to take part in a decision-making experiment, in which your earnings will depend partly on your decisions, partly on the decision of others, and partly on chance.

The entire session will take place through computer terminals, and all interactions between participants will be conducted through the computers. Please do not talk or in any way try to communicate with other participants during the session.

The main task of the experiment consists of 12 matches. Before the main task, you will be asked to complete some comprehension questions. If you have any questions, please raise your hand and the question will be answered so that everyone can hear.

In this experiment, you will earn “points” in each match. Your earnings will be determined by the total points you earn in the 12 matches. Each point has a value of \$0.02. That is, every 100 points generates \$2 in earnings for you. In addition to your earnings from decisions, you will receive a show-up fee of \$10. At the end of the experiment, your earnings will be rounded up to the nearest dollar amount. All your earnings will be paid in cash privately at the end of the experiment.

Main Task

1. In this experiment, you will be asked to make decisions in 12 matches. You will be randomly matched with another participant into a group for every separate match. This random pairing changes in every match.
2. Each match in this experiment corresponds to a game with the following rules.

- At the beginning of each match, each of you and the other participant will be randomly assigned a “color” (either Red or White). After the colors are assigned, you will be able to see the color of the other participant who is paired with you. However, you cannot see your own color!
- There are 3 possible situations, and the probabilities of these situations are summarized in the following table.

Situations	Probabilities
You are Red and the other participant is White.	p
You are White and the other participant is Red.	p
Both of you are Red.	$1 - 2p$

In other words, there is always at least one Red participant among each group.

- Each match is played in rounds. In each match, there are at most 5 rounds. Your color and the other participant’s color are fixed in the match. In each round, you and the other participant will simultaneously choose either “I’m Red” or “wait.” If both participants choose “wait,” then the match will continue to the next round. The match will end:
 - (1) after round 5; or
 - (2) after some round where there is at least one participant choosing “I’m Red.”

This round is called the “terminal round.” Your payoff for this match depends on which round the terminal round is, your action in the terminal round, and your color. Important: your payoff does not depend on the other participant’s color.

- Payoffs:
 - (1) If you choose “wait” in the terminal round, you will get 0 points for this match regardless of your color.
 - (2) If you choose “I’m Red” in the terminal round, your payoff for this match depends on which round the terminal round is and your own color. The payoffs are summarized in the following table. Notice that in each match, you and the other participant will face the same payoff table.

Terminal Round	1	2	3	4	5
Your payoff if your color is Red	p_1	p_2	p_3	p_4	p_5
Your payoff if your color is White	$-w_1$	$-w_2$	$-w_3$	$-w_4$	$-w_5$

(3) Examples:

- a. If you choose “I’m Red” in round 3, you will earn p_3 points if your color is Red and $-w_3$ if your color is White.
- b. If you choose “wait” in round 4 and the other participant chooses “I’m Red” in the same round, you will get 0 points regardless of your color.

3. Decisions:

- After observing the other participant’s color, you and the other participant matched with you will play the game according to the rules described above.
- Therefore, your payoffs are summarized as below.

		You choose “I’m Red” in the terminal round		You choose “wait” in the terminal round
Terminal Round	Your color	Red	White	Red or White
1		p_1	$-w_1$	0
2		p_2	$-w_2$	0
3		p_3	$-w_3$	0
4		p_4	$-w_4$	0
5		p_5	$-w_5$	0

- Each match starts from Round 1. You will make your decision in the following screen.

This is Round 1

	1	2	3	4	5
You					
The other one					

Your action

Wait I’m Red

After you make your decision, the following would happen: If either you or the other participant chooses “I’m Red,” then this round is the terminal round, and your payoff is determined by your action in this round. However, if both

you and the other participant choose “wait,” the match continues to the next round, and you will make your decision in the following screen.

This is Round 2

	1	2	3	4	5
You	wait				
The other one	wait				

Your action

Wait I'm Red

Like the previous round, if either you or the other participant chooses “I’m Red,” the match will end after this round. Yet if both you and the other participant choose “wait,” the match will proceed to the next round.

- If the game proceeds to round 5, then the match will end after this round and your payoff is determined by your action (and color) in round 5.

4. At the end of each match, there will be a summary of the match which includes both of your colors, actions in each round (whenever applicable) and your own payoff for this match.

5. At the beginning of the experiment, you will start from 900 points. You will get paid in cash based on your total points earned from the 12 matches. If your total point is negative, you will only receive the show-up fee.

6. Important:

- After each match, you will be randomly paired with another participant in the next match.
- Your color and the other participant’s color will also be randomly re-drawn in each match. The colors in each match are independent of the colors in other matches.
- The probability distribution of colors and the payoff table will change in each match.

Please raise your hand if you have any questions. The question will be answered so that everyone can hear.

D.6.2 Simultaneous Treatment

General Instructions

Thank you for participating in the experiment. You are about to take part in a decision-making experiment, in which your earnings will depend partly on your decisions, partly on the decision of others, and partly on chance.

The entire session will take place through computer terminals, and all interactions between participants will be conducted through the computers. Please do not talk or in any way try to communicate with other participants during the session.

The main task of the experiment consists of 12 matches. Before the main task, you will be asked to complete some comprehension questions. If you have any questions, please raise your hand and the question will be answered so that everyone can hear.

In this experiment, you will earn “points” in each match. Your earnings will be determined by the total points you earn in the 12 matches. Each point has a value of \$0.02. That is, every 100 points generates \$2 in earnings for you. In addition to your earnings from decisions, you will receive a show-up fee of \$10. At the end of the experiment, your earnings will be rounded up to the nearest dollar amount. All your earnings will be paid in cash privately at the end of the experiment.

Main Task

1. In this experiment, you will be asked to make decisions in 12 matches. You will be randomly matched with another participant into a group for every separate match. This random pairing changes in every match.
2. Each match in this experiment corresponds to a game with the following rules.
 - At the beginning of each match, each of you and the other participant will be randomly assigned a “color” (either Red or White). After the colors are assigned, you will be able to see the color of the other participant who is paired with you. However, you cannot see your own color!
 - There are 3 possible situations, and the probabilities of these situations are summarized in the following table.

Situations	Probabilities
You are Red and the other participant is White.	p
You are White and the other participant is Red.	p
Both of you are Red.	$1 - 2p$

In other words, there is always at least one Red participant among each group.

- Each match is played in rounds. In each match, there are at most 5 rounds. Your color and the other participant's color are fixed in the match. In each round, you and the other participant will simultaneously choose either "I'm Red" or "wait." If both participants choose "wait," then the match will continue to the next round. The match will end:

- (1) after round 5; or
- (2) after some round where there is at least one participant choosing "I'm Red."

This round is called the "terminal round." Your payoff for this match depends on which round the terminal round is, your action in the terminal round, and your color. Important: your payoff does not depend on the other participant's color.

- Payoffs:
 - (1) If you choose "wait" in the terminal round, you will get 0 points for this match regardless of your color.
 - (2) If you choose "I'm Red" in the terminal round, your payoff for this match depends on which round the terminal round is and your own color. The payoffs are summarized in the following table. Notice that in each match, you and the other participant will face the same payoff table.

Terminal Round	1	2	3	4	5
Your payoff if your color is Red	p_1	p_2	p_3	p_4	p_5
Your payoff if your color is White	$-w_1$	$-w_2$	$-w_3$	$-w_4$	$-w_5$

- (3) Examples:

- a. If you choose "I'm Red" in round 3, you will earn p_3 points if your color is Red and $-w_3$ if your color is White.

- b. If you choose “wait” in round 4 and the other participant chooses “I’m Red” in the same round, you will get 0 points regardless of your color.

3. Decisions:

- Instead of playing the game round by round, after observing the other participant’s color, you and the other participant are asked to simultaneously choose a “plan” which describes how you would commit to play the game if the game were played round by round. After you and the other participant both submit the plans, the computer will implement the plans and your payoff is determined accordingly.
- Since the game ends after some participant chooses “I’m Red,” there are six possible plans corresponding to the earliest round you intend to choose “I’m Red” or “always wait.” Specifically, the six plans are listed below.
 - **“I’m Red in Round 1”** means you plan to choose “I’m Red” in Round 1.
 - **“I’m Red in Round 2”** means you plan to choose “wait” in Round 1 and choose “I’m Red” in Round 2.
 - **“I’m Red in Round 3”** means you plan to choose “wait” in Round 1 and Round 2 and choose “I’m Red” in Round 3.
 - **“I’m Red in Round 4”** means you plan to choose “wait” in Round 1 to Round 3 and choose “I’m Red” in Round 4.
 - **“I’m Red in Round 5”** means you plan to choose “wait” in Round 1 to Round 4 and choose “I’m Red” in Round 5.
 - **“Always wait”** means you plan to choose “wait” in Round 1 to Round 5.
- In each match, you will be asked to choose your plan in the following screen.

Your plan: I’m Red in Round 1 I’m Red in Round 2 I’m Red in Round 3
 I’m Red in Round 4 I’m Red in Round 5 always wait

- Therefore, your payoffs are summarized as below.

Terminal Round	Your color	You choose “I’m Red” no later than the other participant		You choose “I’m Red” later or choose “ always wait ”
		Red	White	Red or White
1		p_1	$-w_1$	0
2		p_2	$-w_2$	0
3		p_3	$-w_3$	0
4		p_4	$-w_4$	0
5		p_5	$-w_5$	0

4. At the end of each match, there will be a summary of the match which includes both of your colors, the terminal round, your action, and your own payoff for this match. If you choose “I’m Red” later or at the same round as the other participant, you will be informed the other participant’s exact plan. Otherwise, you will be told that the other participant is “later than you.”

5. At the beginning of the experiment, you will start from 900 points. You will get paid in cash based on your total points earned from the 12 matches. If your total point is negative, you will only receive the show-up fee.

6. Important:

- a. After each match, you will be randomly paired with another participant in the next match.
- b. Your color and the other participant’s color will also be randomly re-drawn in each match. The colors in each match are independent of the colors in other matches.
- c. The probability distribution of colors and the payoff table will change in each match.

Please raise your hand if you have any questions. The question will be answered so that everyone can hear.

D.6.3 Screenshots

Figures D.4 and D.5 show the actual screenshots of the sequential treatment, and Figures D.6 to D.8 display the actual screenshots of the simultaneous treatment. Notice that Figure D.7 represents the feedback screen of a player who selects “I’m

Red” earlier than the opponent, and Figure D.8 provides the perspective from the other player.

Match 1

	You	The Other Participant
You see this	??	Red

Attention:
 1. There is at least one **Red** participant in this group.
 2. The probabilities and the payoff table in each match are different.

	Your color: Red The other's: White	Your color: White The other's: Red	Your color: Red The other's: Red
Probability	1/4	1/4	1/2

Payoffs (in points):

Terminal Round	Your Color	You choose "I'm Red" in the terminal round		You choose "wait" in the terminal round
		Red	White	Red or White
1		100	-444	0
2		95	-422	0
3		90	-401	0
4		86	-381	0
5		81	-362	0

This is Round 1

	1	2	3	4	5
You					
The other one					

Your Action

Wait I'm Red [Confirm](#)

Figure D.4: The Decision Stage of the Sequential Treatment.

End of Match 1

	You	The Other Participant
The true colors	Red	Red

Payoffs (in points):

Terminal Round	Your Color	You choose "I'm Red" in the terminal round		You choose "wait" in the terminal round
		Red	White	Red or White
1		100	-444	0
2		95	-422	0
3		90	-401	0
4		86	-381	0
5		81	-362	0

Round	1	2	3	4	5	Colors
You	Wait	Wait	Red	--	--	Red
The other participant	Wait	Wait	Wait	--	--	Red

Your payoff in this match: **90 points** [Confirm](#)

Figure D.5: The Feedback Stage of the Sequential Treatment.

Match 1

	You	The Other Participant
You see this	??	Red

Attention:
 1. There is at least one **Red** participant in this group.
 2. The probabilities and the payoff table in each match are different.

	Your color: Red The other's: White	Your color: White The other's: Red	Your color: Red The other's: Red
Probability	1/4	1/4	1/2

Payoffs (in points):

Terminal Round	Your Color	You choose "I'm Red " no later than the other participant		You choose "I'm Red " later or choose "always wait"
		Red	White	Red or White
1		100	-444	0
2		95	-422	0
3		90	-401	0
4		86	-381	0
5		81	-362	0

Your plan: I'm **Red** in Round 1 I'm **Red** in Round 2 I'm **Red** in Round 3
 I'm **Red** in Round 4 I'm **Red** in Round 5 always wait

Confirm

Figure D.6: The Decision Stage of the Simultaneous Treatment.

End of Match 1

	You	The Other Participant
The true colors	Red	White

Payoffs (in points):

Terminal Round	Your Color	You choose "I'm Red " no later than the other participant		You choose "I'm Red " later or choose "always wait"
		Red	White	Red or White
1		100	-444	0
2		95	-422	0
3		90	-401	0
4		86	-381	0
5		81	-362	0

Terminal Round	The other's action	Your action	The other's color	Your color	Your payoff
3	Later than you	I'm Red in Round 3	White	Red	90

Confirm

Figure D.7: The Feedback Stage of the Simultaneous Treatment (1).

End of Match 1

	You	The Other Participant
The true colors	White	Red

Payoffs (in points):

Terminal Round	Your Color	You choose "I'm Red " no later than the other participant		You choose "I'm Red " later or choose "always wait"
		Red	White	Red or White
1		100	-444	0
2		95	-422	0
3		90	-401	0
4		86	-381	0
5		81	-362	0

Terminal Round	The other's action	Your action	The other's color	Your color	Your payoff
3	I'm Red in Round 3	always wait	Red	White	0

Confirm

Figure D.8: The Feedback Stage of the Simultaneous Treatment (2).

References

- Bayer, Ralph C and Mickey Chan (2007). *The dirty faces game revisited*. Tech. rep. University of Adelaide, School of Economics.
- Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong (2004). "A cognitive hierarchy model of games". In: *Quarterly Journal of Economics* 119.3, pp. 861–898.
- Fong, Meng-Jhang, Po-Hsuan Lin, and Thomas R Palfrey (2023). "Cursed sequential equilibrium". In: *arXiv preprint arXiv:2301.11971*.
- McKelvey, Richard D and Thomas R Palfrey (1998). "Quantal response equilibria for extensive form games". In: *Experimental Economics* 1.1, pp. 9–41.