

Three Essays on Information Economics

Thesis by
Qiaoxi Zhang

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

2016
Defended May 4, 2016

© 2016

Qiaoxi Zhang

ORCID: orcid.org/0000-0002-3139-7659

All rights reserved

ACKNOWLEDGEMENTS

ABSTRACT

The main theme of my thesis is how uncertainty affects behaviors. I explore how agents seek to resolve uncertainty in different environments. In Chapter 1, agents learn from the messages of informed experts in a signaling game. In Chapter 2, an agent learns about a fixed and uncertain physical environment through dynamic experimentation. In the last chapter, agents learn about others' preferences through the outcome of a central matching mechanism.

Motivated by the question of how opposing political candidates who are policy experts can communicate to voters in a way that helps them win the election, I study a delegation problem with two informed, self-interested agents. Agents make proposals before the decision maker decides to whom to delegate a task. The innovation is that there are multiple issues that the principal and agents care about, and the agents can be vague about any issue in their proposals. Intuition says that agents should be specific about the issues that they are trusted on and vague about other issues. I find the opposite: an agent is disadvantaged by revealing information about certain issues to the decision maker, those on which he is trusted by the principal on. The reason is that doing so enables his opponent to take advantage of this revealed information and undercut him. Essentially, when the principal is on an agent's side for some issue, that agent does not want to be specific, because it creates a visible target for his opponent to react to. He wants to be vague, because that allows the principal's ignorance about the optimal action create an insurmountable obstacle for his opponent. As a result, it is to an agent's advantage to be vague about the issue that he is trusted on.

The second chapter investigates the implication of biased updating in dynamic experimentation such as a firm's R&D process. People exhibit near miss effect during gambling. For example, if the first two wheels of a slot machine indicate a potential final outcome of jackpot but the last wheel indicates a loss, people are motivated to gamble more. An outcome that is close to a success but is still a failure is called a "near miss." In this chapter, I explain the near miss effect in a firm's repeated R&D process. There are two factors that sequentially affect the profitability of R&D, both of which are uncertain. First is whether the R&D team is skilled enough to make a technical breakthrough. If a breakthrough occurs, then a second factor comes into play, which is whether the market demand is high enough to make the product profitable. Moreover, good news for the first stage

is a prerequisite for learning about the second stage. In each one of the infinite periods, the decision maker of the firm decides whether to involve in risky R&D and observe whether the outcome is a failure (no breakthrough), a success (with breakthrough and high market demand), or a near miss (with breakthrough but low market demand). I assume that the decision maker of the firm learns about the skill of the team properly, but when she updates about the market demand, she updates incorrectly and overweighs her prior. In particular, her posterior about the market demand is a convex combination of her prior and the Bayesian posterior. This bias affects the relative updating of the two factors, which gives rise to the near miss effect: after a near miss is observed, the decision maker values doing R&D more than before although she has received no payoff.

I show that if the decision maker is sufficiently biased and overweighs her prior enough, then she exhibits the near miss effect. I also compare the near miss effect for decision makers with different degrees of biases. As it turns out, the more biased a decision maker is, the more severe she exhibits the near miss effect. However, given the decision maker's belief about the two factors, the more biased she is, the less she values R&D. Consequently, the value of R&D is highest for a Bayesian.

In the last chapter, I study how well a centralized matching mechanism works when agents do not know others' preferences. I consider a standard two-sided marriage matching problem, except that agents only know their own preferences. Roth (1989) proved by an example the non-existence of a mechanism with at least one stable equilibria. In his proof, an agent is allowed to report a preference that is realized with ex ante zero probability, which violates the setup of a Bayesian game. Instead, by restricting agents to report only preferences with positive realization probabilities, I show that Roth's result still holds. More interestingly, as long as agents are allowed to form blocking pairs after a matching outcome is announced, the final outcome is always stable with respect to the true preferences. This means that even when the mechanism fails to produce a stable outcome, it can still release enough information for agents to initialize a blocking pair, which induces a stable outcome.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	iv
Table of Contents	vi
List of Illustrations	vii
List of Tables	viii
Chapter I: Vagueness in Multidimensional Proposals	1
1.1 Introduction	1
1.2 The Model	6
1.3 Main Results	11
1.4 Refinement: Extended Intuitive Criterion	13
1.5 Discussion	16
Chapter II: Biased Belief Updating in Dynamic Experimentation	54
2.1 Introduction	54
2.2 The Model	60
2.3 Main Results	64
2.4 Application: Career and Job Mobility	66
2.5 Conclusion	68
Chapter III: Matching with Incomplete Information	80
3.1 Introduction	80
3.2 The Model	80
3.3 The Result	83
3.4 Discussion	86

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
1.1 Only vague about misaligned dimensions	3

LIST OF TABLES

<i>Number</i>		<i>Page</i>
3.1	State of information	82
3.2	$\omega = a$	85
3.3	$\omega = b$	85
3.4	$\omega = c$	85
3.5	$\omega = d$	85

Chapter 1

VAGUENESS IN MULTIDIMENSIONAL PROPOSALS

1.1 Introduction

This paper studies how competition shapes information revelation. Consider two self-interested agents, Agent 1 and Agent 2, who communicate to a decision maker (DM) about their future actions. Agents have private information about the consequences of actions and can strategically choose to be vague about their actions. When actions are multidimensional, whether and on which dimension to be vague are the subjects of this paper.

Delegation often involves multidimensional decisions. Consider a parent who chooses between two schools for her child. Many aspects of a school affect a child's well-being: the physical activity level, the social environment and the curriculum, etc. A school can reveal these information about itself, but may not credibly convey whether it is optimal for the child. For example, a school may announce that their students have PE classes twice a week. Without knowing the optimal frequency of PE classes for a child, a parent cannot evaluate how well the school does in terms of physical well-being. In this paper, I show that if both schools know the consequence of PE class frequency, then the school that is stronger in developing students' physical well-being has an incentive to be vague about their PE class frequency. The reason is, by making a specific announcement of its PE class frequency, it reveals the optimal PE class level to the parent. The weaker school can in turn promise an appropriate PE class level that makes it slightly better overall. In other words, by being specific, a school's advantage is undermined.

Here is the formal setup. Nature chooses a state of the world. Agent 1 and Agent 2 observe the state, then simultaneously announce proposals. The DM, who does not observe the state, tries to infer the state from the proposals, and selects one agent to implement the decision. The outcome of the decision depends on the state and determines everyone's welfare. Both the decision and the state are multidimensional. For each dimension, one agent has an advantage in that his interest is more closely aligned to the DM's than his opponent's.

Agents can choose their levels of commitment as well as actions to commit to. They are allowed to be vague about any dimension of their future action by sending a

null message. Consequently, they have full freedom to take any action if chosen to make the decision. However, they are bound to any specific, non-vague actions that they propose. Since a commitment is binding, it likely reflects an agent's private information about the state. On the contrary, vagueness gives an agent full freedom to implement his own ideal action without revealing their private information. As a result, commitment and information revelation go together.

I show that vagueness is a natural consequence of competition. Not only is vagueness sustainable in equilibrium, an agent's vagueness appears on the dimension that he has an advantage on. If he is vague about dimension 1, then the DM believes that he will implement his ideal action without herself knowing what that is. Since she is ignorant about the state, she is free to adopt any belief about it. In particular, she is free to believe that any specific action proposed by Agent 2 is Agent 2's own ideal action, which is worse than Agent 1's ideal action. Therefore, the DM's ignorance precludes any possible compromises by Agent 2 on his disadvantaged dimension. On the other hand, if Agent 1 is specific about dimension 1 and reveals the state, Agent 2 can then propose an action on dimension 1 which is less biased than his ideal action. Since the DM learns the optimal action from Agent 1, she realizes that Agent 2 is offering a compromise by comparing Agent 2's ideal action given the optimal action and Agent 2's proposed action. Therefore information about Agent 2's disadvantaged dimension allows him to make credible compromises and demonstrate that he is the better agent.

For the solution concept of this two-sender signaling game, I use a strengthening of the weak perfect Bayesian equilibrium. Apart from some regularity equilibrium conditions that simplify the analysis, the DM's belief needs to satisfy a sensible consistency condition. I define and characterize equilibria in Section 1.2 and 1.3.

Signaling games typically have multiple equilibria. The Intuitive Criterion (Cho and Kreps, 1987) is a standard equilibrium refinement for one-sender signaling games. As a side product, I develop a refinement for two-sender signaling games in the spirit of the Intuitive Criterion. Equilibria with vagueness occurring on agents' aligned dimensions survive this refinement (Section 1.4).

As robustness checks, I show in the Appendix II - IV how the results extend when I vary the state space, number of dimensions and the preference of the DM. I also explore the case in which agents can be partially vague and commit to a strict subset of the action space.

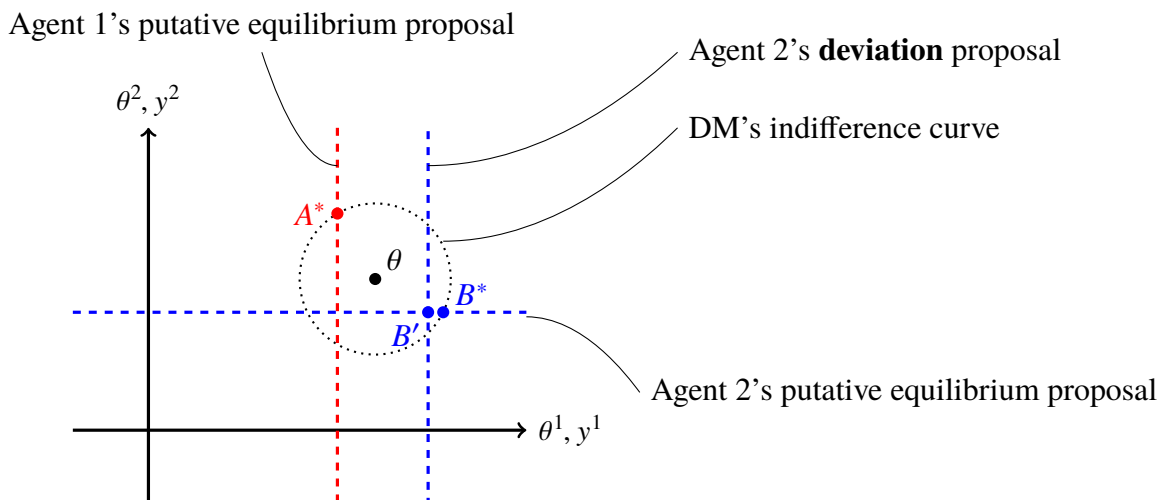


Figure 1.1: Only vague about misaligned dimensions

This mechanism through which competition shapes information revelation can be used to gain insights into electoral competition. There is active research on how parties choose which issues to campaign on as well as invest effort in if elected (Ash, Morelli, and Van Weelden, 2015; Egorov, 2015). In the language of this paper, campaign issue choice is a way for informed political parties to make policy commitments and I draw a connection between this choice and parties' interest alignment with voters. This paper is also related to ambiguity in political campaigns (Meirowitz, 2005; Alesina and Holden, 2008; Kamada and Sugaya, 2014; Kartik, Van Weelden, and Wolton, 2015). This literature focuses on the existence of vagueness when the state space is one-dimensional, whereas I study a multidimensional state space and where vagueness occurs.

An Example

To see why it may be unwise to commit to a specific action, consider a DM choosing an agent to make a two-dimensional decision $y = (y^1, y^2)$. I refer to the vector $\theta - y$ as the outcome, where $\theta \in \mathbb{R}^2$ is the state of the world that is unknown to the DM but known to the two agents. All players have quadratic utilities characterized by their ideal actions given the state. Figure 1.1 illustrates players' ideal actions for any given state θ . The DM's ideal action equals θ . Each agent has a constant bias. A^* and B^* are Agent 1 and Agent 2's ideal actions given θ , respectively. Note that the horizontal distance between A^* and θ is smaller than that between B^* and θ , and the vertical distance between B^* and θ is smaller than that between A^* and θ . Dimension 1 is then called Agent 1's *aligned dimension* and Agent 2's *misaligned*

dimension. The opposite is true for dimension 2.

Consider the putative equilibrium proposals by the agents as shown in Figure 1.1. Agent i 's proposal has a specific commitment to his ideal action on dimension i , but is vague about dimension j . Since he is free to implement any action for dimension j if chosen, the rational choice is to implement his ideal action. So overall he will choose his ideal action on each dimension. Since A^* and B^* are of equal distance to θ , the DM is indifferent between agents' ideal actions and assumed to randomize 50-50 between the agents. Lastly, notice that since each agent commits on his ideal action on dimension i , the DM learns θ^i from Agent i .

However, either agent has an incentive to deviate. Suppose that Agent 2 deviates as follows: instead of being vague about his misaligned dimension, now he is vague about his aligned dimension. The DM is surprised by the proposal profile she observes and her belief over Θ is unspecified. Moreover, for some beliefs, she prefers Agent 2 while the opposite is true for other beliefs. The key to decide which agent is better is to determine the sensible beliefs.

The key idea is that when Agent 2 deviates to make a surprising proposal, the DM should still believe in the information content of Agent 1's proposal. That is, she should believe that Agent 1 has not deviated and learns θ^1 from his proposal. Given this belief, which agent is better? Agent 2 is compromising on dimension 1 by committing to a less biased action. For dimension 2, since he is vague he implements his ideal action. So B' is the action Agent 2 has deviated to. As in equilibrium, Agent 1 will implement A^* . The DM then strictly prefers Agent 2. Since in equilibrium Agent 2 gets A^* half the time and B^* half the time, Agent 2 has made a profitable deviation.

Hence, revealing the state for the opponent's misaligned dimension creates a visible target for him to react to. Vagueness, however, allows the DM's belief to create an insurmountable obstacle for the opponent. Suppose instead that in equilibrium, Agent 1 is vague about both dimensions. Then any concession by Agent 2 is incredible, since nothing stops the DM from believing that he is proposing his own ideal action. For convenience, here I focus on a special case in which no agents have an overall advantage. In Section 1.3, I discuss a general-bias case while still retaining the feature that each agent has a advantage over a dimension. I also provide conditions for the existence of equilibria in which agents are vague about their aligned dimensions only.

Relation to Existing Literature

To my knowledge, this is the first paper to address competition for delegation in a multidimensional action and state space. Ambrus, Baranovskyi, and Kolb (2015), the first study of competition for delegation, uses a one-dimensional setup. Battaglini (2002) studies a multidimensional setting, but communication takes the form of cheap talk. Neither studies the role of vagueness in information disclosure.

My paper is connected to the disclosure literature (see Dranove and Jin (2010) for a survey). Here I briefly discuss the connection of my paper with Grossman (1981). He studies an environment in which a seller chooses whether to credibly disclose his quality and shows that all types of sellers fully disclose. The idea is that suppose two sellers of different qualities are pooled together. Then the buyer is only willing to pay the price as if the quality is an average of the two. Now the higher-quality of the two sellers can profit by disclosing his quality and receive a higher price, since the buyer now has a higher willingness to pay. The argument is based on the assumption that the seller could always have disclosed his exact quality. The counterpart of exact disclosure in my model is vagueness. An agent can always be vague and let the DM know the payoff from choosing him. The difference is that vagueness serves an additional purpose, which is masking the state of the world. This helps sustaining vagueness in equilibrium.

This paper is broadly related to work on competition in information revelation, which takes two forms: verifiable information revelation and cheap talk communication. The former (Gul and Pesendorfer, 2012; Gentzkow and Kamenica, 2015) assumes that informed parties cannot distort information; they can only decide how much information to provide. Cheap talk communication with multiple senders is investigated by Krishna and Morgan (2001) and Battaglini (2002). In particular, informativeness of cheap talk in electoral campaigns is studied by Schnakenberg (2014), Kartik, Squintani, and Tinn (2015) and Kartik and Van Weelden (2014). Blume, O. J. Board, and Kawamura (2007) and Blume and O. Board (2014) study vagueness in one-sender cheap talk in the form of noisy messages.

Another strand of related literature is delegation to informed, biased agents. Alonso and Matouschek (2008) study the DM's problem of how to optimally restrict the set of actions that an agent can take. Li and Suen (2004) discuss when a DM should delegate, and how biased the agent should be in order to be delegated. I focus on the communication from agent to the DM and assume that the only choice the DM can make is choosing between two agents. Ambrus, Baranovskyi, and Kolb

(2015)'s question is probably closest to mine. They investigate agents' proposals when a coarsely informed DM lacks the ability to measure the difference between two proposed actions. They demonstrate the existence of equilibrium in which agents distort their private signals. Distortion comes from the fact that the agents are imperfectly informed, while I study perfectly informed agents.

Lastly, this paper contributes to the equilibrium refinement literature. Refinement for one-sender signaling games is well-studied (Cho and Kreps, 1987; Banks and Sobel, 1987). Vida and Honryo (2015) study the implication of strategic stability (Kohlberg and Mertens, 1986) in general multi-sender signaling games. They find that stability implies a condition for out-of-equilibrium beliefs which states that the receiver rationalizes an out-of-equilibrium signal with the minimal number of deviations (Bagwell and Ramey, 1991). Schultz (1996) applies this belief restriction in a game where two informed parties decide how much public good to provide.

1.2 The Model

A decision affecting a decision maker (DM) and two agents needs to be made. The DM is unable to implement actions and has to delegate it to one of the two agents. Only the agents have the relevant information that pertains to the decision making — they observe a state of the world that determines the consequence of any given action. Afterwards, agents simultaneously announce proposals. The DM then chooses an agent to implement his proposal.

The state, denoted by $\theta = (\theta^1, \theta^2)$, concerns two dimensions: θ^1 is the state on dimension 1 and θ^2 on dimension 2. θ is distributed over $\Theta \equiv \mathbb{R}^2$ according to some continuous distribution function F with density f . f has full support on \mathbb{R}^2 . The space of proposals is $M = (\mathbb{R} \cup \{\emptyset\})^2$, where \mathbb{R} is the action space and \emptyset denotes a message indicating vagueness. For $i = 1, 2$, $m_i = (m_i^1, m_i^2)$ denotes Agent i 's proposal and $m = (m_1, m_2)$ denotes the proposal profile.

If Agent i is chosen, he implements an action $y_i = (y_i^1, y_i^2) \in \mathbb{R}^2$. For each dimension $k \in \{1, 2\}$, the set of feasible actions depends on m_i^k as follows. If Agent i is specific about dimension k and announces $m_i^k \in \mathbb{R}$, then he is committed to m_i^k and must choose $y_i^k = m_i^k$. If Agent i is vague about dimension k and announces $m_i^k = \emptyset$, then he is free to choose any action $y_i^k \in \mathbb{R}$.

The DM and the agents have conflicting interests. When the state is θ and the action is y , the DM's payoff is

$$u_d(\theta, y) = -(\theta^1 - y^1)^2 - (\theta^2 - y^2)^2.$$

Agent 1 and 2's payoffs are

$$\begin{aligned} u_1(\theta, y) &= -(\theta^1 - y^1 - b_1^1)^2 - (\theta^2 - y^2 - b_1^2)^2, \\ u_2(\theta, y) &= -(\theta^1 - y^1 - b_2^1)^2 - (\theta^2 - y^2 - b_2^2)^2, \end{aligned}$$

where $b_1 = (b_1^1, b_1^2)$ and $b_2 = (b_2^1, b_2^2)$ are agents' biases. Given the state θ , the DM's ideal action equals to the state. Agent i 's ideal action is equal to $\theta - b_i$. Throughout the paper I assume that $|b_1^1| < |b_2^1|$ and $|b_2^2| < |b_1^2|$. Since for dimension 1 Agent 1 is less biased than Agent 2, I call dimension 1 Agent 1's *aligned dimension* and Agent 2's *misaligned dimension*. The opposite is true for dimension 2. The rule of the game, all players' utility functions (and therefore their ideal actions given the state), and the state distribution are common knowledge.

An agent's proposal is sufficient for the DM to evaluate her expected payoff from choosing him. To see this, note that if an agent is specific on a dimension, he implements his proposed action. If he is vague on a dimension, he rationally takes his ideal action contingent on the state observed. Since the DM knows the biases of agents, her payoff from choosing this agent is also fixed. Therefore we can ignore agents' eventual action choices and focus on their proposals. A pure strategy for Agent i is a function $s_i : \Theta \rightarrow M$ mapping states into proposals. A strategy for the DM is a function $\beta : M \times M \rightarrow [0, 1]$ mapping proposal profiles into the probabilities that Agent 1 is chosen. The DM's posterior belief $\mu : M \times M \rightarrow \Delta(\Theta)$ maps each possible proposal profile m into a probability measure $\mu(\cdot | m)$ over Θ .

I use the solution concept of weak perfect Bayesian equilibrium (weak PBE) (s_1, s_2, β, μ) satisfying *single-deviation consistency*. This new consistency notion limits how the DM updates her belief about agents' strategies when she is surprised. After observing a surprising move by the agents, she updates her belief under the assumption that agents make strategic choices independently. As a result, knowing that one agent has deviated does not impact her belief about the other agent's strategy. To simplify the analysis, I focus on equilibria in which agents play pure strategies on the equilibrium path, the DM randomizes 50–50 whenever indifferent, agents' choices of vagueness do not depend on the state realization (the invariance condition) and are symmetric (the symmetry condition). The invariance condition says that if Agent i is vague (specific) about dimension k at some state $\bar{\theta}$, then he is vague (specific) about dimension k at all states. The symmetry condition says that if Agent i is vague about his misaligned (aligned) dimension, then Agent j is also vague about his misaligned (aligned) dimension, which is Agent i 's aligned (misaligned) dimension. Therefore two agents' choices of vagueness depend on

their comparative advantage in the same way. Note that I do not assume agents play symmetric strategies, only that the choice of vagueness is symmetric. Agents can vary their specific proposals differently according to the state. From now on, a weak PBE satisfying above requirements is abbreviated as “an equilibrium.” An equilibrium with vagueness is one in which vagueness is supported in equilibrium on some dimension by some agent.

Before formally defining the equilibrium, I first define single-deviation consistency. Given a weak PBE (s_1, s_2, β, μ) , the prior $F \in \Delta(\Theta)$ induces a probability distribution $G_1 \in \Delta(M)$ through s_1 , $G_2 \in \Delta(M)$ through s_2 , and $G_0 \in \Delta(M \times M)$ through s_1 and s_2 . Let $\text{supp}(G_k)$ denote the support of G_k , $k = 0, 1, 2$. Agent i 's proposal m_i is *consistent with equilibrium* if $m_i \in \text{supp}(G_i)$, *inconsistent with equilibrium* if otherwise. A proposal profile m is *on-path* if $m \in \text{supp}(G_0)$, *off-path* if otherwise.

Definition 1 (Single-deviation consistency) *Let (s_1, s_2, β, μ) be a weak PBE and $F \in \Delta(\Theta)$ the prior which induces $G_1, G_2 \in \Delta(M)$ and $G_0 \in \Delta(M \times M)$. The DM's belief μ satisfies single-deviation consistency if, for all $m = (m_1, m_2) \notin \text{supp}(G_0)$,*

$$\mu(m) \in \Delta(\{\theta \in \Theta \mid s_1(\theta) = m_1\} \cup \{\theta \in \Theta \mid s_2(\theta) = m_2\})$$

when $\{\theta \in \Theta \mid s_1(\theta) = m_1\} \cup \{\theta \in \Theta \mid s_2(\theta) = m_2\}$ is nonempty. Otherwise $\mu(m) \in \Delta(\Theta)$.

In other words, a DM faced with an off-path (m_1, m_2) believes that either Agent 1 or Agent 2 has not deviated whenever believing so is possible.

To understand the definition, let's divide the possible off-path $m = (m_1, m_2)$ into the following cases:

1. Neither m_1 or m_2 is consistent with equilibrium;
2. m_i is consistent with equilibrium but m_j is inconsistent with equilibrium;
3. Both m_1 and m_2 are consistent with equilibrium;

In Case 1, the DM learns that both agents have deviated. Since neither m_1 nor m_2 is consistent with equilibrium, both $\{\theta \in \Theta \mid s_1(\theta) = m_1\}$ and $\{\theta \in \Theta \mid s_2(\theta) = m_2\}$ are empty sets and $\mu(m)$ is unrestricted according to Definition 1. Since the DM's action following a bilateral deviation is irrelevant in sustaining the equilibrium, so is her belief.

In Case 2, only Agent j is inconsistent with equilibrium. The DM learns that Agent j has deviated. But the DM has the freedom to either believe that Agent i has also deviated or that Agent i is playing according to equilibrium. Since $\{\theta \in \Theta \mid s_i(\theta) = m_i\}$ is nonempty while $\{\theta \in \Theta \mid s_j(\theta) = m_j\}$ is empty, the DM believes that Agent i has not deviated.

In Case 3, both agents are consistent with equilibrium. To see how this can be possible for an off-path m , consider a putative equilibrium in which $s_1(\theta) = s_2(\theta) = \theta$ for all θ . Now suppose that at some state $\bar{\theta}$, Agent 1 deviates to the proposal $m'_1 = \underline{\theta} \neq \bar{\theta}$. The DM then observes a proposal profile $(\underline{\theta}, \bar{\theta})$. This is an off-path proposal profile because in equilibrium agents announce the same proposal. Since m is off-path, by definition at no θ are $s_1(\theta) = m_1$ and $s_2(\theta) = m_2$ both satisfied. So at least one agent has deviated. In other words, if the DM believes that Agent 1 has not deviated, then she believes that Agent 2 has deviated; if she believes that Agent 2 has not deviated, then she believes that Agent 1 has deviated. Since both $\{\theta \in \Theta \mid s_i(\theta) = m_i\}$ and $\{\theta \in \Theta \mid s_j(\theta) = m_j\}$ are nonempty, she can indeed believe that Agent 1 has not deviated (and so Agent 2 has deviated), or that Agent 2 has not deviated (and so Agent 1 has deviated). Since she cannot decide who the deviator is, her belief allows both cases.

To summarize, for all possible off-path proposal profiles that the DM may face following a unilateral deviation of an agent, the DM identifies the deviator according to single-deviation consistency. It is important to identify the deviator and non-deviator because the DM can rely on the information contained in the non-deviator's proposal and his equilibrium strategy to infer the state.

The notion of single-deviation consistency applies to general signaling games. Its intuition is as follows. If the DM believes that agents make independent strategic choices, then whether an agent has deviated should depend on that agent's proposal only and not his opponent's. Suppose that Agent i is inconsistent with equilibrium while Agent j is consistent with equilibrium. The DM learns that Agent i has deviated, but she should not use this fact as an excuse to change her belief about Agent j 's strategy. She should continue to believe that Agent j is playing his equilibrium strategy s_j . The single-deviation consistency shuts down the channel in which one agent's deviation impacts the DM's belief about the other agent. On the other hand, suppose that both agents are consistent with equilibrium. The single-deviation consistency implies that the DM imposes minimal departure from rationality to rationalize deviations and believes that only one agent is the deviator

whenever possible.

The notion that agents' strategic choices are independent is not new. Battigalli (1996) defines the independence property for conditional systems over the strategy profiles. An implication of independence is that the marginal conditional probabilities about player i 's strategies are independent of information which exclusively concerns player j 's strategies. It is shown that the independence property of conditional systems is necessary for an equivalent assessment to satisfy the consistency notion of sequential equilibrium (Kreps and Wilson, 1982).

Watson (2015) first formally defines perfect Bayesian equilibrium for infinite games without nature moves. The definition retains sequential rationality and puts forward a new notion for consistency, called "plain consistency." Under this framework, the DM's belief has two components: her belief over the strategies of agents and over the state. Since the latter is determined by the former and the proposal profile observed, we can focus on the belief over strategies. According to plain consistency, the DM assigns probability 1 to Agents playing their equilibrium strategies. When she reaches an off-path $m = (m_1, m_2)$ where m_2 is inconsistent with equilibrium and m_1 is consistent with equilibrium, she only alters her belief about Agent 2's strategy. Her belief about Agent 1 should remain as before and so concentrate on Agent 1's equilibrium strategy.

Now we have all the ingredients for the equilibrium definition:

Definition 2 (Equilibrium) *An equilibrium (s_1, s_2, β, μ) is a weak PBE in which*

1. $s_i : \Theta \rightarrow M, i = 1, 2;$
2. $\beta : M \times M \rightarrow \{0, \frac{1}{2}, 1\};$
3. μ satisfies single-deviation consistency;
4. (Invariance) For $i \in \{1, 2\}$ and $k \in \{1, 2\}$,
 - a) if $s_i^k(\bar{\theta}) = \emptyset$ for some $\bar{\theta}$, then $s_i^k(\theta) = \emptyset$ for all θ ;
 - b) if $s_i^k(\bar{\theta}) \neq \emptyset$ for some $\bar{\theta}$, then $s_i^k(\theta) \neq \emptyset$ for all θ .
5. (Symmetry) $\forall \theta$,
 - a) if $s_1(\theta) = (\emptyset, \emptyset)$, then $s_2(\theta) = (\emptyset, \emptyset)$;
 - b) if $s_1(\theta) = (\emptyset, w)$ for some $w \in \mathbb{R}$, then $s_2(\theta) = (z, \emptyset)$ for some $z \in \mathbb{R}$;

c) if $s_1(\theta) = (w, \emptyset)$ for some $w \in \mathbb{R}$, then $s_2(\theta) = (\emptyset, z)$ for some $z \in \mathbb{R}$;

Just to clarify, 4 and 5 are conditions on agents' equilibrium strategies, not the set of proposals that agents can deviate to. In particular, at any state, an agent is free to deviate to a proposal that is vague about any dimension regardless of his bias and the state realization. I investigate the consequence of dropping the symmetry condition in a smaller state space in Appendix II.

1.3 Main Results

What should agents be vague about? Common wisdom may suggest that agents should focus on the dimension that they have advantages on and ignore others. As a result, an agent should commit on their aligned dimension and be vague about their misaligned dimension. The model suggests otherwise. If Agent 1 is vague about dimension 1, then being vague is beneficial since the DM already trusts him on this dimension. More importantly, for any commitment made by Agent 2, the DM is free to believe that Agent 2 is committing on his own ideal action since she has arrived at an off-path information set. On the other hand, if Agent 1 is specific about dimension 1 and reveals θ^1 to the DM, Agent 2 can anchor on this revealed information and deviates by offering a compromise. The DM believes that only Agent 2 has deviated and continues to trust the information about θ^1 revealed by Agent 1. This way Agent 2 is able to credibly compromise on his misaligned dimension.

The results characterize equilibria in terms of where vagueness occurs in agents' proposals and construct an equilibrium. I first consider the case in which agents have zero biases on their aligned dimensions.

Proposition 1 *Suppose that $b_1^1 = b_2^2 = 0$. In all equilibria with vagueness, vagueness occurs on agents' aligned dimensions. An equilibrium in which vagueness only occurs on agents' aligned dimensions exists.*

All proofs are relegated to Appendix I. Here I demonstrate the intuition behind the proof. By Definition 2, in any equilibrium with vagueness, the agents' proposal profile must take one of the following forms:

1. $s_1(\theta) = s_2(\theta) = (\emptyset, \emptyset), \forall \theta$.
2. $s_1(\theta) \in \{\emptyset\} \times \mathbb{R}, s_2(\theta) \in \mathbb{R} \times \{\emptyset\}, \forall \theta$.

3. $s_1(\theta) \in \mathbb{R} \times \{\emptyset\}$, $s_2(\theta) \in \{\emptyset\} \times \mathbb{R}$, $\forall \theta$.

In both 1 and 2, vagueness occurs on agents' aligned dimensions. In 3, vagueness only occurs on agents' misaligned dimensions. I show that 3 cannot be an equilibrium proposal profile. Suppose first that both agents commit to their ideal actions on their aligned dimensions. The DM then learns θ^i from Agent i , $i = 1, 2$. Now suppose that some Agent i deviates to compromise on his misaligned dimension, i.e. propose an action preferred by the DM to his ideal action. By single-deviation consistency, the DM continues to believe the information revealed by m_j and therefore realizes that Agent i is making a compromise. This deviation is profitable as long as the compromise is small enough. The rest of the argument establishes that agents commit to their ideal actions on their aligned dimensions given that they have zero bias.

Although agents cannot be vague only about their misaligned dimensions, they can be vague only about their aligned dimensions. Consider the following putative equilibrium proposals for the agents: $s_1^*(\theta) = (\emptyset, \theta^2 - b_1^2)$, $s_2^*(\theta) = (\theta^1 - b_2^1, \emptyset)$ for all θ . In this equilibrium, each agent reveals the state for his misaligned dimension and implements his ideal action on each dimension. Consider a potential deviation by Agent i to make compromises on dimension i and propose an action strictly better for the DM than his ideal action. This deviation leads to an on-path proposal profile and the DM chooses his equilibrium action. Since Agent i deviates to a worse action but retains his equilibrium probability of winning, the deviation is unprofitable. The rest of the argument shows that even if an agent deviates to a compromise which is inconsistent with equilibrium (so that the DM realizes that he has deviated), the DM is free to believe that his action is his own ideal action and will not change his probability of winning.

Proposition 1 characterizes equilibria with vagueness and establishes its existence when agents have zero biases on their aligned dimensions. The next proposition does so under a more general bias structure and illustrates that the undercutting intuition still applies. With arbitrary bias, it is hard to demonstrate that the equilibrium proposals must be fully revealing regarding the dimension that it is specific about. However, starting with a fully-revealing equilibrium such as one in the following, it is easy to see that such strategies cannot be sustained in equilibrium.

Proposition 2 *Suppose that $|b_1^1| < |b_2^1|$ and $|b_2^2| < |b_1^2|$. Then both agents committing to their ideal actions on the aligned dimensions and being vague about their*

misaligned dimensions cannot be sustained in equilibrium. Both agents being vague about both dimensions can be sustained in equilibrium.

Consider a putative equilibrium proposal profile in which both agents are vague about all dimensions in all states: $s_1^*(\theta) = s_2^*(\theta) = (\emptyset, \emptyset), \forall \theta$. Suppose at some $\bar{\theta}$, Agent 1 deviates to be specific about any given dimension(s). The deviation leads to an off-path proposal profile, at which point the DM should maintain the belief that Agent 2 has not deviated but is free to choose any beliefs regarding Agent 1's strategy. One of the beliefs that can deter this deviation and therefore support (s_1^*, s_2^*) as part of an equilibrium is the belief that Agent 1 has deviated to his own ideal action. Under this belief, the DM maintains his equilibrium action. Since Agent 1's deviation action cannot be strictly preferred to his own ideal action and he gets the same probability of winning as in equilibrium, such a deviation is unprofitable.

1.4 Refinement: Extended Intuitive Criterion

In the last section, I showed that there exist equilibria in which agents are vague about their aligned dimensions. But this should not come as a surprise. In signaling games, there typically exist multiple equilibria supported by unreasonable beliefs. In this section, I show that the beliefs supporting the equilibria I identified are actually very reasonable in the sense that they satisfy a refinement similar to the Intuitive Criterion, adapted for two-sender games. The refinement combines the Intuitive Criterion for one-sender games with single-deviation consistency. For an off-path proposal profile, single-deviation consistency identifies the deviator. The Intuitive Criterion restricts what types of deviations the DM is allowed to believe in.

One-sender: Intuitive Criterion

The Intuitive Criterion for one-sender signaling games places restrictions on the receiver's beliefs after an unexpected message by the sender. The receiver is required to believe that the sender's private information is such that the highest payoff that the sender can get by deviating to the observed message is weakly higher than the sender's equilibrium payoff, given that the receiver does not react to the message with dominated actions.

First I review the idea of Intuitive Criterion. For consistency, I will keep using the terminologies and notations from Section 2.2 to describe both one-sender and two-sender signaling games and the equilibrium refinements. In a one-sender signaling game, given the state θ , the Agent's proposal $m \in M$, the DM's action $\beta \in B$, the

Agent's utility is denoted by $u_a(\theta, m, \beta)$ and the DM's utility $u_d(\theta, m, \beta)$. The DM's belief over Θ given m is denoted by $\mu(\cdot | m)$. Given m and μ , the set of the DM's best responses is:

$$\widetilde{\text{BR}}(\mu, m) = \arg \max_{\beta \in B} \int_{\Theta} u_d(\theta, m, \beta) d\mu(\theta | m).$$

For any nonempty $T \subset \Theta$ and $m \in M$,

$$\text{BR}(T, m) = \bigcup_{\mu: \mu(T|m)=1} \widetilde{\text{BR}}(\mu, m)$$

denotes the set of the DM's best responses according to beliefs supported on T .

When T is empty, let $\text{BR}(T, m) = \text{BR}(\Theta, m)$.

Given an equilibrium (s^*, β^*, μ^*) , the Agent's equilibrium payoff at θ is denoted by $u_a^*(\theta)$. For an off-path proposal m' , the set of states at which the Agent's highest payoff from deviating to m' given that the DM best responds to some belief is higher than the Agent's equilibrium payoff is

$$\Theta(m') = \{\theta | \max_{\beta \in \text{BR}(\Theta, m')} u_a(\theta, m', \beta) \geq u_a^*(\theta)\}.$$

Finally, an equilibrium (s^*, β^*, μ^*) *fails the Intuitive Criterion* if there exists $\theta \in \Theta$ and off-path $m' \in M$ such that

$$u_a^*(\theta) < \min_{\beta \in \text{BR}(\Theta(m'), m')} u_a(\theta, m', \beta).$$

When the DM observes an unexpected proposal m' , the support of her belief is restricted to be states at which m' is potentially profitable. That is, the highest payoff that the Agent can get given that the DM does not play dominated actions is weakly higher than his equilibrium payoff. The Agent then contemplates deviations given that the DM best responds to beliefs thus restricted. If for some θ and off-path proposal m' , any such best response makes the Agent strictly better off compared to in equilibrium when the state is θ , then the equilibrium fails the Intuitive Criterion.

Two-sender: Extended Intuitive Criterion

I extend the Intuitive Criterion to the two-sender case by combining single-deviation consistency with the Intuitive Criterion. Single-deviation consistency identifies the deviator. The Intuitive Criterion identifies the kind of deviations that the receiver is allowed to believe in.

In a two-sender signaling game, given the state $\theta \in \Theta$, agents' proposal profile $m = (m_1, m_2)$, and the DM's action $\beta \in B$, Agent i 's utility is denoted by $u_i(\theta, m, \beta)$ and the DM's utility $u_d(\theta, m, \beta)$ where $m = (m_1, m_2)$. Let $\mu(\cdot | m)$ denote the DM's belief over Θ conditional on m . Given an equilibrium $(s_1^*, s_2^*, \beta^*, \mu^*)$, Agent i 's

equilibrium payoff at θ is denoted by $u_i^*(\theta)$. Similar as before, G_0 is the distribution over $M \times M$ induced by the equilibrium strategies (s_1^*, s_2^*) and the state distribution F over Θ . G_i ($i = 1, 2$) is the distribution over M induced by s_i^* and F . I now introduce the equilibrium refinement.

Definition 3 (The Extended Intuitive Criterion) *Let $(s_1^*, s_2^*, \beta^*, \mu^*)$ be a weak PBE of a two-sender signaling game. For each $m = (m_1, m_2) \notin \text{supp}(G_0)$ and $i \in 1, 2$, form the set*

$$\Theta_i(m) = \left\{ \theta \mid s_j^*(\theta) = m_j, u_i^*(\theta) \leq \max_{\beta \in BR(\Theta, m)} u_i(\theta, s_j(\theta), m_i, \beta) \right\}.$$

$(s_1^, s_2^*, \beta^*, \mu^*)$ fails the Extended Intuitive Criterion if there exists $\theta \in \Theta$, $i \in \{1, 2\}$, $m'_i \in M$ such that*

$$u_i^*(\theta) < \min_{\beta \in BR(\Theta_1(m') \cup \Theta_2(m'), m')} u_i(\theta, m', \beta),$$

where $m' = (m'_i, s_j(\theta))$. An equilibrium satisfying the Extended Intuitive Criterion is called an intuitive equilibrium.

Whenever the DM faces an off-path proposal profile m' , her belief is restricted to be supported on $\Theta_1(m') \cup \Theta_2(m')$. $\Theta_i(m')$ is the set of states at which only Agent i is the deviator and his deviation is potentially profitable. In calculating the payoff from deviating, apart from assuming that the DM does not play dominated actions, the Agent also assumes that the other Agent plays according to equilibrium. Note that $\Theta_i(m)$ may be empty because $m_j \notin \text{supp}(G_j)$ or m_i is not potentially profitable. However, as long as m results from a unilateral deviation and the deviation is potentially profitable, $\Theta_1(m) \cup \Theta_2(m)$ is nonempty. An equilibrium fails the Extended Intuitive Criterion if at some state θ , some Agent i can profitably deviate to some m'_i as long as the DM best responds to restricted beliefs upon observing $(m'_i, s_j(\theta))$.

Robustness of Equilibria

Previously I have identified some equilibria in which agents are vague about their aligned dimensions. When $|b_1^1| < |b_2^1|$ and $|b_2^2| < |b_1^2|$,

$$s_1(\theta) = s_2(\theta) = (\emptyset, \emptyset), \forall \theta$$

can be sustained in equilibrium. As a special case, when $b_1^1 = b_2^2 = 0$

$$s_1(\theta) = (\emptyset, \theta^2 - b_1^2), s_2(\theta) = (\theta^1 - b_2^1, \emptyset), \forall \theta$$

can also be sustained in equilibrium. As it turns out, both of them are robust to the Extended Intuitive Criterion refinement.

Proposition 3 *Suppose that $|b_1^1| < |b_2^1|$ and $|b_2^2| < |b_1^2|$. Intuitive equilibria in which both agents are vague about their aligned dimensions exist. Furthermore if $b_1^1 = b_2^2 = 0$, in all intuitive equilibria with vagueness, vagueness occurs on agents' aligned dimensions.*

The equilibrium in which both agents are vague about both dimensions is supported by the following belief: whenever one of the agents deviates, he is committing to his own ideal action for the dimensions he is specific about. To show that this belief satisfies single-deviation consistency, I first show that for some $\beta \in \{0, \frac{1}{2}, 1\}$, committing to own ideal action gives the deviator higher payoff compared to equilibrium. Then I show that there is some belief $\mu \in \Delta(\Theta)$ according to which β is a best response, and hence β is not dominated.

Furthermore when $b_1^1 = b_2^2 = 0$, I show that the putative equilibria in which both agents are only vague about their misaligned dimensions do not satisfy the Extended Intuitive Criterion. In fact, if we require the DM's belief to satisfy single-deviation consistency, then any agent has a profitable deviation. For this deviation, all the beliefs satisfying single-deviation consistency in fact concentrate on the event that the deviator has made a potentially profitable deviation. Essentially, the Extended Intuitive Criterion does not put further restriction on beliefs apart from those imposed by single-deviation consistency. So the equilibria in which agents are only vague about their misaligned dimensions are not intuitive equilibria. This establishes that all the intuitive equilibria with vagueness has vagueness occurring on agents' aligned dimensions.

1.5 Discussion

One may be interested in knowing how allowing agents' vagueness impacts the DM's payoff. To illustrate this, consider a game which is identical to the setup so far except that agents are not allowed to be vague. That is, $M = \mathbb{R}^2$. There are two consequences of such a restriction. First, it is hard to sustain a simple fully-revealing equilibrium. For example, consider a putative equilibrium in which $s_1(\theta) = s_2(\theta) = \theta$, for all θ . This is in fact not a weak PBE because at any $\bar{\theta} \in \mathbb{R}^2$, there is $\tilde{\theta} \in \mathbb{R}^2$ such that exactly one of the two following events happens: either $\tilde{\theta}$ is a profitable deviation for Agent 1 at $\bar{\theta}$, or $\bar{\theta}$ is a profitable deviation for Agent 2 at $\tilde{\theta}$. The key reason is, whenever an agent deviates, the DM cannot tell who the deviator is. Therefore she must take the same action after seeing Agent 1 proposes $\tilde{\theta}$ and Agent 2 proposes $\bar{\theta}$. If her action is favorable towards Agent 1, then Agent 1

has an incentive to deviate; otherwise Agent 2 has an incentive to deviate.

Second, the DM can potentially get worse-off due to not allowing agents' vagueness. To see this, consider the following equilibrium: $s_1(\theta) = s_2(\theta) = (\bar{\theta}^1, \bar{\theta}^2)$ for all θ . The reason that both agents making the same constant proposal can be sustained in equilibrium is that, whenever an agent deviates to an alternative action on either dimension, the DM is free to believe that the optimal action is closer to the equilibrium action than the deviator's action. In this equilibrium, the DM gets a payoff that depends on the distance from the state realization and the equilibrium proposal.

Consider a pooling equilibrium in the with-vagueness setup: $s_1(\theta) = s_2(\theta) = (\emptyset, \emptyset)$ for all θ . This equilibrium is very similar to the constant equilibrium above since it is also supported by the belief that any deviation is no better than equilibrium for the DM. However in this equilibrium, the DM gets the less-biased agent's ideal action. If his bias is sufficiently small, then the all-vague equilibrium is better for the DM. This is the case because when agents are always vague, undercutting is essentially ruled out. Therefore agents can secure their ideal actions without worrying that their opponent will offer a more favorable action to the DM. The undercutting leads to an equilibrium in which both agents offer the optimal action, which turns out to be non-equilibrium as shown above. Vagueness protects agents from making concessions that lead to non-existence of equilibrium.

Appendix I.

Throughout the proof, I use $u_i^{dev}(\theta, m_i^{dev}, m_j)$ to denote Agent i 's payoff after deviating to m_i^{dev} at state θ , at which Agent j 's equilibrium proposal is m_j . π_i denotes the DM's expected payoff from choosing Agent i given her beliefs and agents' proposals. Given Agent i 's proposal m_i , $y_i(m_i)$ denotes Agent i 's policy. $u_i^*(\theta)$ denotes Agent i 's equilibrium payoff at state θ .

Lemma 1 *Suppose that the agents' strategies s_1, s_2 are as follows:*

$$s_1(\theta) = (w(\theta), \emptyset),$$

$$s_2(\theta) = (\emptyset, z(\theta)),$$

where $w, z : \Theta \rightarrow \mathbb{R}$. Then $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$ for all θ .

Proof. Suppose that $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 0$ for some $\bar{\theta}$. Then Agent 1 has a profitable deviation $s_2(\bar{\theta})$. To see this, first note that Agent 1's equilibrium payoff at $\bar{\theta}$ is

$$u_1^*(\bar{\theta}) = u_1(\bar{\theta}, y_2(\emptyset, z(\bar{\theta}))) = u_1((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1 - b_2^1, z(\bar{\theta}))).$$

Now let Agent 1 deviate to $s_2(\bar{\theta})$. Given $m = (m_1^{dev}, s_2(\bar{\theta})) = ((\emptyset, z(\bar{\theta})), (\emptyset, z(\bar{\theta})))$ and any $\mu \in \Delta(\Theta)$, $\pi_1 > \pi_2$. This is because $|b_1^1| < |b_2^1|$ and both agents propose the same action on dimension 2. Therefore $\beta(m) = 1$ and Agent 1's deviation payoff is

$$u_1^{dev}(\bar{\theta}, m_1^{dev}, s_2(\bar{\theta})) = u_1(\bar{\theta}, y_1(\emptyset, z(\bar{\theta}))) = u_1((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1 - b_1^1, z(\bar{\theta}))) > u_1^*(\bar{\theta}).$$

Similarly, if $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 1$ for some $\bar{\theta}$, then Agent 2 has a profitable deviation $s_1(\bar{\theta})$. To summarize, agents win with equal probability in equilibrium. ■

Proof of Proposition 1

The first part of the proof rules out agents' strategies s_1, s_2 as follows:

$$\begin{aligned} s_1(\theta) &= (w(\theta), \emptyset), \\ s_2(\theta) &= (\emptyset, z(\theta)), \end{aligned}$$

where $w, z : \Theta \rightarrow \mathbb{R}$. I prove this by contradiction by first characterizing w and z and then demonstrating incentives to deviate. Without loss of generality, I assume that $b_1 = (0, b_1^2)$ and $b_2 = (b_2^1, 0)$ where $\|b_1\| \leq \|b_2\|$.

Step 1. $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$ for all θ .

This follows from Lemma 1.

Step 2. $w(\theta) = \theta^1, z(\theta) = \theta^2, \forall \theta$.

Suppose that $w(\bar{\theta}) \neq \bar{\theta}^1$ for some $\bar{\theta}$. Then Agent 1 has a profitable deviation (\emptyset, \emptyset) .

To see this, first note that

$$u_1^*(\bar{\theta}) = \frac{1}{2}u_1(\bar{\theta}, y_1(w(\bar{\theta}), \emptyset)) + \frac{1}{2}u_1(\bar{\theta}, y_2(\emptyset, z(\bar{\theta}))) < \frac{1}{2}u_1(\bar{\theta}, y_2(\emptyset, z(\bar{\theta}))).$$

Now let Agent 1 deviate to $m_1^{dev} = (\emptyset, \emptyset)$. Given $m = (m_1^{dev}, s_2(\bar{\theta})) = ((\emptyset, \emptyset), (\emptyset, z(\bar{\theta})))$ and $\mu \in \Delta(\Theta)$,

$$\pi_1 = \int_{\Theta} u_d(\theta, \theta - b_1) \mu d\theta = -\|b_1\|^2,$$

$$\pi_2 = \int_{\Theta} u_d(\theta, (\theta^1 - b_2^1, z(\bar{\theta}))) \mu d\theta \leq \int_{\Theta} u_d(\theta, (\theta^1 - b_2^1, \theta^2)) \mu d\theta = -\|b_2\|^2 \leq \pi_1.$$

Therefore $\beta(m) \geq \frac{1}{2}$. Moreover, because for any θ and $z(\theta)$,

$$u_1(\theta, y_1(\emptyset, \emptyset)) > u_1(\theta, y_2(\emptyset, z(\theta))),$$

we have

$$\begin{aligned} u_1^{dev}(\bar{\theta}, m_1^{dev}, s_2(\bar{\theta})) &\geq \frac{1}{2}u_1(\bar{\theta}, y_1(\emptyset, \emptyset)) + \frac{1}{2}u_1(\bar{\theta}, y_2(\emptyset, z(\bar{\theta}))) \\ &= \frac{1}{2}u_1(\bar{\theta}, y_2(\emptyset, z(\bar{\theta}))) \\ &> u_1^*(\bar{\theta}). \end{aligned}$$

(\emptyset, \emptyset) is then a profitable deviation for Agent 1. Therefore $w(\theta) = \theta^1$ for all θ and $\pi_1 = -\|b_1\|^2$. Since $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$ for all θ , $\pi_2 = \pi_1 = -\|b_1\|^2$. This immediately rules out $\|b_1\| < \|b_2\|$ since

$$\pi_2 = \int_{\Theta} u_d(\theta, (\theta^1 - b_2^1, z))\mu d\theta \leq \int_{\Theta} u_d(\theta, (\theta^1 - b_2^1, \theta^2))\mu d\theta = -\|b_2\|^2 < -\|b_1\|^2.$$

Therefore s_1, s_2 can only be sustained in equilibrium when $\|b_1\| = \|b_2\|$. In this case, I use the same argument as above to show $z(\theta) = \theta^2$ for all θ .

We have finished characterizing w and z . Given that $w(\theta) = \theta^1, z(\theta) = \theta^2$, for all θ ,

$$u_1^*(\theta) = \frac{1}{2}u_1(\theta, \theta - b_1) + \frac{1}{2}u_1(\theta, \theta - b_2) = -\frac{1}{2}\|b_1 - b_2\|^2.$$

Step 3. At any $\bar{\theta}$, Agent 1 has a profitable deviation $(\emptyset, \bar{\theta}^2 - (1 - \epsilon)b_1^2)$, where $\epsilon > 0$ is small.

Suppose that Agent 1 deviates to said proposal at $\bar{\theta}$. The DM observes $m = (m_1^{dev}, s_2(\bar{\theta})) = ((\emptyset, \bar{\theta}^2 - (1 - \epsilon)b_1^2), (\emptyset, \bar{\theta}^2))$ and believes that only Agent 1 has deviated. Therefore $\mu(m) \in \Delta(T(m))$, where

$$T(m) = \{\theta \mid s_2(\theta) = (\emptyset, \bar{\theta}^2)\} = \{\theta \mid \theta^2 = \bar{\theta}^2\}.$$

For any $\mu \in \Delta(T(m))$,

$$\begin{aligned} \pi_2 &= \int_{\Theta} u_d(\theta, \theta - b_2)\mu d\theta = -\|b_2\|^2 = -\|b_1\|^2, \\ \pi_1 &= \int_{\Theta} u_d(\theta, (\theta^1, \bar{\theta}^2 - (1 - \epsilon)b_1^2))\mu d(\theta) \\ &= \int_{\Theta} u_d(\theta, (\theta^1, \theta^2 - 1 + \epsilon))\mu d\theta = -\|(1 - \epsilon)b_1\|^2 > \pi_2. \end{aligned}$$

Therefore $\beta(m) = 1$ and for ϵ small enough.

$u_1^{dev}(\bar{\theta}, m_1^{dev}, s_1(\bar{\theta})) = u_1(\bar{\theta}, (\bar{\theta}^1, \bar{\theta}^2 - (1 - \epsilon)b_1^2)) = \|(0, \epsilon b_1^2)\|^2 = -\epsilon^2(b_1^2)^2 > u_1^*(\bar{\theta})$, making $(\emptyset, \bar{\theta}^2 - (1 - \epsilon)b_1^2)$ a profitable deviation.

The second part of the proof establishes

$$\begin{aligned} s_1(\theta) &= (\emptyset, \theta^2 - b_1^2), \\ s_2(\theta) &= (\theta^1 - b_2^1, \emptyset), \end{aligned}$$

as an equilibrium. The DM's belief and strategy are as follows:

1. For any $m = ((\emptyset, w), (z, \emptyset))$ where $w, z \in \mathbb{R}$, $\mu(z + b_2^1, w + b_1^2) = 1$.

$$\beta(m) = \begin{cases} \frac{1}{2} & \text{if } \|b_1\| = \|b_2\|, \\ 0 & \text{if } \|b_1\| > \|b_2\|, \\ 1 & \text{if } \|b_1\| < \|b_2\|. \end{cases}$$

2. For any $m = ((\emptyset, \emptyset), (z, \emptyset))$ where $z \in \mathbb{R}$, $\mu\{\tilde{\theta} \mid \tilde{\theta}^1 = z + b_2^1\} = 1$ and $\beta(m)$ is same as 1.

3. For any $m = ((w, \emptyset), (z, \emptyset))$ where $w, z \in \mathbb{R}$, $\mu\{\tilde{\theta} \mid \tilde{\theta}^1 = z + b_2^1\} = 1$.

$$\beta(m) = \begin{cases} \frac{1}{2} & \text{if } \|(z + b_2^1 - w, b_1^2)\| = \|b_2\|, \\ 0 & \text{if } \|(z + b_2^1 - w, b_1^2)\| > \|b_2\|, \\ 1 & \text{if } \|(z + b_2^1 - w, b_1^2)\| < \|b_2\|. \end{cases}$$

4. For any $m = ((q, w), (z, \emptyset))$ where $q, w, z \in \mathbb{R}$, $\mu\{\tilde{\theta} \mid \tilde{\theta}^1 = z + b_2^1, \tilde{\theta}^2 = w + b_1^2\} = 1$.

$$\beta(m) = \begin{cases} \frac{1}{2} & \text{if } \|(z + b_2^1 - q, b_1^2)\| = \|b_2\|, \\ 0 & \text{if } \|(z + b_2^1 - q, b_1^2)\| > \|b_2\|, \\ 1 & \text{if } \|(z + b_2^1 - q, b_1^2)\| < \|b_2\|. \end{cases}$$

5. For any $m = ((\emptyset, w), (\emptyset, \emptyset))$ where $w \in \mathbb{R}$, $\mu\{\tilde{\theta} \mid \tilde{\theta}^2 = w + b_1^2\} = 1$ and $\beta(m)$ is same as 1.

6. For any $m = ((\emptyset, w), (\emptyset, z))$ where $w, z \in \mathbb{R}$, $\mu\{\tilde{\theta} \mid \tilde{\theta}^2 = w + b_1^2\} = 1$.

$$\beta(m) = \begin{cases} \frac{1}{2} & \text{if } \|b_1\| = \|(b_2^1, w + b_1^2 - z)\|, \\ 0 & \text{if } \|b_1\| > \|(b_2^1, w + b_1^2 - z)\|, \\ 1 & \text{if } \|b_1\| < \|(b_2^1, w + b_1^2 - z)\|. \end{cases}$$

7. For any $m = ((\emptyset, w), (q, z))$ where $w, q, z \in \mathbb{R}$, $\mu(q + b_2^1, w + b_1^2) = 1$.

$$\beta(m) = \begin{cases} \frac{1}{2} & \text{if } \|b_1\| = \|(b_2^1, w + b_1^2 - z)\|, \\ 0 & \text{if } \|b_1\| > \|(b_2^1, w + b_1^2 - z)\|, \\ 1 & \text{if } \|b_1\| < \|(b_2^1, w + b_1^2 - z)\|. \end{cases}$$

8. For any other m , $\mu \in \Delta(\Theta)$ and $\beta(m) \in \widetilde{\text{BR}}(\mu, m)$.

$\mu(\cdot \mid m)$ satisfies single-deviation consistency for all m . For m as the result of a unilateral deviation by Agent i such that $m_i \notin \text{supp}(G_i)$, the DM believes that Agent

j has not deviated. m for which $m_i \in \text{supp}(G_i)$ for $i = 1, 2$ while $m \notin \text{supp}(G_0)$ does not exist in this equilibrium since $\Theta = \mathbb{R}^2$.

No agent has an incentive to deviate. Suppose $\|b_1\| > \|b_2\|$. In equilibrium Agent 2 wins with his ideal action, so he has no incentive to deviate. For Agent 1, the possible deviations are as follows:

- (a) If Agent 1 deviates to (\emptyset, w) where $w \neq \theta^2 - b_1^2$, the DM does not detect the deviation and continues to randomize 50-50. Therefore his probability of winning remains unchanged and he gets a worse outcome than his winning outcome in equilibrium when he wins.
- (b) If Agent 1 deviates to (\emptyset, \emptyset) , his probability of winning remains unchanged and he gets the same winning outcome as that in equilibrium.

- (c) If Agent 1 deviates to (q, \emptyset) where $q \in \mathbb{R}$. Since

$$\pi_1 = \int_{\Theta} u_d(\theta, (q, \theta^2 - b_1^2)) \mu d\theta \leq \int_{\Theta} u_d(\theta, (\theta - b_1)) \mu d\theta = -\|b_1\|^2 < \pi_2,$$

he still loses.

- (d) Suppose that Agent 1 deviates to (q, w) where $q, w \in \mathbb{R}$. The DM correctly believes that Agent 2 is proposing his ideal action and that $\theta^1 = z + b_2^1$. Also she believes that Agent 1 is proposing his own ideal action for dimension 2. Therefore, for any $z \in \mathbb{R}$, $\pi_1 < \pi_2$ and $\beta(m) = 0$. So (q, w) is unprofitable.

None of the deviations render a higher-than-equilibrium payoff for Agent 1. Reverse the role of Agent 1 and Agent 2 in the cases above and we get the case in which $\|b_1\| < \|b_2\|$. When $\|b_1\| = \|b_2\|$, in equilibrium agents win with equal probability. The only change from the argument above is:

- (c) If Agent 1 deviates to (q, \emptyset) , then $\pi_1 \leq -\|b_1\|^2 = \pi_2$, so Agent 1's probability of winning is lower compared with equilibrium.
- (d) For any $z \in \mathbb{R}$, $\pi_1 \leq \pi_2$, with equality reached when $z = \theta^1$. So Agent 1's probability of winning is lower compared with equilibrium.

Since in equilibrium, Agent 1 wins with probability $\frac{1}{2}$ and gets his ideal action conditional on winning, neither (q, \emptyset) nor (q, w) is profitable. We have exhausted the different cases for biases and shown that the above strategies constitute an equilibrium.

Proof of Proposition 2

Consider the proposal profile

$$\begin{aligned} s_1(\theta) &= (\theta^1 - b_1^1, \emptyset), \\ s_2(\theta) &= (\emptyset, \theta^2 - b_2^2) \end{aligned}$$

for all θ . The first part of the proof shows that for arbitrary biases b_1, b_2 such that $|b_1^1| < |b_2^1|$ and $|b_2^2| < |b_1^2|$, s_1, s_2 as above cannot be sustained in equilibrium. I show the proof in two cases: $\|b_1\| < \|b_2\|$ and $\|b_1\| = \|b_2\|$. The $\|b_1\| > \|b_2\|$ case is symmetric to $\|b_1\| < \|b_2\|$.

First note that from Lemma 1, $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$ for all θ . Therefore for all θ ,

$$u_1^*(\theta) = \frac{1}{2}u_1(\theta, \theta - b_1) + \frac{1}{2}u_1(\theta, \theta - b_2) = -\frac{1}{2}\|b_1 - b_2\|^2.$$

Case 1. $\|b_1\| < \|b_2\|$.

Agent 1 has a profitable deviation (\emptyset, \emptyset) at any state $\bar{\theta}$. Given $m = (m_1^{dev}, s_2(\bar{\theta})) = ((\emptyset, \emptyset), (\emptyset, \bar{\theta}^2 - b_2^2))$, $\mu\{\tilde{\theta} \mid \tilde{\theta}^2 = s_2(\bar{\theta}) + b_2^2\} = 1$. Then

$$\begin{aligned} \pi_1 &= \int_{\Theta} u_d(\theta, \theta - b_1) \mu d\theta = -\|b_1\|^2, \\ \pi_2 &= \int_{\Theta} u_d(\theta, \theta - b_2) \mu d\theta = -\|b_2\|^2 < -\|b_1\|^2. \end{aligned}$$

Therefore $\beta(m) = 1$ and

$$u_1^{dev}(\bar{\theta}, m_1^{dev}, s_2(\bar{\theta})) = u_1(\bar{\theta}, \bar{\theta} - b_1) = 0 > u_1^*(\bar{\theta}).$$

Case 2. $\|b_1\| = \|b_2\|$.

Agent 1 has a profitable deviation $(\emptyset, \bar{\theta}^2 - (1 - \epsilon)b_1^2)$ for small $\epsilon > 0$. Given $m = (m_1^{dev}, s_2(\bar{\theta})) = ((\emptyset, \bar{\theta}^2 - (1 - \epsilon)b_1^2), (\emptyset, \bar{\theta}^2 - b_2^2))$, $\mu\{\tilde{\theta} \mid \tilde{\theta}^2 = s_2(\bar{\theta}) + b_2^2\} = 1$.

Then

$$\begin{aligned} \pi_1 &= \int_{\Theta} u_d(\theta, (\theta^1 - b_1^1, \theta^2 - (1 - \epsilon)b_1^2)) \mu d\theta = -\|(b_1^1, (1 - \epsilon)b_1^2)\|^2, \\ \pi_2 &= \int_{\Theta} u_d(\theta, \theta - b_2) \mu d\theta = -\|b_2\|^2 = -\|b_1\|^2 < -\|(b_1^1, (1 - \epsilon)b_1^2)\|^2. \end{aligned}$$

Therefore $\beta(m) = 1$ and for $\epsilon > 0$ small enough,

$$u_1^{dev}(\bar{\theta}, m_1^{dev}, s_2(\bar{\theta})) = u_1(\bar{\theta}, (\bar{\theta}^1 - b_1^1, \bar{\theta}^2 - (1 - \epsilon)b_1^2)) = -\epsilon^2 |b_1^2|^2 > u_1^*(\bar{\theta}).$$

The second part of the proof shows that the following is an equilibrium proposal profile:

$$s_1(\theta) = s_2(\theta) = (\emptyset, \emptyset), \forall \theta.$$

The DM's belief and strategy is as follows: for any m such that there is exactly one agent who is specific for any dimension(s), the DM believes that he is proposing

his own ideal action. In other words, if m is such that $m_j = (\emptyset, \emptyset)$ and $m_i \in (\mathbb{R} \times \{\emptyset\}) \cup (\{\emptyset\} \times \mathbb{R}) \cup (\mathbb{R}^2)$, $\mu\{\tilde{\theta} \mid \tilde{\theta}^j = m_i^j + b_i^j, \forall j \text{ s.t. } m_i^j \in \mathbb{R}\} = 1$. Any other m is the result of bilateral deviation and $\mu(m)$ is not restricted by single-deviation consistency or consequential in sustaining the equilibrium. Therefore $\mu(m) \in \Delta(\Theta)$ and $\beta(m) \in \widetilde{\text{BR}}(\mu, m)$.

To see that no agent has an incentive to deviate, suppose that at $\bar{\theta}$ Agent i deviates to m_i^{dev} . Given $m = (m_i^{dev}, m_j)$, note that

$$\pi_i = \int_{\Theta} u_d(\theta, \theta - b_i) \mu d\theta.$$

Therefore deviating does not change how favorable the DM perceives the deviator's proposal. So $\beta(m) = \beta(s_1(\bar{\theta}), s_2(\bar{\theta}))$ and m_i^{dev} is unprofitable.

Proof of Proposition 3

First I show that the equilibrium

$$s_1(\theta) = s_2(\theta) = (\emptyset, \emptyset), \forall \theta$$

satisfies the Extended Intuitive Criterion. The proof is divided into two cases: $\|b_1\| > \|b_2\|$ and $\|b_1\| = \|b_2\|$. $\|b_1\| < \|b_2\|$ is symmetric.

Case 1. $\|b_1\| > \|b_2\|$.

In equilibrium, Agent 2 wins and the outcome is b_2 . Therefore for any θ ,

$$u_1^*(\theta) = u_1(\theta, \theta - b_2) = -\|b_1 - b_2\|^2$$

$$u_2^*(\theta) = u_2(\theta, \theta - b_2) = 0.$$

For any $m \neq ((\emptyset, \emptyset), (\emptyset, \emptyset))$ that results from a unilateral deviation, $\mu(\cdot \mid m)$ is as follows: if m is such that $m_j = (\emptyset, \emptyset)$ and $m_i \in (\mathbb{R} \times \{\emptyset\}) \cup (\{\emptyset\} \times \mathbb{R}) \cup (\mathbb{R}^2)$,

$$\mu\{\tilde{\theta} \mid \tilde{\theta}^j = m_i^j + b_i^j, \forall j \text{ s.t. } m_i^j \in \mathbb{R}\} = 1.$$

For m such that $m_1 = (\emptyset, \emptyset)$ and $m_2 \in (\mathbb{R} \times \{\emptyset\}) \cup (\{\emptyset\} \times \mathbb{R}) \cup (\mathbb{R}^2)$, $\mu(\cdot \mid m)$ assigns probability 1 to the event that Agent 2 has deviated to committing to his ideal action.

1. If $\beta(m) = 0$, then for any θ , $u_2^{dev}(\theta, s_1(\theta), m_2, \beta) = u_2(\theta, \theta - b_2) = 0$, so the highest payoff Agent 2 gets by deviating is no less than his equilibrium payoff.
2. $\beta(m) = 0$ is not dominated because it is the best response to $\mu(\cdot \mid m)$.

For m such that $m_2 = (\emptyset, \emptyset)$ and $m_1 \in (\mathbb{R} \times \{\emptyset\}) \cup (\{\emptyset\} \times \mathbb{R}) \cup (\mathbb{R}^2)$, $\mu(\cdot \mid m)$ assigns probability 1 to the event that Agent 1 has deviated to committing to his ideal action.

1. If $\beta(m) = 0$, then for any θ , $u_1^{dev}(\theta, s_2(\theta), m_1, \beta) = u_1(\theta, \theta - b_2) = u_1^*(\theta)$, so the highest payoff Agent 1 gets by deviating is no less than his equilibrium payoff.
2. $\beta(m) = 0$ is not dominated because it is the best response to $\mu(\cdot | m)$.

Case 2. $\|b_1\| = \|b_2\|$.

In equilibrium, both agents win with probability $\frac{1}{2}$. Therefore for any θ ,

$$u_1^*(\theta) = \frac{1}{2}u_1(\theta, \theta - b_1) + \frac{1}{2}u_1(\theta, \theta - b_2) = -\frac{1}{2}\|b_1 - b_2\|^2,$$

$$u_2^*(\theta) = \frac{1}{2}u_2(\theta, \theta - b_1) + \frac{1}{2}u_2(\theta, \theta - b_2) = -\frac{1}{2}\|b_1 - b_2\|^2.$$

For m such that $m_1 = (\emptyset, \emptyset)$ and $m_2 \in (\mathbb{R} \times \{\emptyset\}) \cup (\{\emptyset\} \times \mathbb{R}) \cup (\mathbb{R}^2)$, $\mu(\cdot | m)$ assigns probability 1 to the event that Agent 2 has deviated to committing to his ideal action.

1. If $\beta(m) = \frac{1}{2}$, then for any θ , $u_2^{dev}(\theta, s_1(\theta), m_2, \beta) = u_2^*(\theta)$, so the highest payoff Agent 2 gets by deviating is no less than his equilibrium payoff.
2. $\beta(m) = \frac{1}{2}$ is not dominated because it is the best response to $\mu(\cdot | m)$.

For m such that $m_2 = (\emptyset, \emptyset)$ and $m_1 \in (\mathbb{R} \times \{\emptyset\}) \cup (\{\emptyset\} \times \mathbb{R}) \cup (\mathbb{R}^2)$, $\mu(\cdot | m)$ assigns probability 1 to the event that Agent 1 has deviated to committing to his ideal action.

1. If $\beta(m) = \frac{1}{2}$, then for any θ , $u_1^{dev}(\theta, s_2(\theta), m_1, \beta) = u_1^*(\theta)$, so the highest payoff Agent 1 gets by deviating is no less than his equilibrium payoff.
2. $\beta(m) = \frac{1}{2}$ is not dominated because it is the best response to $\mu(\cdot | m)$, the actual belief.

When $\|b_1^1\| = \|b_2^2\| = 0$,

$$s_1(\theta) = s_2(\theta) = (\emptyset, \emptyset), \forall \theta$$

continues to be an intuitive equilibrium.

Moreover, from Proposition 1,

$$s_1(\theta) = (w(\theta), \emptyset),$$

$$s_2(\theta) = (\emptyset, z(\theta)),$$

where $w, z : \Theta \rightarrow \mathbb{R}$ is not an equilibrium. To see it is not an intuitive equilibrium, note that in Step 2, after Agent 1's deviation (\emptyset, \emptyset) , the DM's belief about the profitability of the deviation is independent of her belief about the state. This is because for any state, such a deviation is profitable when $\beta = 1$. Moreover, $\beta = 1$ is not dominated because if the DM believes that Agent 2 has not deviated, then it is a best response since $\|b_1\| \leq \|b_2\|$. And this is required by both the single-deviation consistency and the Extended Intuitive Criterion.

In Step 3, Agent 1's undercutting goes through because the restricted beliefs under both single-deviation consistency and the Extended Intuitive Criterion assigns probability 1 to Agent 2 has not deviated. In addition, the Extended Intuitive Criterion assigns probability 1 to the event that Agent 1's deviation is potentially profitable. Since Agent 1's deviation is indeed profitable if $\beta = 1$ and $\beta = 1$ is a best response to the belief that only Agent 1 has deviated, this deviation is profitable even if we require the equilibrium to be intuitive.

To be precise, the three steps are as follows (assuming without loss of generality $\|b_1\| \leq \|b_2\|$):

Step 1. $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}, \forall \theta$.

Step 2. $w(\theta) = \theta^1, z(\theta) = \theta^2, \forall \theta$.

Steps 1 and 2 are identical to the those in the proof for Proposition 1.

Step 3. At any $\bar{\theta}$, Agent 1 has a profitable deviation $(\emptyset, \bar{\theta}^2 - (1 - \epsilon)b_1^2)$, where $\epsilon > 0$ is small.

Suppose that Agent 1 deviates to said proposal at $\bar{\theta}$. The DM observes $m = (m_1^{dev}, s_2(\bar{\theta})) = ((\emptyset, \bar{\theta}^2 - (1 - \epsilon)b_1^2), (\emptyset, \bar{\theta}^2))$. Then

$$\Theta_1(m) = \{\tilde{\theta} \mid s_2(\tilde{\theta}) = (\emptyset, \bar{\theta}^2), u_1^*(\tilde{\theta}) \leq \max_{\beta \in \text{BR}(\Theta, m)} u_1(\tilde{\theta}, s_2(\tilde{\theta}), m_1^{dev}, \beta)\} = \{\tilde{\theta} \mid \bar{\theta}^2 = \bar{\theta}^2\}.$$

whereas $\Theta_2(m)$ is empty. This is because s_2 pins down θ^2 and for any θ such that $\theta^2 = \bar{\theta}^2$, $u_1(\theta, s_2(\theta), m_1^{dev}, 1) = u_1(\bar{\theta}, (\bar{\theta}^1 - b_1^1, \bar{\theta}^2 - (1 - \epsilon)b_1^2)) = -\epsilon^2|b_1^2|^2 > u_1^*(\bar{\theta})$ for small enough ϵ . Moreover, $1 \in \text{BR}(\Theta_1(m), m)$ and therefore is not dominated. $\Theta_2(m)$ is empty because $m_1^{dev} \in \text{supp}(G_1)$ given s_1 .

For all $\mu \in \Delta(\Theta_1(m))$, $\pi_1 > \pi_2$. Therefore $\text{BR}(\Theta_1(m), m) = \{1\}$ and $u_1^{dev}(\bar{\theta}, s_2(\bar{\theta}), m_1^{dev}) > u_1^*(\bar{\theta})$, making m_1^{dev} a profitable deviation.

Appendix II. General Strategies in a Binary State Space

In order to relax the assumption that agents' equilibrium choices of vagueness are symmetric, I study the following setup as a special case. θ is uniformly distributed over $\Theta = \{(0, 0), (0, 1), (1, 0), (1, 1)\}$. In order to allow arbitrary compromises by the agents and ensure both agents' ideal actions are implementable, I let action space to be the interval $[-1, 1]$. Therefore the proposal space is $M = ([-1, 1] \cup \{\emptyset\})^2$. The agents' biases are $b_1 = (0, 1)$ and $b_2 = (1, 0)$. The next proposition illustrates where vagueness is likely to occur.

Proposition 4 1. *An intuitive equilibrium in which each agent is vague on his aligned dimension and commits on his misaligned dimension exists.*

2. *There does not exist a weak PBE in which, at some state, each agent is vague on his misaligned dimension and commits on his aligned dimension.*

For part 1, I construct a semi-revealing equilibrium in which the DM learns the state at $(1, 1)$, but assigns equal probabilities to the rest of the states when the state is not $(1, 1)$. The construction relies on the following property of any equilibrium: if both agents propose the same action at a state, then any deviation is potentially profitable. Even if Agent i has deviated to a less-preferred action, as long as the DM selects Agent i with zero probability, he gets the same payoff as in equilibrium.

Part 2 first shows that specific commitments must be agents' own ideal actions. I then show that there is a profitable deviation for Agent 1 regardless of the off-path beliefs of the DM. In other words, there is no need to resort to equilibrium refinement to rule out these strategies.

Proof of Proposition 4.

(1) Consider a proposal profile

$$s_1(\theta) = (\emptyset, q(\theta)),$$

$$s_2(\theta) = (z(\theta), \emptyset)$$

in which $q, z : \Theta \rightarrow [-1, 1]$ are as follows:

$$q(\theta) = z(\theta) = \begin{cases} 1 & \text{if } \theta = (1, 1), \\ 0 & \text{otherwise.} \end{cases}$$

Let $\beta : M \times M \rightarrow [-1, 1]$ and $\mu : M \times M \rightarrow \Delta(\Theta)$ be as follows:

- If m is on-path, $\beta(m) = \frac{1}{2}$.

$$\mu((1, 1) \mid (\emptyset, 1), (1, \emptyset)) = 1;$$

$$\begin{aligned} \mu((0, 0) \mid (\emptyset, 0), (0, \emptyset)) &= \mu((0, 1) \mid (\emptyset, 0), (0, \emptyset)) \\ &= \mu((1, 0) \mid (\emptyset, 0), (0, \emptyset)) = \frac{1}{3}. \end{aligned}$$

- If m is off-path and there is at least one Agent i with m_i consistent with equilibrium, $\mu(\cdot \mid m)$ is defined as follows:

If $m_1 \neq (\emptyset, 1)$ and $m_2 \neq (1, \emptyset)$, $\mu((0, 0) \mid m) = 1$; otherwise $\mu((1, 1) \mid m) = 1$.

- If $m = (m_1, m_2)$ where both m_1 and m_2 are inconsistent with equilibrium, $\mu(m)$ is unrestricted.

To be precise, $\mu(\cdot \mid m)$ and $\beta(m)$ for all off-path m are defined as follows:

- a) For $m = (m_1^{dev}, (0, \emptyset))$ where m_1^{dev} is inconsistent with equilibrium:
 $\mu((0, 0) \mid m) = 1$. If $m_1^{dev} = (0, 0)$, then $\beta(m) = \frac{1}{2}$; otherwise $\beta(m) = 0$.
- b) For $m = (m_1^{dev}, (1, \emptyset))$ where m_1^{dev} is inconsistent with equilibrium:
 $\mu((1, 1) \mid m) = 1$. If $m_1^{dev} = (1, 1)$, then $\beta(m) = \frac{1}{2}$; otherwise $\beta(m) = 0$.
- c) For $m = ((\emptyset, 0), m_2^{dev})$ where m_2^{dev} is inconsistent with equilibrium:
 $\mu((0, 0) \mid m) = 1$. If $m_2^{dev} = (0, 0)$, then $\beta(m) = \frac{1}{2}$; otherwise $\beta(m) = 1$.
- d) For $m = ((\emptyset, 1), m_2^{dev})$ where m_2^{dev} is inconsistent with equilibrium:
 $\mu((1, 1) \mid m) = 1$. If $m_2^{dev} = (1, 1)$, then $\beta(m) = \frac{1}{2}$; otherwise $\beta(m) = 1$.
- e) For $m = (m_1^{dev}, m_2^{dev})$ where both m_1^{dev} and m_2^{dev} are inconsistent with equilibrium, $\mu(\cdot \mid m)$ is unrestricted. $\beta(m) \in \widetilde{\text{BR}}(\mu, m)$.
- f) For $m = ((\emptyset, 0), (1, \emptyset))$, $\mu((1, 1) \mid m) = 1$, $\beta(m) = 0$.
- g) For $m = ((\emptyset, 1), (0, \emptyset))$, $\mu((1, 1) \mid m) = 1$, $\beta(m) = 1$.

It is easy to verify that for all m , $\beta(m) \in \widetilde{\text{BR}}(\mu(m), m)$. To see that the equilibrium satisfies the Extended Intuitive Criterion, note that for all m in Case (1)a, $(0, 0) \in \Theta_1(m)$. To see this, note that

$$\begin{aligned} \Theta_1(m) &= \{\theta \mid s_2(\theta) = (0, \emptyset), u_1^*(\theta) \leq \max_{\beta \in \text{BR}(\Theta, m)} \beta u_1(\theta, y_1(m_1)) + (1 \\ &\quad - \beta) u_1(\theta, y_2(m_2))\}. \end{aligned}$$

Since $u_1^*(0, 0) = -1 = u_1((0, 0), y_2(m_2))$, the above inequality always holds for $(0, 0)$, with equality reached by $\beta = 0$. Similarly, $(0, 0) \in \Theta_2(m)$ for all m in Case (1)c.

For all m in Case (1)b, the DM believes that Agent 2 has not deviated and $(1, \emptyset)$ pins down the state $(1, 1)$. Moreover, if Agent 1's deviation is potentially profitable, then $(1, 1) \in \Theta_1(m)$; otherwise, $\Theta_1(m)$ is the empty set and $\mu(\cdot | m)$ is unrestricted. Therefore $\mu((1, 1) | m)$ is supported on the set of states for which no agent has made a deviation that is not potentially profitable. Similarly for all m in Case (1)d.

For m in Case (1)f, the DM believes that Agent 1 has deviated and therefore $\theta = (1, 1)$. At $(1, 1)$, Agent 1's deviation $(\emptyset, 0)$ is potentially profitable because $u_1((1, 1), y_1(\emptyset, 0)) = u_1((1, 1), (1, 0)) = 0 > -1 = u_1^*(1, 1)$. Therefore $(1, 1) \in \Theta_1(m)$. Similarly, $(1, 1) \in \Theta_2(m)$ for m in Case (1)g.

Now I show that Agent 1 has no incentives to deviate; Agent 2 is symmetric.

- At state $(0, 0)$, $u_1^*(0, 0) = -1$.

$$(i) \text{ If } m_1^{dev} = (\emptyset, 1), u_1^{dev}((0, 0), m_1^{dev}, m_2) = u_1((0, 0), y_1(\emptyset, 1)) = u_1((0, 0), (0, 1)) = -4 < -1 = u_1^*(0, 0).$$

$$(ii) \text{ If } m_1^{dev} = (0, 0), u_1^{dev}((0, 0), m_1^{dev}, m_2) = \frac{1}{2}u_1((0, 0), y_1(0, 0)) + \frac{1}{2}u_1((0, 0), y_2(0, \emptyset)) = -1 = u_1^*(0, 0).$$

$$(iii) \text{ For any other } m_1^{dev} \text{ inconsistent with equilibrium, } u_1^{dev}((0, 0), m_1^{dev}, m_2) = u_1((0, 0), y_2(0, \emptyset)) = u_1((0, 0), (0, 0)) = -1 = u_1^*(0, 0).$$

- At state $(0, 1)$, $u_1^*(0, 1) = -\frac{1}{2}$.

$$(i) \text{ If } m_1^{dev} = (\emptyset, 1), u_1^{dev}((0, 1), m_1^{dev}, m_2) = u_1((0, 1), y_1(\emptyset, 1)) = u_1((0, 1), (0, 1)) = -1 < -\frac{1}{2} = u_1^*(0, 1).$$

$$(ii) \text{ If } m_1^{dev} = (0, 0), u_1^{dev}((0, 1), m_1^{dev}, m_2) = \frac{1}{2}u_1((0, 1), y_1(0, 0)) + \frac{1}{2}u_1((0, 1), y_2(0, \emptyset)) = -\frac{1}{2} = u_1^*(0, 1).$$

$$(iii) \text{ For any other } m_1^{dev} \text{ inconsistent with equilibrium, } u_1^{dev}((0, 1), m_1^{dev}, m_2) = u_1((0, 1), y_2(0, \emptyset)) = -1 < -\frac{1}{2} = u_1^*(0, 1).$$

- At state $(1, 0)$, $u_1^*(1, 0) = -\frac{3}{2}$.

$$(i) \text{ If } m_1^{dev} = (\emptyset, 1), u_1^{dev}((1, 0), m_1^{dev}, m_2) = -4 < -\frac{3}{2} = u_1^*(1, 0).$$

$$(ii) \text{ If } m_1^{dev} = (0, 0), u_1^{dev}((1, 0), m_1^{dev}, m_2) = \frac{1}{2}u_1((1, 0), y_1(0, 0)) + \frac{1}{2}u_1((1, 0), y_2(0, \emptyset)) = -2 < -\frac{3}{2} = u_1^*(1, 0).$$

$$(iii) \text{ For any other } m_1^{dev} \text{ inconsistent with equilibrium, } u_1^{dev}((1, 0), m_1^{dev}, m_2) = u_1((1, 0), y_2(0, \emptyset)) = -2 < -\frac{3}{2} = u_1^*(1, 0).$$

- At state $(1, 1)$, $u_1^*(1, 1) = -1$.

(i) If $m_1^{dev} = (\emptyset, 0)$, $u_1^{dev}((1, 1), m_1^{dev}, m_2) = u_1((1, 1), y_2(1, \emptyset)) = -1 = u_1^*(1, 1)$.

(ii) If $m_1^{dev} = (1, 1)$, $u_1^{dev}((1, 1), m_1^{dev}, m_2) = \frac{1}{2}u_1((1, 1), y_1(1, 1)) + \frac{1}{2}u_1((1, 1), y_2(1, \emptyset)) = -1 = u_1^*(1, 1)$.

(iii) For any other m_1^{dev} inconsistent with equilibrium, $u_1^{dev}((1, 1), m_1^{dev}, m_2) = u_1((1, 1), y_2(1, \emptyset)) = -1 = u_1^*(1, 1)$.

Therefore Agent 1 has no incentives to deviate at any state; this is similar for Agent 2. I have established that above is an equilibrium.

(2) I prove this by showing that any strategy profile in which each agent is vague on his misaligned dimension and commits on his aligned dimension at some state is not a weak PBE strategy profile. I prove by contradiction: suppose that there exists $\bar{\theta}$ and $w, z \in [-1, 1]$ such that

$$s_1(\bar{\theta}) = (w, \emptyset),$$

$$s_2(\bar{\theta}) = (\emptyset, z)$$

is a weak PBE proposal profile when state is $\bar{\theta}$.

Step 1. $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = \frac{1}{2}$.

If $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 0$, then Agent 1 has a profitable deviation $s_2(\bar{\theta})$. To see this, first note that

$$u_1^*(\bar{\theta}) = u_1((\bar{\theta}^1, \bar{\theta}^2), y_2(\emptyset, z)) = u_1((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1 - 1, z)).$$

Now let Agent 1 deviate to $s_2(\bar{\theta})$. Given $m = (m_1^{dev}, s_2(\bar{\theta})) = ((\emptyset, z), (\emptyset, z))$ and any $\mu \in \Delta(\Theta)$, $\pi_1 > \pi_2$. Therefore $\beta(m) = 1$ and $u_1^{dev}(\bar{\theta}, m_1^{dev}, s_2(\bar{\theta})) = u_1(\bar{\theta}, y_1(\emptyset, z)) = u_1((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1, z)) > u_1^*(\bar{\theta})$.

Similarly, if $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 1$, then Agent 2 has a profitable deviation $s_1(\bar{\theta})$.

Step 2. $w = \bar{\theta}^1, z = \bar{\theta}^2$.

Suppose $w \neq \bar{\theta}^1$. Then Agent 1 has a profitable deviation (\emptyset, \emptyset) . To see this, first note that

$$\begin{aligned} u_1^*(\bar{\theta}) &= \frac{1}{2}u_1(\bar{\theta}, y_1(w, \emptyset)) + \frac{1}{2}u_1(\bar{\theta}, y_2(\emptyset, z)) \\ &< \frac{1}{2}u_1((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1, \bar{\theta}^2 - 1)) + \frac{1}{2}u_1(\bar{\theta}, y_2(\emptyset, z)). \end{aligned}$$

Now let Agent 1 deviate to $m_1^{dev} = (\emptyset, \emptyset)$. Given $m = (m_1^{dev}, s_2(\bar{\theta})) = ((\emptyset, \emptyset), (\emptyset, z))$ and any $\mu \in \Delta(\Theta)$, $\pi_1 \geq \pi_2$. Therefore $\beta(m) \geq \frac{1}{2}$ and

$$\begin{aligned} u_1^{dev}(\bar{\theta}, m_1^{dev}, s_2(\bar{\theta})) &\geq \frac{1}{2}u_1(\bar{\theta}, y_1(\emptyset, \emptyset)) + \frac{1}{2}u_1(\bar{\theta}, y_2(\emptyset, z)) \\ &= \frac{1}{2}u_1((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1, \bar{\theta}^2 - 1)) + \frac{1}{2}u_1(\bar{\theta}, y_2(\emptyset, z)) \\ &> u_1^*(\bar{\theta}). \end{aligned}$$

Similarly if $z \neq \bar{\theta}^2$, then Agent 2 has a profitable deviation (\emptyset, \emptyset) .

Step 3. $u_1^*(\bar{\theta}) = -1$.

From the steps above,

$$\begin{aligned} u_1^*(\bar{\theta}) &= \frac{1}{2}u_1(\bar{\theta}, y_1(w, \emptyset)) + \frac{1}{2}u_2(\bar{\theta}, y_2(\emptyset, z)) \\ &= \frac{1}{2}u_1((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1, \bar{\theta}^2 - 1)) + \frac{1}{2}u_1((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1 - 1, \bar{\theta}^2)) \\ &= -1. \end{aligned}$$

Step 4. If $\bar{\theta}^2 = 0$, then Agent 1 has a profitable deviation $(\emptyset, -0.4)$. If $\bar{\theta}^2 = 1$, then Agent 1 has a profitable deviation $(\emptyset, 0.6)$.

I first consider the case in which $\bar{\theta}^2 = 0$. By Step 2, $z = 0$. Given $m = (m_1^{dev}, s_2(\bar{\theta})) = ((\emptyset, -0.4), (\emptyset, 0))$ and any $\mu \in \Delta(\Theta)$, $\pi_1 > \pi_2$. Therefore $\beta(m) = 1$. $u_1^{dev}(\bar{\theta}, m_1^{dev}, s_2(\bar{\theta})) = u_1(\bar{\theta}, y_1(\emptyset, -0.4)) = u_1((\bar{\theta}^1, 0), (\bar{\theta}^1, -0.4)) = -0.36 > -1 = u_1^*(\bar{\theta})$.

Now I consider the case in which $\bar{\theta}^2 = 1$. By Step 2, $z = 1$. Given $m = (m_1^{dev}, s_2(\bar{\theta})) = ((\emptyset, 0.6), (\emptyset, 1))$ and any $\mu \in \Delta(\Theta)$, $\pi_1 > \pi_2$. Therefore $\beta(m) = 1$. $u_1^{dev}(\bar{\theta}, m_1^{dev}, s_2(\bar{\theta})) = u_1(\bar{\theta}, y_1(\emptyset, 0.6)) = u_1((\bar{\theta}^1, 1), (\bar{\theta}^1, 0.6)) = -0.36 > -1 = u_1^*(\bar{\theta})$.

I have established that Agent 1 has a profitable deviation at $\bar{\theta}$. Therefore for no $\theta \in \Theta$ and $w, z \in [-1, 1]$ is

$$\begin{aligned} s_1(\theta) &= (w, \emptyset), \\ s_2(\theta) &= (\emptyset, z), \end{aligned}$$

a weak PBE proposal profile.

■

The next result characterizes intuitive equilibria in the binary-state environment. In particular, in no equilibria is any Agent i specific about dimension i and vague about dimension j . As one can see, the simpler state space renders a stronger result.

Proposition 5 *In any intuitive equilibrium with symmetry condition removed, an agent who is vague on some dimension must be vague on his aligned dimension.*

I prove this by first supposing that Agent 1 is vague about his misaligned dimension and commits on his aligned dimension. Then I show that the proposal profile given by Agent 1's strategy and any strategy by Agent 2 cannot be sustained in equilibrium. The case in which Agent 2 is vague on his misaligned dimension and commits on his aligned dimension is symmetric. The binary state space helps in two ways. First, because the state space is smaller than the action space, some actions are credible compromises regardless of the state. Second, since there are altogether only four different states, each agent can only make four possibly different proposals. This makes it easy to pin down agents' equilibrium strategies based on the posterior beliefs of the DM.

Proof of Proposition 5. Let $s_1(\theta) = (q(\theta), \emptyset)$ for all θ , where $q : \Theta \rightarrow [-1, 1]$. I show that s_1 is not a strategy in any intuitive equilibrium. The arguments for Agent 2 is symmetric. There are 4 cases:

1. $s_2(\theta) = (\emptyset, z(\theta))$ for all θ , where $z : \Theta \rightarrow [-1, 1]$.
2. $s_2(\theta) = (\emptyset, \emptyset)$ for all θ .
3. $s_2(\theta) = (z(\theta), w(\theta))$ for all θ , where $z, w : \Theta \rightarrow [-1, 1]$.
4. $s_2(\theta) = (z(\theta), \emptyset)$ for all θ , where $z : \Theta \rightarrow [-1, 1]$.

For each of the 4 cases, I show that (s_1, s_2) is not an equilibrium proposal profile.

1. This is shown in Proposition 4.
2. I prove this by contradiction. Suppose the proposal profile is a weak PBE.

Step 1. $\forall \theta, \beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$.

If $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 0$ for some $\bar{\theta}$, then Agent 1 has a profitable deviation (\emptyset, \emptyset) . To see this, note that

$$u_1^*(\bar{\theta}) = u_1(\bar{\theta}, y_2(\emptyset, \emptyset)) = u_1((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1 - 1, \bar{\theta}^2)).$$

Now let Agent 1 deviate to (\emptyset, \emptyset) . Given $m = (m_1^{dev}, s_2(\bar{\theta})) = ((\emptyset, \emptyset), (\emptyset, \emptyset))$, $\beta(m) = \frac{1}{2}$ and

$$\begin{aligned} u_1^{dev}(\bar{\theta}, m_1^{dev}, s_2(\bar{\theta})) &= \frac{1}{2}u_1(\bar{\theta}, y_1(\emptyset, \emptyset)) + \frac{1}{2}u_1(\bar{\theta}, y_2(\emptyset, \emptyset)) \\ &= \frac{1}{2}u_1((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1, \bar{\theta}^2 - 1)) + \frac{1}{2}u_1((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1 - 1, \bar{\theta}^2)) \\ &> u_1^*(\bar{\theta}). \end{aligned}$$

If $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 1$ for some $\bar{\theta}$, then Agent 2 has a profitable deviation $(q(\bar{\theta}), \emptyset)$. The argument is identical to the one used in Step 1 of the proof for part (2) of the Proposition 4.

Step 2. $q(\theta) = \theta^1, \forall \theta$.

Fix any $\bar{\theta}$. Let μ denote the DM's belief at $\bar{\theta}$. Since $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = \frac{1}{2}$, $\pi_1 = \pi_2 = u_d((\bar{\theta}^1, \bar{\theta}^2), (\bar{\theta}^1 - 1, \bar{\theta}^2)) = -1$ given μ . Moreover, given any $\tilde{\mu} \in \Delta(\Theta)$, $\pi_1 \leq -1$ with equality reached when $\tilde{\mu}\{\tilde{\theta} \mid \tilde{\theta}^1 = q(\bar{\theta})\} = 1$. Therefore $\mu\{\tilde{\theta} \mid \tilde{\theta}^1 = q(\bar{\theta})\} = 1$. Since μ is consistent with s_1 and s_2 , $q(\bar{\theta}) = \bar{\theta}^1$.

Step 3. Agent 2 has a profitable deviation $(\frac{1}{2}, \emptyset)$ at state $(1, 0)$.

From Step 2, $(s_1(1, 0), s_2(1, 0)) = ((1, \emptyset), (\emptyset, \emptyset))$. From Step 1 and 2, $u_2^*(1, 0) = \frac{1}{2}u_2((1, 0), y_1(1, \emptyset)) + \frac{1}{2}u_2((1, 0), y_2(\emptyset, \emptyset)) = -1$.

Now let Agent 2 deviate to $(\frac{1}{2}, \emptyset)$. Given $m = (s_1(1, 0), m_2^{dev}) = ((1, \emptyset), (\frac{1}{2}, \emptyset))$, for all $\mu \in \Delta(\Theta)$, $\pi_1 < \pi_2$. To see why, note that if $\theta^1 = 0$,

$$u_d(\theta, y_2(\frac{1}{2}, \emptyset)) = u_d((0, \theta^2), (\frac{1}{2}, \theta^2)) > u_d((0, \theta^2), (1, \theta^2 - 1)) = u_d(\theta, y_1(1, \emptyset)).$$

On the other hand, if $\theta^1 = 1$,

$$u_d(\theta, y_2(\frac{1}{2}, \emptyset)) = u_d((1, \theta^2), (\frac{1}{2}, \theta^2)) > u_d((1, \theta^2), (1, \theta^2 - 1)) = u_d(\theta, y_1(1, \emptyset)).$$

So, given any $\mu \in \Delta(\Theta)$, $\pi_1 < \pi_2$ and $\beta(m) = 0$. Therefore $u_2^{dev}((1, 0), s_1(1, 0), m_2^{dev}) = u_2((1, 0), y_2(\frac{1}{2}, \emptyset)) = u_2((1, 0), (\frac{1}{2}, 0)) = -\frac{1}{4} > -1 = u_2^*(1, 0)$.

3. I prove this by contradiction. Suppose the proposal profile is an equilibrium proposal profile. First I prove the following 4 claims.

Claim 1 $\forall \theta, \beta(s_1(\theta), s_2(\theta)) \in \{0, \frac{1}{2}\}$.

Proof of Claim 1. If for some $\bar{\theta}$, $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 1$, then Agent 2 has a profitable deviation $s_1(\bar{\theta})$. The argument is identical to the one used in Step 1 of the proof for part (2) of the Proposition 4. ■

Claim 2 *If for some θ , $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$, then $z(\theta) = \theta^1 - 1$ and $w(\theta) = \theta^2$.*

Proof of Claim 2. Suppose on the contrary that either $z(\theta) \neq \theta^1 - 1$ or $w(\theta) \neq \theta^2$. Then

$$\begin{aligned} u_2^*(\theta) &= \frac{1}{2}u_2(\theta, y_2(z(\theta), w(\theta))) + \frac{1}{2}u_2(\theta, y_1(q(\theta), \emptyset)) \\ &< \frac{1}{2}u_2((\theta^1, \theta^2), (\theta^1 - 1, \theta^2)) + \frac{1}{2}u_2(\theta, y_1(q(\theta), \emptyset)). \end{aligned}$$

Now let Agent 2 deviate to (\emptyset, \emptyset) . Given $m = (s_1(\theta), m_2^{dev}) = ((q(\theta), \emptyset), (\emptyset, \emptyset))$ and any $\mu \in \Delta(\Theta)$, $\pi_1 \leq \pi_2$. Therefore $\beta(m) \leq \frac{1}{2}$ and

$$\begin{aligned} u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) &\geq \frac{1}{2}u_2(\theta, y_2(\emptyset, \emptyset)) + \frac{1}{2}u_2(\theta, y_1(q(\theta), \emptyset)) \\ &= \frac{1}{2}u_2((\theta^1, \theta^2), (\theta^1 - 1, \theta^2)) + \frac{1}{2}u_2(\theta, y_1(q(\theta), \emptyset)) \\ &> u_2^*(\theta). \end{aligned}$$

■

Claim 3 *If for some θ , $\beta(s_1(\theta), s_2(\theta)) = 0$, then $z(\theta) = \theta^1$.*

Proof of Claim 3. Suppose on the contrary that $z(\theta) \neq \theta^1$. Then

$$\begin{aligned} u_1^*(\theta) &= u_1(\theta, y_2(z(\theta), w(\theta))) \\ &= u_1((\theta^1, \theta^2), (z(\theta), w(\theta))) \\ &< u_1((\theta^1, \theta^2), (\theta^1, w(\theta))). \end{aligned}$$

Now let Agent 1 deviate to $(\emptyset, w(\theta))$. Given $m = (m_1^{dev}, s_2(\theta)) = ((\emptyset, w(\theta)), (z(\theta), w(\theta)))$ and any $\mu \in \Delta(\Theta)$, $\pi_1 \geq \pi_2$. Therefore $\beta(m) \geq \frac{1}{2}$ and

$$\begin{aligned} u_1^{dev}(\theta, m_1^{dev}, s_2(\theta)) &\geq \frac{1}{2}u_1(\theta, y_1(\emptyset, w(\theta))) + \frac{1}{2}u_1(\theta, y_2(z(\theta), w(\theta))) \\ &= \frac{1}{2}u_1((\theta^1, \theta^2), (\theta^1, w(\theta))) + \frac{1}{2}u_1((\theta^1, \theta^2), (z(\theta), w(\theta))) \\ &> u_1((\theta^1, \theta^2), (z(\theta), w(\theta))) \\ &= u_1^*(\theta). \end{aligned}$$

■

Claim 4 *If for some θ , $\beta(s_1(\theta), s_2(\theta)) = 0$, then $u_2^*(\theta) \leq -1$.*

Proof of Claim 4. From Claim 3,

$$\begin{aligned}
u_2^*(\theta) &= u_2(\theta, y_2(s_2(\theta))) \\
&= u_2((\theta^1, \theta^2), (z(\theta), w(\theta))) \\
&= u_2((\theta^1, \theta^2), (\theta^1, w(\theta))) \\
&\leq u_2((\theta^1, \theta^2), (\theta^1, \theta^2)) \\
&= -1.
\end{aligned}$$

■

For the rest of the proof, I fix any θ and show that $(s_1(\theta), s_2(\theta)) = ((q(\theta), \emptyset), (z(\theta), w(\theta)))$ cannot be sustained in equilibrium for any $q, z, w : \Theta \rightarrow [-1, 1]$. We divide the discussion into 3 cases.

Case 1 For all $\tilde{\theta}$ such that $s_1(\tilde{\theta}) = s_1(\theta) = (q(\theta), \emptyset)$, $\tilde{\theta}^1 = q(\theta)$.

If $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$, then $u_2^*(\theta) = \frac{1}{2}u_2(\theta, y_1(s_1(\theta))) + \frac{1}{2}u_2(\theta, y_2(s_2(\theta))) = \frac{1}{2}u_2((\theta^1, \theta^2), (\theta^1, \theta^2 - 1)) + \frac{1}{2}u_2((\theta^1, \theta^2), (\theta^1 - 1, \theta^2)) = -1$. On the other hand, if $\beta(s_1(\theta), s_2(\theta)) = 0$, then $u_2^*(\theta) \leq -1$ by Claim 4. Therefore in both situations we have $u_2^*(\theta) \leq -1$.

Now let Agent 2 deviate to $(q(\theta) - 1 + \epsilon, \emptyset)$, where ϵ is a very small positive number. Given $m = (s_1(\theta), m_2^{dev}) = ((q(\theta), \emptyset), (q(\theta) - 1 + \epsilon, \emptyset))$, we have that for all $\tilde{\theta} \in \Theta_2(m)$, $\tilde{\theta}^1 = q(\theta)$. Since for any $\mu \in \Delta(\Theta_2(m))$,

$$\begin{aligned}
\pi_1 &= u_d((\theta^1, \theta^2), y_1(q(\theta), \emptyset)) \\
&= u_d((\theta^1, \theta^2), (\theta^1, \theta^2 - 1)) \\
&< u_d((\theta^1, \theta^2), (\theta^1 - 1 + \epsilon, \theta^2)) \\
&= u_d((\theta^1, \theta^2), y_2(q(\theta) - 1 + \epsilon, \emptyset)) \\
&= \pi_2,
\end{aligned}$$

we have $\beta(m) = 0$. Therefore $u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) = u_2(\theta, y_2(m_2^{dev})) = u_2((\theta^1, \theta^2), (\theta^1 - 1 + \epsilon, \theta^2)) = -\epsilon^2 > -1$.

Therefore any $s_1(\theta)$ that belongs to Case 1 cannot be sustained in equilibrium.

Case 2 For all $\tilde{\theta}$ such that $s_1(\tilde{\theta}) = s_1(\theta) = (q(\theta), \emptyset)$, $\tilde{\theta}^1 \neq q(\theta)$.

If $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$, then $s_2(\theta) = (\theta^1 - 1, \theta^2)$. Therefore given DM's belief μ at θ , $\pi_2 = u_d((\theta^1, \theta^2), (\theta^1 - 1, \theta^2)) > \pi_1$, contradicting $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$. Therefore $\beta(s_1(\theta), s_2(\theta)) = 0$ and $u_2^*(\theta) \leq -1$.

Now let Agent 2 deviate to (\emptyset, \emptyset) . Given $m = (s_1(\theta), m_2^{dev}) = ((q(\theta), \emptyset), (\emptyset, \emptyset))$ and any $\mu \in \Delta(\Theta_2(m))$, $\pi_1 < -1 = \pi_2$. Therefore $\beta(m) = 0$ and $u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) = u_2(\theta, y_2(m_2^{dev})) = u_2((\theta^1, \theta^2), (\theta^1 - 1, \theta^2)) = 0 > u_2^*(\theta)$.

Therefore any $s_1(\theta)$ that belongs to Case 2 cannot be sustained in equilibrium.

Case 3 *There exists $\tilde{\theta}$ such that $s_1(\tilde{\theta}) = s_1(\theta) = (q(\theta), \emptyset)$ and $\tilde{\theta}^1 = q(\theta)$, and there exists $\tilde{\theta}$ such that $s_1(\tilde{\theta}) = s_1(\theta) = (q(\theta), \emptyset)$ and $\tilde{\theta}^1 \neq q(\theta)$.*

In this case, $q(\theta) \in \{0, 1\}$. We further divide this case into 6 subcases.

Subcase 1 $q(\theta) = 1$. *For all $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((1, \emptyset), s_2(\theta))$, $\tilde{\theta}^1 = 1$.*

If $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$, then $u_2^*(\theta) = \frac{1}{2}u_2(\theta, y_1(s_1(\theta))) + \frac{1}{2}u_2(\theta, y_2(s_2(\theta))) = -1$. If $\beta(s_1(\theta), s_2(\theta)) = 0$, then $u_2^*(\theta) \leq -1$. Therefore for both situations we have $u_2^*(\theta) \leq -1$.

Now let Agent 2 deviate to $(1 - \epsilon, \emptyset)$, where ϵ is a very small positive number. Given $m = (s_1(\theta), m_2^{dev}) = ((1, \emptyset), (1 - \epsilon, \emptyset))$ and any $\mu \in \Delta(\Theta)$, $\pi_2 > \pi_1$. To see why, note that when $\theta^1 = 1$,

$$u_d(\theta, y_1(1, \emptyset)) = u_d((1, \theta^2), (1, \theta^2 - 1)) = -1,$$

while

$$u_d(\theta, y_2(1 - \epsilon, \emptyset)) = u_d((1, \theta^2), (1 - \epsilon, \theta^2)) = -\epsilon^2 < -1.$$

On the other hand, when $\theta^1 = 0$,

$$u_d(\theta, y_1(1, \emptyset)) = u_d((0, \theta^2), (1, \theta^2 - 1)) = -2,$$

and

$$u_d(\theta, y_2(1 - \epsilon, \emptyset)) = u_d((0, \theta^2), (1 - \epsilon, \theta^2)) = -(1 - \epsilon)^2 > -2.$$

So $\beta(m) = 0$ and

$$\begin{aligned} u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) &= u_2(\theta, y_2(m_2^{dev})) \\ &= u_2((1, \theta^2), (1 - \epsilon, \theta^2)) \\ &= -(1 - \epsilon)^2 \\ &> -1 = u_2^*(\theta). \end{aligned}$$

Therefore any $s_1(\theta)$ that belongs to Subcase 1 cannot be sustained in equilibrium.

Subcase 2 $q(\theta) = 1$. For all $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((1, \emptyset), s_2(\theta))$, $\tilde{\theta}^1 = 0$.

If $s_1(\theta)$ falls into this subcase, then there exists $\theta_1 = (1, \theta_1^b)$ such that $(s_1(\theta_1), s_2(\theta_1)) = ((1, \emptyset), s_2(\theta_1))$ and $s_2(\theta_1) \neq s_2(\theta)$; otherwise it by definition contradicts $(q(\theta), \emptyset)$ belongs to Case 3.

Moreover, the DM's belief μ at θ_1 must assign probability $0 < p < 1$ to $\{\tilde{\theta} \mid \tilde{\theta}^1 = 1\}$. The reason is that, if $p = 0$, then μ is inconsistent with equilibrium strategy; if $p = 1$, then $s_1(\theta_1)$ belongs to Subcase 1, which we have established as unsustainable in equilibrium.

Since μ is consistent with the equilibrium strategies, there exists $\theta_2 = (0, 1 - \theta_2^2)$ such that $(s_1(\theta_2), s_2(\theta_2)) = ((1, \emptyset), s_2(\theta_1))$.

The DM elects Agent 2 with the same probability at states θ_1 and θ_2 since $(s_1(\theta_1), s_2(\theta_1)) = (s_1(\theta_2), s_2(\theta_2))$. If $\beta(s_1(\theta_1), s_2(\theta_1)) = \frac{1}{2}$, then $z(\theta_1) = \theta_1^1 - 1 = \theta_2^1 - 1$. This is impossible since $\theta_1^1 = 1$ and $\theta_2^1 = 0$. On the other hand, if $\beta(s_1(\theta_1), s_2(\theta_1)) = 0$, then $z(\theta_1) = \theta_1^1 = \theta_2^1$ which is also impossible. Now we have reached a contradiction. Therefore any $s_1(\theta)$ that belongs to Subcase 2 cannot be sustained in equilibrium.

Subcase 3 $q(\theta) = 1$. There exists $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((1, \emptyset), s_2(\theta))$ and $\tilde{\theta}^1 = 1$. There exists $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((1, \emptyset), s_2(\theta))$ and $\tilde{\theta}^1 = 0$.

This subcase is ruled out using the argument in the last paragraph of Subcase 2.

Subcase 4 $q(\theta) = 0$. For all $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((0, \emptyset), s_2(\theta))$, $\tilde{\theta}^1 = 1$.

We first prove the following claim:

Claim 5 In any equilibrium with $(s_1(\theta), s_2(\theta)) = ((q(\theta), \emptyset), (z(\theta), w(\theta)))$, there does not exist θ such that $\theta^1 = 1$ and $s_1(\theta) = (0, \emptyset)$.

Proof of Claim 5. Suppose on the contrary that $s_1(\theta) = (0, \emptyset)$ for $\theta = (1, \theta^2)$. First note that $u_2^*(\theta) = 0$, otherwise let Agent 2 deviate to $(0, \emptyset)$. Given

$m = ((0, \emptyset), (0, \emptyset))$, since for any $\mu \in \Delta(\Theta)$, $\pi_1 < \pi_2$, we have that $\beta(m) = 0$. Therefore $u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) = u_2(\theta, y_2(m_2^{dev})) = u_2(1, 0) = 0$.

Since $u_2(\theta, y_1(s_1(\theta))) < 0$, we must have $\beta(s_1(\theta), s_2(\theta)) = 0$ and $s_2(\theta) = (\theta^1 - 1, \theta^2)$. This contradicts Claim 3. ■

Now let's come back to Subcase 4. By definition of Subcase 4, $\theta^1 = 1$. By Claim 5, $s_1(\theta)$ that falls into Subcase 4 cannot be sustained in equilibrium.

Subcase 5 $q(\theta) = 0$. For all $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((0, \emptyset), s_2(\theta))$, $\tilde{\theta}^1 = 0$.

The argument in first paragraph for Case 1 shows that $u_2^*(\theta) \leq -1$.

Now let Agent 2 deviate to $(-\frac{1}{2}, \emptyset)$. Given $m = (s_1(\theta), m_2^{dev}) = ((0, \emptyset), (-\frac{1}{2}, \emptyset))$, by Claim 5, for any $\tilde{\theta} \in \Theta_2(m)$, $\tilde{\theta}^1 = 0$. So for any $\mu \in \Delta(\Theta_2(m))$, $\pi_1 = -1 < -\frac{1}{4} = \pi_2$ and $\beta(m) = 0$. So $u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) = u_2(\theta, y_2(m_2^{dev})) = u_2((0, \theta^2), (\frac{1}{2}, \theta^2)) = -\frac{1}{4} > u_2^*(\theta)$. Therefore any $s_1(\theta)$ that belongs to Subcase 5 cannot be sustained in equilibrium.

Subcase 6 $q(\theta) = 0$. There exists $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((0, \emptyset), s_2(\theta))$ and $\tilde{\theta}^1 = 1$. There exists $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((0, \emptyset), s_2(\theta))$ and $\tilde{\theta}^1 = 0$.

This subcase is ruled out using the argument in the last paragraph of Subcase 2.

Therefore all cases have been shown to be unsustainable in equilibrium.

4. I prove this by contradiction. Suppose the proposal profile is an equilibrium proposal profile. I first prove the following 3 claims.

Claim 1 $\beta(s_1(\theta), s_2(\theta)) \leq \frac{1}{2}, \forall \theta$.

Proof of Claim 1. Suppose on the contrary that for some $\bar{\theta}$, $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 1$. Then $u_2^*(\bar{\theta}) = u_2(\bar{\theta}, y_1(s_1(\bar{\theta}))) = u_2((\bar{\theta}^1, \bar{\theta}^2), (q(\bar{\theta}), \bar{\theta}^2 - 1))$.

Now let Agent 2 deviate to $s_1(\bar{\theta})$. Given $m = (s_1(\bar{\theta}), m_2^{dev}) = ((q(\bar{\theta}), \emptyset), (q(\bar{\theta}), \emptyset))$ and any $\mu \in \Delta(\Theta)$, $\pi_2 > \pi_1$. Therefore $\beta(m) = 0$ and $u_2^{dev}(\bar{\theta}, s_1(\bar{\theta}), m_2^{dev}) = u_2(\bar{\theta}, y_2(m_2^{dev})) = u_2((\bar{\theta}^1, \bar{\theta}^2), (q(\bar{\theta}), \bar{\theta}^2)) > u_2^*(\bar{\theta})$.

■

Claim 2 *If for some θ , $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$, then $z(\theta) = \theta^1 - 1$.*

Proof of Claim 2. Suppose on the contrary that $z(\theta) \neq \theta^1 - 1$. Therefore

$$\begin{aligned} u_2^*(\theta) &= \frac{1}{2}u_2(\theta, y_1(s_1(\theta))) + \frac{1}{2}u_2(\theta, y_2(s_2(\theta))) \\ &< \frac{1}{2}u_2(\theta, y_1(s_1(\theta))) + \frac{1}{2}u_2((\theta^1, \theta^2), (\theta^1 - 1, \theta^2)) \\ &= \frac{1}{2}u_2(\theta, y_1(s_1(\theta))). \end{aligned}$$

Now let Agent 2 deviate to (\emptyset, \emptyset) . Given $m = (s_1(\theta), m_2^{dev}) = ((q(\theta), \emptyset), (\emptyset, \emptyset))$ and any $\mu \in \Delta(\Theta)$, $\pi_1 \leq \pi_2$. Therefore $\beta(m) \leq \frac{1}{2}$ and

$$\begin{aligned} u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) &\geq \frac{1}{2}u_2(\theta, y_1(s_1(\theta))) + \frac{1}{2}u_2(\theta, y_2(m_2^{dev})) \\ &= \frac{1}{2}u_2(\theta, y_1(s_1(\theta))) \\ &> u_2^*(\theta). \end{aligned}$$

■

Claim 3 $\beta((1, \emptyset), (\delta, \emptyset)) = 0, \forall \delta \in (0, 1)$.

Proof of Claim 3. Let $m = ((1, \emptyset), (\delta, \emptyset))$. If $\theta^1 = 1$, then

$$u_d(\theta, y_1(m_1)) = u_d((1, \theta^2), (1, \theta^2 - 1)) = -1,$$

while

$$u_d(\theta - y_2(m_2)) = u_d((1, \theta^2), (\delta, \theta^2)) = -(1 - \delta)^2 > -1.$$

On the other hand, if $\theta^1 = 0$, then

$$u_d(\theta - y_1(m_1)) = u_d((0, \theta^2), (1, \theta^2 - 1)) = -2,$$

while

$$u_d(\theta - y_2(m_2)) = u_d((0, \theta^2), (\delta, \theta^2)) = -\delta^2 > -2.$$

Therefore for any $\mu \in \Delta(\Theta)$, $\pi_1 < \pi_2$ and $\beta(m) = 0$. ■

For the rest of the proof, we fix any θ and divide the discussion into 3 cases to show $(s_1(\theta), s_2(\theta)) = ((q(\theta), \emptyset), (z(\theta), \emptyset))$ cannot be sustained in equilibrium for any $q, z : \Theta \rightarrow [-1, 1]$.

Case 1 *For all $\tilde{\theta}$ such that $s_1(\tilde{\theta}) = s_1(\theta) = (q(\theta), \emptyset)$, $\tilde{\theta}^1 = q(\theta)$.*

There are two possibilities: $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$ and $\beta(s_1(\theta), s_2(\theta)) = 1$.

- If $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$, then by Claim 2,

$$\begin{aligned} u_2^*(\theta) &= \frac{1}{2}u_2(\theta, y_1(s_1(\theta))) + \frac{1}{2}u_2(\theta, y_2(s_2(\theta))) \\ &= \frac{1}{2}u_2((\theta^1, \theta^2), (\theta^1, \theta^2 - 1)) + \frac{1}{2}u_2((\theta^1, \theta^2), (\theta^1 - 1, \theta^2)) \\ &= -1. \end{aligned}$$

Now let Agent 2 deviate to $(q(\theta) - 1 + \epsilon, \emptyset)$ inconsistent with equilibrium, where ϵ is a very small positive number. This can be done since Agent 2 makes at most 4 different proposals in equilibrium. Given $m = (s_1(\theta), m_2^{dev}) = ((q(\theta), \emptyset), (q(\theta) - 1 + \epsilon, \emptyset))$ and any $\tilde{\theta} \in \Theta_2(m)$, $\tilde{\theta}^1 = q(\theta)$. Therefore for any $\mu \in \Delta(\Theta_2(m))$, $\pi_1 = -1 < -(1 - \epsilon)^2 = u_d((\theta^1, \theta^2), (\theta^1 - 1 + \epsilon, \theta^2)) = \pi_2$. Therefore $\beta(m) = 0$ and $u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) = u_2(\theta, y_2(m_2^{dev})) = u_2((\theta^1, \theta^2), (\theta^1 - 1 + \epsilon, \theta^2)) = -\epsilon^2 > -1 = u_2^*(\theta)$.

- If $\beta(s_1(\theta), s_2(\theta)) = 0$, then

$$u_2^*(\theta) = u_2(\theta, y_2(s_2(\theta))) = u_2(\theta, y_2(z(\theta), \emptyset)) = -(\theta^1 - z(\theta) - 1)^2.$$

Moreover, given the DM's belief μ at θ , $\pi_1 < \pi_2$. Since $(q(\theta), \emptyset)$ belongs to Case 1, the DM learns $\theta^1 = q(\theta)$ at θ . Therefore $\pi_1 = -1 < \pi_2 = u_d((\theta^1, \theta^2), (z(\theta), \theta^2))$. Thus $|\theta^1 - z(\theta)| < 1$.

Now let Agent 2 deviate to $(z(\theta) - \epsilon, \emptyset)$ inconsistent with equilibrium, where $z(\theta) - \epsilon > \theta^1 - 1$. Therefore $\theta^1 - z(\theta) + \epsilon \in (-1, 1)$.

Given $m = (s_1(\theta), m_2^{dev}) = ((q(\theta), \emptyset), (z(\theta) - \epsilon, \emptyset))$ and any $\tilde{\theta} \in \Theta_2(m)$, $\tilde{\theta}^1 = q(\theta)$. Therefore for any $\mu \in \Delta(\Theta_2(m))$, $\pi_1 = -1 < -(\theta^1 - z(\theta) + \epsilon)^2 = u_d((\theta^1, \theta^2), (z(\theta) - \epsilon, \theta^2)) = \pi_2$. Therefore $\beta(m) = 0$ and

$$\begin{aligned} u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) &= u_2(\theta, y_2(m_2^{dev})) \\ &= u_2((\theta^1, \theta^2), (z(\theta) - \epsilon, \theta^2)) \\ &= -(\theta^1 - z(\theta) + \epsilon - 1)^2 \\ &> -(\theta^1 - z(\theta) - 1)^2 \\ &= u_2^*(\theta). \end{aligned}$$

Therefore $s_1(\theta)$ that falls into Case 1 cannot be sustained in equilibrium.

Case 2 For all $\tilde{\theta}$ such that $s_1(\tilde{\theta}) = s_1(\theta)$, $\tilde{\theta}^1 \neq q(\theta)$.

I show that if $s_1(\theta)$ falls into this case, then $z(\theta) = \theta^1 - 1$ and $\beta(s_1(\theta), s_2(\theta)) = 0$. To see this, suppose the contrary. Then $u_2^*(\theta) < 0$. Now let Agent 2 deviate to (\emptyset, \emptyset) . Given $m = (s_1(\theta), m_2^{dev}) = ((q(\theta), \emptyset), (\emptyset, \emptyset))$ and any

$\mu \in \Delta(\Theta_2(m)), \pi_1 < -1 = \pi_2$. Therefore $\beta(m) = 0$ and $u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) = u_2(\theta, y_2(m_2^{dev})) = 0 > u_2^*(\theta)$.

Therefore for any $s_1(\theta)$ that falls into Case 2, Agent 2 is proposing his ideal action $(\theta^1 - 1, \emptyset)$ and wins with probability 1.

Case 3 *There exists $\tilde{\theta}$ such that $s_1(\tilde{\theta}) = s_1(\theta) = (q(\theta), \emptyset)$ and $\tilde{\theta}^1 = q(\theta)$, and there exists $\tilde{\theta}$ such that $s_1(\tilde{\theta}) = s_1(\theta) = (q(\theta), \emptyset)$ and $\tilde{\theta}^1 \neq q(\theta)$.*

In this case, $q(\theta) \in \{0, 1\}$. I further divide this case into 6 subcases.

Subcase 1 $q(\theta) = 1$. *For all $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((1, \emptyset), s_2(\theta))$, $\tilde{\theta}^1 = 1$.*

There are two possibilities: $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$ and $\beta(s_1(\theta), s_2(\theta)) = 1$.

- If $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$, then by Claim 2, $z(\theta) = \theta^1 - 1$ and $u_2^*(\theta) = \frac{1}{2}u_2(\theta, y_1(s_1(\theta))) + \frac{1}{2}u_2(\theta, y_2(s_2(\theta))) = -1$.

Now let Agent 2 deviate to an off path $(1 - \epsilon, \emptyset)$, where ϵ is a very small positive number. Given $m = (s_1(\theta), m_2^{dev}) = ((1, \emptyset), (1 - \epsilon, \emptyset))$, by Claim 3, $\beta(m) = 0$. Therefore $u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) = u_2(\theta, y_2(m_2^{dev})) = u_2((1, \theta^2), (1 - \epsilon, \theta^2)) = -(1 - \epsilon)^2 > -1 = u_2^*(\theta)$.

- If $\beta(s_1(\theta), s_2(\theta)) = 1$, then $u_2^*(\theta) = u_2(\theta, y_2(s_2(\theta))) = u_2((1, \theta^2), (z(\theta), \theta^2)) = -(z(\theta))^2$.

Moreover, given the DM's belief μ at θ , $\pi_1 < \pi_2$. Since $s_1(\theta)$ belongs to Subcase 1, the DM learns $\theta^1 = q(\theta) = 1$ at θ . Therefore $\pi_1 = -1$ and $\pi_2 = u_d((1, \theta^2), (z(\theta), \theta^2)) = -(1 - z(\theta))^2$, where $|1 - z(\theta)| < 1$. Therefore $z(\theta) \in (0, 1]$.

Now let Agent 2 deviate to an $(z(\theta) - \epsilon, \emptyset)$ inconsistent with equilibrium, where $z(\theta) - \epsilon \in (0, 1)$. Given $m = (s_1(\theta), m_2^{dev}) = ((1, \emptyset), (z(\theta) - \epsilon, \emptyset))$, by Claim 3, $\beta(m) = 0$. Therefore

$$\begin{aligned} u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) &= u_2(\theta, y_2(m_2^{dev})) \\ &= u_2((1, \theta^2), (z(\theta) - \epsilon, \theta^2)) \\ &= -(z(\theta) - \epsilon)^2 \\ &> -(z(\theta))^2 \\ &= u_2^*(\theta). \end{aligned}$$

Therefore $s_1(\theta)$ that falls into Subcase 1 cannot be sustained in equilibrium.

Subcase 2 $q(\theta) = 1$. For all $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((1, \emptyset), s_2(\theta))$, $\tilde{\theta}^1 = 0$.

If $s_1(\theta)$ falls into this subcase, then there exists $\theta_1 = (1, \theta_1^b)$ such that $(s_1(\theta_1), s_2(\theta_1)) = ((1, \emptyset), s_2(\theta_1))$ and $s_2(\theta_1) \neq s_2(\theta)$; otherwise it contradicts the definition of Case 3 and Subcase 2.

Moreover, the DM's belief μ at θ_1 must assign probability $0 < p < 1$ to $\{\tilde{\theta} \mid \tilde{\theta}^1 = 1\}$. The reason is that, if $p = 0$, then μ is inconsistent with equilibrium strategies; if $p = 1$, then $s_1(\theta_1)$ belongs to Subcase 1, which we have established as unsustainable in equilibrium.

Since μ is consistent with equilibrium strategy, there exists $\theta_2 = (0, 1 - \theta^2)$ such that $(s_1(\theta_2), s_2(\theta_2)) = (s_1(\theta_1), s_2(\theta_1)) = ((1, \emptyset), s_2(\theta_1))$.

The DM elects Agent 2 with the same probability at states θ_1 and θ_2 since $(s_1(\theta_1), s_2(\theta_1)) = (s_1(\theta_2), s_2(\theta_2))$. If $\beta(s_1(\theta_1), s_2(\theta_1)) = \frac{1}{2}$, then by Claim 2 we have $z(\theta_1) = \theta_1^a - 1 = \theta_2^a - 1$. But this is impossible since $\theta_1^1 = 1$ and $\theta_2^1 = 0$. Therefore $\beta(s_1(\theta_1), s_2(\theta_1)) = 0$ and $u_2^*(\theta_1) = u_2((1, \theta_1^2), (z(\theta_1), \theta_1^2)) = -(z(\theta_1))^2$.

I argue that $z(\theta_1) = z(\theta_2) = 0$; otherwise at state θ_1 , let Agent 2 deviate to (ϵ, \emptyset) inconsistent with equilibrium such that $0 < \epsilon < |z(\theta_1)|$. Then given $\tilde{m} = ((1, \emptyset), (\epsilon, \emptyset))$, $\beta(\tilde{m}) = 0$ by Claim 3. Therefore $u_2^{dev}(\theta_1, (1, \emptyset), (\epsilon, \emptyset)) = u_2(\theta, y_2(m_2^{dev})) = -\epsilon^2 > -(z(\theta_1))^2 = u_2^*(\theta_1)$.

Therefore we have that $\beta((1, \emptyset), (0, \emptyset)) = 0$ since $(s_1(\theta_1), s_2(\theta_1)) = ((1, \emptyset), (0, \emptyset))$ and $\beta(s_1(\theta_1), s_2(\theta_1)) = 0$.

This means for any $\tilde{\theta}$ such that $\tilde{\theta}^1 = 1$ and $s_1(\tilde{\theta}) = (1, \emptyset)$, $s_2(\tilde{\theta}) = (0, \emptyset)$ and $u_2^*(\tilde{\theta}) = 0$; otherwise, Agent 2 has incentive to deviate to $(0, \emptyset)$.

Now I show that Agent 2 has a profitable deviation at θ_2 . First note that $u_2^*(\theta_2) = u_2(\theta_2, y_2(s_2(\theta_2))) = u_2((0, \theta_2^2), (0, \theta_2^2)) = -1$. Now let Agent 2 deviate to $(\epsilon - 1, \emptyset)$ inconsistent with equilibrium where ϵ is a very small positive number. Then given $m = (s_1(\theta), m_2^{dev}) = ((1, \emptyset), (\epsilon - 1, \emptyset))$,

$$\Theta_2(m) = \{\tilde{\theta} \mid s_1(\tilde{\theta}) = (1, \emptyset),$$

$$u_2^*(\tilde{\theta}) \leq \max_{\beta \in \text{BR}(\Theta, m)} (1 - \beta)u_2(\tilde{\theta}, y_2(\epsilon - 1, \emptyset)) + \beta u_2(\tilde{\theta}, y_1(1, \emptyset))\}.$$

For $\tilde{\theta}$ such that $\tilde{\theta}^1 = 1$ and $s_1(\tilde{\theta}) = (1, \emptyset)$, since

$$u_2^*(\tilde{\theta}) = 0 > u_2(\tilde{\theta}, y_1(1, \emptyset))$$

and

$$u_2^*(\tilde{\theta}) = 0 > u_2(\tilde{\theta}, y_2(\epsilon - 1, \emptyset)),$$

we have

$$u_2^*(\tilde{\theta}) > \max_{\beta \in \text{BR}(\Theta, m)} (1 - \beta)u_2(\tilde{\theta}, y_2(\epsilon - 1, \emptyset)) + \beta u_2(\tilde{\theta}, y_1(1, \emptyset)).$$

Therefore for any $\tilde{\theta} \in \Theta_2(m)$, $\tilde{\theta}^1 = 0$. Thus for any $\mu \in \Delta(\Theta_2(m))$,

$$\pi_1 = u_d((0, \tilde{\theta}^2), y_1(1, \emptyset)) = u_d(0, \tilde{\theta}^2), (1, \tilde{\theta}^2 - 1)) = -2$$

while

$$\pi_2 = u_d((0, \tilde{\theta}^2), y_2(\epsilon - 1, \emptyset)) = u_d((0, \tilde{\theta}^2), (\epsilon - 1, \tilde{\theta}^2)) = -(1 - \epsilon)^2 > -2.$$

So $\beta(m) = 0$ and

$$\begin{aligned} u_2^{dev}(\theta_2, (1, \emptyset), (\epsilon - 1, \emptyset)) &= u_2((0, \theta_2^b), y_2(\epsilon - 1, \emptyset)) \\ &= u_2((0, \theta_2^b), (\epsilon - 1, \theta_2^b)) \\ &= -\epsilon^2 \\ &> -1 \\ &= u_2^*(\theta_2). \end{aligned}$$

Therefore $s_1(\theta)$ that falls into Subcase 2 cannot be sustained in equilibrium.

Subcase 3 $q(\theta) = 1$. There exists $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((1, \emptyset), s_2(\theta))$ and $\tilde{\theta}^1 = 1$. There exists $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((1, \emptyset), s_2(\theta))$ and $\tilde{\theta}^1 = 0$.

Let $\theta_1 = (0, \theta_1^2)$ and $\theta_2 = (1, \theta_2^2)$ be such that $(s_1(\theta_1), s_2(\theta_1)) = (s_1(\theta_2), s_2(\theta_2))$. The same argument from Subcase 2 shows that Agent 2 has incentive to deviate at θ_2 . Therefore $s_1(\theta)$ that falls into Subcase 3 cannot be sustained in equilibrium.

Subcase 4 $q(\theta) = 0$. For all $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((0, \emptyset), s_2(\theta))$, $\tilde{\theta}^1 = 0$.

I first prove the following claim:

Claim 4 If for some $\theta = (1, \theta^2)$, $s_1(\theta) = (0, \emptyset)$, then $s_2(\theta) = (0, \emptyset)$ and $\beta(s_1(\theta), s_2(\theta)) = 0$.

Proof of Claim 4. $u_2^*(\theta) = 0$; otherwise, Agent 2 has a profitable deviation $(0, \emptyset)$. To see this, note that given $m = ((0, \emptyset), (0, \emptyset))$, for any $\mu \in \Delta(\Theta)$, $\pi_1 < \pi_2$ and therefore $\beta(m) = 0$. Therefore $u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) = u_2(\theta, y_2(m_2^{dev})) = 0$.

Since $u_2(\theta, y_1(s_1(\theta))) < 0$, we have that $\beta(s_1(\theta), s_2(\theta)) = 0$ and that $s_2(\theta) = (0, \emptyset)$. ■

Now let's come back to Subcase 4. There are two possibilities: $z(\theta) > -1$ and $z(\theta) = -1$.

- If $z(\theta) > -1$, then Agent 2 has a profitable deviation to $(z(\theta) - \epsilon, \emptyset)$ inconsistent with equilibrium where $z(\theta) - \epsilon > -1$. To see this, note that given $m = (s_1(\theta), m_2^{dev}) = ((0, \emptyset), (z(\theta) - \epsilon, \emptyset))$, by Claim 4, any $\tilde{\theta}$ with $\tilde{\theta}^1 = 1$ does not belong to $\Theta_2(m)$. Therefore if $\tilde{\theta} \in \Theta_2(m)$, then $\tilde{\theta}^1 = 0$. Thus for any $\mu \in \Delta(\Theta_2(m))$, $\pi_1 < \pi_2$ and $\beta(m) = 0$. So

$$\begin{aligned} u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) &= u_2(\theta, y_2(m_2^{dev})) \\ &= u_2((0, \theta^2), y_2(z(\theta) - \epsilon, \emptyset)) \\ &= -(z(\theta) + 1 - \epsilon)^2 \\ &> -(z(\theta) + 1)^2 \\ &= u_2(\theta, y_2(s_2(\theta))) \\ &\geq u_2^*(\theta). \end{aligned}$$

The last inequality comes from the following observation: if an agent wins with positive probability, then he weakly prefers the outcome from his own proposal to that from his opponent's proposal. To see this, fix any θ . Let $(s_1(\theta), s_2(\theta))$ be denoted by m . Suppose that $\beta(m) > 0$ and $u_1(\theta, y_1(m_1)) < u_1(\theta, y_2(m_2))$. Therefore $u_1^*(\theta) < u_1(\theta, y_2(m_2))$. Now let Agent 1 deviate to $m_1^{dev} = (\emptyset, \emptyset)$. Then $u_1^{dev}(\theta, m_1^{dev}, m_2) \geq u_1(\theta, y_2(m_2)) > u_1^*(\theta)$, making m_1^{dev} a profitable deviation. This means that $\beta(m) > 0$ implies $u_1(\theta, y_1(m_1)) \leq u_1(\theta, y_2(m_2))$. The case for Agent 2 is similar.

- $z(\theta) = -1$. Let μ be the DM's belief at θ . Then by the definition of Subcase 4, $\mu\{\tilde{\theta} \mid \tilde{\theta}^1 = 0\} = 1$. Given $m = (s_1(\theta), s_2(\theta)) = ((0, \emptyset), (-1, \emptyset))$ and μ , $\pi_1 = \pi_2$. Therefore $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$ and $u_2^*(\theta) = \frac{1}{2}u_2(\theta, y_1(s_1(\theta))) + \frac{1}{2}u_2(\theta, y_2(s_2(\theta))) = -1$.

Now let Agent 2 deviate to $(\epsilon - 1, \emptyset)$ inconsistent with equilibrium. Given $m = (s_1(\theta), m_2^{dev}) = ((0, \emptyset), (\epsilon - 1, \emptyset))$, same argument for the $z(\theta) > -1$ case above shows that for any $\tilde{\theta} \in \Theta(m)$, $\tilde{\theta}^1 = 0$. Therefore $\beta(m) = 0$

and $u_2^{dev}(\theta, s_1(\theta), m_2^{dev}) = u_2(\theta, y_2(m_2^{dev})) = u_2((0, \theta^2), (\epsilon - 1, \theta^2)) = -\epsilon^2 > u_2^*(\theta)$.

Therefore if $s_1(\theta)$ falls into Subcase 4, then it cannot be sustained in equilibrium.

Subcase 5 $q(\theta) = 0$. For all $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((0, \emptyset), s_2(\theta))$, $\tilde{\theta}^1 = 1$.

By the definition of Case 3 and Subcase 5, there exists $\theta_1 = (0, \theta_1^b)$ such that $s_1(\theta_1) = (0, \emptyset)$ and $s_2(\theta_1) \neq s_2(\theta)$. Then the DM's belief μ at θ must assign probability $0 < p < 1$ to $\{\tilde{\theta} \mid \tilde{\theta}^1 = 0\}$: if $p = 0$, then μ is inconsistent with equilibrium strategies; if $p = 1$, then $s_1(\theta_1)$ falls into Subcase 4, which we have shown to be unsustainable in equilibrium.

Therefore there exists $\theta_2 = (1, \theta_2^b)$ such that $(s_1(\theta_2), s_2(\theta_2)) = (s_1(\theta_1), s_2(\theta_1))$. The same argument in Subcase 2 shows that Agent 2 has a profitable deviation at θ_1 .

Therefore if $s_1(\theta)$ falls into Subcase 5, then it cannot be sustained in equilibrium.

Subcase 6 $q(\theta) = 0$. There exists $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((0, \emptyset), s_2(\theta))$ and $\tilde{\theta}^1 = 1$. There exists $\tilde{\theta}$ such that $(s_1(\tilde{\theta}), s_2(\tilde{\theta})) = ((0, \emptyset), s_2(\theta))$ and $\tilde{\theta}^1 = 0$.

Let $\theta_1 = (0, \theta_1^b)$ and $\theta_2 = (1, \theta_2^b)$ be such that $(s_1(\theta_1), s_2(\theta_1)) = (s_1(\theta_2), s_2(\theta_2))$. Same argument for Subcase 5 shows that Agent 2 has a profitable deviation at θ_1 .

Therefore if $s_1(\theta)$ falls into Subcase 6, then it cannot be sustained in equilibrium.

So the only sustainable case is Case 2. Therefore in any equilibrium in which $(s_1(\theta), s_2(\theta)) = ((q(\theta), \emptyset), (z(\theta), \emptyset))$ for all θ , $z(\theta) = \theta^1 - 1$, $q(\theta) \neq \theta^1$, and $\beta(s_1(\theta), s_2(\theta)) = 0$.

Now let Agent 1 deviate to (\emptyset, \emptyset) at any state θ . Given $m = (m_1^{dev}, s_2(\theta)) = ((\emptyset, \emptyset), (\theta^1 - 1, \emptyset))$, for any $\tilde{\theta} \in \Theta_1(m)$, $\tilde{\theta}^1 = 1 + z(\theta)$. Therefore for any $\mu \in \Delta(\Theta_1(m))$, $\pi_1 = -1 = \pi_2$ and $\beta(m) = \frac{1}{2}$. Hence, $u_1^{dev}(\theta, m_1^{dev}, s_2(\theta)) = -1 > -2 = u_1^*(\theta)$. (\emptyset, \emptyset) is a profitable deviation for Agent 1.

■

Appendix III. N dimensions

I focus on the case in which for each dimension $k \in N$, there is $i \in \{1, 2\}$ such that $|b_i^k| < |b_j^k|$. That is, for each dimension, one agent is more aligned with the DM than his opponent. I call such a dimension j Agent i 's aligned dimension and Agent j 's misaligned dimension. The set of Agent i 's aligned dimensions is denoted by N_i . The following result shows that in this environment, agents should not be vague on their misaligned dimensions and commit on their aligned dimensions.

Proposition 6 *For $i \in \{1, 2\}$, suppose that $\forall k \in N_i, s_i^j(\theta) = \theta^j - b_i^j$; otherwise $s_i^k(\theta) = \emptyset$. Then (s_1, s_2) is not part of an equilibrium strategy profile. Furthermore, suppose that for each i , whenever $k \in N_i, |b_i^k| = 0$. For all $k \in N_i, s_i^k(\theta) \in \mathbb{R}$; otherwise $s_i^k(\theta) = \emptyset$. Then (s_1, s_2) is not part of an equilibrium strategy profile.*

The intuition of the result is similar as the previously-analyzed environment with two dimensions and completely (mis)aligned agents. First, if such (s_1, s_2) is part of an intuitive equilibrium, then the outcome must be a tie, otherwise the losing agent can mimic the winning agent's proposal and win. Second, each specific commitment must be its proposer's ideal action. To see this, suppose that Agent i commits on dimension k and that his commitment is not equal to $\theta^k - b_i^k$. Then let Agent i deviate to be vague on all dimensions. The DM then believes that Agent i will implement his own ideal action, whereas Agent j will at best implement his own ideal action. Therefore Agent i 's probability of winning after the deviation is no worse than in equilibrium and his winning proposal is strictly better than in equilibrium. Lastly, given that specific commitment is revealing about the state, the undercutting argument destroys the equilibrium.

Proof of Proposition 6. Suppose that for all θ , for $\forall k \in N_i, s_i^k(\theta) = \theta^k - b_i^k$; otherwise $s_i^k(\theta) = \emptyset$. In equilibrium, both agents propose their ideal actions.

Step 1. $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}, \forall \theta$.

Suppose on the contrary that $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 1$ for some $\bar{\theta}$. Then

$$u_2^*(\bar{\theta}) = u_2(\bar{\theta}, y_1(s_1(\bar{\theta}))) = u_2(\bar{\theta}, \bar{\theta} - b_1) = -\|b_1 - b_2\|^2.$$

Let Agent 2 deviate to $s_1(\bar{\theta})$. Given $m = (s_1(\bar{\theta}), s_1(\bar{\theta}))$, the DM believes that only Agent 2 has deviated. Let $y_2 = (y_2^1, y_2^2)$ be Agent 2's policy. For $k \in N_1$,

$s_2^k(\bar{\theta}) = \bar{\theta}^k - b_1^k$ so $y_2^k = \bar{\theta}^k - b_1^k$. For $k \in N_2$, $s_2^k(\bar{\theta}) = \emptyset$ so $y_2^k = \bar{\theta}^k - b_2^k$.

$$\pi_1 = \int_{\Theta} u_d(\theta, \theta - b_1) \mu d\theta = -\|b_1\|^2,$$

$$\pi_2 = \int_{\Theta} u_d(\theta, y_2(s_1(\bar{\theta}))) \mu d\theta = - \sum_{k \in N_1} |b_1^k|^2 - \sum_{k \in N_2} |b_2^k|^2 > -\|b_1\|^2.$$

So $\beta(m) = 0$ and

$$u_2^{dev}(\bar{\theta}, m_2^{dev}, s_1(\bar{\theta})) = u_2(\bar{\theta}, y_2(s_1(\bar{\theta}))) = - \sum_{k \in N_1} |b_1^k - b_2^k|^2 > u_2^*(\bar{\theta}).$$

$s_1(\bar{\theta})$ is then a profitable deviation for Agent 2. The proof is similar for when $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 0$ for some $\bar{\theta}$. Since $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2} \forall \theta$, the DM is indifferent between agents. Since both agents propose their ideal actions, we can only have $\|b_1\| = \|b_2\|$. Therefore for all θ ,

$$u_1^*(\theta) = -\frac{1}{2} \|b_1 - b_2\|^2.$$

Step 2. Agent 1 has a profitable deviation.

At $\bar{\theta}$, let Agent 1 deviate to m_1^{dev} such that: for some $n \in N_2$, $m_1^{dev, n} = \bar{\theta}^n - (1 - \epsilon)b_1^n$. $m_1^{dev, k} = \emptyset$ for all $k \neq n$. Given $m = (m_1^{dev}, s_2(\bar{\theta}))$, the DM believes that only Agent 1 has deviated since $m_1^{dev} \notin \text{supp}(G_1)$. Therefore $\mu\{\bar{\theta} \mid \bar{\theta}^n = s_2^n(\bar{\theta}) + b_2^n\} = 1$.

$$\pi_2 = -\|b_2\|^2,$$

$$\pi_1 = - \sum_{k \neq n} |b_1^k|^2 - |(1 - \epsilon)b_1^n|^2 > -\|b_1\|^2 = -\|b_2\|^2.$$

So $\beta(m) = 1$ and Agent 1 gets arbitrarily close to outcome b_1 , which makes m_1^{dev} a profitable deviation.

Now suppose that for each i , $|b_i^k| = 0$ for all $k \in N_i$. Consider putative equilibrium strategies for agents as follows:

$$s_1^k(\theta) = w^k(\theta), \forall k \in N_1, \forall \theta; s_1^k(\theta) = \emptyset, \forall k \in N_2, \forall \theta,$$

$$s_2^l(\theta) = z^l(\theta), \forall l \in N_2, \forall \theta; s_2^l(\theta) = \emptyset, \forall l \in N_1, \forall \theta,$$

where $w^k : \Theta \rightarrow \mathbb{R}$ for $k \in N_1$ and $z^l : \Theta \rightarrow \mathbb{R}$ for $l \in N_2$.

Step 1. $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$ for all θ .

Suppose on the contrary that $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 1$ for some $\bar{\theta}$. Then

$$u_2^*(\bar{\theta}) = u_2(\bar{\theta}, y_1(s_1(\bar{\theta}))) = - \sum_{k \in N_1} |\theta^k - b_2^k - s_1^k(\bar{\theta})|^2 - \sum_{k \in N_2} |b_1^k - b_2^k|^2.$$

Now let Agent 2 deviate to $s_1(\bar{\theta})$. Given $m = (s_1(\bar{\theta}), s_1(\bar{\theta}))$, for any $\mu \in \Delta(\Theta)$,

$$\begin{aligned}\pi_1 &= \int_{\Theta} u_d(\theta, y_1(s_1(\bar{\theta}))) \mu d\theta = \int_{\Theta} - \sum_{k \in N_1} |\theta^k - s_1^k(\bar{\theta})|^2 \mu d\theta - \sum_{k \in N_2} |b_1^k|^2, \\ \pi_2 &= \int_{\Theta} u_d(\theta, y_2(s_1(\bar{\theta}))) \mu d\theta = \int_{\Theta} - \sum_{k \in N_1} |\theta^k - s_1^k(\bar{\theta})|^2 \mu d\theta - \sum_{k \in N_2} |b_2^k|^2.\end{aligned}$$

Since for any $k \in N_2$, $|b_1^k| > |b_2^k|$, $-\sum_{k \in N_2} |b_1^k|^2 < -\sum_{k \in N_2} |b_2^k|^2$. So $\pi_1 < \pi_2$ and $\beta(m) = 0$.

$$u_2^{dev}(\bar{\theta}) = - \sum_{k \in N_1} |\theta^k - b_2^k - s_1^k(\bar{\theta})|^2 - \sum_{k \in N_2} |b_2^k - b_2^k|^2 > u_2^*(\bar{\theta}),$$

making $s_1(\bar{\theta})$ a profitable deviation for Agent 2.

Step 2. For all $k \in N_1$, $w^k(\theta) = \theta^k - b_1^k$ for all θ ; for all $l \in N_2$, $z^l(\theta) = \theta^l - b_2^l$ for all θ .

Without loss of generality, I assume that $\|b_1\| \leq \|b_2\|$. Suppose for some $\bar{\theta}$ and some $\bar{k} \in N_1$, $w^{\bar{k}}(\bar{\theta}) \neq \bar{\theta}^{\bar{k}} - b_1^{\bar{k}}$. Then

$$\begin{aligned}u_1^*(\bar{\theta}) &= \frac{1}{2} u_1(\bar{\theta}, y_1(s_1(\bar{\theta}))) + \frac{1}{2} u_1(\bar{\theta}, y_2(s_2(\bar{\theta}))) \\ &= \frac{1}{2} \left[\sum_{k \neq \bar{k} \in N_1} -|\bar{\theta}^k - s_1^k(\bar{\theta}) - b_1^k|^2 - |\bar{\theta}^{\bar{k}} - w^{\bar{k}}(\bar{\theta}) - b_1^{\bar{k}}|^2 \right] + \frac{1}{2} u_1(\bar{\theta}, y_2(s_2(\bar{\theta}))) \\ &< \frac{1}{2} \left[\sum_{k \neq \bar{k} \in N_1} -|\bar{\theta}^k - s_1^k(\bar{\theta}) - b_1^k|^2 \right] + \frac{1}{2} u_1(\bar{\theta}, y_2(s_2(\bar{\theta}))).\end{aligned}$$

Let Agent 1 deviate to $m_1^{dev} = (\emptyset, \emptyset, \dots, \emptyset)$ at $\bar{\theta}$. Given $m = (m_1^{dev}, s_2(\bar{\theta}))$, the DM believes that only Agent 1 has deviated and

$$\begin{aligned}\pi_1 &= -\|b_1\|^2 \geq -\|b_2\|^2, \\ \pi_2 &= - \int_{\Theta} u_d(\theta, y_2(s_2(\bar{\theta}))) \mu d\theta = \int_{\Theta} - \sum_{l \in N_2} |\theta^l - z^l(\bar{\theta})|^2 \mu d\theta - \sum_{k \in N_1} |b_2^k|^2 \leq -\|b_2\|^2.\end{aligned}$$

Therefore $\beta(m) \geq \frac{1}{2}$ and

$$u_1^{dev}(\bar{\theta}) \geq \frac{1}{2} u_1(\bar{\theta}, y_1(m_1^{dev})) + \frac{1}{2} u_1(\bar{\theta}, y_2(s_2(\bar{\theta}))) = \frac{1}{2} u_1(\bar{\theta}, y_2(s_2(\bar{\theta}))) > u_1^*(\bar{\theta}),$$

making m_1^{dev} a profitable deviation. Therefore for all $k \in N_1$, $w^k(\theta) = \theta^k - b_1^k$ for all θ . This implies that for all θ , $\pi_1 = u_d(\theta, \theta - b_1) = -\|b_1\|^2$. Since $\beta(m) = \frac{1}{2}$ for all θ and for all $\mu \in \Delta(\Theta)$, $\pi_2 \leq -\|b_2\|^2 \leq \pi_1$, above s_1 and s_2 can only be sustained in equilibrium if $\|b_1\| = \|b_2\|$. Given that $\|b_1\| = \|b_2\|$, I can use the argument above to show that for all $l \in N_2$, $z^l(\theta) = \theta^l - b_2^l$.

Step 3. Agent 1 has a profitable deviation.

This step is identical to the Step 2 in the first part of the proof. ■

Appendix IV. General Preferences of the DM

Now we focus back on the case in which there are two dimensions and agents' ideal outcomes are $(0, 1)$ and $(1, 0)$. Let the DM's payoff function be $u_d(\theta, y) = -(\theta^1 - y^a)^2 - \alpha(\theta^2 - y^b)^2$, where $0 < \alpha < 1$. Agents' payoff functions are as before:

$$u_1(\theta, y) = -(\theta^1 - y^1)^2 - (\theta^2 - y^2 - 1)^2,$$

$$u_2(\theta, y) = -(\theta^1 - y^1 - 1)^2 - (\theta^2 - y^2)^2.$$

Since $u_d(\theta, y_1(\emptyset, \emptyset)) = -\alpha$, $u_v(\theta, y_2(\emptyset, \emptyset)) = -1$, $\forall \theta$, Agent 1 has an overall advantage. The next result shows that in this case, agents should not be vague on their misaligned dimensions and commit on their aligned dimensions.

Proposition 7 *Suppose that $b_1 = (0, 1)$ and $b_2 = (1, 0)$ and the DM weighs dimension 1 more than dimension 2. In all equilibria with vagueness, vagueness occurs on agents' aligned dimensions.*

The intuition is simple. The same argument as before show that each agent wins with probability $\frac{1}{2}$, which means that Agent 1 cannot get his own ideal action with probability 1. Now let Agent 1 deviate to be vague on both dimensions. Then the DM believes that Agent 1 will implement his own ideal action, while Agent 2 will at best implement his own ideal action, which is strictly less preferred to Agent 1's. Therefore Agent 1 wins with probability 1 and gets his own ideal action.

Proof of Proposition 7. Suppose on the contrary that for all θ ,

$$s_1(\theta) = (w(\theta), \emptyset),$$

$$s_2(\theta) = (\emptyset, z(\theta)).$$

From Lemma 1, $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}$ for all θ . Therefore $u_1^*(\theta) < 0$ for all θ .

Let Agent 1 deviate to (\emptyset, \emptyset) at $\bar{\theta}$. Given $m = (m_1^{dev}, s_2(\bar{\theta}))$ and for any $\mu \in \Delta(\Theta)$,

$$\pi_1 = -\alpha,$$

$$\pi_2 = \int_{\Theta} u_d(\theta, y_2(\emptyset, z(\bar{\theta}))) \mu d\theta \leq -\|b_2\|^2 = -1 < \pi_1$$

So $\beta(m_1^{dev}, s_2(\bar{\theta})) = 1$ and

$$u_1^{dev}(\bar{\theta}, (\emptyset, \emptyset), s_2(\bar{\theta})) = u_1(\bar{\theta}, y_1(\emptyset, \emptyset)) = 0 > u_1^*(\bar{\theta}),$$

making (\emptyset, \emptyset) a profitable deviation for Agent 1. ■

Appendix V. Committing to Subsets of Actions

So far, the definition of vagueness is a complete lack of commitment, which is equivalent to committing to the entire action space. One may also think of a less extreme kind of vagueness and let agents commit to strict subsets of actions. Formally, let $M_i = (2^{\mathbb{R}})^2$ and $s_i : \Theta \rightarrow (2^{\mathbb{R}})^2$. So for each dimension, an agent can choose an arbitrary subset of the action space to commit to. Note that the previous notion of vagueness is committing to \mathbb{R} under this new definition.

The next proposition shows that it is still not part of an equilibrium in which both agents are specific about their aligned dimensions, as long as vagueness on their misaligned dimensions is equivalent to committing to the entire action set.

Proposition 8 *Suppose that $b_1 = (0, 1)$, $b_2 = (1, 0)$. Let $y, z : \Theta \rightarrow \mathbb{R}$ and $A, B : \Theta \rightarrow 2^{\mathbb{R}}$ be such that $\theta^2 - b_1^2 \in A(\theta)$ and $\theta^1 - b_2^1 \in B(\theta)$ for all θ . Then $s_1(\theta) = (w(\theta), A(\theta))$, $s_2(\theta) = (B(\theta), z(\theta))$ cannot be supported in an equilibrium.*

In other words, there does not exist an equilibrium in which each agent is specific about his aligned dimension and committing to a subset containing his ideal action on his misaligned dimension.

The idea of the proof is very similar to the previous case, where the agents can only choose between being precise and complete vagueness. For simplicity I focus on the case in which $b_1 = (0, 1)$ and $b_2 = (1, 0)$. Note first that as long as a subset contains the proposer's ideal action, the action that will be implemented is his ideal action. The proof consists of three steps. First, each agent wins with probability $\frac{1}{2}$. The reason is, if some agent wins, the losing agent can deviate to mimic the winner's proposed action on his specific dimension and commit to \mathbb{R} on the winner's vague dimension. Now two agents are both specific on the loser's misaligned dimension. On the loser's aligned dimension, both implement their ideal actions. So now the loser is preferred by the DM and wins with a strictly better outcome on his aligned dimension. Next, $w(\theta) = \theta^1 - b_1^1$ and $z(\theta) = \theta^2 - b_2^2$. That is, on their specific dimensions, both agents commit to their ideal actions. If an agent commits to an alternative action at some state, then he can deviate to (\mathbb{R}, \mathbb{R}) . The DM believes that the deviator will implement his ideal action, while the non-deviator will implement his ideal action at best. Since the DM is indifferent between two agents' ideal actions, the deviator wins with probability at least $\frac{1}{2}$ with a strictly better winning outcome. Lastly, now any agent can undercut as before by committing to \mathbb{R} on his aligned dimension and promise an aligned action on his misaligned dimension.

With the option to choose arbitrary levels of vagueness, the previous equilibria remain. In the current notation, they are written as:

$$s_1(\theta) = s_2(\theta) = (\mathbb{R}, \mathbb{R}), \forall \theta$$

and when $b_1 = (0, 1)$, $b_2 = (1, 0)$,

$$s_1(\theta) = (\mathbb{R}, \{\theta^2 - b_1^2\}), s_2(\theta) = (\{\theta^1 - b_2^1\}, \mathbb{R}), \forall \theta.$$

Complete vagueness precludes undercutting by opponents while preserving own ideal action. However, there is difficulty extending these results to partial vagueness. For example, in order to sustain

$$s_1(\theta) = (A(\theta), w(\theta)), s_2(\theta) = (z(\theta), B(\theta)), \forall \theta,$$

we need to make sure that $A(\theta)$ does not reveal θ^1 and $B(\theta)$ does not reveal θ^2 . As one can see, the key reason for undercutting to carry through is revelation of the state. With complete vagueness, this can always be done because as long as an agent is always completely vague, his proposal does not reveal the state. Moreover, he has incentive to do so because being completely vague guarantees his ideal action. However, if an agent is only allowed to be partially vague, i.e. committing to $A(\theta) \subsetneq \mathbb{R}$, then it is harder to satisfy both conditions. In order not to reveal the state, the agent will need to be pooling, so $A(\theta)$ cannot reveal the state perfectly. Moreover, to ensure that the agent is willing to commit to $A(\theta)$ and not \mathbb{R} , $A(\theta)$ must be sufficiently close to the agent's ideal action, if not containing it.

Proof of Proposition 8. I show that if s_1, s_2 is as follows, then some agent has an incentive to deviate:

$$s_1(\theta) = (w(\theta), A(\theta)),$$

$$s_2(\theta) = (B(\theta), z(\theta)),$$

where $w, z : \Theta \rightarrow \mathbb{R}$ and $A, B : \Theta \rightarrow 2^{\mathbb{R}}$ and $\theta^2 - 1 \in A(\theta)$, $\theta^1 - 1 \in B(\theta)$ for all θ .

Step 1. $\beta(s_1(\theta), s_2(\theta)) = \frac{1}{2}, \forall \theta$.

Suppose for some $\bar{\theta}$, $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 1$. (The case in which $\beta(s_1(\bar{\theta}), s_2(\bar{\theta})) = 0$ is symmetric.) Then

$$u_2^*(\bar{\theta}) = u_2(\bar{\theta}, y_1(w(\bar{\theta}), A(\bar{\theta}))) = -|\bar{\theta}^1 - w(\bar{\theta}) - 1|^2 - 1.$$

Let Agent 2 deviate to $m_2^{dev} = (w(\bar{\theta}), \mathbb{R})$ at $\bar{\theta}$. Then given $m = ((w(\bar{\theta}), A(\bar{\theta})), (w(\bar{\theta}), \mathbb{R}))$, for any $\mu \in \Delta(\Theta)$,

$$\begin{aligned} \pi_1 &= \int_{\Theta} -(\theta^1 - w(\bar{\theta}))^2 \mu d\theta - 1, \\ \pi_2 &= \int_{\Theta} -(\theta^1 - w(\bar{\theta}))^2 \mu d\theta > \pi_1. \end{aligned}$$

Therefore $\beta(m) = 0$ and

$$u_2^{dev}(\bar{\theta}, s_1(\bar{\theta}), m_2^{dev}) = u_2(\bar{\theta}, y_2(w(\bar{\theta}), \mathbb{R})) = -|\bar{\theta}^1 - w(\bar{\theta}) - 1|^2 > u_2^*(\bar{\theta}),$$

making m_2^{dev} a profitable deviation.

Step 2. $w(\theta) = \theta^1$, $z(\theta) = \theta^2$ for all θ .

Suppose for some $\bar{\theta}$, $w(\bar{\theta}) \neq \bar{\theta}^1$ (the case for which $z(\bar{\theta}) \neq \bar{\theta}^2$ is symmetric). Then $u_1(\bar{\theta}, y_1(s_1(\bar{\theta}))) < 0$. Let Agent 1 deviate to (\mathbb{R}, \mathbb{R}) . Given $m = ((\mathbb{R}, \mathbb{R}), (B(\bar{\theta}), z(\bar{\theta})))$, The DM believes that only Agent 1 has deviated. Therefore

$$\pi_1 = -\|b_1\|^2 = -1,$$

$$\pi_2 = -|b_2^1|^2 - \int_{\Theta} |z(\bar{\theta}) - \theta^2|^2 \mu d\theta \leq -|b_2^1|^2 = -1.$$

Since $\pi_2 \leq \pi_1$, $\beta(m) \geq \frac{1}{2}$ and

$$\begin{aligned} u_1^{dev}(\bar{\theta}, m_1^{dev}, s_2(\bar{\theta})) &\geq \frac{1}{2}u_1(\bar{\theta}, y_1(\mathbb{R}, \mathbb{R})) + \frac{1}{2}u_1(\bar{\theta}, y_2(s_2(\bar{\theta}))) \\ &> \frac{1}{2}u_1(\bar{\theta}, y_1(s_1(\bar{\theta}))) + \frac{1}{2}u_1(\bar{\theta}, y_2(s_2(\bar{\theta}))), \end{aligned}$$

making m_1^{dev} a profitable deviation.

So for all θ , $w(\theta) = \theta^1$ and $z(\theta) = \theta^2$.

$$u_2^*(\theta) = \frac{1}{2}(0) - \frac{1}{2}(-2) = -1.$$

Step 3. Agent 2 has a profitable deviation.

Let Agent 2 deviate to $(\bar{\theta}^1 - 1 + \epsilon, \mathbb{R})$ at $\bar{\theta}$. Given $m = ((w(\bar{\theta}), A(\bar{\theta})), (\bar{\theta}^1 - 1 + \epsilon, \mathbb{R}))$, the DM believes that only Agent 2 has deviated. Therefore $\mu\{\tilde{\theta}^1 \mid \tilde{\theta}^1 = w(\bar{\theta})\} = 1$.

$$\pi_1 = -1,$$

$$\pi_2 = \int_{\Theta} [\theta^1 - (\bar{\theta}^1 - 1 + \epsilon)]^2 \mu d\theta = -(1 - \epsilon)^2 > \pi_1.$$

So $\beta(m) = 0$ and

$$u_2^{dev}(\bar{\theta}, s_1(\bar{\theta}), m_2^{dev}) = u_2(\bar{\theta}, y_2(m_2^{dev})) = -\epsilon^2 > u_2^*(\bar{\theta})$$

for small enough $\epsilon > 0$, making m_2^{dev} a profitable deviation for Agent 2. ■

References

- Alesina, Alberto F and Richard Holden (2008). "Ambiguity and Extremism in Elections". In: *Harvard Institute of Economic Research Discussion Paper 2157* 2157.
- Alonso, Ricardo and Niko Matouschek (2008). "Optimal delegation". In: *The Review of Economic Studies* 75.1, pp. 259–293.

- Ambrus, Attila, Volodymyr Baranovskyi, and Aaron Kolb (2015). “A Delegation-Based Theory of Expertise”. In: *Economic Research Initiatives at Duke (ERID) Working Paper No. 193*. Available at SSRN: <http://ssrn.com/abstract=2643672>.
- Ash, Elliott, Massimo Morelli, and Richard Van Weelden (2015). “Elections and Divisiveness: Theory and Evidence”. In: *Available at SSRN: http://ssrn.com/abstract=2571667*.
- Bagwell, Kyle and Garey Ramey (1991). “Oligopoly Limit Pricing”. In: *The Rand Journal of Economics*, pp. 155–172.
- Banks, Jeffrey S and Joel Sobel (1987). “Equilibrium selection in signaling games”. In: *Econometrica: Journal of the Econometric Society*, pp. 647–661.
- Battaglini, Marco (2002). “Multiple Referrals and Multidimensional Cheap Talk”. In: *Econometrica*, pp. 1379–1401.
- Battigalli, Pierpaolo (1996). “Strategic independence and perfect Bayesian equilibria”. In: *Journal of Economic Theory* 70.1, pp. 201–234.
- Blume, Andreas and Oliver Board (2014). “Intentional vagueness”. In: *Erkenntnis* 79.4, pp. 855–899.
- Blume, Andreas, Oliver J Board, and Kohei Kawamura (2007). “Noisy talk”. In: *Theoretical Economics* 2.4, pp. 395–440.
- Cho, In-Koo and David M Kreps (1987). “Signaling Games and Stable Equilibria”. In: *The Quarterly Journal of Economics*, pp. 179–221.
- Dranove, David and Ginger Zhe Jin (2010). “Quality Disclosure and Certification: Theory and Practice”. In: *Journal of Economic Literature* 48.4, pp. 935–63.
- Egorov, Georgy (2015). “Single-Issue Campaigns and Multidimensional Politics”. In: *Working Paper No. 21265, National Bureau of Economic Research*.
- Gentzkow, Matthew and Emir Kamenica (2015). “Competition in Persuasion”. In: *Available at SSRN: http://ssrn.com/abstract=2643672*.
- Grossman, Sanford J (1981). “The Informational Role of Warranties and Private Disclosure about Product Quality”. In: *Journal of Law and Economics*, pp. 461–483.
- Gul, Faruk and Wolfgang Pesendorfer (2012). “The War of Information”. In: *The Review of Economic Studies* 79.2, pp. 707–734.
- Kamada, Yuichiro and Takuo Sugaya (2014). “Valence Candidates and Ambiguous Platforms in Policy Announcement Games”. In: *Mimeo*.
- Kartik, Navin, Francesco Squintani, and Katrin Tinn (2015). “Information Revelation and Pandering in Elections”. In: *Mimeo*.
- Kartik, Navin and Richard Van Weelden (2014). “Informative Cheap Talk in Elections”. In: *Mimeo*.

- Kartik, Navin, Richard Van Weelden, and Stephane Wolton (2015). “Electoral Ambiguity and Political Representation”. In: *Available at SSRN 2682691*.
- Kohlberg, Elon and Jean-Francois Mertens (1986). “On the Strategic Stability of Equilibria”. In: *Econometrica: Journal of the Econometric Society*, pp. 1003–1037.
- Kreps, David M and Robert Wilson (1982). “Sequential equilibria”. In: *Econometrica: Journal of the Econometric Society*, pp. 863–894.
- Krishna, Vijay and John Morgan (2001). “A Model of Expertise”. In: *The Quarterly Journal of Economics* 116.2, pp. 747–775.
- Li, Hao and Wing Suen (2004). “Delegating decisions to experts”. In: *Journal of Political Economy* 112.S1, S311–S335.
- Meirowitz, Adam (2005). “Informational Party Primaries and Strategic Ambiguity”. In: *Journal of Theoretical Politics* 17.1, pp. 107–136.
- Schnakenberg, Keith E (2014). “Directional Cheap Talk in Electoral Campaigns”. In: *Available at SSRN: <http://ssrn.com/abstract=2412316>*.
- Schultz, Christian (1996). “Polarization and Inefficient Policies”. In: *Review of Economic Studies*, pp. 331–343.
- Vida, Péter and Takakazu Honryo (2015). “Iterative Forward Induction and Strategic Stability versus Unprejudiced Beliefs”. In: *Mimeo*.
- Watson, Joel (2015). “Perfect Bayesian Equilibrium: General Definitions and Illustrations”. In:

*Chapter 2***BIASED BELIEF UPDATING IN DYNAMIC
EXPERIMENTATION****2.1 Introduction**

One repeatedly observed phenomenon in gambling studies is the “near miss effect,” which is the gambling-reinforcing effect of a failure perceived as close to a win. I formally define the near miss effect and offer an theoretical explanation. The explanation is not new and is based on evidence and theories in psychology. The framework is general and can be applied to other situations in which the decision maker learns how much control she has over a stochastic outcome through experimentation.

Gambling is popular in the United States. In 2007 alone, 55 million people made 376 million trips to casinos (Barberis, 2012). This paper focuses on a robust and particularly puzzling aspect of gambling behavior. When encountering a loss that is perceived as close to a win, gamblers increase their desire to play. In an experiment by Clark et al. (2009), subjects complete a two-reel slot machine task in which they select the icon on the first reel and wait for the icon on the second reel to be realized. A win in this task is a match between the two icons located at the payline on both reels. A near miss is a mismatch in which a matching icon is right above or below the payline. Other outcomes are called full misses. In half the trials subjects choose personally the icon on the first reel, while the computer chooses for them for the rest of the trials. After choice on the first reel is made, subjects are asked “how do you rate your chance of winning?” After the icon on the second reel is realized, they are asked “how pleased are you with the result?” and “how much do you want to continue playing?” Two results stand out. First, when the choice of the icon on the first reel is made personally, subjects report higher ratings of chance of winning. Second, although both near misses and full misses give no payoffs, compared with full misses (i.e. mismatches which are not near misses), near misses increase subjects’ desire to play. However, this increase is restricted to trials in which subjects have personal control over the first reel.

The first result indicates that subjects seem to believe that their actions matter and that they have certain skills in picking the “right” icon - they believe the chance

of winning is higher when they choose the icon on the first reel themselves. The increased desire to play even when no payoff is realized conditional on placing the gamble themselves seems to indicate that subjects derive favorable information from a near miss, which has to do with their skill. I discuss other evidence in Subsection 2.1.

The story I give is as follows. The gambler believes that she has a fixed but unknown skill, which may or may not affect her probability of success. If it does, then a higher skill combined with a luck factor leads to a higher success rate. If it does not, then a higher skill is canceled out by a luck factor and as a result the success rate is unchanged. As she plays, she updates her belief about her skill and whether it matters to the outcome. Although a near miss does not yield any payoff, it increases the gambler's perceived success rate through a feedback of high skill but low chance that her skill matters.

To be more specific, there are three possible outcomes from each round of play: a success, a near miss and a failure. A success yields a fixed, positive payoff while both a failure and a near miss yield zero payoff. So the gambler only cares about the success rate.

She has two models of how the gambling task operates in mind and is unsure which one is true: the skill model or the chance model. In both models, outcomes are realized in two sequential stages. A signal is generated in each stage which can be either good or bad. If a bad signal is generated in the first stage, the outcome is a failure. If good signals are generated in both stages, the outcome is a success. If a good signal is generated in the first stage and a bad signal in the second stage, the outcome is a near miss.

The signal realization of the first stage is interpreted as the result of skill, which is the probability that the prerequisite of success is reached. If not, then the outcome is a failure. If yes, then luck comes in at the second stage which determines whether the outcome is a success or a near miss. The skill factor is common for both models, so the probability of the prerequisite being met is the same. The models differ in terms of their chance factors. In the skill model, the probability of good luck is positive while that in the chance model is zero. As a result, conditional on the prerequisite being met the skill model generates more successes than the chance model. Moreover, a high skill translates into a high success rate only in the skill model.

The gambler knows the luck factor for each model, but is not sure about which model is true. She is also uncertain about the skill factor which is common for both models. Given this setup, it is easy to see how a near miss may increase the gambler's valuation for gambling. If she knows for sure that the skill model is true, then a near miss indicates high skill which translates into a high success rate. If she knows for sure that the chance model is true, then a near miss is meaningless because although it also indicates high skill, it does not affect the success rate. Given a fixed skill, a near miss is more likely to come from a chance model, decreasing the gambler's belief that her skill matters. Overall, a near miss increases the perceived success rate if it signals high skill more than it signals a chance model.

One may question the assumption that a gambler is uncertain about whether her skill matters in an "obvious" chance model such as gambling. There is ample experimental evidence indicating that it is very easy to induce subjects to believe that their action matters in purely chance-driven tasks (Langer, 1975). By introducing into lotteries factors which would affect the success rate in a task in which skill matters, such as competition, familiarity with the task and involvement, she induce the subjects to believe that their skill matters. In other words, the subjects exhibit "the illusion of control." Therefore it is reasonable to assume that at least when the gambler first starts playing, her prior assigns a positive probability to the skill model being the true one.

I model the gambler's problem as a infinite, discrete period two-armed bandit. In each period, she decides whether to pull a risky arm or a safe arm. The risky arm generates three possible outcomes: success, failure and near miss. Outcomes are observed instantaneously after the gambler pulls the risky arm. The safe arm gives a fixed, known payoff. Given the gambler's belief, there is an index function representing her willingness to play. The gambler exhibits the near miss effect if the index function increases after a near miss is observed.

Instead of characterizing beliefs for which the gambler exhibits the near miss effect, I parametrize a gambler by her bias λ and show the existence of near miss effect for large enough λ for arbitrary beliefs. The advantage of this approach lies in that it does not rely on the specification of the gambler's prior beliefs; nor does it require an explicit formula of the index function.

The bias is defined as follows. The gambler's belief about whether the task is a skill model or a chance model is sticky and updated more slowly than Bayesian. In particular, her posterior belief that the task is a skill model is a convex combination

of her prior and the Bayesian posterior, with the weight assigned to the prior being $\lambda \in [0, 1]$. When $\lambda = 0$, the gambler is Bayesian. When $\lambda = 1$, the gambler is fully biased and does not update.

I now discuss the implication of biased updating. Since the gambler updates about her skill in a Bayesian way but updates slowly about what kind of task she is faced with, she essentially over-attributes an outcome towards her skill. In other words, she believes that she has more control over the outcome than she actually does. Apart from the illusion of control, this assumption is also reminiscent of the fundamental attribution error, which is a tendency to attribute an outcome more towards the personal factor rather than the situational factor.

As an example of the fundamental attribution error, in E. E. Jones and V. A. Harris (1967) subjects were presented essays and were asked to infer from them the attitude of the essay writers. Even when told that the essays were written when the writers had little choice about what attitude to express, subjects still believed that the writers' attitude is consistent with that expressed in their essays. Therefore when an outcome is the result of both a personal factor and a situational factor, people believe that the personal factor is responsible.

To see the connection between the experiment and this paper, notice that the skill takes the place of the personal factor and the type of task the gambler is faced with (i.e. skill or chance model) takes the place of the situational factor. We learn from the experiment that subjects, knowing that the task is a chance model in which the personal factor does not matter (or matter very little), still attribute the outcome towards the personal factor. In this paper, for any non-degenerate belief about the personal and situational factor, the gambler relatively underestimates the effect of the situational factor and overestimates that of the personal factor.

Using this approach, I give some conditions in terms of the bias level under which the gambler exhibits the near miss effect. I show that for any given belief about the gambling task, a Bayesian gambler values gambling more than a biased gambler. Moreover for any belief, if the probability of a success conditional on a non-failure is high enough in a skill model, then whenever a Bayesian gambler exhibits the near miss effect, then so does a biased gambler. Lastly, for any belief there is a cutoff bias level such that a more biased gambler exhibits the near miss effect.

I want to emphasize that it is not my intention to simply rationalize the data by incorporating biases. If one is only concerned about rationalizing as much data as

possible, then arguably there are simpler models which do not require the two-stage construction. My intention is to depict the mechanism through which people learn two things at once: how skillful they are, and whether their skill matters. In this situation, one common mistake is to over-estimate one's control over the outcome - both the illusion of control and fundamental attribution error are evidence of this mistake. They lead to the modeling choice that (1) the gambler believes that the gambling task is a skill model with positive probability, and (2) the gambler over-attributes the outcome towards her skill. One can certainly make the gambler under-attribute the outcome towards her skill by updating aggressively about the type of task she is faced with, but that would not be appropriate for the situation I am trying to model.

A direct application of the model is a technology firm's innovation process. R&D involves multiple layers of uncertainty. The skill of the R&D team determines whether a breakthrough occurs in the product development stage. Conditional on a breakthrough, the market demand then determines whether the product will be profitable. Therefore both factors determine the profitability of R&D and the decision maker of the firm learns about them through dynamic experimentation.

In this context, the decision maker's bias lies in over-estimation of how much control the R&D team has over the outcome. For example, when the market demand is low because the line of products the firm produces is dated, the decision maker attributes it to the R&D team's failure to make the product appealing. In the event of this bias, a decision maker values R&D less compared to the case with no bias. Moreover, she is more likely to continue R&D even under the feedback that the market demand is low for her line of products. This is because she treats it as good news about something she has control over - her team - more than she treats it as bad news about something she does not have control over - the market.

As an indirect application, I also use this experimentation framework to study a worker's occupation switching. When a worker considers what career to choose a job from as well as what job to choose, often there is uncertainty over how well-equipped she is for a career as well as whether a job is a good fit. When the outcome from experimenting with a job can be a success, a failure or a near miss, I study what outcomes trigger a job or career switch.

Related Literature

The near miss effect is widely studied in psychology and gambling. It is long recognized in psychology literature that near misses increase gambling propensity (Reid, 1986). Experiments manipulating the frequency of near misses have shown effects on slot machine gambling persistence (Kassinove and Schare, 2001; Côté et al., 2003; MacLin et al., 2007).¹ People differ in their physiological responses to a near miss and a full miss (Qi et al., 2011).

This paper is also related to biased information processing. Experimental evidence consistently suggests that people process information in a biased manner. Möbius et al. (2014) find that subjects update too little in response to both positive and negative feedback. Moreover, they over-weight positive feedback relative to negative. Eil and Rao (2011) find that subjects deviate more from Bayes' rule in response to negative feedback than positive feedback. Ertac (2011) compares subjects' deviation from Bayes' rule when processing information with and without self-relevance and finds evidence of pessimism in response to negative feedback in the self-relevance context. A series of evidence of underreaction of stock prices to news announcements is surveyed in Bernard (1992).

Some theoretical models depart from classical models by modifying Bayes' rule. Compte and Postlewaite (2004) is probably the most closely related to this paper. Their model confidence-enhanced performance and make two assumptions: performance is affected by perceived empirical frequency of success rate in the past, and the perceived frequency differs from the actual frequency. It is shown that when confidence affects performance, biased perception may be welfare enhancing. Although the framework is very similar to mine, they assume that there is no experimentation and the agent engages in the risky activity iff the expected payoff equals the cost. Jehiel (2005) and Eyster and Rabin (2005) each put forward an equilibrium notion in which players best response according to non-standard beliefs. In Jehiel (2005), instead of learning other players' behavior at each decision node, players bundle their decision nodes into analogy classes and best respond to the others' average behavior within each class. In Eyster and Rabin (2005), players underestimate information contained in other people's actions and therefore do not update sufficiently.

¹Depending on the experiment setup, there are different operational definitions of near misses. For example, near miss is defined as 2 identical symbols on the first two reels followed by a third different symbol on the third reel in Côté et al. (2003); it is defined as the case when three out of the four slots are of the same number, with the last number being different, e.g. 3 3 3 7 in Kassinove and Schare (2001). MacLin et al. (2007) specify near miss to be two winning symbols on the payline and the final winning symbol directly above or below the payline.

There is a vast literature on dynamic experimentation focusing on rational experimenters. Rothschild (1974) is the first known model of a seller learning about the demand by charging different prices and observing the results. The main conclusion is that firms do not acquire perfect information in learning about the demand. As a result, some firms charge incorrect prices, which explains the persistence of price diversity. Keller and Rady (1999) studies the seller's problem in an environment subject to random changes. Grossman, Kihlstrom, and Mirman (1977) studies a problem in which a buyer decides whether to buy from seller whose value is uncertain and shows that learning induces a buyer to buy more of the drug compared with when no learning takes place. Bar-Isaac (2003) considers a situation in which both the buyer and the seller are uncertain about the value of the seller's product. Another series of studies look at many-agent experimentation and studies the strategic interaction between experimenters (Bergemann and Välimäki, 1996; Bolton and C. Harris, 1999; Keller, Rady, and Cripps, 2005; Keller and Rady, 2010; Klein and Rady, 2011).

2.2 The Model

Consider an infinitely-lived player, who chooses between undertaking a risky activity ($s_t = 1$) or not ($s_t = 0$) in period $t = 1, 2, \dots$. Undertaking the activity in period t generates a stochastic outcome Y_t carrying a reward. $Y_t \in \{0, 1, 1'\}$ where 0 represents a failure, 1 represents a success, and $1'$ represents a near miss. Y_t is drawn from a distribution contingent on the realization of ω . ω can be either g (good) or b (bad). If $\omega = g$,

$$Y_t = \begin{cases} 0 & \text{with probability } 1 - q, \\ 1 & \text{with probability } q\epsilon, \\ 1' & \text{with probability } q(1 - \epsilon). \end{cases}$$

and if $\omega = b$,

$$Y_t = \begin{cases} 0 & \text{with probability } 1 - q, \\ 1 & \text{with probability } q\epsilon', \\ 1' & \text{with probability } q(1 - \epsilon'), \end{cases}$$

where $\epsilon > \epsilon'$. Therefore in both states of the world, an outcome has the same probability to be a failure. Conditional on the outcome not being a failure, more successes are generated when $\omega = g$. The player gets a payoff of 1 from a success and 0 from a near miss or a failure. Not undertaking the activity results in a payoff of c_0 in the current period.

Nature selects $\omega \in \{g, b\}$ and $q \in [0, 1]$, which are fixed throughout the time. The player knows ϵ and ϵ' , but is unsure about ω and q . Her initial belief about ω is characterized by a prior probability $\xi^0 \in (0, 1)$ that $\omega = g$. Her initial belief about q is denoted by a distribution F^0 with the density f^0 . I assume that both ω and q are independent and have non-degenerate distributions.

After each period of undertaking the risky activity, the player observes the realization of the outcome and updates her beliefs about ω and q . Specifically, given any prior belief (ξ, f) and newly realized outcome $Y = y$, a Bayesian player's updated belief is as follows. She updates her belief f about q upwards to f^+ after observing a success or a near miss, downwards to f^- after observing a failure. Letting $\mu_f = \int_0^1 qf(q)dq$, we have

$$f^+(q) = \frac{qf(q)}{\int_0^1 \tilde{q}f(\tilde{q}) d\tilde{q}} = \frac{qf(q)}{\mu_f},$$

$$f^-(q) = \frac{(1-q)f(q)}{\int_0^1 (1-\tilde{q})f(\tilde{q}) d\tilde{q}} = \frac{(1-q)f(q)}{1-\mu_f}.$$

Since a failure is generated with the same probability when ω is good and bad, the player updates her belief ξ about ω only if the outcome is not a failure. Since the good state generates successes with higher probability than the bad state, she updates ξ upwards to ξ^+ after observing a success and downwards to ξ^- after observing a near miss. Letting $g(\xi) = \xi\epsilon + (1-\xi)\epsilon'$ denote the probability of an outcome being a success given that it is not a failure, we have

$$\xi^+ = \frac{\xi\epsilon}{\xi\epsilon + (1-\xi)\epsilon'} = \frac{\xi\epsilon}{g(\xi)},$$

$$\xi^- = \frac{\xi(1-\epsilon)}{\xi(1-\epsilon) + (1-\xi)(1-\epsilon')} = \frac{\xi(1-\epsilon)}{1-g(\xi)}.$$

A biased player updates about q in the same way as a Bayesian, but underreacts to information pertaining to updating about ω . Let $\chi \in [0, 1]$ be the bias of the player and $\xi^+(\chi)$ and $\xi^-(\chi)$ be the updated beliefs after observing a success and a near miss, respectively. We have

$$\xi^+(\chi) = \chi\xi + (1-\chi)\frac{\xi\epsilon}{\xi\epsilon + (1-\xi)\epsilon'} = \chi\xi + (1-\chi)\frac{\xi\epsilon}{g(\xi)},$$

$$\xi^-(\chi) = \chi\xi + (1-\chi)\frac{\xi(1-\epsilon)}{\xi(1-\epsilon) + (1-\xi)(1-\epsilon')} = \chi\xi + (1-\chi)\frac{\xi(1-\epsilon)}{1-g(\xi)}.$$

Therefore the posterior of a player with bias χ is the sum of her prior weighted by χ and the Bayesian posterior weighted by $(1-\chi)$. A player with $\chi = 0$ is called a Bayesian player. A player with $\chi = 1$ is called a fully biased player. To summarize,

let $x = (\xi, f, \chi)$ be the state of the following dynamic programming problem:

$$\hat{V}(x) = \max \left\{ \frac{c_0}{1-\delta}; r(x) + \delta \mathbb{E}[\hat{V}(\tilde{x})|x] \right\},$$

where δ is the discount factor and $r(x) = g(\xi)\mu_f$ is the expected utility from a single-period reward given x . If $x_t = (\xi, f, \chi)$, then

$$x_{t+1} = \begin{cases} (\xi^+(\chi), f^+, \chi) & \text{with probability } g(\xi)\mu_f, \\ (\xi^-(\chi), f^+, \chi) & \text{with probability } [1 - g(\xi)]\mu_f, \\ (\xi, f^-, \chi) & \text{with probability } (1 - \mu_f). \end{cases}$$

To summarize, the player faces a two-armed bandit problem and chooses between a safe arm and a risky arm in each period t . The safe arm generates a fixed payoff of c_0 . The risky arm generates an expected payoff of $r(x)$. Then based on the outcome realization, x_t evolves into x_{t+1} , which determines the expected value for choosing the risky arm in the next period. The solution for this problem is provided by the Gittins Index (D. M. Jones and Gittins, 1972) and later summarized by Kumar (1985):

Theorem 1 (Gittins and Jones) *There is a real-valued function $v : X \rightarrow \mathbb{R}$ such that for all x , the optimal action is to play the risky arm iff $v(x) \geq c_0$. Moreover, for each x , v is defined as follows: let $\tilde{V}(x, c) := \max \left\{ \frac{c}{1-\delta}; r(x) + \delta \mathbb{E}[\tilde{V}(\tilde{x}, c)|x] \right\}$. Then $v(x) := \inf \left\{ c : \tilde{V}(x, c) = \frac{c}{1-\delta} \right\}$.*

For each x , the value of the risky arm is calibrated by an arm giving fixed payoff $v(x)$. The optimal strategy is then to play the risky arm iff $v(x) \geq c_0$. Instead of being just the per-period expected reward payoff $r(x)$, the calibration also takes into account the value for experimentation. The calibrated value is the smallest fixed reward such that the player weakly prefers getting the fixed reward of the calibrating arm each period from now on, to playing the risky arm for one more period and facing the same problem next period.

To see that $v(x)$ is well-defined, note that for $c > 1$, $\tilde{V}(x, c) = \frac{c}{1-\delta}$. In other words, when the reward from the calibrating arm exceeds the best possible reward from the risky arm, the optimal action is to play the calibrating arm. Therefore $\{c : \tilde{V}(x, c) = \frac{c}{1-\delta}\}$ is nonempty. When $c < 0$, the optimal action is to play the risky arm, so $\{c : \tilde{V}(x, c) = \frac{c}{1-\delta}\}$ is bounded below by 0.

To simplify later exposition, define

$$V(x, c) := \tilde{V}(x, c) - \frac{c}{1-\delta} = \max\{0; r(x) - c + \delta \mathbb{E}[V(\tilde{x}, c)|x]\} := \max\{0; W(x, c)\}.$$

So

$$v(x) = \inf\{c : W(x, c) \leq 0\}.$$

Proposition 9 *For each x , $c = v(x)$ is the unique solution to the equation $W(x, c) = 0$.*

The proof closely follows that of Lemma 1 in Rothschild (1974). I present it here since the proof strategy will be used throughout the paper. In order to show that $W(x, c) = 0$ has a unique solution, I show that it is continuous and strictly decreasing in c . In proving the continuity of W , I define W^t recursively as the expected value of experimenting for one more period when there are t periods left. The actual W is obtained by taking t to infinity. The continuity of W follows from the fact that the uniform limit of continuous functions is continuous.

Proof. I define $W(x, c)$ and $V(x, c)$ recursively as follows:

$$\begin{aligned} V^0(x, c) &= 0, \\ W^t(x, c) &= r(x) - c + \delta E[V^{t-1}(\tilde{x}, c)|x], \\ V^t(x, c) &= \max\{0, W^t(x, c)\}, \\ V(x, c) &= \lim_{t \rightarrow \infty} V^t(x, c), \quad W(x, c) = \lim_{t \rightarrow \infty} W^t(x, c). \end{aligned}$$

To see that $\lim_{t \rightarrow \infty} W^t(x, c)$ and $\lim_{t \rightarrow \infty} V^t(x, c)$ exist, first note that for each (x, c) , $V^t(x, c)$ and $W^t(x, c)$ increase in t . This is because the player can always imitate what she would do when there are fewer periods left and stop early. Second, for each t and (x, c) , $V^t(x, c)$ and $W^t(x, c)$ are bounded above by $\frac{1}{1-\delta}$. This is because the player cannot get a payoff higher than 1 per period. Therefore $V(x, c) = \lim_{t \rightarrow \infty} V^t(x, c)$ and $W(x, c) = \lim_{t \rightarrow \infty} W^t(x, c)$ exist.

Now I show that $W(x, c)$ is continuous in c by showing that it is continuous in (ξ, χ, c) . I will use this result in the future.

First I show by induction that $W^t(x, c)$ is continuous in (ξ, χ, c) for each t . For the case of $t = 1$, $W^1(x, c) = g(\xi)\mu_f - c$ is clearly continuous in (ξ, χ, c) . Now suppose $W^t(x, c)$ is continuous in (ξ, χ, c) . Since the function $\max\{\cdot, \cdot\}$ is continuous, $V^t(x, c)$ is also continuous in (ξ, χ, c) .

$$\begin{aligned} W^{t+1}(x, c) &= g(\xi)\mu_f \\ &\quad - c + \delta\mu_f \left[g(\xi)V^t(\xi^+(\chi), f^+, \chi, c) + [1 - g(\xi)]V^t(\xi^-(\chi), f^+, \chi, c) \right] + \delta(1 \\ &\quad - \mu_f)V^t(\xi, f^-, \chi, c), \end{aligned}$$

and $\xi^+(\chi)$ and $\xi^-(\chi)$ are both continuous in (ξ, χ, c) , so is $W^{t+1}(x, c)$. So $W^t(x, c)$ is continuous in (ξ, χ, c) for all t .

Now I show that $W^t(x, c)$ uniformly converges to $W(x, c)$. Let $\bar{W}^t(x, c)$ denote the discounted sum of expected payoff for t periods, using the optimal policy intended for infinite periods. Then we have

$$W^t(x, c) \geq \bar{W}^t(x, c),$$

$$W(x, c) - \bar{W}^t(x, c) \leq \frac{\delta^t}{1 - \delta}.$$

Therefore,

$$W(x, c) \leq \bar{W}^t(x, c) + \frac{\delta^t}{1 - \delta} \leq W^t(x, c) + \frac{\delta^t}{1 - \delta}.$$

So

$$|W^t(x, c) - W(x, c)| \leq \frac{\delta^t}{1 - \delta}.$$

Since the above inequality does not depend on (ξ, χ, c) , convergence is uniform.

Uniform limit of continuous functions is continuous, so $W(x, c)$ is continuous in (ξ, χ, c) .

Since $W(x, c)$ is continuous in c , $W(x, 0) \geq 0$, and $W(x, 1) \leq 0$ for all x , by the intermediate value theorem there exists $c \in [0, 1]$ such that $W(x, c) = 0$ for each x . Since $r(x) - c$ is strictly decreasing in c and $V(x, c)$ is weakly decreasing in c , $W(x, c)$ is strictly decreasing in c . So the solution to $W(x, c) = 0$ exists and is unique for each x . ■

2.3 Main Results

I show two implications of biased updating on experimentation. First, bias decreases a player's value for experimentation. When a Bayesian player and a player with arbitrary bias have the same beliefs, the Bayesian player values experimentation more. Second, bias more likely leads to a stronger near miss effect. For any belief, as long as the conditional probability of success is high enough, the near miss effect is stronger for a biased player than for a Bayesian player. As for its existence, for any belief there is a cutoff bias level such that a player above this level exhibits the near miss effect.

Proposition 10 *Suppose that $\xi > 0$ and $\chi \in [0, 1]$. Then for any (ξ, f) , $v(\xi, f, 0) \geq v(\xi, f, \chi)$.*

With the same belief, the Gittins index of the Bayesian player is higher than that of a

biased player. The proof has two main steps. The first step shows that the Bayesian player's value function $V(\xi, f, 0, c)$ is convex in ξ . The second step shows that, conditional on getting a non-failure, the expected value function is higher when the arguments are the Bayesian posteriors than the biased posteriors. Intuitively, since Bayesian posterior is more dispersed than the biased posterior, the expected value of a convex function of the Bayesian posterior is higher than that for the biased posterior. The full proof is presented in the Appendix.

Now we consider how bias leads to the near miss effect. Since the value for experimentation can be characterized by the Gittins index, a player exhibits the near miss effect if the Gittins index increases after a near miss is observed.

Definition 4 *A player at state (ξ, f, χ) exhibits the near miss effect if $N(\xi, f, \chi) = v(\xi^-(\chi), f^+, \chi) - v(\xi, f, \chi) > 0$. For $x_1 = (\xi_1, f_1, \chi_1)$ and $x_2 = (\xi_2, f_2, \chi_2)$, the near miss effect is stronger for x_1 than x_2 if $N(x_1) > N(x_2)$.*

Proposition 10 facilitates comparison of the near miss effect between players. Since we know that if the Gittins index is higher for the Bayesian player, it suffices to know the relationship of the Gittins index for both players after a near miss is observed. As it turns out, as long as the conditional probability of success when $\omega = g$ is high enough, the Gittins index for a Bayesian player is lower than that of a biased player.

Proposition 11 *Given $\xi > 0, f, \chi \in (0, 1]$, there exists ϵ^* such that for all $\epsilon > \epsilon^*$, $N(\xi, f, 0) < N(\xi, f, \chi)$. That is, if successes are generated with high enough probability when $\omega = g$, then all else equal, an arbitrarily biased player exhibits a stronger near miss effect than a Bayesian player.*

The intuition is as follows. Suppose the conditional probability of success when $\omega = g$ is 1. In other words, near miss occurs with probability 0. Then whenever a Bayesian player observes a near miss, her belief ξ is updated to 0. Since she will not update ξ any more, she is equivalent to a fully biased player with belief $\xi = 0$. At this point, her value function is lower than that of a fully biased player with belief $\xi > 0$. The same relationship holds for the Gittins index of both players. I then use the continuity of the Gittins index in ϵ to show that for ϵ close to 1, the Gittins index of the biased player is higher than that of a Bayesian player.

The next result gives conditions for the near miss effect. Near miss effect is frequently observed in slot machine gambling. In this stylized experimentation, a player learns

about the realization of each wheel and update her belief about the profitability of gambling. A near miss will lead her to increase her belief about the wheels with good realizations and decrease that about the wheels with bad realizations. It is easy to see that if the player does not update her belief about the bad wheels sufficiently, than a near miss will be a good news overall. The next result formalizes this intuition.

Proposition 12 *For all (ξ, f) where f is not a one-point distribution, there exists $\chi^*(\xi, f)$ such that $\chi > \chi^*(\xi, f)$ implies that $N(\xi, f, \chi) > 0$.*

The proof first shows that the near miss effect is stronger for a fully biased player than a Bayesian player. Since the Gittins index is continuous in bias, the same relationship holds for a sufficiently biased player.

2.4 Application: Career and Job Mobility

Consider a worker trying to choose a job and a career. Each career consists of different jobs. Both the fit of the worker for a career and a job are uncertain. In order to learn about these uncertainties, she need to experiment with a job. The outcome from experimentation contains both job- and career-specific information. For example, a worker may choose to become an economist by first experimenting with research in macroeconomics. In the process, she may learn about her ability or fit for a career in economics as well as for a job in macroeconomics. While it is easy to learn whether she has the basic required skill for the career, learning about job-specific information is conditional on her having acquired the basic skills. Inability to produce a publishable result in macroeconomics may lead her to doubt her ability to engage in economic research in general, but not whether macroeconomics is indeed an appropriate field for her. The latter can only be learnt after she has produced a publishable result and sent it to a journal.

In this environment, learning is conditional on the worker's decision to experiment as well as the outcome of experimentation. Based on the result of learning, the worker may choose to continue experimenting with the current job, or to switch between jobs and careers. The question I am interested in is what causes a worker to switch between jobs or careers. Below I show that total failure leads to career switching but not job switching, while failure which is perceived close to success may lead to both.

A worker is faced with M possible career choices. Career $m \in M$ consists of $|K(m)|$ jobs to choose from. Let (m, k) denote job k of career m . Each career has a specific

set of quantifiable skills. Skills for career m is quantified by $q_m \in [0, 1]$, interpreted as the probability that the worker satisfies the prerequisites for a success for any job within the career. In the event that she does not satisfy the prerequisites, the outcome is declared a failure. Otherwise, she is able to further learn about the current job, which can be good or bad. If the outcome is a success, she gets a fixed payoff; otherwise she gets zero payoff. a failure leads to update only about her belief about the career-specific skill, while a non-failure also leads to update about job-specific information. The conditional probability of success is ϵ when a job is good and $\epsilon' < \epsilon$ when a job is bad. In each period $t = 1, 2, \dots$, the worker chooses a job to work on. I assume that the worker's skill for any career remains fixed regardless of which job she chooses to work on, and that there is no switching cost.

Let $(\xi_{(m,k)}, f_m)$ denote the worker's current belief about the quality of job k of career m and her skill about career m . The value of the job is measured by its Gittins index $v(\xi_{(m,k)}, f_m)$. Let's consider how $v(\xi_{(m,k)}, f_m)$ evolves based on the outcome observed. First, note that the current job the worker is experimenting with must be the one with the highest Gittins index. That is, if the worker is currently experimenting with (m^*, k^*) , then $v(\xi_{(m^*,k^*)}, f_{m^*}) \geq v(\xi_{(m',k')}, f_{m'})$ for all $m' \in M$ and $k' \in K(m')$. Now consider the case in which the worker gets a success. This leads to an increase in both $\xi_{(m^*,k^*)}$ and f_{m^*} . So for any job (m^*, k) , its Gittins index becomes $v(\xi_{(m^*,k)}, f_{m^*}^+)$ and job (m^*, k^*) 's Gittins index becomes $v(\xi_{(m^*,k^*)}^+, f_{m^*}^+)$. Since $v(\xi_{(m^*,k^*)}, f_{m^*}) \geq v(\xi_{(m^*,k)}, f_m)$, $v(\xi_{(m^*,k^*)}^+, f_{m^*}^+) \geq v(\xi_{(m^*,k)}, f_{m^*}^+)$. For jobs outside of career m^* , their Gittins indices remain unchanged. Therefore the current job (m^*, k^*) remains the one with the highest Gittins index and the worker does not make switches.

Now suppose that the worker satisfies the prerequisite for a success but eventually is met with a failure. Then her belief about job (m^*, k^*) becomes $(\xi_{(m^*,k^*)}^-, f_{m^*}^+)$ and her belief about any other job in the career (m^*, k) becomes $(\xi_{(m^*,k)}, f_{m^*}^+)$. Here three things may happen. Job (m^*, k^*) may continue to be the most lucrative job in career m^* , in which case she makes no switches. Otherwise, she could switch to a different job or career.

Lastly, suppose that the worker is met with a failure without satisfying the prerequisite for a success. Then her belief about job (m^*, k^*) becomes $(\xi_{(m^*,k^*)}^-, f_{m^*}^-)$ and her belief about any other job in career m^* (m^*, k) becomes $(\xi_{(m^*,k)}, f_{m^*}^-)$. In this case, all jobs in career m^* has become less attractive, but (m^*, k^*) remains the best within career m^* . The worker may switch, but she would not make a within-career switch.

If in addition the worker underreacts in learning about job-specific information, then a near miss makes the worker insufficiently update her belief about job (m^*, k^*) . As a result, she is less likely to make a switch of any kind after a near miss. Since the worker only makes switches within a career after a near miss, underreaction leads to fewer job changes within a career.

2.5 Conclusion

I study a dynamic experimentation problem in which the experimenter is biased in her belief updating. The experimenter updates about two unknown parameters and under-reacts to information pertaining to one parameter. As a result of conservatism, she over-attribute the outcome observed to the other parameter. I show that attribution bias leads to undervaluation of experimentation. Moreover, a biased player is more susceptible of the near miss effect frequently observed in slot machine gambling. As an application of the model, I analyze a worker's career and job switching choice.

Appendix

I first prove a useful lemma.

Lemma 2 For all (ξ, f, c) , $g(\xi) \max \{0, g(\xi^+) \mu_{f^+} - c\} + [1 - g(\xi)] \max \{0, g(\xi^-) \mu_{f^+} - c\} \geq \max \{0, g(\xi) \mu_{f^+} - c\}$.

Proof. First note that $g(\xi)g(\xi^+) + [1 - g(\xi)]g(\xi^-) = g(\xi)$. Indeed, the RHS is $\Pr(1 \mid \{1, 1'\})$ in the current period and the LHS is the expected value $\Pr(1 \mid \{1, 1'\})$ in the next period.

Second, note that $0 \leq g(\xi^+) > g(\xi) > g(\xi^-) \leq 1$. There are four cases.

Case 1. $g(\xi^+) \mu_{f^+} - c \leq 0$. Both LHS and RHS equal 0.

Case 2. $g(\xi) \mu_{f^+} - c \leq 0 < g(\xi^+) \mu_{f^+} - c$. LHS = $g(\xi)[g(\xi^+) \mu_{f^+} - c] \geq 0 \geq$ RHS.

Case 3. $g(\xi^-) \mu_{f^+} - c \leq 0 < g(\xi) \mu_{f^+} - c$.

$$\begin{aligned} \text{LHS} &= g(\xi)[g(\xi^+) \mu_{f^+} - c] \\ &\geq g(\xi)[g(\xi^+) \mu_{f^+} - c] + [1 - g(\xi)][g(\xi^-) \mu_{f^+} - c] \\ &= g(\xi) \mu_{f^+} - c = \text{RHS}. \end{aligned}$$

Case 4. $g(\xi^-)\mu_f - c > 0$.

$$\begin{aligned} \text{LHS} &= g(\xi)[g(\xi^+)\mu_{f^+} - c] + [1 - g(\xi)][g(\xi^-)\mu_{f^+} - c] \\ &= g(\xi)\mu_{f^+} - c = \text{RHS}. \end{aligned}$$

■

Now I introduce some notation. For any $K \in \mathbb{N}$, $\xi^{a_1 a_2 \dots a_K}$ where $a_i \in \{+, -\}$ is the Bayesian posterior obtained by the operation a_1 on ξ , then a_2 on ξ^{a_1} , then a_3 on $\xi^{a_1 a_2}$, \dots , and finally a_K on $\xi^{a_1 \dots a_{K-1}}$. Similarly for $f^{a_1 a_2 \dots a_K}$.

Lemma 3 For each (ξ, f, c) ,

$$g(\xi)V^t(\xi^+, f^+, 0, c) + [1 - g(\xi)]V^t(\xi^-, f^+, 0, c) \geq V^t(\xi, f^+, 0, c).$$

Proof. I prove by induction. When $t = 1$, the statement is exactly Lemma 2. Now suppose the statement holds for t . First note that

$$\begin{aligned} &g(\xi)V^{t+1}(\xi^+, f^+, 0, c) + [1 \\ &- g(\xi)]V^{t+1}(\xi^-, f^+, 0, c) = g(\xi) \max\{0, W^{t+1}(\xi^+, f^+, 0, c)\} + [1 \\ &- g(\xi)] \max\{0, W^{t+1}(\xi^-, f^+, 0, c)\} = \max\{0, g(\xi)W^{t+1}(\xi^+, f^+, 0, c)\} + \max\{0, [1 \\ &- g(\xi)]W^{t+1}(\xi^-, f^+, 0, c)\} \geq \max\{0, g(\xi)W^{t+1}(\xi^+, f^+, 0, c) + [1 \\ &- g(\xi)]W^{t+1}(\xi^-, f^+, 0, c)\}. \end{aligned}$$

By the definition of W^{t+1} ,

$$\begin{aligned} &g(\xi)W^{t+1}(\xi^+, f^+, 0, c) + [1 - g(\xi)]W^{t+1}(\xi^-, f^+, 0, c) = g(\xi) \left\{ g(\xi^+)\mu_{f^+} \right. \\ &- c + \delta\mu_{f^+}g(\xi^+)V^t(\xi^{++}, f^{++}, 0, c) + \delta\mu_{f^+}[1 - g(\xi^+)]V^t(\xi^{+-}, f^{++}, 0, c) + \delta(1 \\ &- \mu_{f^+})V^t(\xi^+, f^{+-}, 0, c) \left. \right\} + [1 - g(\xi)] \left\{ g(\xi^-)\mu_{f^+} \right. \\ &- c + \delta\mu_{f^+}g(\xi^-)V^t(\xi^{-+}, f^{++}, 0, c) + \delta\mu_{f^+}[1 - g(\xi^-)]V^t(\xi^{--}, f^{++}, 0, c) + \delta(1 \\ &- \mu_{f^+})V^t(\xi^-, f^{+-}, 0, c) \left. \right\}. \end{aligned}$$

Using the induction hypothesis,

$$\begin{aligned} &\geq g(\xi)[g(\xi^+)\mu_{f^+} - c] + g(\xi)\delta\mu_{f^+}V^t(\xi^+, f^{++}, 0, c) + \delta g(\xi)(1 \\ &- \mu_{f^+})V^t(\xi^+, f^{+-}, 0, c) + [1 - g(\xi)][g(\xi^-)\mu_{f^+} - c] + [1 \\ &- g(\xi)]\delta\mu_{f^+}V^t(\xi^-, f^{++}, 0, c) + [1 - g(\xi)]\delta(1 - \mu_{f^+})V^t(\xi^-, f^{+-}, 0, c). \end{aligned}$$

Using the first step of the proof of Lemma 2,

$$= g(\xi)\mu_{f^+} - c + \delta g(\xi)[\mu_{f^+}V^t(\xi^+, f^{++}, 0, c) + (1 - \mu_{f^+})V^t(\xi^+, f^{+-}, 0, c)] + \delta[1 - g(\xi)][\mu_{f^+}V^t(\xi^-, f^{++}, 0, c) + (1 - \mu_{f^+})V^t(\xi^-, f^{+-}, 0, c)].$$

Regrouping the terms,

$$= g(\xi)\mu_{f^+} - c + \delta\mu_{f^+} \left\{ g(\xi)V^t(\xi^+, f^{++}, 0, c) + [1 - g(\xi)]V^t(\xi^-, f^{++}, 0, c) \right\} + \delta(1 - \mu_{f^+}) \left\{ g(\xi)V^t(\xi^+, f^{+-}, 0, c) + [1 - g(\xi)]V^t(\xi^-, f^{+-}, 0, c) \right\}.$$

Using the induction hypothesis again on the $\delta(1 - \mu_{f^+})$ branch,

$$\geq g(\xi)\mu_{f^+} - c + \delta\mu_{f^+} \left\{ g(\xi)V^t(\xi^+, f^{++}, 0, c) + [1 - g(\xi)]V^t(\xi^-, f^{++}, 0, c) \right\} + \delta(1 - \mu_{f^+})V^t(\xi, f^{+-}, 0, c).$$

By definition,

$$= W^{t+1}(\xi, f^+, 0, c).$$

To summarize, we have shown that

$$g(\xi)W^{t+1}(\xi^+, f^+, 0, c) + [1 - g(\xi)]W^{t+1}(\xi^-, f^+, 0, c) \geq W^{t+1}(\xi, f^+, 0, c),$$

which implies

$$g(\xi)V^{t+1}(\xi^+, f^+, 0, c) + [1 - g(\xi)]V^{t+1}(\xi^-, f^+, 0, c) \geq V^{t+1}(\xi, f^+, 0, c).$$

This concludes the induction procedure.

■

Proof of Proposition 10

Lemma 4 - 7 establish Proposition 10. Given a sequence a_1, a_2, \dots , I define the sequence $\{z_i\}$ as follows: let $\xi^{a_0} = \xi$. For $i \geq 1$:

$$z_i = \begin{cases} g(\xi^{a_1 a_2 \dots a_{i-1}}) & \text{if } a_i = +, \\ 1 - g(\xi^{a_1 a_2 \dots a_{i-1}}) & \text{if } a_i = -. \end{cases}$$

Lemma 4 Fix ϵ and ϵ' . For each $K \in \mathbb{N}$, $z_1 \cdot z_2 \dots \cdot z_K$ is linear in $g(\xi)$.

Proof. I prove by induction on K . When $a_1 = +$, $z_1 = g(\xi)$. When $a_1 = -$, $z_1 = 1 - g(\xi)$. Either way, z_1 is linear in $g(\xi)$. Suppose $z_1 \cdot z_2 \dots \cdot z_{K-1}$ is linear in $g(\xi)$. Then

$$z_1 \cdot z_2 \dots \cdot z_K = \begin{cases} g(\xi)(z_2 \cdot z_3 \dots \cdot z_K) & \text{if } a_1 = +, \\ [1 - g(\xi)](z_2 \cdot z_3 \dots \cdot z_K) & \text{if } a_1 = -. \end{cases}$$

Since $z_2 \cdot z_3 \dots \cdot z_K$ is linear in $g(\xi^{a_1})$ by the induction hypothesis and

$$g(\xi^{a_1}) = \begin{cases} g(\xi^+) & \text{if } a_1 = +, \\ g(\xi^-) & \text{if } a_1 = -, \end{cases}$$

it suffices to show that $g(\xi)g(\xi^+)$ and $[1 - g(\xi)]g(\xi^-)$ are both linear in $g(\xi)$. First,

$$\begin{aligned} g(\xi)g(\xi^+) &= g(\xi)g\left(\frac{\xi\epsilon}{g(\xi)}\right) \\ &= g(\xi) \left[(\epsilon - \epsilon') \frac{\xi\epsilon}{g(\xi)} + \epsilon' \right] \\ &= (\epsilon - \epsilon')\xi\epsilon + \epsilon'g(\xi) \\ &= \epsilon g(\xi) - \epsilon\epsilon' + \epsilon'g(\xi) \end{aligned}$$

and therefore is linear in $g(\xi)$. Moreover, by the proof of Lemma 2, we know that

$$g(\xi)g(\xi^+) + [1 - g(\xi)]g(\xi^-) = g(\xi).$$

Since $[1 - g(\xi)]g(\xi^-) = g(\xi) - g(\xi)g(\xi^+)$, it is also linear in $g(\xi)$. So $z_1 \cdot z_2 \dots \cdot z_K$ is linear in $g(\xi)$. ■

Lemma 5 For each $K \geq 2$, $a_1, a_2, \dots, a_{K-1} \in \{+, -\}$, $1 \leq i \leq K - 1$,

$$z_1 \cdot z_2 \dots \cdot z_i \cdot W^{K-i}(\xi^{a_1 a_2 \dots a_i}, f, 0, c)$$

is convex in ξ .

Proof. I prove by induction. When $K = 2$, it suffices to show that $z_1 \cdot W^1(\xi^{a_1}, f, 0, c)$ is linear in ξ .

$$z_1 \cdot W^1(\xi^{a_1}, f, 0, c) = z_1[g(\xi^{a_1})\mu_f - c] = \begin{cases} g(\xi)[g(\xi^+)\mu_f - c] & \text{if } a_1 = +, \\ [1 - g(\xi)][g(\xi^-)\mu_f - c] & \text{if } a_1 = -, \end{cases}$$

both are linear in $g(\xi)$ by Lemma 4. Since $g(\xi)$ is linear in ξ , both are linear in ξ also.

Now suppose that for all $K \in \mathbb{N}$ and $1 \leq i \leq K - 1$,

$$z_1 \cdot z_2 \dots \cdot z_i \cdot W^{K-i}(\xi^{a_1 a_2 \dots a_i}, f, 0, c)$$

is convex in ξ . I want to show that for all $K + 1 \in \mathbb{N}$ and $1 \leq i \leq K$,

$$z_1 \cdot z_2 \dots \cdot z_i \cdot W^{K+1-i}(\xi^{a_1 a_2 \dots a_i}, f, 0, c)$$

is also convex in ξ . Notice that

$$\begin{aligned} z_1 \cdot W^K(\xi^{a_1}, f, 0, c) &= z_1 g(\xi^{a_1}) \mu_f - c z_1 + z_1 \delta \mu_f g(\xi^{a_1}) \max\{0, W^{K-1}(\xi^{a_1+}, f^+, 0, c)\} \\ &\quad + z_1 \delta \mu_f [1 - g(\xi^{a_1})] \max\{0, W^{K-1}(\xi^{a_1-}, f^+, 0, c)\} \\ &\quad + z_1 \delta (1 - \mu_f) \max\{0, W^{K-1}(\xi^{a_1}, f^-, 0, c)\} \end{aligned}$$

is guaranteed to be convex if the followings hold:

1. $z_1 g(\xi^{a_1}) \mu_f - c z_1$ is linear in ξ ;
2. $z_1 g(\xi^{a_1}) W^{K-1}(\xi^{a_1+}, f^+, 0, c)$ is convex in ξ ;
3. $z_1 [1 - g(\xi^{a_1})] W^{K-1}(\xi^{a_1-}, f^+, 0, c)$ is convex in ξ ;
4. $z_1 W^{K-1}(\xi^{a_1}, f^-, 0, c)$ is convex in ξ .

1 is shown above in the $K = 2$ case. 4 comes from the induction hypothesis that $z_1 \cdot W^{K-1}(\xi^{a_1}, f, 0, c)$ is convex in ξ . 2 and 3 will hold if for all $a_1, a_2 \in \{+, -\}$, $z_1 \cdot z_2 \cdot W^{K-1}(\xi^{a_1 a_2}, f, 0, c)$ is convex. Therefore to show that $z_1 \cdot W^K(\xi^{a_1}, f, 0, c)$ is convex, it suffices to show that $z_1 \cdot z_2 \cdot W^{K-1}(\xi^{a_1 a_2}, f, 0, c)$ is convex. Notice that

$$\begin{aligned} z_1 z_2 W^{K-1}(\xi^{a_1 a_2}, f, 0, c) &= z_1 z_2 [g(\xi^{a_1 a_2}) \mu_f - c] \\ &\quad + \delta \mu_f \left\{ g(\xi^{a_1 a_2}) z_1 z_2 \max\{0, W^{K-2}(\xi^{a_1 a_2+}, f^+, 0, c)\} \right. \\ &\quad \left. + [1 - g(\xi^{a_1 a_2})] z_1 z_2 \max\{0, W^{K-2}(\xi^{a_1 a_2-}, f^+, 0, c)\} \right\} \\ &\quad + \delta (1 - \mu_f) z_1 z_2 \max\{0, W^{K-2}(\xi^{a_1 a_2}, f^-, 0, c)\} \end{aligned}$$

is convex if we can show that $z_1 \cdot z_2 \cdot z_3 \cdot W^{K-2}(\xi^{a_1 a_2 a_3}, f, 0, c)$ is convex for all $a_1, a_2, a_3 \in \{+, -\}$ (since $z_1 \cdot z_2 W^{K-2}(\xi^{a_1 a_2}, f^-, 0, c)$ is convex from the induction hypothesis). Following this logic, it suffices to show the convexity of

$$z_1 \cdot z_2 \dots \cdot z_K W^1(\xi^{a_1 a_2 \dots a_K}, f, 0, c),$$

which is equal to

$$z_1 \cdot z_2 \dots \cdot z_K [g(\xi^{a_1 a_2 \dots a_K}) \mu_f - c] = z_1 \cdot z_2 \dots \cdot z_K \cdot z_{K+1} \mu_f - c z_1 \cdot z_2 \dots \cdot z_K$$

when $a_{K+1} = +$. It is linear in $g(\xi)$ by Lemma 4, and hence linear (and convex) in ξ . ■

Lemma 6 $V^t(\xi, f, 0, c)$ is convex in ξ .

Proof. Since $V^t(\xi, f, 0, c) = \max\{0, W^t(\xi, f, 0, c)\}$, it suffices to show $W^t(\xi, f, 0, c)$ is convex in ξ . I prove by induction.

$W^1(\xi, f, 0, c) = g(\xi)\mu_f - c$ is linear in ξ and therefore convex in ξ . Suppose that $W^t(\xi, f, 0, c)$ is convex in ξ . Then

$$\begin{aligned} W^{t+1}(\xi, f, 0, c) &= g(\xi)\mu_f - c + \delta\mu_f \{g(\xi) \max\{0, W^t(\xi^+, f^+, 0, c)\} \\ &\quad + [1 - g(\xi)] \max\{0, W^t(\xi^-, f^+, 0, c)\}\} \\ &\quad + \delta(1 - \mu_f) \max\{0, W^t(\xi, f^-, 0, c)\} \end{aligned}$$

is guaranteed to be convex if:

1. $g(\xi)W^t(\xi^+, f^+, 0, c)$ is convex;
2. $[1 - g(\xi)]W^t(\xi^-, f^+, 0, c)$ is convex;
3. $W^t(\xi, f^-, 0, c)$ is convex.

3 follows from induction hypothesis. 1 and 2 follow from Lemma 5. ■

Lemma 7 *Let $\chi > 0$. Then for any (ξ, f, c) , $g(\xi)V^t(\xi^+, f, 0, c) + [1 - g(\xi)]V^t(\xi^-, f, 0, c) \geq g(\xi)V^t(\xi^+(\chi), f, 0, c) + [1 - g(\xi)]V^t(\xi^-(\chi), f, 0, c)$.*

Proof. For each ξ , define two random variables Y and Z as follows:

$$Y = \begin{cases} \xi^+ & \text{with probability } g(\xi), \\ \xi^- & \text{with probability } 1 - g(\xi). \end{cases}$$

$$Z = \begin{cases} \xi^+(\chi) & \text{with probability } g(\xi), \\ \xi^-(\chi) & \text{with probability } 1 - g(\xi). \end{cases}$$

Simple calculations show that $E[Y] = E[Z]$. Since $\xi^- < \xi^-(\chi) < \xi^+(\chi) < \xi^+$, Z is Y 's mean-preserving spread. Therefore Z second-order stochastically dominates Y . Since $V^t(\xi, f, 0, c)$ is convex in ξ , $g(\xi)V^t(\xi^+, f, 0, c) + [1 - g(\xi)]V^t(\xi^-, f, 0, c) = E[V^t(Y, f, 0, c)] \geq E[V^t(Z, f, 0, c)] = g(\xi)V^t(\xi^+(\chi), f, 0, c) + [1 - g(\xi)]V^t(\xi^-(\chi), f, 0, c)$. ■

Proof of Proposition 10. First, I prove by induction that $W^t(\xi, f, 0, c) \geq W^t(\xi, f, \chi, c)$ for all t . $W^1(\xi, f, 0, c) = W^1(\xi, f, \chi, c) = g(\xi)\mu_f - c$. Now suppose that $W^t(\xi, f, 0, c) \geq W^t(\xi, f, \chi, c)$. Then

$$\begin{aligned}
W^{t+1}(\xi, f, 0, c) &= g(\xi)\mu_f \\
&\quad - c + \delta\mu_f \left\{ g(\xi)V^t(\xi^+, f^+, 0, c) + [1 - g(\xi)]V^t(\xi^-, f^+, 0, c) \right\} + \delta(1 \\
&\quad - \mu_f)V^t(\xi, f^-, 0, c) \\
&\geq g(\xi)\mu_f - c + \delta\mu_f \left\{ g(\xi)V^t(\xi^+(\chi), f^+, 0, c) + [1 - g(\xi)]V^t(\xi^-(\chi), f^+, 0, c) \right\} + \delta(1 \\
&\quad - \mu_f)V^t(\xi, f^-, 0, c) \\
&\geq g(\xi)\mu_f - c + \delta\mu_f \left\{ g(\xi)V^t(\xi^+(\chi), f^+, \chi, c) + [1 - g(\xi)]V^t(\xi^-(\chi), f^+, \chi, c) \right\} + \delta(1 \\
&\quad - \mu_f)V^t(\xi, f^-, \chi, c) \\
&= W^{t+1}(\xi, f, \chi, c).
\end{aligned}$$

The first inequality follows from Lemma 7. The second inequality follows from the induction hypothesis.

This concludes the induction. Since $W^t(\xi, f, 0, c) \geq W^t(\xi, f, \chi, c)$ for all t , $W(\xi, f, 0, c) = \lim_{t \rightarrow \infty} W^t(\xi, f, 0, c) \geq \lim_{t \rightarrow \infty} W^t(\xi, f, \chi, c) = W(\xi, f, \chi, c)$. Finally, note that both $W(\xi, f, 0, c)$ and $W(\xi, f, \chi, c)$ strictly decrease in c . $v(\xi, f, 0)$ and $v(\xi, f, \chi)$ solves $W(\xi, f, 0, c) = 0$ and $W(\xi, f, \chi, c) = 0$, respectively. Therefore $v(\xi, f, 0) \geq v(\xi, f, \chi)$. ■

Proof of Proposition 11

Lemma 8 $v(\xi, f, \chi) \geq v(\xi, f, 1)$ for $\chi \in [0, 1]$.

Proof. I prove by induction that $W^t(\xi, f, \chi, c) \geq W^t(\xi, f, 1, c)$. First note that $W^1(\xi, f, \chi, c) = W^1(\xi, f, 1, c) = g(\xi)\mu_f - c$. Now suppose $W^t(\xi, f, \chi, c) \geq W^t(\xi, f, 1, c)$. A consequence of this is $V^t(\xi, f, \chi, c) \geq V^t(\xi, f, 1, c)$.

$$\begin{aligned}
W^{t+1}(\xi, f, \chi, c) &= g(\xi)\mu_f \\
&\quad - c + \delta\mu_f \left[g(\xi)V^t(\xi^+(\chi), f^+, \chi, c) + [1 - g(\xi)]V^t(\xi^-(\chi), f^+, \chi, c) \right] + \delta(1 \\
&\quad - \mu_f)V^t(\xi, f^-, \chi, c) \\
&\geq g(\xi)\mu_f - c + \delta\mu_f V^t(\xi, f^+, \chi, c) + \delta(1 - \mu_f)V^t(\xi, f^-, \chi, c) \\
&\geq g(\xi)\mu_f - c + \delta\mu_f V^t(\xi, f^+, 1, c) + \delta(1 - \mu_f)V^t(\xi, f^-, 1, c) \\
&= W^{t+1}(\xi, f, 1, c).
\end{aligned}$$

The first inequality follows from Lemma 3. Note that its proof does not rely on $\chi = 0$. The second inequality follows from the induction hypothesis. This concludes the induction. Therefore $W(\xi, f, \chi, c) \geq W(\xi, f, 1, c)$, which implies that $v(\xi, f, \chi) \geq v(\xi, f, 1)$. ■

For the next lemma, we need the following implicit function theorem for non-differentiable functions:

Theorem 2 (Kumagai, 1980) *Let $F : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a continuous mapping. Suppose that*

$$F(x_0, y_0) = 0.$$

There exist open neighborhoods $A_0 \subset \mathbb{R}^n$ and $B_0 \subset \mathbb{R}^m$ of x_0 and y_0 , respectively, such that, for all $y \in B_0$, the equation

$$F(x, y) = 0$$

has a unique solution

$$x = Hy \in A_0,$$

where H is a continuous mapping from B_0 into A_0 iff there exist open neighborhoods $A \subset \mathbb{R}^n$ and $B \subset \mathbb{R}^m$ of x_0 and y_0 , respectively, such that, for all $y \in B$, $F(\cdot, y) : A \rightarrow \mathbb{R}^n$ is locally one-to-one.

Lemma 9 *Fix $\xi > 0, f, \epsilon'$. Then there exists ϵ^* such that for $\epsilon > \epsilon^*$, $v(\xi^-, f^+, 0) < v(\xi, f^+, \chi)$.*

Proof. When $\epsilon = 1$, $v(\xi^-, f^+, 0) = v(0, f^+, 1) < v(\xi^-(\chi), f^+, 1) \leq v(\xi^-(\chi), f^+, \chi)$. The first equality follows because when $\epsilon = 1$, observing a near miss makes a Bayesian player believe that $\omega = b$ with probability 1. Since she will not update her belief about ω further, she is essentially a fully biased player with $\xi = 0$. The first inequality follows because $w(\xi, f, 1, c)$ is strictly increasing in ξ . The second inequality follows from Lemma 8. Therefore when $\epsilon = 1$, $v(\xi^-, f^+, 0) - v(\xi, f^+, \chi) < 0$.

Now I show that $v(\xi^-, f^+, 0) - v(\xi, f^+, \chi)$ is continuous in ϵ . First note that given ξ, ϵ' and f , $W : (c, \epsilon) \rightarrow \mathbb{R}$ is a continuous mapping. Moreover, for any ϵ , $W(\xi^-, f^+, 0, c)$ is strictly decreasing in c , so it is one-to-one. Also for each ϵ , $c = v(\xi^-, f^+, 0)$ solves $W(\xi^-, f^+, 0, c) = 0$. By Theorem Kumagai (1980), $v(\xi^-, f^+, 0)$ is continuous in ϵ . Similarly, $v(\xi, f^+, \chi)$ is continuous in ϵ . Therefore

$\nu(\xi^-, f^+, 0) - \nu(\xi, f^+, 1)$ is continuous in ϵ . , $\nu(\xi^-, f^+, 0) - \nu(\xi, f^+, \chi)$ is continuous in ϵ . Therefore there is ϵ^* such that $\nu(\xi^-, f^+, 0) - \nu(\xi, f^+, \chi) < 0$ for $\epsilon > \epsilon^*$. ■

Proof of Proposition 11. From Proposition 10, $\nu(\xi, f, 0) \geq \nu(\xi, f, \chi)$. From Lemma 9, there exists ϵ^* such that for $\epsilon > \epsilon^*$, $\nu(\xi^-, f^+, 0) < \nu(\xi, f^+, \chi)$. Therefore there exists ϵ^* such that for $\epsilon > \epsilon^*$, $N(\xi, f, 0) < N(\xi, f, \chi)$. ■

Proof of Proposition 12

Proof. The proof contains 3 steps. 1) $N(\xi, f, 1) > 0$. 2) For each ξ, f , $\nu(\xi, f, \chi)$ is continuous in χ . 3) For each ξ, f , $\nu(\xi^-(\chi), f, \chi)$ is continuous in χ .

Step 1 $N(\xi, f, 1) > 0$.

Since $\xi^-(1) = \xi$, it suffices to show that $\nu(\xi^-(1), f^+, 1) = \nu(\xi, f^+, 1) > \nu(\xi, f, 1)$. First, I prove by induction that $W^t(\xi, f^+, 1, c) > W^t(\xi, f, 1, c) > W^t(\xi, f^-, 1, c)$ for all t . Since

$$g(\xi)\mu_{f^+} - c > g(\xi)\mu_f - c > g(\xi)\mu_{f^-} - c,$$

we have

$$W^1(\xi, f^+, 1, c) > W^1(\xi, f, 1, c) > W^1(\xi, f^-, 1, c).$$

Now suppose that $W^t(\xi, f^+, 1, c) > W^t(\xi, f, 1, c) > W^t(\xi, f^-, 1, c)$.

$$\begin{aligned} W^{t+1}(\xi, f^+, 1, c) &= g(\xi)\mu_{f^+} - c + \delta \left[\mu_{f^+} V^t(\xi, f^{++}, 1, c) + (1 - \mu_{f^+}) V^t(\xi, f^{+-}, 1, c) \right] \\ &> g(\xi)\mu_f - c + \delta \left[\mu_{f^+} V^t(\xi, f^+, 1, c) + (1 - \mu_{f^+}) V^t(\xi, f^-, 1, c) \right] \\ &\geq g(\xi)\mu_f - c + \delta \left[\mu_f V^t(\xi, f^+, 1, c) + (1 - \mu_f) V^t(\xi, f^-, 1, c) \right] \\ &= W^{t+1}(\xi, f, 1, c). \end{aligned}$$

To understand the first inequality, note that $\mu_{f^+} > \mu_f$ so $g(\xi)\mu_{f^+} - c > g(\xi)\mu_f - c$. Moreover, $W^t(\xi, f^{++}, 1, c) > W^t(\xi, f^+, 1, c)$ by the induction hypothesis. So $V^t(\xi, f^{++}, 1, c) \geq V^t(\xi, f^+, 1, c)$. Lastly, since $f^{+-} = f^{-+}$, $W^t(\xi, f^{+-}, 1, c) = W^t(\xi, f^{-+}, 1, c) > W^t(\xi, f^-, 1, c)$ by the induction hypothesis, which implies that $V^t(\xi, f^{+-}, 1, c) \geq V^t(\xi, f^-, 1, c)$. The second inequality follows from the fact that $V^t(\xi, f^+, 1, c) > V^t(\xi, f^-, 1, c)$ and that $\mu_{f^+} > \mu_f$.

This concludes the induction procedure for showing that $W^t(\xi, f^+, 1, c) > W^t(\xi, f, 1, c)$. $W^t(\xi, f, 1, c) > W^t(\xi, f^-, 1, c)$ can be shown similarly. Therefore, $\lim_{t \rightarrow \infty} W^t(\xi, f^+, 1, c) > \lim_{t \rightarrow \infty} W^t(\xi, f, 1, c) > \lim_{t \rightarrow \infty} W^t(\xi, f^-, 1, c)$, so we have

$$W(\xi, f^+, 1, c) \geq W(\xi, f, 1, c) \geq W(\xi, f^-, 1, c)$$

and

$$V(\xi, f^+, 1, c) \geq V(\xi, f, 1, c) \geq V(\xi, f^-, 1, c).$$

Now,

$$\begin{aligned} W(\xi, f^+, 1, c) &= g(\xi)\mu_{f^+} - c + \delta \left[\mu_{f^+}V(\xi, f^{++}, 1, c) + (1 - \mu_{f^+})V(\xi, f^{+-}, 1, c) \right] \\ &> g(\xi)\mu_f - c + \delta \left[\mu_fV(\xi, f^+, 1, c) + (1 - \mu_f)V(\xi, f^-, 1, c) \right] \\ &= W(\xi, f, 1, c). \end{aligned}$$

The inequality comes from the following facts: $\mu_{f^+} > \mu_f$; $V(\xi, f^{++}, 1, c) \geq V(\xi, f^+, 1, c)$, $V(\xi, f^{+-}, 1, c) \geq V(\xi, f^-, 1, c)$. Therefore

$$v(\xi, f^+, 1) > v(\xi, f, 1)$$

and $N(\xi, f, 1) > 0$.

Step 2 $v(\xi, f, \chi)$ is continuous in χ .

Since $v(\xi, f, \chi)$ is the solution to $W(\xi, f, \chi, c) = 0$, from Theorem 2, we need only to verify:

1. $W(\xi, f, \chi, c)$ is continuous in (χ, c) .
2. $W(\xi, f, \chi, c)$ is one-to-one in c .

1 is shown in the proof for Proposition 9. 2 follows from the fact that $W(\xi, f, \chi, c)$ is strictly decreasing in c .

Step 3 $v(\xi^-(\chi), f, \chi)$ is continuous in χ .

Since $\xi^-(\chi)$ is continuous in χ , we need to show that $v(\xi, f, \chi)$ is continuous in (ξ, χ) . Similar to last step, we need to verify that:

1. $W(\xi, f, \chi, c)$ is continuous in (ξ, χ, c) .
2. $W(\xi, f, \chi, c)$ one-to-one in c .

Both hold for the same reasons as in the last step.

From Step 2 and 3, $N(\xi, f, \chi) = v(\xi^-(\chi), f, \chi) - v(\xi, f, \chi)$ is continuous in χ . Since $N(\xi, f, 1) > 0$, we have that there exists a neighborhood C of $\chi = 1$ such that for all $\chi \in C$, $N(\xi, f, \chi) > 0$. This establishes the result. ■

References

- Barberis, Nicholas (2012). “A Model of Casino Gambling”. In: *Management Science* 58.1, pp. 35–51.
- Bar-Isaac, Heski (2003). “Reputation and Survival: Learning in a Dynamic Signaling Model”. In: *The Review of Economic Studies* 70.2, pp. 231–251.
- Bergemann, Dirk and Juuso Välimäki (1996). “Learning and Strategic Pricing”. In: *Econometrica: Journal of the Econometric Society*, pp. 1125–1149.
- Bernard, Victor Lewis (1992). *Stock Price Reactions to Earnings Announcements: A Summary of Recent Anomalous Evidence and Possible Explanations*. Division of Research, Graduate School of Business Administration, University of Michigan.
- Bolton, Patrick and Christopher Harris (1999). “Strategic Experimentation”. In: *Econometrica* 67.2, pp. 349–374.
- Clark, Luke et al. (2009). “Gambling Near-misses Enhance Motivation to Gamble and Recruit Win-related Brain Circuitry”. In: *Neuron* 61.3, pp. 481–490.
- Compte, Olivier and Andrew Postlewaite (2004). “Confidence-Enhanced Performance”. In: *American Economic Review*, pp. 1536–1557.
- Côté, Denis et al. (2003). “Near Wins Prolong Gambling on a Video Lottery Terminal”. In: *Journal of Gambling Studies* 19.4, pp. 433–438.
- Eil, David and Justin M Rao (2011). “The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself”. In: *American Economic Journal: Microeconomics* 3.2, pp. 114–138.
- Ertac, Seda (2011). “Does Self-Relevance Affect Information Processing? Experimental Evidence on the Response to Performance and Non-Performance Feedback”. In: *Journal of Economic Behavior & Organization* 80.3, pp. 532–545.
- Eyster, Erik and Matthew Rabin (2005). “Cursed Equilibrium”. In: *Econometrica* 73.5, pp. 1623–1672.
- Grossman, Sanford J, Richard E Kihlstrom, and Leonard J Mirman (1977). “A Bayesian Approach to the Production of Information and Learning by Doing”. In: *The Review of Economic Studies* 44.3, pp. 533–547.
- Jehiel, Philippe (2005). “Analogy-based Expectation Equilibrium”. In: *Journal of Economic theory* 123.2, pp. 81–104.
- Jones, David Morian and John C Gittins (1972). *A Dynamic Allocation Index for the Sequential Design of Experiments*. University of Cambridge, Department of Engineering.
- Jones, Edward E and Victor A Harris (1967). “The Attribution of Attitudes”. In: *Journal of Experimental Social Psychology* 3.1, pp. 1–24.

- Kassinove, Jeffrey I and Mitchell L Schare (2001). "Effects of the "Near Miss" and the "Big Win" on Persistence at Slot Machine Gambling." In: *Psychology of Addictive Behaviors* 15.2, p. 155.
- Keller, Godfrey and Sven Rady (1999). "Optimal Experimentation in a Changing Environment". In: *The Review of Economic Studies* 66.3, pp. 475–507.
- (2010). "Strategic Experimentation with Poisson Bandits". In: *Theoretical Economics* 5.2, pp. 275–311.
- Keller, Godfrey, Sven Rady, and Martin Cripps (2005). "Strategic Experimentation with Exponential Bandits". In: *Econometrica* 73.1, pp. 39–68.
- Klein, Nicolas and Sven Rady (2011). "Negatively Correlated Bandits". In: *The Review of Economic Studies*, rdq025.
- Kumagai, Sadatoshi (1980). "An Implicit Function Theorem: Comment". In: *Journal of Optimization Theory and Applications* 31.2, pp. 285–288.
- Kumar, PR (1985). "A Survey of Some Results in Stochastic Adaptive Control". In: *SIAM Journal on Control and Optimization* 23.3, pp. 329–380.
- Langer, Ellen J (1975). "The illusion of Control." In: *Journal of personality and social psychology* 32.2, p. 311.
- MacLin, Otto H et al. (2007). "Using a Computer Simulation of Three Slot Machines to Investigate a Gambler's Preference Among Varying Densities of Near-miss Alternatives". In: *Behavior Research Methods* 39.2, pp. 237–241.
- Möbius, Markus M et al. (2014). "Managing Self-Confidence". In:
- Qi, Senqing et al. (2011). "Neural Correlates of Near-misses Effect in Gambling". In: *Neuroscience Letters* 493.3, pp. 80–85. ISSN: 0304-3940.
- Reid, RL (1986). "The Psychology of the Near Miss". In: *Journal of Gambling Behavior* 2.1, pp. 32–39.
- Rothschild, Michael (1974). "A Two-armed Bandit Theory of Market Pricing". In: *Journal of Economic Theory* 9.2, pp. 185–202.

Chapter 3

MATCHING WITH INCOMPLETE INFORMATION

3.1 Introduction

Consider a two-sided matching game in which agents only know their own preferences with certainty and have a common prior over the possible preference profiles. Can you design a mechanism with an equilibrium that is always stable with respect to the true preferences for any realization of the preference profiles? The answer is no. Roth (1989) demonstrates this by way of an example. For a particular distribution of preference profiles, no mechanism has an equilibrium which is stable at all realized states. The argument is, if such a mechanism exists, then there is a corresponding stable direct revelation mechanism with a truth-telling equilibrium. However, given that everyone else is reporting truthfully, some agent has a profitable deviation by lying about her preference.

In the proof, the agent deviates to report a preference that is realized with zero probability *ex ante*. Implicitly, this approach assumes that the mechanism designer does not know the prior distribution of agents' preferences or their support. In this paper, I show that even if the mechanism designer does learn the common prior and only allows agents to report preferences realized with positive probabilities, profitable deviations still exist. That is, the incentive compatibility constraint is in fact harsher than what is previously discovered.

After confirming the result from Roth (1989) under a stricter requirement for the deviations allowed, I go on to illustrate another aspect of the incentive compatibility constraints. Not only is the matching outcome inevitably unstable sometimes, it is "truly unstable" in the sense that the blocking agents realize their ability to block after observing the matching outcome. Using the same example used in the proof, I show that any mechanism sometimes produces a truly unstable match. Lastly, I discuss the potential to reach a stable matching with blocking pairs alone, free from equilibrium forces.

3.2 The Model

The market is composed of a group of men and women. Each agent has a preference ordering over the agents from the other group as well as the prospect of being

single. Each agent only knows his or her own preference and has a common prior distribution of the preference profiles. In a general matching game, each agent has a set of actions available irrespective of his or her own preference. According to the rule of the game, an action profile leads to a matching outcome. Agents derive utility from their matched partners.

Formally, a general two-sided matching game with incomplete information about preferences is denoted by

$$\Gamma = (N = M \cup W, \{D_i\}_{i \in N}, g, U = \prod_{i \in N} U_i, F).$$

$M = \{m_1, m_2, \dots, m_{|M|}\}$ is a finite set of men and $W = \{w_1, w_2, \dots, w_{|W|}\}$ is a finite set of women with $|M| \geq 2$ and $|W| \geq 2$. D_i is the set of actions available for Agent $i \in N$ after each possible history in this general matching game. For each $i \in M$, $U_i = \{u_i : W \cup \{i\} \rightarrow \mathbb{R}\}$ is the set of all utility functions defined over the possible matches for i and the prospect of remaining single, similarly for each $i \in W$. I assume that agents have utility functions of the following form: the least preferred partner (including staying single, for which the agent's partner is him/herself) is assigned utility 0. The second least preferred partner is assigned utility 1, the third least preferred assigned 2, etc. \mathcal{M} denotes the set of all matchings between M and W . Given $u \in U$, let $S(u)$ be the set of stable matchings with respect to u . $g : \prod_{i \in N} D_i \rightarrow \Delta(\mathcal{M})$ maps an action profile to a lottery over matchings. A pure strategy for Agent i is a function $\sigma_i : U_i \rightarrow D_i$. F is a discrete probability distribution over U , where $F(u)$ is the probability assigned to $u \in U$. F has support U . $[U, F]$ is called the state of information of the game Γ . $[(D_i)_{i \in N}, g]$ is called a mechanism. Γ is common knowledge. In addition, Agent i learns u_i prior to the play.

A revelation game Γ_R is a game in which the mechanism is of the form (U, ϕ) , where $\phi : U \rightarrow \Delta(\mathcal{M})$ maps stated preferences into lotteries over matchings. Therefore in a revelation game each agent reports a preference. Given any equilibrium σ^* of Γ , let $\Gamma_R(\sigma^*)$ be the corresponding revelation game such that $\phi(u) = g(\sigma^*(u))$ for each $u \in U$. By the revelation principle, truth telling is an equilibrium in $\Gamma_R(\sigma^*)$ and when all agents tell the truth in $\Gamma_R(\sigma^*)$, the resulting matching is the same as that when the agents play σ^* in Γ .

Theorem 3 (Roth, 1989) *If there are at least two agents on each side of the market, then for any general mechanism $[(D_i)_{i \in N}, g]$ there exist states of information $[U, F]$ for which every equilibrium σ of the resulting game Γ has the property that $g(\sigma(u)) \notin$*

$\Delta(S(u))$ for some $u \in U$. (And the set of such u with $g(\sigma(u)) \notin \Delta(S(u))$ has positive probability under F .) That is, there exists no mechanism with the property that at least one of its equilibria is always stable with respect to the true preferences at every realization of a game.

Theorem 3 is proved by giving a state of information $[U, F]$ that causes every stable revelation mechanism (and hence every mechanism) to fail. The state of information is shown in Table 1. $M = \{m_1, m_2\}$ and $W = \{w_1, w_2\}$. m_2 and w_2 have constant preference across all states. For each state $\omega \in \Omega \equiv \{a, b, c, d\}$, $u(\omega)$ is the corresponding $u \in U$. $\Pr(\omega)$ is the probability that $u(\omega)$ is realized. The case with larger sets of agents follows from the fact that the four agents who play a role in the proof can be embedded in any larger set of agents without affecting the conclusion, as long as they do not consider any additional agents to be acceptable matches.

When $\omega = c$, without loss of generality assume the matching $(m_1, w_1), (m_2, w_2)$ is given probability higher than $\frac{1}{2}$. The proof shows that no truth-telling equilibria exists. The argument is as follows. If q is sufficiently small, in which case c becomes the most probable, then w_2 can report her preference to be m_1 and ensure the matching $(m_1, w_2), (m_2, w_1)$ be carried out. Therefore given that everyone else is reporting truthfully, w_2 has an incentive to deviate.

However, w_2 has deviated to report a preference that is realized with probability 0 according to the common prior. The proof therefore implicitly assumes that the mechanism designer does not know the support of agents' preferences. In the next section, I show that if the prior distribution is indeed known to the mechanism designer, still no truth-telling equilibrium exists.

Table 3.1: State of information

State ω	Probability $\Pr(\omega)$	Preference $u(\omega)$	Stable matching $S(u(\omega))$			
a	q^2	$m_1 : w_1$ $w_1 : m_2$ $m_2 : w_2, w_1$ $w_2 : m_1, m_2$	m_2			
b	$q(1 - q)$	$m_1 : w_1, w_2$ $w_1 : m_2$ $m_2 : w_2, w_1$ $w_2 : m_1, m_2$	m_1	m_2		
c	$(1 - q)^2$	$m_1 : w_1, w_2$ $w_1 : m_2, m_1$ $m_2 : w_2, w_1$ $w_2 : m_1, m_2$	m_1	m_2	m_1	m_2
d	$q(1 - q)$	$m_1 : w_1$ $w_1 : m_2, m_1$ $m_2 : w_2, w_1$ $w_2 : m_1, m_2$	m_1	m_2		

3.3 The Result

Consider a revelation game $\Gamma_R = (N = M \cup W, U, F, \phi)$ in which $\phi(u) \in \Delta(S(u))$ for any u ; that is, the mechanism is stable with respect to the reported preferences. A pure strategy for Agent i is $\sigma_i : U_i \rightarrow U_i$. I focus on pure strategy equilibria throughout the paper. Let $u_{-i} = \prod_{j \neq i} u_j$. For a matching μ and Agent i , $\mu(i)$ denotes Agent i 's matching partner. $F(u_{-i} | u_i)$ denotes Agent i 's posterior probability assigned to $u = (u_i, u_{-i})$ given that i 's preference is u_i . The equilibrium is defined as follows.

Definition 5 A Bayesian Nash equilibrium $\sigma^* = (\sigma_i^*)_{i \in N}$ of the game $\Gamma_R = (N, U, f, \phi)$ is as follows: $\forall i, \forall u_i \in U_i, \forall u'_i \in U_i$,

$$\sum_{u_{-i}} F(u_{-i} | u_i) \cdot \sum_{\mu \in \mathcal{M}} \phi(\sigma_{-i}^*(u_{-i}), \sigma_i^*(u_i))(\mu) \cdot u_i(\mu(i)) \geq \sum_{u_{-i}} F(u_{-i} | u_i) \cdot \sum_{\mu \in \mathcal{M}} \phi(\sigma_{-i}^*(u_{-i}), u'_i)(\mu) \cdot u_i(\mu(i)).$$

Given an equilibrium σ^* and Agent j 's equilibrium action u_j^* , let $\sigma_j^{*-1}(u_j^*) = \{\tilde{u}_j \in U_j \mid \sigma_j^*(\tilde{u}_j) = u_j^*\}$. Agent i 's posterior probability assigned to Agent j 's preference being u_j , given that Agent i 's preference is

$$\Pi_i(u_j \mid u_i, \sigma^*, u_j^*) = \begin{cases} \frac{\sum_{\{\tilde{u}: \tilde{u}_j = u_j, \tilde{u}_i = u_i\}} F(\tilde{u})}{\sum_{\{\tilde{u}: \tilde{u}_j \in \sigma_j^{*-1}(u_j^*), \tilde{u}_i = u_i\}} F(\tilde{u})} & \text{if } u_j \in \sigma_j^{*-1}(u_j^*) \\ & \text{and } \sum_{\{\tilde{u}: \tilde{u}_j \in \sigma_j^{*-1}(u_j^*), \tilde{u}_i = u_i\}} F(\tilde{u}) \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Next, I define an instability notion for a revelation game with incomplete information. Apart from requiring that a blocking pair exists, I require at least one of the agents in the blocking pair realizes their ability to block at some preference realization.

Definition 6 A matching μ is truly unstable if the following two conditions hold:

1. μ is unstable. That is, there exists (i, j) where $i \in M, j \in W$ such that $u_i(j) > u(\mu(i))$ and $u_j(i) > u(\mu(j))$. Such (i, j) is called a blocking pair. Denote the set of blocking pairs given a matching μ by $B(\mu)$.
2. There exists a blocking pair of which at least one agent realizes his or her ability to block after observing the equilibrium actions at some realization u . That is, there exists $(i, j) \in B(\mu)$ and $u \in U$ such that either

$$\forall u_j \text{ s.t. } \Pi_i(u_j | u_i, \sigma^*, u_j^*) > 0, u_j(i) > u_j(\mu(j)),$$

or

$$\forall u_i \text{ s.t. } \Pi_j(u_i|u_j, \sigma^*, u_i^*) > 0, u_i(j) > u_i(\mu(i)).$$

The first condition is the usual stability notion when there is complete information over agents' preferences. The second condition is new. It states that for some blocking pair, at least one of the blocking agents, say Agent $i \in M$, realizes his ability to block. In particular, given the other blocking Agent $j \in W$'s message and the equilibrium strategies, for every preference of Agent j , she strictly prefers Agent i to her current matching partner. Agent i then realizes his ability to block.

Now I introduce the main result:

Theorem 4 *If there are at least two agents on each side of the market, then for any general mechanism $[\{D_i\}_{i \in N}, g]$ there exist states of information $[U, F]$ for which every equilibrium σ of the resulting game Γ has the property that for some $u \in U$ and μ truly unstable, $g(\sigma(u))(\mu) > 0$. (And the set of such u with $g(\sigma(u)) \notin \Delta(S(u))$ has positive probability under F .) That is, for any mechanism, at least one of its equilibria is truly unstable with respect to the true preferences at some realization of the game.*

Proof. Let ϕ be any stable revelation matching mechanism. Consider the state of information in Table 1 from Roth (1989). Since agents can only observe their own preferences, the information partition of each agent is as follows:

$$P(m_1) = \{\{a, d\}, \{b, c\}\},$$

$$P(w_1) = \{\{a, b\}, \{c, d\}\},$$

$$P(m_2) = \{a, b, c, d\},$$

$$P(w_2) = \{a, b, c, d\}.$$

By reporting preferences, agents report the information partition element that they find themselves in. Since the mechanism designer knows that m_2 and w_2 cannot distinguish any states, there is no need for them to report anything. Therefore only m_1 and w_1 are the active agents. When $\omega = c$, let the matching $((m_1, w_1), (m_2, w_2))$ be realized with probability s .

I say that an agent "reveals" if the agent truthfully reports his or her own preference, i.e., the element of partition this agent observes, and "lies" if the agent does not truthfully report his or her own preference, i.e., the element of partition this agent does not observe. For example, when state a is realized, m_1 will be lying if he

reports $\{b, c\}$, and revealing if he reports $\{a, d\}$. m_1 and w_1 's payoffs in each state given their actions and are summarized in Table 3.2 - 3.5:

	w_1 lies	w_1 reveals
m_1 lies	$2s, 2s$	$0, 2$
m_1 reveals	$2, 0$	$1, 1$

Table 3.2: $\omega = a$

	w_1 lies	w_1 reveals
m_1 lies	$2, 0$	$0, 1$
m_1 reveals	$1 + s, 2s$	$1, 2$

Table 3.3: $\omega = b$

	w_1 lies	w_1 reveals
m_1 lies	$0, 0$	$2, 1$
m_1 reveals	$1, 2$	$1 + s, 2 + s$

Table 3.4: $\omega = c$

	w_1 lies	w_1 reveals
m_1 lies	$0, 2$	$2s, 1 + s$
m_1 reveals	$1, 0$	$2, 1$

Table 3.5: $\omega = d$

From Table 3.2 and 3.5, if m_1 is in the partition set $\{a, d\}$, he will reveal regardless of w_1 's report. Similarly, from Table 3.2 and 3.3 if w_1 is in the partition set $\{a, b\}$, she will reveal regardless of m_1 's report. Therefore the key in characterizing equilibrium is to pin down m_1 's report at partition $\{b, c\}$ and w_1 's report at partition $\{c, d\}$.

Simple calculations show that there are three equilibria:

1. If $\frac{1-2q}{1-q} \leq s \leq \frac{q}{1-q}$, both m_1 and w_1 always reveal.
2. If $s \leq \frac{1-2q}{1-q}$, m_1 always reports $\{a, d\}$ and w_1 always reveals.
3. If $s \geq \frac{q}{1-q}$, w_1 always reports $\{a, b\}$ and m_1 always reveals.

In order for $\frac{1-2q}{1-q} \leq s \leq \frac{q}{1-q}$, we need $q \geq \frac{1}{3}$. Therefore as long as $q < \frac{1}{3}$, in every equilibrium someone lies. This implies that under the current state of information, truth-telling can never be equilibrium in any game induced by a stable revelation mechanism. Therefore any equilibrium in the original matching game cannot be always stable. This proves Roth's theorem. Now we need to find for each equilibrium the state that makes the matching truly unstable.

Equilibrium 2 When $\omega = b$, since m_1 reports $\{a, d\}$ and w_1 reports $\{a, b\}$, the matching outcome will be $S(u(a))$ where m_1 is single. Therefore m_1 , who prefers w_2 than being single and knows w_2 prefer himself to m_2 in all states, will realize he can form a blocking pair with w_2 .

Equilibrium 3 When $\omega = d$, since w_1 reports $\{a, b\}$ and m_1 reports $\{a, d\}$, the matching outcome will be $S(u(a))$ where w_1 is single. Therefore w_1 , who prefers m_1 than being single and knows m_1 prefer herself to being single in all states, will realize she can form a blocking pair with m_1 . ■

3.4 Discussion

An alternative way to fix Roth (1989)'s proof is by starting with an alternative state of information: with probability ϵ , the preference profile is as those in Table 1; with probability $1 - \epsilon$, the profile is uniformly drawn from the set of all possible combinations of all agents' possible preferences. As long as ϵ is small enough, state $\omega = c$ is realized with high enough probability such that for w_2 can profitably deviate to report m_1 .

It is easy to check that when Equilibrium 1 is played, only when $\omega = b$ is the outcome unstable with respect to true preference. In this case, if m_1 and w_2 form a blocking pair which is publicly observed, w_1 will realize that m_1 was lying and that $\omega \in \{a, b\} \cap \{b, c\} = b$ and form a blocking pair with m_2 , thus reaching the stable matching outcome. Similarly for Equilibrium 2. Therefore, information released through blocking pair formation can lead to more blocking pairs forming, leading the matching outcome to stability. This implies that blocking pairs may be able to serve an informational role and lead to stability, without the help of equilibrium.

To illustrate, suppose that m_1 does not know w_1 's strategy and w_1 does not know m_1 's strategy. That is, no equilibrium condition is assumed. However, I assume that agents do not play dominated strategies. Recall that this means when $\omega \in \{a, d\}$, m_1 reveals and reports $\{a, d\}$; when $\omega \in \{a, b\}$, w_1 reveals and reports $\{a, b\}$. There are

three possible outcomes: (m_2, w_2) , $((m_1, w_2), (m_2, w_1))$ and $((m_1, w_1), (m_2, w_2))$.

First note that when the matching outcome is $((m_1, w_2), (m_2, w_1))$, which is the stable matching when $\omega \in \{b, c\}$, m_1 must have reported $\{b, c\}$. Since $\{b, c\}$ is dominated when $\omega \in \{a, d\}$, this implies that $\omega \in \{b, c\}$. However, m_1 does not know whether $\omega = b$ or $\omega = c$ whereas w_1 knows this. Either way, the matching outcome is stable with respect to true preferences and no agent can form a blocking pair. A symmetric argument applies when the matching outcome is $((m_1, w_1), (m_2, w_2))$.

Now suppose that the matching outcome is (m_2, w_2) , which is the stable matching when $\omega = a$. Therefore m_1 must have reported $\{a, d\}$ and w_1 must have reported $\{a, b\}$. Without the knowledge that agents are playing according to an equilibrium, $\omega \in \{a, b, c, d\}$. Next I show that for each ω , the outcome after blocking pair formation is stable. Moreover, m_1 can manipulate such that the outcome is men-optimal in every state.

Suppose $\omega = a$. Given m_1 's information partition, he learns that $\omega \in \{a, d\}$. If $\omega = a$, then w_1 truly finds m_1 unacceptable and m_1 cannot form a blocking pair with her. If $\omega = d$, then m_1 expects w_1 to form a blocking pair with him. This is because w_1 cannot form a blocking pair with her preferred partner m_2 since he is already matched with his favorite partner. Either way, m_1 does not need to initiate a blocking pair. On the other hand, w_1 will not block either since her only acceptable partner is matched with his favorite partner. In this case, the matching outcome is "truly stable" in the sense that agents will not form blocking pairs, even given the information released through matching outcome and the lack of formation of blocking pairs.

Suppose $\omega = b$. Given m_1 's information partition, he learns that $\omega \in \{b, c\}$. In either case, he will be better off matched with either women. When $\omega = b$, he can only form a blocking pair with w_2 because w_1 finds him unacceptable. If $\omega = c$, he can form a blocking pair with either w_1 or w_2 . On the other hand, if $\omega = c$, w_1 will form a blocking pair with m_1 . Therefore m_1 does not need to initiate a blocking pair with w_1 and will block with w_2 instead. However, when $\omega = b$, w_1 will not form a blocking pair because he finds m_1 unacceptable and m_2 is already matched with his favorite partner. Therefore m_1 forms a blocking pair with w_2 , after which w_1 may form a blocking pair with m_2 , leading to the outcome $((m_1, w_2), (m_2, w_1))$, which is stable with respect to the true preference.

Suppose $\omega = c$. Given m_1 's information partition, he learns that $\omega \in \{b, c\}$. From

the last paragraph, m_1 forms a blocking pair with w_2 . However, w_1 learns that $\omega \in \{c, d\}$ and will indeed form a blocking pair with m_1 , leading to the outcome $((m_1, w_1), (m_2, w_2))$, stable with respect to the true preference.

Suppose $\omega = d$. Again, w_1 forms a blocking pair with m_1 , leading to the stable match $((m_1, w_1), (m_2, w_2))$.

Therefore, given that (m_2, w_2) is the matching outcome, at each ω , blocking pairs lead to a match stable with respect to the true preference. For other outcomes, no blocking pairs are formed and the outcome is stable with respect to true preferences. Therefore, allowing blocking pairs leads to a stable outcome under all circumstances.

When $\omega \in \{a, b, d\}$, there is only one stable match available. When $\omega = c$, however, there are two stable matches. Therefore the only possibility for manipulating to get a more preferred stable match lies in when $\omega = c$. Now I show that m_1 can get the men-optimal matching at $\omega = c$. From the argument before, if the matching outcome is (m_2, w_2) when $\omega = c$, the eventual match after blocking pairs are formed is the men-optimal matching $((m_1, w_1), (m_2, w_2))$. If the matching outcome is $((m_1, w_1), (m_2, w_2))$ or $((m_1, w_2), (m_2, w_1))$, the outcome is finalized. Therefore in order to avoid the women-optimal matching $((m_1, w_2), (m_2, w_1))$, m_1 should report $\{a, d\}$ when he learns that $\omega \in \{b, c\}$. This way, the outcome is men-optimal at every state.

References

Roth, Alvin E (1989). "Two-sided matching with incomplete information about others' preferences". In: *Games and Economic Behavior* 1.2, pp. 191–209.