Investigating Psychology-Influenced Economic Models in Lab, Field, and Theory

Thesis by

Alexander L. Brown

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2008

(Defended May 15, 2008)

Acknowledgements

First and foremost, I thank my mom and dad. They have provided me with tremendous support throughout graduate school (and a tremendous amount of encouragement before graduate school!). Even though they have entirely different styles of parenting, they both made me feel like I always had someone there if I needed it. Yet, at the same time, they knew how to talk to me so I felt stronger and more self-reliant.

Second, I would not have attended Caltech if not for John Kagel. He introduced me to experimental economics at Ohio State as an undergraduate, which began to my interest in the field. At Caltech, Colin Camerer, my advisor, provided a great deal of support, creativity, and influence on my research. He is the most responsible for developing my skills as an economist.

Surviving at Caltech is as much a social matter as an academic one, for that reason I am tremendously grateful for the support of the pigeonholers, Rafael Verduzco, Justin Bois, and Aditya Khair. There were times when I would see them for every lunch and every weekend; I miss you guys. Akira Villar, the one pigeonholer who remains, deserves special thanks for putting up with me during this last year on the job market.

I am also thankful to have had the opportunity to work with Charlie Plott. I thank Leeat Yariv for her initiative in helping me prepare for the job market. Dave Grether provided valuable insight as my fourth committee member.

For the rest of HSS, I must thank Sera Linardi, especially; she has the power to cheer me up on even the bleakest days. Morgan Llewellyn and Mike Alton are the type of friends I hope I can have in every department. Laurent Mathevet deserves credit for finally teaching me Kuhn-Tucker. I cherish the lunch conversations I had with Kyle Mattes. I thank Julian Romero for always inviting me to run a barefoot marathon with him, regardless of my limitations in fitness.

I wish to thank the TGIF regulars, John Ledyard, Kim Border, P.J. Healy, and Stuart McDonnell, our time together not only helped me through my first few years, it also gave me inside story on the department.

My sisters, Katie Brown and Anjuli Agarwal also deserve credit; knowing they were out there gave me a little extra fire. And without Suzan Brown, I would not have had the ability to contact them as often as I did.

Yoshie Narui, Nick Brunelli, and Tony Roy deserve praise for watching nearly every Ohio State football game with me over the last five years. Go Bucks!

I thank miscellaneously Mandy Agler, Ariana Arredondo, Meghana Bhatt, Jo Ducey, Will Ford, Julie Goldstein, Heidi Kamp, Ian Krajbich, Laurel Auchampaugh, Edith Liu, Howard Marvel, Mitchell Muessen, and Merrielle Spain.

Finally, I want to thank the people I met on the Graduate Student Council, Ernie's AF, the Two Baggers, the Laugh out Louds, HSS Softball, Union Station Foundation, the Budapest Behavioral Economics Workshop, and the Undergraduate Economics Society.

I also have chapter-specific aknowlegements:

My first chapter "Learning And Visceral Temptation In Dynamic Optimization Experiments" was coauthored with Colin F. Camerer and Zhikang Eric Chua. I thank them both for their contributions. Zhikang ran most of the first study experiments himself, I ran most of the second study. Colin helped with much of the writing. I ran the second set of experiments and did most of the data analysis. We all thank the support of NSF grant SES-0078911, Chris Carroll, Daniel Houser, Paul Kattman, George Loewenstein, Tanga McDaniel, John Hey, Nat Wilcox, three anonymous referees, and editor Ed Glaeser for helpful comments. We also thank Julie Malmquist of the SSEL Caltech lab, Chong Juin Kuan (NUS), Hackjin Kim, Tony Bruguier, and especially Min Jeong Kang (who ran several of the beverage-condition subjects herself) for help in doing the experiments.

My second chapter "To Review or Not Review? Limited Strategic Thinking at the Box Office" was coauthored with Colin F. Camerer and Dan Lovallo. I thank them both for their contributions. Colin helped with much of the writing and Dan arranged for the initial data and organized interviews with movie insiders. We all thank audiences at Caltech, especially Tom Palfrey, Leeat Yariv, Stuart McDonald, and Robert Sherman. We also thank audiences at Chicago GSB, Berkeley, Yale Graduate Student Conference on Behavioral Science, 2007 North American ESA, SJDM, and SEA. Thanks to Sera Linardi for a suggestion that lead to an 8-fold improvement in CH runtimes; to Esther Hwang, Ferdinand Dubin, and especially Jonathan Garrity for help with initial data collection. We also thank Carmina Clarke and Ferdinand Dubin for additional data collection.

My third chapter "Endogenous Time Preference and Personal Rules" is my own. I thank Robert Östling, Morgan Kouser, Stuart McDonnell, Dražen Prelec, and Kim Border for helpful comments.

Abstract

This thesis demonstrates the effectiveness of novel, psychology-influenced models of economics on traditional economic structures through the lab, field, and theory.

Chapter 2 observes how subjects are able to solve the computationally difficult buffer stock savings model first for monetary earnings and then in terms of cola when thirsty. The first experiments suggested that subjects saved too little initially, but learned to save optimally within four repeated lifecycles, or 1–2 lifecycles when examining the behavior of others. The second experiment, the first of its kind to combine savings models with visceral temptation in a laboratory, found evidence that subjects when receiving rewards immediately did worse than with a ten-minute delay, consistent with the quasi-hyperbolic discounting models and several other studies.

Chapter 3 examines the decision of film distributors to deliberately withhold from critics low-quality movies. In equilibrium, through iterative reasoning, moviegoers should correctly infer quality and a cold opening should not be profitable. Therefore, cold openings provide a natural field setting to test models of limited strategic thinking as well as the rational-actor, quantal response equilibrium model. In a data set of 856 widely released movies, cold opening produces a significant, 14–17%, increase in domestic box office revenue. Parameter estimates of moviegoers behavior fit those observed in experiments. However, distributor parameters imply they overestimate their consumers and could earn more by increasing the frequency of cold openings.

Chapter 4 examines two types of "personal rules" through a model where immediacy preference changes with decisions. That is, choosing (or not choosing) a tempting alternative makes it more (less) tempting in the future. "Descriptive" rules are the backward-induction solution to the problem. With finite periods, agents may avoid the tempting alternative if their choice is going to be repeated, exhibiting the precedent effect, but they also may exhibit procrastination knowing that in the future, they will avoid temptation anyway. "Prescriptive" rules, involving an agent changing his belief structure in order to bring about a more preferred outcome, can eliminate this procrastination effect, but lose their power under an infinite time horizon.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis demonstrates the effectiveness of novel, psychology-influenced models of economics on traditional economic structures through the lab, field, and theory. Traditional economic theory views individuals as consistent selves with one motivational drive that maximizes utility and, by assumption, individual welfare. However, empirical evidence has begun to refute that assumption. This thesis adds to that rebuttal in a constructive manner: rather than just pronounce traditional theory wrong, it aims to find a specific theory to support as an alternative to ultimately inform economic theory.

Specifically, this thesis focuses on two of the most recognized areas of behavioral economics: immediacy preference and limited iterated strategic thinking. Immediacy preference represents the idea that individuals do not have uniform time preferences: an individual who prefers $10 today vs. $12 tomorrow may not have the same preference when thinking about $10 in 10 days from now vs. $12 in 11 days, though standard economic theory would suggest just that. This property can lead to very peculiar behavior on the part of economic agents as two chapters investigate. Chapter 2 finds evidence of dynamically inconsistent preferences as thirsty subjects make simulated savings decisions for cola rewards. Chapter 4 theoretically examines how individuals might overcome dynamic inconsistency by making personal rules to enforce habits.

Limited strategic thinking is the idea that one cannot continue the process of thinking about another thinking about one thinking about another... etc. For instance when a producer refuses to disclose the quality of his product when he can costlessly do so, the equilibrium result is that his good is the worst quality possible. But this inference takes several steps of strategic thinking and most people to not make the equilibrium inference. Chapter 3 examines the box office of movies that distributors did not show to the critics and finds they do much better than their counterparts.

But there is a caveat to those unfamiliar with these types of models. The intent of these studies is not even to disprove traditional economic theory with another established theory that can take on any number of values. It is more than that. It is to validate the theory within a range of established parameters for a global class of problems (both in the field, and in the lab). For these theories to be integrated into economic theory, they must give specific predictions. Currently, both the cognitive hierarchy model and quasi-hyperbolic models give around 1.5 and 0.6, respectively, in both the lab and the field (Chapters 2 and 3 provide more evidence of that).

The last chapter also uses quasi-hyperbolic discounting; its aim is to design new theory to explain a dilemma created by immediacy preference. Personal rules have been suggested to solve that dilemma (Ainslie, 1975, 1992), and the chapter attempts to explain those rules in a manner both consistent with game theory and psychology. Ultimately by bringing those two fields together, it would create a model that could be tested and calibrated in the future like the previous two aforementioned models.

This thesis proceeds as follows:

Chapter 2 observes how subjects are able to solve the computationally difficult buffer stock savings model first for monetary earnings and then in terms of cola when thirsty. The first experiments suggested that subjects saved too little initially, but learned to save optimally within four repeated lifecycles, or 1–2 lifecycles when examining the behavior of others. The second experiment, the first of its kind to combine savings models with visceral temptation in a laboratory, found evidence that subjects when receiving rewards immediately did worse than with a ten-minute delay, consistent with the quasi-hyperbolic discounting models and several other studies.

Chapter 3 examines the decision of film distributors to deliberately withhold from critics low-quality movies. In equilibrium, through iterative reasoning, moviegoers should correctly infer quality and a cold opening should not be profitable. Therefore, cold openings provide a natural field setting to test models of limited strategic thinking as well as the rational-actor, quantal response equilibrium model. In a data set of 856 widely released movies, cold opening produces a significant, 14–17%, increase in domestic box office revenue. Parameter estimates of moviegoers behavior fit those observed in experiments. However, distributor parameters imply they overestimate their consumers and could earn more by increasing the frequency of cold openings.

Chapter 4 examines two types of "personal rules" through a model where immediacy preference changes with decisions. That is, choosing (or not choosing) a tempting alternative makes it more (less) tempting in the future. "Descriptive" rules are the backward-induction solution to the problem. With finite periods, agents may avoid the tempting alternative if their choice is going to be repeated, exhibiting the precedent effect, but they also may exhibit procrastination knowing that in the future, they will avoid temptation anyway. "Prescriptive" rules, involving an agent changing his belief structure in order to bring about a more preferred outcome, can eliminate this procrastination effect, but lose their power under an infinite time horizon.

Chapter 2

Learning And Visceral Temptation In Dynamic Savings Experiments

2.1 Introduction

Dynamic optimization of sequential choices is central to many different economic analyses. In the most interesting cases, current choices affect state variables, which either constrain future choices or influence future utility. Decisions of this type include extensive-form games with type updating, job search, fertility timing, purchases of durables and equipment replacement, investment with learning-by-doing, and many diet and health choices (including addictive consumption).

Our paper explores how well people make these types of decisions in a complex experimental environment. The essential design parameters were taken from models of consumer savings with income uncertainty and habit formation, which created a two-state finite dynamic program. Optimal saving required subjects to save a lot in early periods to buffer against bad income shocks and to avoid creating an early consumption "internality" from habit that reduces utility from future consumption. Field evidence on whether people save optimally is mixed. The earliest sophisticated analyses assumed certain (or certainty-equivalent) income in order to solve the models, and rejected many of the predictions of lifecycle theory. With the advent of better computing power and relaxation of restriction on the income process, Zeldes (1989) and Carroll (1992, 1997) were able to explain many aspects of consumption found in data using Friedman's 1957 original ideas.

Some recent studies suggested saving is optimal (Scholtz et al., 2006; Lusardi et al., 2001; Darlin, 2007). Still other research has argued that consumers make fundamental mistakes regarding savings (Haveman et al., 2006; Choi et al., 2003, 2005) and typically undersave relative to optimal levels (though see Rick et al., 2008). The fact that minor changes in details of savings plans induced higher saving is consistent with the hypothesis that people save too little, left to their own devices, but can be "nudged" to save more. (e.g., Thaler and Benartzi, 2004; Choi et al., forthcoming) Finally, the research that is most parametrically related to our own is on "tightwads" who overspend (Rick et al., 2008). Angeletos et al. (2001) modeled undersaving using a quasi-hyperbolic (β-δ) consumption model. They found that a model in which agents have an immediate preference for consumption ($\beta < 1$) was better calibrated to aggregate data than a model with only exponential discounting.[1]

Laboratory experiments may be of some use in this area of research. The experiments inevitably reflect the classic trade-off in generalizing from stylized lab experiments to naturally occurring choice: Experiments have high "internal validity" because the maintained assumptions of a particular theory (e.g., about utility functions and beliefs about the income process) can be clearly implemented with experimental control. Experimental comparison of different treatments can also shed some light on competing explanations (that is the potential advantage of the immediate-delayed consumption comparison in our second study). However, the generalizability of the experiments to most actual savings decisions is debatable because experiments necessarily take place over a very short horizon and the savings lifecycle is long. These experiments, much like a time-lapse photograph, show a process in a short amount of time (a few hours) that usually spans a much longer time interval (many years). Therefore, while we discuss the results in terms of their relevance to debates about lifecycle savings, readers are entitled to think of the data as more generalizable about much shorter-term dynamic decisions like consumption of addictive substances, or taking up an exercise regimen or diet.

Two classes of explanations for apparent evidence of undersaving are: Bounded rationality, and a preference for immediacy (or present-bias). Each was addressed in a separate experiment. Reporting the two experiments together enables direct judgment of which explanation is generally better (or whether both have some merit) and allows some parametric comparison.

In the first experiment, subjects had an oppurtunity to learn privately over seven experimental

[1]As the differences in these studies indicate, it is difficult to conclusively reject or accept the basic premise of lifecycle saving, which is that current saving correctly anticipates future needs and income variation, and smoothes consumption (Browning and Lusardi, 1996; Venti, 2006). The difficulty stems from the fact that econometric tests of the lifecycle model typically depend on many auxiliary assumptions about utility functions, separability across time, income expectations, retirement and other institutional rules, sorting, and credit market constraints. Apparent statistical evidence of undersaving in any particular study might be due to one or more econometric misspecifications or to mismeasurement of capital gains, educational returns, or durable consumption flows (Gale et al., 1999). For example, using the economic surprise of German reunification, Fuchs-Schundeln and Schundeln (2005) find that evidence of buffer-stock savings is sensitive to self-selection of risk-averse workers into low-risk professions.

lifecycles or to learn "socially" from the choices of other subjects. If learning created movement toward optimal choice, that was prima facie evidence of bounded rationality in initial choices (because highly rational subjects would not need to learn). The goal in this study was to see how close subjects were to optimal, how well they learned, and whether deviations from optimality resemble those in field data (as calibrated by β and δ) in quasi-hyperbolic models. The results of the first experiment were that subjects undersaved at first but were able to learn very quickly from social information and less quickly from private learning.

The approximation to optimality observed after learning was surprisingly close, which turned our attention to the second basic explanation for undersaving, which was a preference for immediacy (perhaps reflecting visceral temptation). Even if people fully understand optimal savings rules, optimal saving might be hard to implement if the choices that generate current utility are viscerally tempting (as in addictions).

Therefore, the second experiment used thirsty subjects who chose how much beverage to consume, as an experimental-scale model of more dramatic types of visceral temptation. Some subjects received beverages immediately (i.e., their period t decisions led to physical consumption in period t) and others received beverages with a 10-period delay (as if ordering from a catalog for later delivery). This immediate-delayed contrast was a simple way to detect the preference for immediacy (since the β term only influenced decisions when consumption was immediate). In the second experiment, subjects generally consumed more when rewards were immediate than when rewards were delayed, a difference which is consistent with models of hyperbolic discounting (i.e., Ainslie, 1975; Laibson, 1997) and dual-self conflict (e.g., Bernheim and Rangel, 2004; Fudenberg and Levine, 2006; Loewenstein and O'Donoghue, 2004). Structural parameter estimates of β and δ yielded values comparable to those in other lab and field studies (mean β of 0.6–0.7), albeit over very different time horizons.

2.2 Two Explanations for Undersaving

The experimental design implemented the assumptions of the buffer stock savings model of Carroll et al. (2000). Agents earned income each period, subject to stochastic independent shocks from a distribution they knew. In each period their available cash was the previous buffer stock, plus new income. In each period they chose how much of this available cash to spend on consumption and the rest was saved. Utility in each period depended upon a ratio of current consumption to a habit index. The habit index was a depreciated sum of previous consumption (as in the pioneering design of Fehr and Zych (1998), based on Becker and Murphy (1988)). An entire 30-period lifecycle was repeated several times with different income realizations each time. Two alternative explanations, which have been hypothesized to explain consumer undersaving in the field, suggested why subjects might have saved too little in this experimental environment.

2.2.1 Bounded Rationality

One explanation for apparent undersaving is the bounded rationality of consumers; They may undersave because their rationality is bounded, and solving for optimal saving in the buffer stock model is computationally difficult.[2] As Carroll (2001) put it bluntly

> "One problem is the spectacular contrast between the sophisticated mathematical apparatus required to solve the optimal consumption problem and the mathematical imbecility of most consumers. (p. 41)"

After all, the reason economists used an approximation with certainty-equivalent income for many years was because they were not able to solve the same problem posed in our experiment themselves, before later advances in computing.[3] Allen and Carroll (2001) also showed that learning by simple reinforcement is far too slow to produce convergence to optimal saving in reasonable time scales. It is possible that consumers simply cannot figure out or learn over time from modest experience how to save optimally.

With the exception of Bernasconi and Kirchkamp (2000), previous experimental work with simpler models also found evidence of undersaving and attributed it to bounded rationality. (Kotlikoff et al., 2001; Carbone, 2005; Carbone and Hey, 2004; Hey, 1988; Hey and Dardanoni, 1988; Fehr and Zych, 1998; Ballinger et al., 2003, 2006). Since there was no widely accepted theory of how bounded rationally should be modeled formally in these settings,[4] the presence of rationality bounds was inferred indirectly: if subjects made mistakes in the first lifecycle, but learned over time or from

[2]Another kind of bound on rationality is that consumers are overoptimistic about future income or underestimate the force of habit formation. Note that these possibilities are ruled out by inducing beliefs about the income process and subjects understanding of the degree of habit formation. So if we find that subjects save optimally, but believe that Americans do not, then the experiments suggest that misperceptions about income and habit formation could be the culprit in generating suboptimal saving in the field data.

[3]Friedman (1953) also had another idea about complex models. Perhaps individuals are not able to solve dynamic savings models, but are able to reach near optimal results through trial and error, much the same way a pool player sinks balls in pockets without understanding the physics behind shots. However, Carroll (2001) was also not satisfied with this explanation, his own experiments revealed this process could take hundreds of lifecycles (Allen and Carroll, 2001). He suspected social learning from others might allow individuals to learn to save optimally in far less time.

[4]Ballinger et al. (2006) modeled bounded rationality as individuals only looking ahead a fixed number of periods. They interpreted the results of Ballinger et al. (2003) to suggest most subjects only look ahead two periods.

the social examples[5] (our first study also included social learning (similar to Ballinger et al., 2003)), then we would infer that their initial mistakes resulted from bounds on rationality, because subjects with unbounded rationality would not need to learn from experience. Development of a more precise theory of rationality bounds and learning remains a priority for future research (and is discussed further in the conclusion).

2.2.2 Temptation and Dynamically Inconsistent Preferences

A second explanation for undersavings is that consumers know how to save optimally, but cannot resist short-term temptations to consume for some products. For example, the availability of widespread credit can contribute to overspending if, psychologically, credit cards anesthetize the "pain of paying" (Prelec and Loewenstein, 1998). As Carroll (2001) suggested

> "There certainly seems to be strong evidence that American households are now using credit cards in non-optimal ways. The optimal use of credit cards (at least as implied by solving the final optimizing model discussed above) is as an emergency reserve to be drawn on only rarely, in response to a particularly bad shock or series of shocks. However, the median household with at least one credit card holds about $7000 in debt on all cards combined. (p. 42)"

Laibson et al. (2003) argued that this pattern is explained by consumers who have a powerful preference for immediate consumption, rather than an expression of bounded rationality.

To test this explanation for undersavings, in the second experimental study we converted consumption from numbers to actual sips of beverage (for thirsty subjects). Comparing immediate and delayed delivery of beverage consumption enabled us to study the strength of temptation and dynamic inconsistency that might result. Of course, small amounts of beverage are not as dramatic as temptations like drug addiction, gambling, and credit card spending, but they were feasible in the lab and gave us a first contrast between money rewards and visceral temptations that can guide future research.

[5] Social learning can be considered a form of aggregating several lifetimes of information. For other experiments on information aggregation, see research on markets (Plott and Sunder, 1982, 1988) and information cascades (in markets) (Anderson and Holt, 1997).

2.3 Study 1: Learning with Money Rewards

2.3.1 Experimental Design

Participants were carefully instructed about the basic concepts of the experiment, and how their decisions and the random income draws would determine how much money they would earn (see sections A.1 and A.2 for details and instructional tables). To avoid demand effects and to enhance memorability, economic jargon like "income shocks," "habit stock," and "utility," were translated into plainer language—"adjustment factor," "lifestyle index," and "points," respectively.

Subjects chose C_t in each period from cash-on-hand, which is the sum of previous cash plus new income (Y_t). Income in each period was $Y_t = P_t \eta_t$, the product of P_t permanent income that grew at five percent ($P_t = (1.05)P_{t+1}$, with initial $P_t = 100$) and a multiplicative shock η_t, which was lognormally distributed $\left(\log \eta \sim N\left(\frac{1}{2}, 1\right)\right)$. There was no interest rate and discount factor, and no borrowing or investment. Period-specific utility depended on consumption and on an accumulated level of habit, according to

$$u(C_t, H_{t-1}) = k + \frac{\theta}{1-\rho}\left(\frac{C_t + \hat{\epsilon}}{H_{t-1}^{\gamma}}\right)^{1-\rho} \tag{2.1}$$

with risk-aversion parameter $\rho = 3$ and a habit strength exponent $\gamma = 0.6$.[6] The habit stock grew according to $H_t = \lambda H_t - 1 + C_t$ where $\lambda = 0.7$ is a depreciation rate (as in Fehr and Zych, 1998) and the initial habit $H_0 = 10$. Thus, larger early consumption built up the habit level and depreciated future-period utility. This "internality" implied that optimization requires restrained consumption in early periods.

The subject's problem was to choose the stream of consumption $\tilde{C}$ in each period t to maximize his expected utility,

$$E_t\left[\sum_{s=t}^{T} u\left(\tilde{C}_s, \tilde{H}_{s-1}\right)\right]. \tag{2.2}$$

Because $T = 30$ in the experiments, the problem could be simplified to a dynamic programming problem with two state variables, cash-on-hand C_t and habit H_t (after dividing both variables by the permanent income P_t).

The experimental environment was designed to have some basic empirical features of savings in the modern American economy. The 5% income growth and lognormality of multiplicative

[6] Since $\rho = 3$, the term k is the upper asymptote of utility. θ is a scaling parameter, and $\hat{\epsilon}$ bounds the utility function from below. In the experiments, $\hat{\epsilon} = 2.7$, similar to Ballinger et al. (2003). Scaling factors are $\theta = 750$ and $k = 40$.

shocks were shown by Carroll (1992) to characterize US data. However, we chose $T = 30$ to compress the lifecycles (compared to American annualized lifetimes) in order to create "lifecycles" which were long enough to create a savings challenge and interesting dynamics, but short enough to allow several lifecycles in each experimental session. We also multiplied the standard deviation of multiplicative income shocks η_t by five (creating a standard deviation of 1, rather than Carroll's estimate of 0.2) in order to deliberately produce more income variation.

The goal of experiments like these was *not* to precisely recreate all the empirical properties of naturally occurring decisions in a particular setting. After all, parametric properties of savings problems vary widely across periods of history and across countries so there is no single "real world" to serve as a unique design target. The goal, instead, was to explore a range of environments in which the theory might apply in order to judge when the theory is likely to work and when it is likely to fail. We deliberately chose income shock volatility that is larger than that observed in the modern American economy because higher income variation created a more analytically challenging environment in which deviations from rationality would be more clearly observed. The design also combined uncertain income and habit formation, because (a) previous experiments have already studied each separately and (b) since combining them made the problem much more complicated, if learning would occur then the power of learning would be established with more force.

The instructions explained *all* the details of the structure described above. To make the details easier to understand, we included 30-draw samples from the lognormal distribution to give participants a feel for how much their income could vary and showed the utility functions and habit stock evolution using numerical tables (see A.2 for instructions and descriptions of tables). One table illustrated how the habit stock in each period was determined by the previous periods habit stock and the current spending. A separate table showed how their spending and habit stock in one period determined their utility points in that period. Before participating, subjects took a quiz testing them on how their choices, habit levels, and income shocks would determine utility points. The quiz was designed to satisfy concerns that suboptimal consumption decisions do not arise from confusion about how their decisions map into points (and eventual money earnings).

Consumption decisions were input to an Excel interface which displayed the income obtained, the corresponding cash available, and the habit stock for each consumption choice (see Figure 2.1). The program also calculated and displayed the possible points (i.e., utilities) that could be obtained from different levels of spending, and the corresponding savings available for the next period. Participants could experiment by inputting different consumption amounts and see how much utility

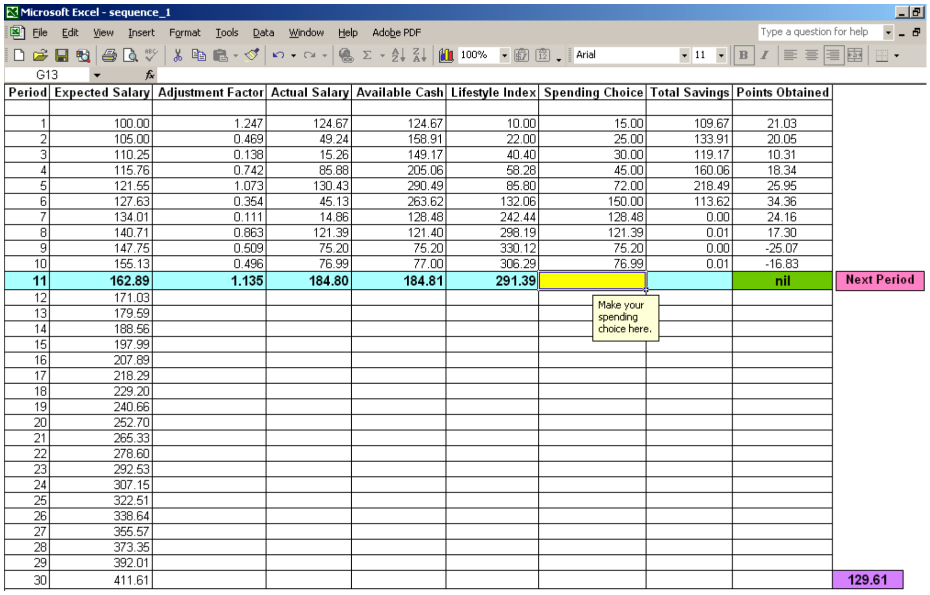

Period	Expected Salary	Adjustment Factor	Actual Salary	Available Cash	Lifestyle Index	Spending Choice	Total Savings	Points Obtained	
1	100.00	1.247	124.67	124.67	10.00	15.00	109.67	21.03	
2	105.00	0.469	49.24	158.91	22.00	25.00	133.91	20.05	
3	110.25	0.138	15.26	149.17	40.40	30.00	119.17	10.31	
4	115.76	0.742	85.88	205.06	58.28	45.00	160.06	18.34	
5	121.55	1.073	130.43	290.49	85.80	72.00	218.49	25.95	
6	127.63	0.354	45.13	263.62	132.06	150.00	113.62	34.36	
7	134.01	0.111	14.86	128.48	242.44	128.48	0.00	24.16	
8	140.71	0.863	121.39	121.40	298.19	121.39	0.01	17.30	
9	147.75	0.509	75.20	75.20	330.12	75.20	0.00	-25.07	
10	155.13	0.496	76.99	77.00	306.29	76.99	0.01	-16.83	
11	**162.89**	**1.135**	**184.80**	**184.81**	**291.39**			**nil**	**Next Period**
12	171.03								
13	179.59								
14	188.56								
15	197.99								
16	207.89								
17	218.29								
18	229.20								
19	240.66								
20	252.70								
21	265.33								
22	278.60								
23	292.53								
24	307.15								
25	322.51								
26	338.64								
27	355.57								
28	373.35								
29	392.01								
30	411.61								**129.61**

Figure 2.1: Screenshot of the Excel interface

they would earn, and how much cash they would have available at the start of the next period. Most participants tried out several spending choices before making a decision (especially in the first couple of lifecycles). This process was repeated until the end of the lifecycle of 30 periods. (The program automatically spent all cash in the final period 30.) There were a total of seven lifecycles, to see how rapidly subjects could learn across lifecycles. Each participant's total payoff was a pre-announced linear function of the total points earned in all lifecycles[7] plus a $5 show-up payment. Subjects earned between $7.50 to $65 with an average of $45.

After thirty-six (36) subjects had participated in the private learning condition described above, thirty-six (36) more participated in a "social learning" condition.[8] In the social learning condition, as part of their initial instructions, subjects were also given samples of what three actual subjects had done in the private learning condition. The three samples were taken from the highest-earning subject, the lowest-earning subject, and from one subject chosen at random from the private condition in their subject pool. The social learning subjects were told exactly how these three samples were chosen.

There were many ways to implement social learning or imitation (e.g., Ballinger et al., 2003, used direct talking). Our method mimicked intergenerational imitation in which a parent points

[7]The exchange rates were US $1.50 for every 100 experimental points in Caltech, and US $2.50 in Singapore (using an exchange rate of US $1 $\approx$ Sing $1.70).

[8]The tables looked like the screens the participants had, showing income each period, cash-on-hand, spending decisions, and points from each period of a 30-period lifecycle.

out three role models—a great success who retires wealthy, a neer-do-well who ends up broke, and a random acquaintance. The high-earning role model might have been a subject who overspent early on (relative to the optimum) but got lucky by receiving high-income draws. In that case, subjects copied the "successful" subject too directly, they would have easily overspent relative to the optimum; so it was not clear whether social learning would have actually helped, hurt, or had no effect.[9]

Participants were 35 undergraduates from the National University of Singapore (NUS) and 37 undergraduates from California Institute of Technology. These students were unusually adept at analytical thinking so they should represent an upper bound on how well average consumers do in these intertemporal optimization problems. The participants were recruited using the universities mail servers. Half the participants (18 from each school) did the experiment with private learning and approximately half (17 NUS, 19 Caltech) did the experiment with social learning. Each group had seven lifecycles of 30 periods of income draws. To simplify data analysis, within each condition all participants received the same income draws (but the draws were different in the two learning conditions).[10] Most participants completed the instruction and seven lifecycles in about 90 minutes.

2.3.2 Basic Results

Under optimality people should act as if they make ex ante optimal savings decisions under uncertainty, discounting future utilities exponentially, given their beliefs about future income and other structural parameters. In our experimental design, subjects should have saved a lot to build up a buffer stock, and then spent roughly their average income once their buffer stock is large enough. The buffer stock would protect against bad future income draws, and high early saving would limit the negative "internality" of current spending on future utility, (which occurs because of the controlled effect of habit formation). Figure 2.2 illustrates an optimal path of consumption, and cash-on-hand, given a particular lifecycle of income shocks (based on parameters used in the experiment, described later). Savings is the gap between the black optimal consumption line and the gray cash-on-hand line. In this example, the optimal consumer should spend less than current income in early periods except 6–7 (when income happened to be unusually low and consumers should dip into their

[9]Ex post we know that the highest scoring subject underconsumed for the first ten periods, but then overconsumed for the remainder. The subject was lucky to draw a high income realization in the later periods so he could still produce positive utility under a high level of lifestyle habit.

[10]The income realizations were different so that the social learning subjects would never have a lifecycle that matched exactly the income realizations seen by the role model subjects (drawn from the private learning condition).

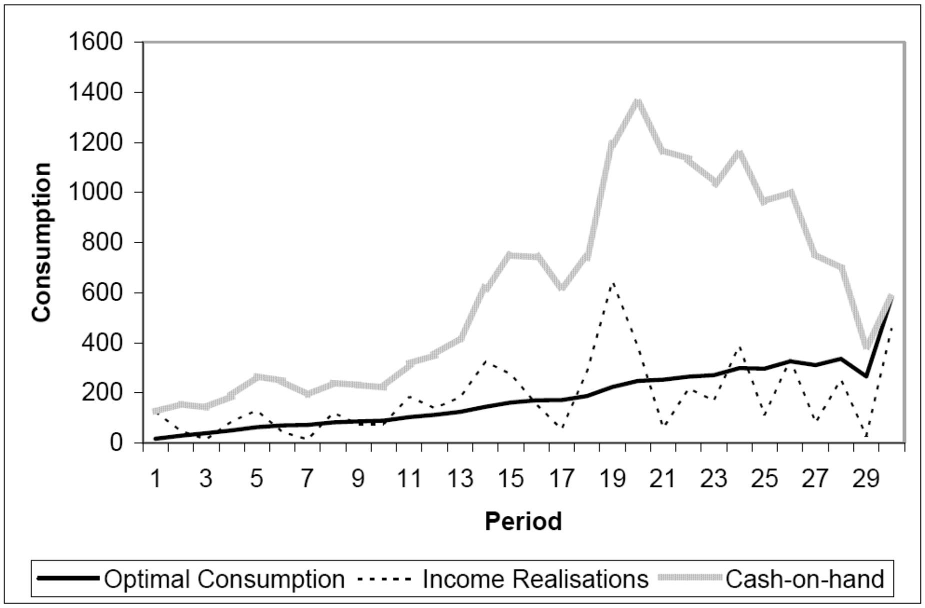

Figure 2.2: An optimal consumption path

savings to earn a reasonable utility from consumption). The optimal cash-on-hand in the example steadily rises to 1500 in period 20, building up a buffer stock which is about six times the annual income at that point. That is, consumers should brace themselves for a rainy day by saving until about period 20. After period 20, they should start to dissave by spending more than their current income and dipping into their cash-on-hand (i.e., the optimal consumption line is usually above the dotted income line after period 20).

Table 2.1 gives summary statistics of actual point outcomes in the two learning conditions. The first and second rows give the average of total lifecycle points in each condition, and the standard deviation across subjects. The third row is the difference between the average point total and the (unconditional) optimal point total.[11] The fifth row is the total income in each lifecycle (which gives an idea of whether deviations from optimality in a particularly lifecycle are due to bad decisions or to bad income luck). With only private learning, performance in the first three lifecycles was well below the unconditional optimum and highly variable across subjects. However, by lifecycle four the average subject earned point totals within 80% of the optimum and the variability across subjects shrank.

Table 2.1 (bottom panel) shows that social learning brought point outcomes close to the optimum

[11]Note that in some cases, the average subject does better than the unconditional or conditional optimum (i.e., the deviation from optimality is positive). This can happen if participants overspend (underspend) but get lucky (unlucky) and have good (bad) income shocks in later periods.

Lifecycle	1	2	3	4	5	6	7
			Private Learning				
Mean Points $\sum_{t=1}^{T} u(C_t, H_{t-1})$ (average over subjects)	**118**	**-53**	**224**	**450**	**-65**	**435**	**440**
Std. dev. Points	**635**	**694**	**498**	**297**	**475**	**255**	**146**
Deviation from Optimum $\sum_{t=1}^{T} u(C_t, H_{t-1}) - \sum_{t=1}^{T} u(C_t^*, H_{t-1}^*)$ (average over subjects)	***-453***	**-628**	**-349**	**-125**	**11**	**-153**	**-149**
Total Income $\sum_{t=s}^{T} X_t$	**5471**	**7083**	**5215**	**6235**	**4300**	**4571**	**4789**
			Social Learning				
Mean Points	**325**	**586**	**559**	**589**	**309**	**539**	**504**
Std. dev. Points	**238**	**54**	**93**	**62**	**255**	**73**	**47**
Deviation from Optimum	**-215**	**-68**	**-69**	**-66**	**-220**	**-66**	**-49**
Total Income	**4342**	**5416**	**5224**	**5901**	**4193**	**5344**	**5050**

Table 2.1: Summary statistics of actual point outcomes

rapidly. The mean and variation of points in the very first lifecycle with social learning are similar to those statistics from lifecycles 4–7 with only private learning.[12]

2.3.3 Behavior Relative to Conditional Optimization

The Table 2.1 statistics compare point totals to *unconditional* optimal level of spending in each period. This can be a misleading comparison because *conditional* optimal spending in each period depended on the participants *actual* cash-on-hand and accumulated habit stock. A subject who had made some bad decisions in early periods, but then wised up and made conditionally optimal decisions in later periods, would look bad in Table 2.1 but may have been close to conditionally optimal overall when those few early mistakes are averaged with the smarter later decisions.

Each subject's average conditional deviation for each period is the difference between their actual spending and the optimum (conditioned on that participants earlier decisions). Figure 2.3 plots the conditional deviation paths for lifecycles 1 and 7 with private learning, along with 95%

[12]It should be noted that in both conditions lifecycle 5 featured the lowest total income (the harshest income draws). In condition 1 it managed to cause the subjects and the ex ante optimal path to have negative utility. In condition 2 it only reduced the utility of the subjects.

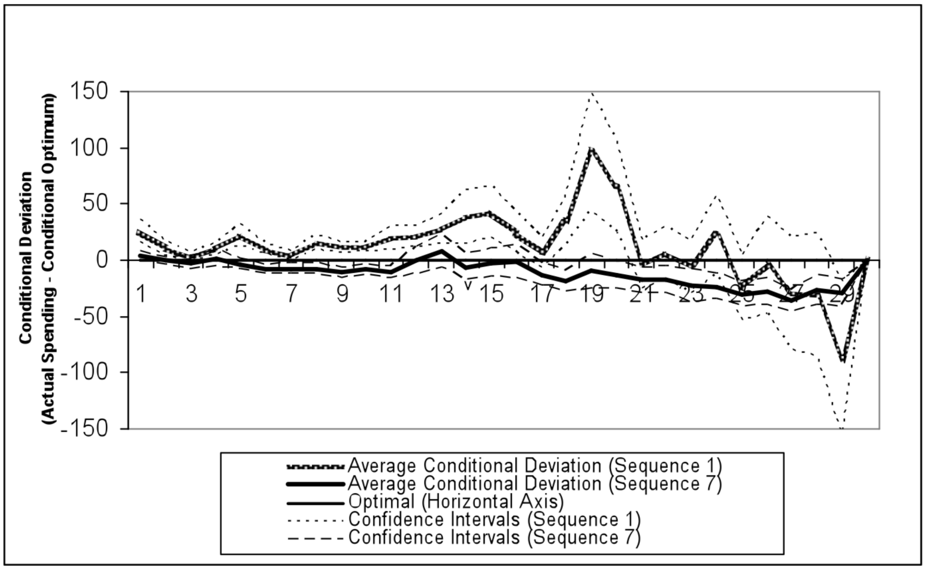

Figure 2.3: Deviations from conditional optima, lifecycle 1 and 7, private learning

confidence intervals (dotted lines). Since the optimal conditional path in Figure 2.3 is the zero-deviation horizontal line, the reader can judge at a glance whether deviations are significant by seeing whether the confidence interval covers the zero line or is far from it.

Figure 2.3 confirms the conclusion from Table 2.1: With only private learning, participants in lifecycle 1 spent significantly more than optimal in early periods, until about period 20 (when they often spent too little). However, the lifecycle 7 conditional deviations are never significantly different from zero, which shows that learning was very effective over the seven lifecycles. In fact, the actual spending path is insignificantly different from the conditional optima by lifecycle 4.

Figure 2.4 shows the analogous data for the social learning condition. These small deviations are deliberately plotted with the same y-axis scale as in Figure 2.3, to show how much smaller the deviations are when there is social learning compared to private learning. Deviations are insignificantly different from zero in most periods. There is also little difference between lifecycles 1 and 7 in the social learning condition. The initial performance is so close to optimal that there is little left to learn over the seven lifestyles.

To measure the effects of private and social learning, we regressed the log of the absolute deviation from the conditional optimum on dummy variables for lifecycles (excluding the first lifecycle), the period number and its square, and dummy variables for social learning condition, gender (Female$=1$, mean$=.43$) and ethnicity (Chinese$=1$, mean$=.50$).[13]

[13]See Chua and Camerer (2004) for details. Ethnicity is of interest because Singaporean Chinese have one of the

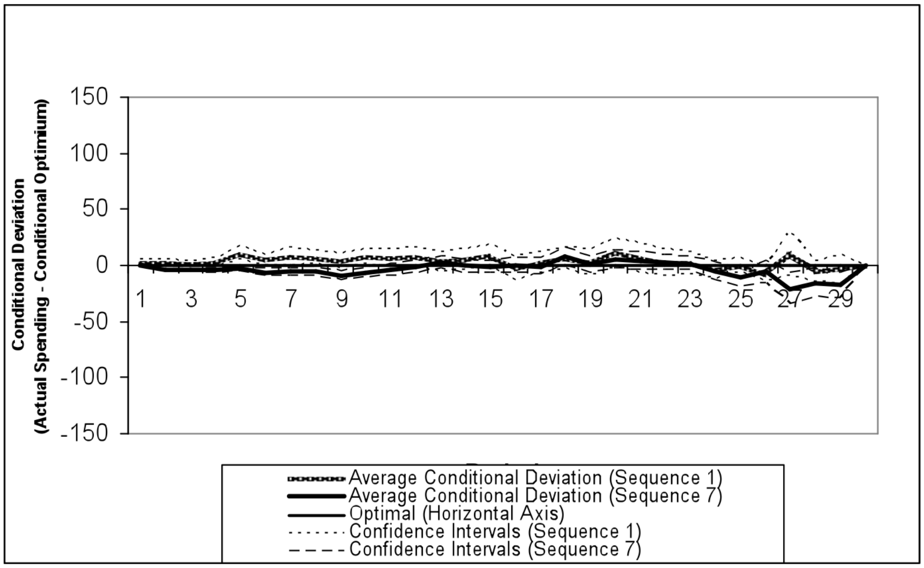

Figure 2.4: Deviations from conditional optima, lifecycle 1 and 7, social learning

Table 2.2 shows the results. The period effect is positive (but nonlinear because the period^2 effect is negative) because the absolute deviations are larger in later periods, when incomes are larger. The social learning main effect is highly significant (it implies a 24% reduction in conditional deviation), as are the dummy variables lifecycles, 5 and 7, reflecting learning across lifecycles. There is no significant effect of ethnicity and a small effect of gender (women deviate about 20% more).

2.4 Study 2: Beverage Rewards and Temptation

2.4.1 Experimental Design

Study 2 was the same as the first study except for one large change.[14] Lifecycles 1, 2, 4 and 5 (with money rewards) were the same as in study 1. However, in lifecycle 3 subjects received a fixed monetary payment for their participation but did not earn any additional money for decisions.

highest savings rates in the world (see Carroll et al., 1999). Participant random effects were also included to control for individual differences, which are substantial. In a broader specification a Caltech dummy variable was also included but is insignificant and is dropped. The Chinese dummy variable is correlated with subject pool, but not strongly. There are many ethnic Chinese students at Caltech, and Singaporean students are not exclusively Chinese.

[14]One reason to keep the complex design with habit formation and stochastic income was because behavioral research suggests that higher cognitive loads make people more likely to succumb to visceral temptation (Shiv and Fedorikhin, 2002). Additionally subjects are more likely to succumb to temptation if they are unaware they are doing so (Baumeister et al., 1994) or if the signals of doing so are noisy (Bodner and Prelec, 2003; Benabou and Tirole, 2004).

	Model (3)
Social Learning	**-0.24***
	(-2.51)
Lifecycle 2	**0.092***
	(2.30)
Lifecycle 3	**-0.027**
	(-0.67)
Lifecycle 4	**0.075**
	(1.86)
Lifecycle 5	**-0.43****
	(-10.69)
Lifecycle 6	**-0.063**
	(-1.58)
Lifecycle 7	**-0.17****
	(-4.21)
Period	**0.084****
	(15.91)
Period Squared	**-0.00034***
	(-2.01)
Female	**0.19***
	(1.99)
Chinese	**0.0006**
	(0.01)
Constant	**0.77****
	(16.39)
R^2	**0.20**

Table 2.2: Regression of log(absolute conditional deviation) (t-statistics in parentheses)

Instead, in each period they drank an amount of a beverage[15] proportional to their consumption decisions each period (1 ml beverage for each 2 points). The Excel interface was modified to show the total milliliters of beverage reward to be obtained, rather than points (utilities). It also displayed the maximum milliliters of beverage reward that could be obtained from spending all available cash immediately. As noted in the introduction, this change was designed to see if savings decisions about abstract money reward were different than viscerally tempting rewards—namely, liquid consumption by thirsty subjects.

To make this reward appealing and limit satiation across the experiment, subjects were asked not to drink for four hours before the experiment began.[16] They also began by eating some salty snacks. Since it took them 45 minutes to read the instructions and to complete two 30-period lifecycles for money before the beverage lifecycle, they were definitely thirsty by the time they reached the beverage lifecycle. It is likely that they did not satiate during the lifecycle because no subject received more than 350 ml of soda (less than a 12 oz. can) of Coca-cola in that lifecycle, subjects would only be able to drink a maximum of 20 ml/period (0.7 oz), and beverage periods were separated by one minute.[17] Subjects were required to drink their entire beverage consumption in that one-minute period (and they always did).

A syringe pump was used to deliver an exact amount of beverage into a cup.[18] If subjects incurred a negative number of points in any period, they incurred a debt of sorts—they would not receive any beverage until that level had been offset by future positive point totals. This debt was "forgiven" at the end of the beverage lifecycle because we could not force subjects to "pay back" the debt by taking away the beverage (as we do in the money lifecycles).

There were two different reward-delivery conditions in the beverage lifecycle. In the immediate condition subjects received their beverage reward right after making their decision. In the delayed condition subjects received their beverage reward chosen in period t ten periods after making their decision, in period $t + 10$.[19] Quasi-hyperbolic or present-bias models of time discounting had

[15] Subjects were given their preference of Coke or Pepsi, and could substitute Diet Coke or Diet Pepsi if they requested it. We used these beverages because they are widely valued, water was as motivating as colas, and because pilot subjects (including the middle coauthor) thought fruit juices that were tried were too filling and might induce satiation which complicates the analysis.

[16] There is no way to know whether all subjects obeyed our request to show up thirsty. However, because assignment to the immediate and delayed conditions did not depend upon apparent thirst, uncontrolled and unmeasured differences in pre-experiment thirst are sources of sampling error in comparing the two groups which lower the power of the test and bias the test against finding a difference between the immediate and delayed conditions.

[17] The concavity of utility and properties of the buffer stock savings model ensure that no subject could earn more than 700 points in any beverage or monetary lifecycle.

[18] See section A.1 for a diagram of the beverage delivery apparatus.

[19] To standardize both conditions completely there were forty periods of one minute each in lifecycle 3. In the imme-

predicted that subjects would drink more beverage in the early periods of the immediate condition because delayed rewards would be heavily discounted (for much more detail on this model and the similar dual-self model, see supplemental section A.3). In the delayed condition, immediate choices did not lead to immediate consumption so the present bias term in β-δ discounting would disappear. Intuitively, the delayed condition would provide external self-control that helps β-δ discounters. Subjects should have drunk more overall in the delayed condition if they were quasi-hyperbolic discounters.

Subjects were $n = 52$ Caltech students.[20] Because a single liquid-delivery apparatus was used, experiments were conducted in a single office rather than a computer lab with one subject at a time. As a result, this study was more laborious than most economics experiments (taking about 130 hours of experimenter-subject contact time).

2.4.2 Results

2.4.2.1 Total Beverage Awarded

The hyperbolic discounting and dual-self models predicted that subjects in the immediate condition would receive less beverage than in the delayed condition, because they would consume relatively more compared to a total-reward-maximizing optimum in early periods.[21] This prediction is empirically correct (see Table 2.3, row 1). The immediate-condition subjects drank less total beverage on average (179 ml, $s = 84.6$) than the delayed-condition subjects (226 ml, $s = 79.0$). There was substantial variation across subjects, but this difference is significant at conventional levels by one-tailed tests (t-test $p = 0.047$, Mann-Whitney rank sum test $p = 0.015$).

diate condition, subjects did nothing in the last ten periods. In the delayed condition, subjects made decisions in the first ten periods of the delayed condition but received no rewards. In the last ten periods of that condition subjects received their rewards from periods 21–30 but made no decisions.

[20] The first 44 subjects were run from April 21 to July 27, 2005. After that, 11 more subjects were run from February 7–16, 2006 to enlarge the sample and check robustness of the result. Two subjects refused to drink during the beverage period and were dropped from the analysis. Another subjects data were lost by mistake.

[21] An alternative explanation is that the first taste of cola primed subjects to consume more. Since immediate subjects first received cola after period 1 and delayed subjects first received cola after period 11, this priming could be responsible for the difference in total rewards. While thirst priming has been studied in psychology through subliminal means (e.g., Strahan et al., 2002), to our knowledge no psychological work has studied of found evidence of this specific type of priming. Further the data finds an average consumption increase (11.7 vs. 14.2 ml) between periods 1 and 2 for the immediate condition, but an average consumption *decrease* (10.4 vs. 8.2 ml) for the delayed. The jump between periods 1 and 2 in the immediate condition is most likely explained by a very high income draw in period 2 ($\eta_t \approx 3.56$) that on average sextupled cash-on-hand (76.6 vs. 429.1).

	Immediate	Delayed	Parametric test	Nonparametric
Total beverage received	176.78 (81.31)	215.65 (82.89)	t=1.71 p=0.047	z=2.09 p=0.018
Total expected losses from optimal (bounded at 350 ml)	171.91 (128.13)	96.98 (104.04)	t=2.35 p=0.011	z=2.34 p=0.010
Average expected loss from overspending	18.36 (28.78)	6.40 (10.91)	t=1.92 p=0.031	z=1.77 p=0.038

Table 2.3: Summary statistics comparing immediate and delayed conditions in the beverage lifecycle. Sample standard deviations are in parentheses below means. All p-values are one-tailed.

2.4.2.2 Adjusting for Skill

Simply comparing total beverages in the immediate and delayed conditions does not control for possible differences in skill or discounting between subjects in those conditions, which could be evidenced by differential performance in the four money lifecycles. To control for these skill differences, we estimate the regression

$$P_{it} = a + b_1 r_1 + b_2 r_2 + b_3 r_4 + b_4 r_5 + b_5 I + e_{it} \tag{2.3}$$

where P_{it} is the point total for subject i in lifecycle t, r is a dummy variable for lifecycle i, and I is a dummy variable for the immediate condition. If immediate consumption triggered overconsumption and poorer savings accumulation, $b_5 < 0$.

Notice that point totals can be negative for the beverage lifecycle, but the total ml of beverage consumed cannot be negative. (Subjects could not have been forced to "pay back" liquid once it is consumed.) This constraint is different than for the money rounds because a monetary point debt accumulated in one lifecycle could be offset by other lifecycles (and subjects are aware of this difference in incentive structures). If a large beverage deficit (> 350 ml) occurred in an earlier period, subjects would have known that no amount of spending could have erased this deficit. As a result, when subjects had large negative point totals they could have become indifferent about future decisions (their marginal incentive disappears) and produced high negative points. These high deviations occurred disproportionately in the immediate condition, which then greatly overstates b_5 when the dependent variable is points (see Table 2.4).[22] In order to reduce the effects of these outliers, two

[22] Since subjects know they will not be forced to pay back previously consumed beverage, it is conceivable that they exploit this design property by deliberately overconsuming in early periods and then running up point debts they do not have to pay. We do not estimate such a model because the period-specific maximum of liquid consumption is 20 ml per period, so the marginal beverage value of increased consumption falls sharply. As a result, subjects who are trying to optimize total liquid would smooth consumption and would never deliberately run up a debt. In terms of our estimation

	Points	Beverage	Sign-preserved Log Points
Immediate condition *(I)*	-28137.36** (7284.97)	-39.58* (20.39)	-4.21** (1.43)
r_1	-171.61 (6264.08)	-20.78 (16.77)	-2.30 (1.20)
r_2	-458.17 (6264.08)	-57.27** (16.77)	-3.73** (1.20)
r_4	105.10 (6264.08)	44.28** (16.77)	1.14 (1.20)
r_5	-656.71 (6264.08)	-26.28 (16.77)	-6.93** (1.20)
constant	282.30 (5151.29)	215.82** (15.14)	4.60** (1.02)
R^2	0.09	0.16	0.22
N	268	268	268

**p<.05 (one-tailed)* ***p<.01 (one-tailed)*

Table 2.4: Regression of periods and condition on subject performance. The standard errors on dummy variables are identical because of the type of covariance matrix.

alternative regressions were run. In the second specification, each lifecycle money point total was calculated as if it were a beverage lifecycle (i.e., periods with negative utility are ignored). In the third specification, extreme point totals were reduced in magnitude by taking the logarithms of their absolute values with their sign preserved (i.e., the dependent variable is $[|P_{it}|/P_{it}]\log(|P_{it}|)$.

Table 2.4 shows the results of a random effects regression run on each model. In all three specifications the sign of b_5, the effect of the immediate condition, is negative and significant at $p < 0.05$. In fact, these results are stronger in significance than the parametric t-tests reported in Table 2.3, which implies that accounting for individual differences in skill by using the money-lifecycle results actually enhances the significance of the immediate-delayed condition difference (by reducing variation from cross-subject differences in skill or patience.)

These analyses use the overall point totals in the lifecycle. As in study 1, it is also useful to examine *conditional* deviations in each period given decisions in previous periods. For each period in the beverage lifecycle we calculated the future expected points for that subject resulting from her decision, compared to the future expected points from a conditionally total-reward-maximizing optimal decision in that period. We then converted these amounts to ml of beverage and totaled these values over all thirty periods. Since no subject received more than 350 ml of beverage in the

below, a deliberate strategy of overconsuming because of anticipated "bankruptcy" would be misclassified as a low value of the discounting parameter δ. There is no a priori reason to think this pattern will be more common in the immediate and delayed conditions if both types of subjects have similar discounting patterns.

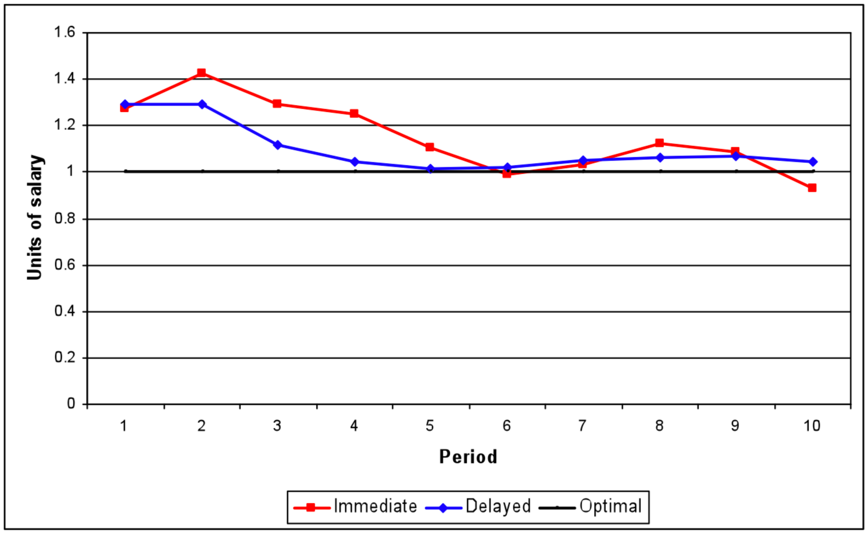

Figure 2.5: Ratio of average consumption to conditional optimal by condition, beverage lifecycle, periods 1–10

lifecycle or less than 0 ml, we bounded all totals at 350 ml. Row 2 of Table 2.3 shows the results. The average total expected beverage loss, in conditional deviation from optimality, was much higher for the immediate condition than for the delayed condition (about twice as high).

2.4.2.3 Exploring the Time Series of Overspending in Early Periods

Figure 2.5 shows the average ratios of spending to conditionally optimal spending. (In the first 10 periods the optimal line is now just a flat line at a ratio of 1.) Figure 2.5 confirms that even when conditioning on past decisions, the immediate-condition subjects were spending more in the first five periods. (After that period the higher number of subjects with beverage deficits and large habits in the immediate condition pushed down their overspending.)[23] Another diagnostic statistic is the average overspending in those periods in which subjects overspent compared to the conditional optimum. The immediate group subjects actually made somewhat fewer overspending decisions than the delayed-condition subjects (41% vs. 51% of decisions),[24] but when they had overspent, the immediate condition subjects spent much more than was optimal (Table 2.3, row 3), which created greater expected losses.

[23] Immediate subjects have more beverage deficits (4 subjects vs. 1 in period 6; 15 vs. 8 by period 10) and higher average habit levels accumulated (218 vs. 185 in period 6) than the delayed condition. It is not the case that the immediate subjects have satiated on soft drinks compared to the delayed group, because the immediate subjects have only drunk about 57 ml (2 oz) on average after five periods.

[24] Periods in which a subject encountered a deficit of 20 ml or greater were omitted in this analysis.

2.5 Estimating Quasi-Hyperbolic Discounting Parameters

The results presented to this point have supported the basic prediction of the hyperbolic discounting and dual self models, that subjects in the immediate condition would consume less overall. Because the hyperbolic model is clearly parameterized, we can also estimate best-fitting values of the parameters $\tilde{\delta}$ and $\tilde{\beta}$ from savings decisions and compare those values to estimates from other studies. The analysis is restricted to observations when subjects did not encounter beverage deficits. When a subject encountered a beverage deficit, their decision could only be made to receive *future* rewards and so $\tilde{\beta}$, the immediate bias term, should not apply even in the immediate-beverage condition.[25]

In the quasi-hyperbolic model, the weights placed on immediate and future rewards are 1, $\beta\delta$, $\beta\delta^2, \ldots \beta\delta^t \ldots$ If δ is close to one the terms δ^t are close in numerical value, so there will be many combinations of (β, δ) values which produce similar sequences of weights and similar choices. It is therefore difficult to estimate the two parameters separately. When (β, δ) were maximized simultaneously, the analysis often yielded toward 0 or above 1. We therefore use a two-stage procedure to calibrate δ and β for each subject.

Since behavior in the delayed condition gave no information about the present bias β, in theory, the delayed-condition data is used to estimate δ. So we first search for best-fitting values of δ_D which explain delayed-condition subject choices as if they were maximizing discounted expected utility of consumption with a discount rate δ_D and $\beta = 1$. These estimates minimize the sum of squared percentage deviations between the actual consumption and the consumption predicted by the model. This estimation gives a distribution of δ_D estimates with a mean of 0.904 and standard deviation, across subjects, of 0.230. This mean value is reasonable but is significantly less than one at the 2% level by a cross-subject t-test (see Table 2.5). (Note that a discount factor around 0.9 is more plausibly interpreted as a reduced-form expression of suboptimal choice rather than true time preference for these short-horizon experiments.)

The next challenge is to estimate β_I values in the immediate condition, using reasonable values of δ.[26] The procedure we use first fixes $\beta = 1$ for each immediate-condition subject and then

[25] Additionally, subjects with high enough beverage deficits knew they would not receive liquid again and have no incentive to choose one spending decision over another. While some subjects never encountered a beverage deficit, and others encountered them early, each subject was given a single parameter value and the results were analyzed so that each subjects value counts as much as any other.

[26] Using the mean of the delayed-condition estimates δ_D and estimating subject-specific β_I works poorly because differences in δ values for those subjects from the mean δ_D leads to implausible variation in estimates of β_I. The problem with using the delayed-condition mean δ_D for the immediate-condition subjects is the following: suppose an immediate-condition subjects δ is smaller than the mean δ_D. Then the best-fitting sequence of weights $1, \beta\delta, \beta\delta^2, \ldots \beta\delta^2$ will overestimate β because the β parameter is forced to pick up the slack for the under-estimated δ. Similarly, if the

Model	Standard	Sophisticated	Naïve
mean $\tilde{\delta}$ of Delayed (std deviation)	1 n/a	0.904** (0.230)	0.904** (0.230)
mean $\tilde{\beta}$ of Immediate (std deviation)	1 n/a	0.619*** (0.211)	0.721*** (0.134)
average squared deviations per period, before deficits, immediate subjects, using mean $\tilde{\beta}$ and $\tilde{\delta}$	0.230	0.189	0.193

*1-tailed t-test of parameter < 1: *p<.05 **p<.02 ***p<0.01*

Table 2.5: Two-stage parameter estimates of $\tilde{\delta}$ and $\tilde{\beta}$

estimates a best-fitting value of δ_I for each of those subjects. These values are shown in Table 2.6; the mean δ_I is 0.85 and the standard deviation is 0.24. Since we are fixing $\beta = 1$, but we believe the actual β_I values might be below 1, we need to adjust the δ_I values in some way that permits more precise estimation of β. We do this by projecting the subject-specific values of δ_I onto the value of the distribution of δ_D estimated from the delayed-condition subjects which has the same standardized deviation. That is, a specific immediate-condition estimate δ_I is adjusted to an estimate δ_I^* where $(\delta_I^* - 0.904)/0.23 = (\delta - 0.85)/0.24$. This procedure permits individual differences in δ_I values, but yokes their distribution to the distribution of δ_D values to permit better identification of β. Using these adjusted values of δ_I^* for each immediate-condition subject, we then estimate β_I for each subject.

There is one further complication. In quasi-hyperbolic models, people can be either sophisticated or naïve (e.g., O'Donoghue and Rabin, 1999). Sophisticated subjects discount delayed payoffs steeply but understand that in the future they will discount steeply too. Naïve subjects discount steeply but believe, mistakenly, that their current discount factors applied to future periods will also be applied to later decisions.

The difference between sophistication and naïveté can be illustrated in a three period example. In the first period, both types of subjects apply weights 1, $\beta\delta$ and $\beta\delta^2$ to the three periods. However, the sophisticated subject knows that the discount rates 1 and $\beta\delta$ will actually be applied to periods 2 and 3 when period 2 decisions are made, and accounts for this weighting in forecasting period 2 and 3 choices. The naïve subject thinks the discount rates $\beta\delta$ and $\beta\delta^2$ will be used in period 2 to weight period 2 and period 3 utilities; since the $\beta\delta$ term will divide out in optimization, the naïve subject therefore thinks the relative weights applied in periods 2 and 3 will be 1 and δ (i.e., the naïve subject thinks he will act like an exponential discounter in the future).

immediate-condition δ is below the mean δ_D, β will be underestimated. Indeed, when we tried this procedure the estimate of β tends to bifurcate to the lower and upper bounds placed on β.

Subject	Best fit $\tilde{\delta}$	Projected $\tilde{\delta}$ Delayed Condition	Soph $\tilde{\beta}$	Naïve $\tilde{\beta}$	Soph SS fit per period	Naïve SS fit per period	Soph fit- Naïve fit
i1	1.07	1.11	1.10	1.19	0.074	0.082	0.008
i2	0.83	0.88	0.32	0.86	0.014	0.022	0.008
i3	0.41	0.48	0.79	0.63	0.393	0.364	-0.029
i4	1.05	1.09	0.89	0.69	0.051	0.051	0.000
i5	1.06	1.10	0.54	0.58	0.005	0.001	-0.004
i6	0.51	0.58	0.49	0.73	0.012	0.010	-0.002
i7	0.87	0.92	0.70	0.74	0.004	0.002	-0.002
i8	1.02	1.07	0.70	0.58	0.492	0.467	-0.025
i9	1.04	1.08	0.57	0.71	0.008	0.012	0.004
i10	0.83	0.88	0.67	0.76	0.004	0.004	0.000
i11	1.01	1.05	0.40	0.82	0.017	0.023	0.006
i12	0.12	0.21	0.08	0.44	0.014	0.001	-0.012
i13	0.70	0.75	0.69	0.75	0.027	0.023	-0.004
i14	0.91	0.96	0.73	0.69	0.021	0.013	-0.008
i15	1.10	1.13	0.61	0.70	0.005	0.010	0.004
i16	1.04	1.08	0.67	0.58	0.063	0.045	-0.018
i17	1.04	1.08	0.77	0.86	0.027	0.033	0.006
i18	0.89	0.94	0.82	0.69	0.056	0.042	-0.014
i19	0.99	1.04	0.19	0.61	0.042	0.038	-0.005
i20	0.97	1.01	0.61	0.69	0.002	0.001	-0.001
i21	0.87	0.92	0.65	0.73	0.077	0.074	-0.004
i22	0.99	1.03	0.54	0.79	0.018	0.023	0.004
i23	0.72	0.78	0.71	0.78	0.016	0.014	-0.002
i24	0.94	0.99	0.76	0.66	0.049	0.036	-0.013
i25	0.45	0.52	0.58	0.72	0.064	0.058	-0.006
i26	0.79	0.85	0.51	0.77	0.007	0.011	0.003
Mean	0.85	0.90	0.62	0.72	0.060	0.056	-0.004
Median	0.93	0.97	0.66	0.71	0.020	0.023	-0.002
St Dev	0.24	0.23	0.21	0.13	0.116	0.109	0.010

Table 2.6: Estimated $\tilde{\beta}$ of individual subjects in immediate condition

In simple choice experiments these two behavioral assumptions are difficult to distinguish empirically, but our 30-period experiment gives some empirical leverage for distinguishing them. We therefore estimate β values (using the adjustment procedure described above) assuming both sophisticated and naïve forecasting of future behavior (see supplemental section A.3 for full details).

The results are summarized in Table 2.5. The estimates of β in both the sophisticated and naïve models are clustered around 0.6–0.7. Table 2.6 shows individual results;[27] all but one subject's estimate is below 1 for both specifications, so the hypothesis that there is no present bias ($\beta = 1$) is strongly rejected.[28] The estimates of $\tilde{\beta}$ are in the ballpark of estimates of Angeletos et al. (2001) ($\tilde{\beta} = 0.55$), Laibson et al. (2007) ($\tilde{\beta} = 0.7$), DellaVigna and Paserman (2005) ($\tilde{\beta} = 0.9$), and Tanaka et al. (2006) ($\tilde{\beta} = 0.74$–0.89) (from macroeconomic calibration, consumption data, unemployment spells, and experiments in Vietnam, respectively). The values also are close to other experiments with a much different design that used juice and water rewards ($\tilde{\beta} = 0.52$, McClure et al. (2007)).

Measured by the sum of squared deviations, the naïve model fits better in 16 of 26 subjects. Since this structure is not deliberately designed to distinguish the two specifications, this is just a clue that both specifications should be taken seriously as explanations of behavior in future work.

2.6 Applying the Quasi-Hyperbolic Model to Bounded Rationality and Social Learning

Study 1 and study 2 both feature decisions made by subjects in the same experimental framework, so it is useful to have a unified approach to compare the results parametrically. In the previous section, we estimated the results in study 2 using the quasi-hyperbolic model because that model has been designed and commonly used to calibrate the tradeoffs between immediate and delayed rewards over time, and used to model temptation. Any reasonable unified model must accommodate the empirical immediate-delayed difference and the β-δ model is one way to do so. However, it has never been used to represent bounded rationality or the effects of private and social learning on decisions

[27]A possible correlate of individual β values are subject values on the Barratt Impulsivity Scale. After subjects had completed their experimental session, they answered a survey measuring their total "impulsivity" on the BIS 11 Barratt Impulsivity scale (Patton et al., 1995). However, these values show little correlation with the individual naïve betas, sophisticated betas, and subject performance (correlations smaller than 0.1 in absolute value).

[28]The correlation of β and δ estimates across subjects is around .35 for both specifications of β, so there is no serious identification problem.

One approach to creating a unified model is to estimate β-δ parameter values for the private and social learning effects. Since learning means, empirically, saving more at the beginning, in the β-δ framework learning is expressed as a change in these preference parameters. While this is clearly a reduced-form approximation, it is the approach we take below. We return to the details after discussing why other approaches are not likely to fit these data any better. Of course, future research should certainly tackle the problem of developing a more sensible and unified approach, and designing the best experiments to test it.

Two other ways to potentially model private and social learning involve limited planning horizons and rules of thumb.

Limited planning horizons: Ballinger et al. (2006) found that their data are reasonably explained by a model in which subjects tend to only think ahead three periods. A more general version of this approach is a model in which people optimize but act as if only K periods remain (Ballinger et al. estimate $K = 3$). Note that full optimality is $K > 30$ (in these 30-period experiments) and in beginning period $T - K$ subjects fully optimize (conditionally). Empirically, this model cannot explain all our data. The period at which consumption becomes conditionally optimal can be used to approximate K, because in period $T - K$ people will begin to (conditionally) optimize. Figure 2.3 suggests K is around 10 because their decisions are conditionally optimal beginning around period 20. But a model with a horizon $K = 10$ predicts that subjects will consume everything in the delayed condition of study 2 (since consumption is delivered 10 periods later in the delayed condition). This appealing model cannot easily account parametrically for both the oversavings in study 1 (which implies K around 10) and the fact that there is a limit on consumption in the 10-period-delayed condition of study 2. To be clear, our view is that the truncated-horizon model is a very plausible one, but it just does not do well in explaining the central empirical features of both of our studies.

Rules of thumb: Another approach to model bounded rationality and observed learning is that consumers use a rule of thumb which is adjusted by experience (e.g., Cochrane, 1989). Two plausible rules are consuming a constant fraction of current income or a constant fraction of accumulated cash-on-hand. Neither model fits our experimental data especially well.[29]

The actual consumption-to-income ratio does not exhibit a trend across periods, but fluctuates

[29] See Chua and Camerer (2004) for a regression of consumption against conditionally optimal consumption and rule-of-thumb spending of a constant percentage of current income. The latter term has no statistical weight. Our interpretation is that while subjects are not exactly optimizing (they clearly undersave in early periods with only private learning), the variation across the 30 periods is much better picked up by variation in optimal consumption than by a constant rule of thumb.

wildly across periods (see supplemental section A.4 for details). For example, in the first lifecycle of the private learning condition, in half the periods the propensity to consume out of current income is less than 1, but in six of 30 periods it is above 2 (i.e., subjects spend all the current income and also dip into savings, because current income is too low to produce an adequate consumption utility). Subjects seem to have some intuitive ability, even in the first lifecycle, to adjust spending from current income to smooth consumption across periods so a simple rule-of-thumb model is strongly rejected. The actual consumption-to-cash-on-hand ratio also exhibits little trend before learning takes place, but also fluctuates substantially (in six periods it is around .4, and in seven periods it is around .7, in the first private learning lifecycle). Learning does change this ratio so it looks like a quadratic polynomial across periods after learning takes place.

The most promising approach to modeling bounded rationality is a propensity to consume policy function, which is a low-order polynomial in the state variables (including cash-on-hand and the number of periods remaining), which adjusts with experience (e.g., Houser et al., 2004). Adjusting from experience is not so straightforward, however, because each lifecycle only provides one observation on performance of a particular cross-period policy. It is not clear how to adjust a polynomial policy across 30 periods from a single observation on the entire dynamic policy's performance rapidly enough to match the human learning we observe (see the well-known "credit assignment" problem in learning of dynamic policies, e.g., Holland, 1985).

Furthermore, even if we had an ideal model of variation and learning across rule-of-thumb policies that could explain the observed learning in experiment 1, these models are not likely to explain the immediate-delayed condition effect in study 2. Therefore, we use the β-δ approach as a benchmark unified model (since it is the most natural way to explain the immediate-delayed difference). We treat estimated changes in those parameters as very reduced-form expressions of learning as expressed through these parameters.

Because of the problem in separately identifying β and δ in this design, ad hoc methods are used to first identify δ, then estimate β given the estimates of δ.

First consider private learning. We assume that $\beta = 1$ in lifecycle 7, estimate δ values from those subjects in lifecycle 7 (the estimate is 1.00), then apply those estimates of δ to estimate β in lifecycles 1 and 4. These numbers indicate the strength of learning, if learning is assumed to only change β and not affect δ. These results are shown in Table 2.7. The estimates of first- and fourth-lifecycle β are 0.415 and 0.778 for the naïve model and 0.273 and 0.585 for the sophisticated model. These numbers are a crude indication of the size of the suboptimality in consumption; they suggest

Condition	Lifetime 1		Lifetime 4		Lifetime 7 [a]
	Naïve	Sophisticated	naïve	sophisticated	
private learning	0.415 (0.339)	0.273 (0.352)	0.778 (0.323)	0.585 (0.450)	1[b] n/a
social learning	0.671 (0.230)	0.421 (0.424)	0.974 (0.192)	1.025 (0.398)	1[c] n/a

Table 2.7: Two-stage parameter estimates of $\tilde{\beta}$ and $\tilde{\delta}$ by learning condition $(N = 36)$

that about half the gap between the first lifecycle β and the value of one is closed by lifecycle 4.

For social learning, we use the same procedure. First $\beta = 1$ is assumed in lifecycle 7 and an estimate of δ is derived from the lifecycle 7 data (the estimates is 1.03). Since the estimate of δ is so close to 1, we use the value of 1 to enhance comparability across the two conditions when later estimating β in lifecycles 1 and 4. The resulting estimates are 0.671 and 0.974 for the naïve model and 0.421 and 1.025 for the sophisticated model.

Together, all these figures give us a crude parametric index of the strength of learning, when learning is parameterized by a change in the immediacy preference β—more appropriately, a reduced-form proxy for undersaving. Assuming $\beta = 1$ in the last lifecycle (to permit identification), social learning increases β estimates substantially in the first lifecycle compared to private learning (from 0.415 to 0.671, assuming naïveté, or 0.273 to 0.421 assuming sophistication). The learning is apparently much more rapid in social learning, as well, since the estimates of β are very close to one even in lifecycle 4 but are still far from one in private learning.

2.7 Conclusion

Dynamic choice models in which current choices influence future constraints or utilities are computationally difficult. Savings in the presence of income uncertainty and habit formation are an example of choice models in this class, which are especially relevant in the economy.

Empirical evidence on savings suggests people are not always saving optimally (though many studies are consistent with some features of optimal savings). However, tests with field data depend sensitively on assumptions about expectations, separability of consumption, and other unobservables. Experiments control for these assumptions. Simple experiments done by others generally show that experimental subjects save too little. Our goal in this paper was to extend this research to a more complicated (and lifelike) environment that combines income uncertainty and habit formation.

The experiments are designed to examine two explanations for undersavings. The first is bounded rationality and the second is an immediacy preference (i.e., even if people knew the optimal saving rules they could not execute them). We test these explanations in two experimental studies using a common design and pool of highly skilled subjects. We use the β-δ model to measure undersaving across both studies.

We find that subjects saved much too little at first, but learned to save close to optimal amounts after three or four lifecycles of direct experience (private learning). Furthermore, subjects who have received social learning, examples of successful, unsuccessful, and average experimental performance, produced savings decisions that are quite close to optimal even in their first lifecycle. Since consumers are limited to one lifecycle of private learning (absent reincarnation with memory), it would be interesting to know what types of social learning are more effective. Does social learning work better when it comes from family and friends, from total strangers, from financial planners, or from training and education? Our data suggest that one type of social learning works well but invites consideration of other forms which can be tested in future experiments and in field data.

The fact that subjects could learn to save optimally for money rewards led us to explore whether they saved optimally when rewards are more immediate and visceral—when thirsty subjects' rewards were immediate sips of a cola beverage. The subjects who sipped the beverage immediately also overspent (i.e., overdrank), compared to group of subjects who made decisions in one period but did not get to sip that period's beverage amount until ten periods later. As a result of their overspending, subjects in the immediate-reward condition earned less total rewards than those in the delayed condition, and received less than the theoretical, total-reward-maximizing optimum. This unique feature of our second study provides a model for future studies of highly tempting decisions like addiction, overeating, and perhaps spending splurges. The difference in the performance of the immediate and delayed conditions is consistent with the predictions of both the quasi-hyperbolic and dual self models, and is not consistent with the standard exponential model. When parameters of the quasi-hyperbolic model are calibrated from subject decisions in the immediate condition, the mean best-fitting (the degree of present bias) is 0.62 for the sophisticated case and 0.72 for the naïve case. These values are close to values observed in some other studies using both calibrations to aggregate data and direct experimental measurement. Parameter estimates using beta as a representation of inexperience are much lower than those observed in field data and experiments which allow immediacy preference. Although this model was not intended to be applied to measure bounded rationality in a reduced-form way, the lower β suggest that if consumers were inexperienced and did not know how much to consume, they would be much more impatient than has been inferred by Angeletos et al. (2001) and Laibson et al. (2007) from aggregate savings and investment.

There are many directions for future research. The experimental paradigm could also be extended by adding more lifelike features, such as stochastic mortality, retirement, and supply-side advice that either tempts subjects more or gives them good advice. The fact that subjects in the delayed condition are able to resist temptation better (and drink more total beverage as a result) corroborates the conclusion of models like Bernheim and Rangel's (2004), that creating a time wedge between "ordering" and consuming may be helpful to people. This observation suggests an experimental way to measure demand for external self-control. The immediate-condition subjects are making a mistake, but they can't help doing so. If they had access to external commitment, sophisticated hyperbolics would seek external commitment. Future experiments could allow subjects in beverage studies the choice between whether they want to participate in the immediate or delayed condition; sophisticated subjects should opt for the imposed delay. Naïve hyperbolics and exponential discounters would be indifferent about both conditions. An alternative theory (Gul and Pesendorfer, 2001) suggests that agents might prefer the delayed condition if it reduces disutility from temptation.

The natural question about experiments of this type is how well their results generalize to naturally occurring savings by different groups of people. While economic agents cannot experience more than one lifecycle, they can learn from the savings success and mistakes of others. Retirement advisors may exist because individuals are unable to make retirement decisions in one lifecycle but can make good decisions after observing multiple lifecycles (and those histories are bottled and sold by advisors) and with formal tools to analyze and explain what to do. The market may have solved the cognitive problem in savings models by producing a supply of helpful retirement advisors. These phenomena can be studied in experiments too, by allowing markets for advice and group-level decisions (e.g., household saving) to see whether these institutions contribute to optimal choice.

Chapter 3

To Review or Not Review? Limited Strategic Thinking at the Box Office

3.1 Introduction

The central principle in Bayesian-Nash equilibrium analysis of games with information asymmetry is that players can correctly infer what other players know from their actions. In contrast, models in which strategic thinking is limited due to cognitive constraints allow the possibility that some players *do not* correctly infer what the actions of other players imply.

Models of limited strategic thinking have been shown to explain data from a wide variety of experimental games[1] and auctions[2] better than equilibrium predictions. In many cases, these models express the idea that players with private information can fool some of the people, some of the time, in contrast to the equilibrium assumption that nobody is ever fooled (see Crawford, 2003).

This paper is one of the first applications of different models of limited strategic thinking,[3] to a novel field setting. The novel setting is the differential box office earned by movies that are "cold opened," i.e., deliberately unavailable for pre-release review by critics, compared to movies that get pre-release reviews. Through an iterative reasoning process, rational moviegoers should correctly infer quality and a cold opening premium should not exist. In equilibrium, therefore, cold openings provide a natural field setting to examine models of limited strategic thinking (cognitive hierarchy and cursed equilibrium) as well as the rational-actor, quantal response equilibrium model.

This study estimates behavioral parameters (which ideally include full rationality as a limiting

[1] See Nagel (1995), Stahl and Wilson (1995), Camerer et al. (2004), Crawford and Iriberri (2007a).

[2] See Crawford and Iriberri (2007b) and Wang (2006).

[3] Two unpublished studies using field data and cognitive hierarchy approaches are Östling et al., (2007) using Swedish lottery choices and experimental analogues, and Goldfarb and Yang (2007) using estimation of firm adoption of 56K modems. The Östling et al. study compares QRE and cognitive hierarchy approaches but Goldfarb and Yang do not compare to QRE, and neither paper estimates the cursed equilibrium model as we do.

case) to see whether parameters have some stability across economic domains and have similar value for consumer (moviegoers) and firm (distributors) behavior.

Ours is one example of a more general class of games in which a producer or person that knows something about a product's quality can choose whether to disclose its quality or not (see Verrecchia (2001, section 3) and Fishman and Hagerty (2003) for surveys). For instance, a car salesman can signal a vehicle's quality by adding a warranty (Grossman, 1981). Online daters can decide whether to post a picture or not (Levitt and Dubner, 2005). Restaurants can voluntarily post health department ratings when not required to by law (Jin and Leslie, 2003). HMOs can choose whether to voluntarily disclose quality by submitting to independent accreditation (Jin, 2005). A regulated firm can selectively report information about its industry to regulators (Milgrom, 1981). A hedge fund can selectively report past earnings (Malkiel and Saha, 2005). In politics and law, the analogous situation is when one can choose to answer a direct question about a fact they might know, or avoid answering the question (e.g., "pleading the fifth" in legal settings).[4]

The interesting empirical question in these settings is what limitedly rational consumers and voters infer from the reluctance to reveal quality when it is easy to do so.

The behavioral explanation of cold opening of movies is straightforward. Suppose moviegoers prefer to see high-quality movies, but some moviegoers do not pay attention to reviews or do not infer anything negative about the quality of cold opened films (even though unreviewed movies are lower in quality, empirically).[5] If film studios know a film's quality in advance, believe that critics will judge quality fairly but think that some moviegoers will not realize that no review is bad news about quality, they will deliberately keep some mediocre movies from being critically reviewed in advance, by opening them cold.[6] The technique has been used in the industry for some time (and is increasingly used recently). As Litwak (1986) notes

> "As a courtesy, and to ensure that reviews are ready by the time a film is released, studios arrange advance screenings for critics. However, if negative reviews are expected, the studio may decide not to screen a picture, hoping to delay the bad news. (p. 241)"

[4] In 1999 while campaigning for President, George W. Bush was the only one of the twelve major candidates to not disclose whether he had used cocaine. He subsequently revealed that he had not used cocaine for the previous seven years, and then later disclosed he had not used cocaine for the 15 years before his father had become President in 1989. (DeFrank, 1999) Although no physical evidence of cocaine use was found, disclosure literature would suggest this implies Bush had used cocaine prior to 1974.

[5] Cold-opened movies have an average metacritic rating (0–100) of 22, while all movies have an average rating of 48. See Section 3.2 for an explanation of metacritic ratings.

[6] As Greg Basser, CEO of Village Roadshow Entertainment Group, told us, "If you screen [a bad movie] for critics all they can do is say something which may prevent someone from going to the movie."

If some moviegoers do not realize that a cold opening is bad news about quality, there will be a cold opening premium: Box office revenues will be higher for movies opened cold (controlling for quality as measured by later critic ratings and other characteristics such as budget, star power, etc., that influence revenue).

A fully rational analysis, due originally to Grossman (1981) and Milgrom (1981), implies no cold opening premium. If moviegoers correctly infer that a cold opened movie is probably bad news, and form conditional rational expectations given that belief, then there will be no cold opening premium and only the worst movies will be opened cold.

The argument can be illustrated numerically. Suppose movie quality is uniformly distributed from 0 to 100. If distributors cold open movies with quality below a cutoff 50, moviegoers with rational expectations will infer that the expected quality of a cold opened movie is 25. But then it would pay to screen movies with qualities between 26 and 100, and only cold open movies with qualities 25 or below. Generally, if the distributors do not screen movies with qualities below q^*, the consumers' conditional expectation if a movie if unscreened is $q^*/2$, so it pays to screen movies with qualities $q \in (q^*/2, 100]$ rather than quality below q^*. The logical conclusion of iterating this reasoning is that only the worst movies (quality 0) are unscreened.[7]

The rational logic does not appear to completely jibe with some basic facts about movies. Seven percent of the movies in our sample are opened cold (and that percentage has increased sharply in recent years, see Figure 3.1). Regressions show that cold opening appears to generate a box office premium (compared to similar-quality movies that are pre-reviewed, and including other controls). The estimated premium is robust to various specifications and is absent in UK and Mexico grosses (which are typically later than US releases, so that information about movie quality is likely to have leaked out across national borders).[8]

Assuming moviegoers infer quality from critics, we estimate specific parametric models of the

[7]There are other models of discretionary disclosure that do not have this full unraveling result. Some models have costly disclosure, which cause distributors to only reveal information up to a certain threshold (Viscusi, 1978; Jovanovic, 1982). Other models have sellers uninformed about the quality of their product but able to learn it with some probability (Dye, 1985; Jung and Kwon, 1988; Dye and Sridhar, 1995) or at a cost (Matthews and Postlewaite, 1985; Farrell, 1986; Shavell, 1994).

We do not believe either assumption fits this particular industry. Fishman and Hagerty (2003) allow a portion of consumers to be unable to interpret revealed information. They find three main equilibria, one where quality is always revealed, one where it is never revealed, and one where high quality is revealed and low quality is not. However, because there are only two quality levels, in the third equilibrium all quality is also revealed. To our knowledge, an explanation involving a proportion of uniformed consumers in the population cannot generate a box office premium or explain why a non-zero percentage of movies were cold opened.

[8]For all practical purposes, there are no cold-opened movies in Mexico or the UK, since only the biggest blockbusters, which are not cold opened in the US, are released simultaneously in those countries.

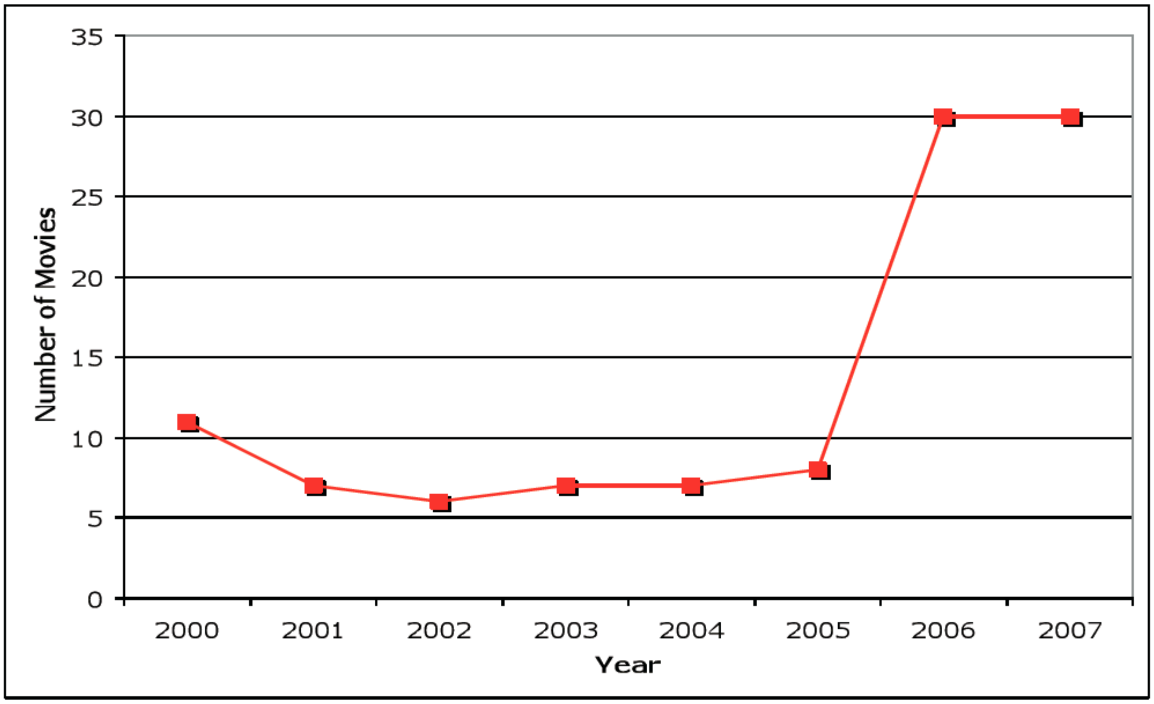

Figure 3.1: Cold openings by year, 2000–2006

degree of limited strategic thinking. A benchmark equilibrium model is quantal reponse equilibrium (QRE). Two models which allow systematic errors in beliefs about the actions of other players, developed to explain experiments, are cursed equilibrium (Eyster and Rabin, 2005) and a cognitive hierarchy (Camerer et al., 2004). It is important to note that fully rational behavior is a special case of those models, because their behavioral parameters have specific numerical values if players are fully rational. Thus, the model estimation allows a sharp comparison of rational limiting cases and more general behavioral specifications, and gives us a numerical measure of the degreed limited rationality. Field applications like these are important in showing whether principles of limited rationality that were inspired and calibrated by experimental data can also explain some basic facts in larger-scale field settings (see DellaVigna, 2007, for many examples).

The degree of limited rationality by consumers shown here is not large. Roughly speaking, the estimates suggest about ten times a year a few million Americans pay \$5–10 in money, and an hour or two in time, to see a movie they might not have seen if they had inferred from the lack of pre-release reviews that the movie was not very good. But this small, synchronized mistake by millions of moviegoers is a multi-million-dollar profit opportunity for movie distributors. This example is a reminder that the industrial organization implications of a behavioral mistake depends on the psychology underlying the mistake and on industrial structure (e.g., Ellison, 2006).

Furthermore, the central conclusion here is that some consumers fail to infer that product quality is bad when objective quality reviews are actively avoided or withheld by producers. This possibility could be explored in many other types of markets and settings where the failure to signal quality should itself be informative.

In fact, three field studies of consumer quality disclosure found results consistent with the hypothesis that not all consumers fully infer quality information from a failure to disclose (a conclusion similar to ours), by comparing firm behavior under voluntary and mandatory disclosure. Mathios (2000) studied nutrition labelling of salad dressing. Most low-fat dressings (less than 9 grams of fat per serving) were voluntarily labelled for fat content before mandatory disclosure, while only 15% of high-fat dressings were labelled. After mandatory disclosure, the share of the high-fat dressings fell by about 20%. Jin and Leslie (2003) studied the effects of a shift from voluntary to mandatory posting of health-rating cards in Los Angeles restaurants. They find that mandatory disclosure increases hygiene scores by 5.3%, which is about half a standard deviation of the distribution, and which is modestly significantly higher than under voluntary disclosure.[9] Jin (2005) shows that HMOs do not voluntarily disclose quality (via NCQA accreditation) in markets that are the least competitive. She also finds that HMOs which voluntarily disclose by seeking accreditation tend to serve areas with large employers, which suggests that HMOs are responding to more sophisticated customers (big firms with more savvy benefits managers).

All of these studies (and ours) are inconsistent with the strong hypothesis that customer strategic thinking leads to complete disclosure.

This chapter is organized as follows: Section 3.2 describes the dataset, the cold-opening premium, and refutes possible alternative explanations for the premium. Section 3.3 introduces a general model for moviegoer and distributor behavior; Section 3.4 calibrates it from the data to the QRE, Cursed, and CH models, and compares those models to data observed in the labratory. Section 3.5 concludes, examines this empirical work with others, and provides suggestions for future research.

[9] Their test probably understates the effects of a shift from voluntary to mandatory disclosure because some of the voluntary-disclosure cities were expected to adopt mandatory disclosure in the near future. Restaurants might have begun complying early during the last parts of the voluntary regime, and earlier than they would have if they did not expect a shift to mandatory disclosure.

3.2 Data

The data set contains all 890 movies widely released[10] in the U.S. in their first weekend, over the $6\frac{1}{2}$ year period from January 1, 2000 to June 30, 2006.[11]

Critic and moviegoer ratings are used to measure quality. Metacritic.com ratings are used to measure critic ratings. Metacritic.com quantifies and averages ratings from over 30 movie critics from newspapers, magazines, and websites. The metacritic rating is available for all non-cold-opened movies on the day they are released and available on Monday for cold opened movies. In this way it is generally exogenous from box office revenue measures.[12]

A natural question to ask is whether metacritic.com ratings accurately capture the quality of movies as perceived by moviegoers and revealed by demand. Our analysis indicates they do. Figure 3.2 indicates that metacritic ratings are highly correlated with the logarithm of total US box office revenue. This result is also found in studies of critic influence (Eliashberg and Shugan, 1997; Reinstein and Snyder, 2005).

We also examine the aggregated user ratings on imdb.com, the largest internet site for user movie reviews. We find a high correlation (.76) between metacritc scores and imdb user reviews (see Figure 3.3). The result is not specific to genre (see Table 3.3 below). Metacritic scores align with two clear indicators of movie popularity (imdb and box office).[13]

The squares in Figure 3.2 represent the cold-opened movies in our sample. No cold opened movie has a metacritic rating higher than 55, and the average rating for those movies is 25. However, the graph does not conclusively show whether cold opened movies do better than non-cold-opened

[10] Attention is restricted to movies initially released in over 300 theaters. Movies in more limited release have much less box office impact (they are usually art house movies that use a platform strategy of starting on a few screens, then expanding). It is also likely that information about quality leaks out more rapidly for these movies if they later go into wide release, even when they are initially opened cold.

[11] Movies before 2000 are excluded because Metacritic.com's records did not cover every movie from before 2000.

[12] Ratings such as the imdb.com user rating are determined by the people who see the movie and who give reviews afterwards. We treat that variable as a measure of popularity, but consider it endogenous to box office.

[13] Of course the assumption that we are making is that critic reviews influence moviegoers. Alternatively, (i) critics have different sensibilities than moviegoers and have no correlation with actual popularity; (ii) critics have the same sensibilities, but moviegoers ignore them so they have predictive, but no influencing power.

We find a strong correlation between box office revenue and critic rating as well as ex post moviegoer ratings of movies and critic ratings across genres. With this result and the findings of both Eliashberg and Shugan (1997) and Reinstein and Snyder (2005) (who both studied this very issue) we assume that critics generally have the same quality beliefs as moviegoers. The evidence is weaker that critics influence moviegoers. While survey evidence (Simmons, 1994) suggests one third of moviegoers use critical reviews to make decisions, Eliashberg and Shugan reached no definitive conclusion on this issue, and Reinstein and Snyder found evidence that critic ratings matter for specific genres. However, that study only examined the effect of a specific two critics (i.e., Siskel and Ebert) delaying their review. A cold opening delays *all* reviews and thus might have a greater effect. Because this evidence is somewhat inconclusive, we will use several different tests to check our hypothesis that it is indeed the cold opening increasing box office and thus the critic reviews (or lack thereof) influencing moviegoers.

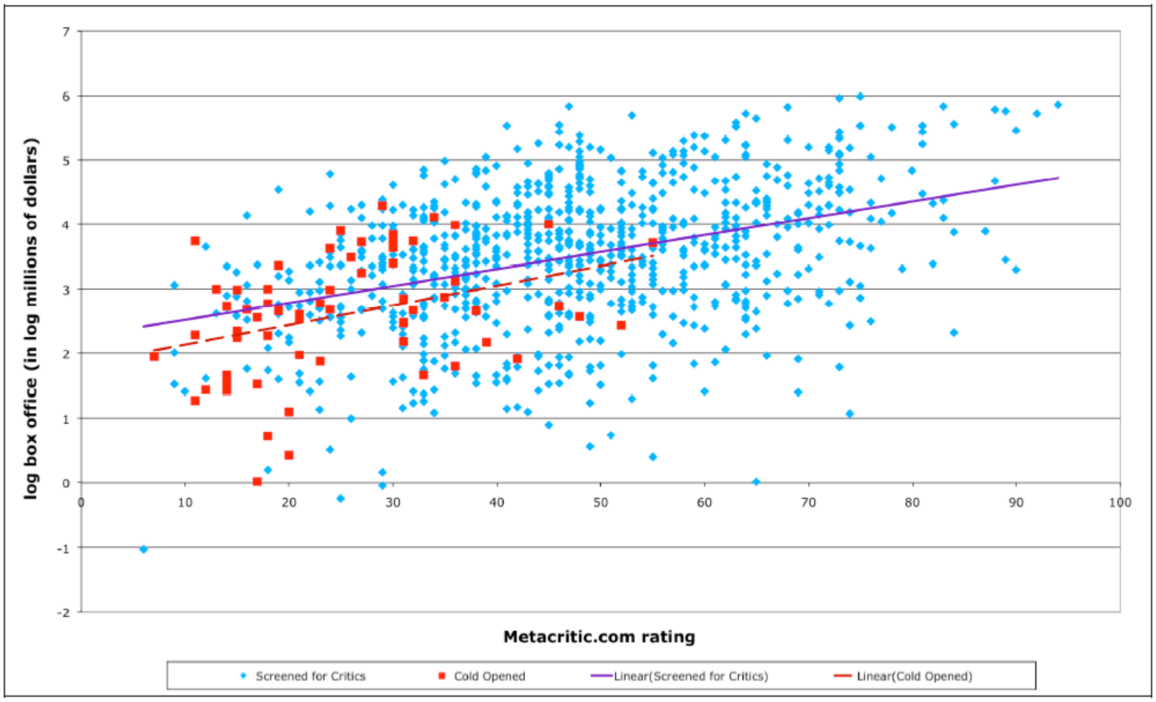

Figure 3.2: Scatterplot of metacritic.com quality ratings and log box office revenue

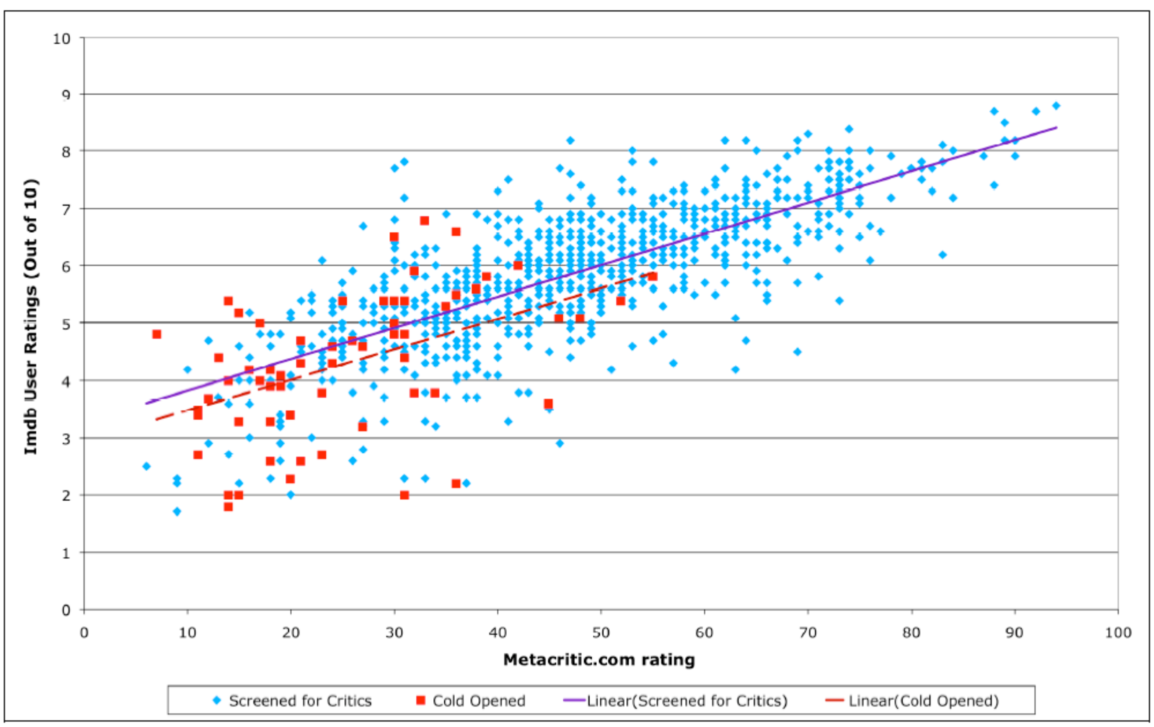

Figure 3.3: Scatter plot of metacritic.com quality ratings and imdb user ratings

variable\regression	mean	median	standard dev
cold	1.242	1.260	0.973
log total box office revenue (logcubo)	3.443	3.510	1.092
log 1st weekend box office revenue (logwkd)	2.354	2.390	0.942
metacritic rating (crit)	45.793	46.000	16.813
theaters opened (in thousands) (thtr)	2.435	2.550	0.787
production budget (in millions) (bud)	42.299	33.360	33.484
average competitor budget (in millions) (comp)	42.061	35.110	27.465
average log star ranking (star)	4.607	4.500	1.672
summer open (1=Jun, Jul, Aug) (sum)	0.249	6.130	4.514
sequel or adaptation (1=yes) (sq/adpt)	0.389	0.000	0.488
opening days bef fri (1=Thurs, etc.) (beffri)	0.224	0.000	0.658
opening wkd length (days) (wkdlen)	0.109	0.000	0.336
early foreign open (days) (forbef)	11.804	0.000	99.724
action/ adventure (1) (act/adv)	0.164	0.000	0.371
animated (1) (ani)	0.060	0.000	0.237
comedy (1) (com)	0.380	0.000	0.486
documentary (1) (doc)	0.006	0.000	0.075
fantasy/scifi (1) (fant/sci)	0.062	0.000	0.241
supense/ horror (1) (susp/hor)	0.157	0.000	0.364
year of release (2003=0) (year)	-0.166	0.000	1.902
PG (1) (pg)	0.158	0.000	0.365
PG-13 (1) (pg13)	0.478	0.000	0.500
R (1) (r)	0.326	0.000	0.469

Table 3.1: Summary statistics for variables. There are only 856 observations for production budget; all other variables have 890 values.

movies because there are other variables that are not included in Figure 3.2 that correlate with box office revenues (e.g., the initial number of screens on which the movie is shown).

Cold opening, box office revenues, movie genres and ratings, production budgets, and star power ratings are collected from various data sources (see Supplemental Section B.1 for a more detailed description). All these variables were used in a regression model to test if movies that are cold opened have significantly greater logged opening weekend and total US box office. Table 3.1 provides summary statistics for all variables.

Each movie, j, has a metacritic.com rating, q_j, a dummy variable for whether a movie was cold opened, c_j, (=1 if cold) and a vector X_j of other variables. The regression model is

$$\log y_j = aX_j + bq_j + dc_j + \epsilon_j \tag{3.1}$$

where y_j is logged opening weekend or total US box office for movie j in 2003 dollars, standardized

variable\regression	logcubo	logcubo	logwkd	logwkd
cold	0.154**	0.140**	0.147**	0.171***
	(0.090)	(0.088)	(0.073)	(0.073)
crit	0.021***	0.020***	0.013***	0.013***
	(0.001)	(0.001)	(0.001)	(0.001)
thtr	0.864***	0.853***	0.848***	0.820***
	(0.039)	(0.035)	(0.032)	(0.029)
bud	0.003***	0.002***	0.002***	0.002***
	(0.001)	(0.001)	(0.001)	(0.001)
comp	0.022***	0.003***	0.001**	0.001**
	(0.001)	(0.001)	(0.001)	(0.001)
star	-0.045**	-0.054***	-0.029**	-0.048***
	(0.015)	(0.014)	(0.013)	(0.011)
sum	0.052*	-	0.028	-
	(0.050)		(0.041)	
sq/adpt	0.124***	0.116***	0.119***	0.111***
	(0.045)	(0.045)	(0.037)	(0.037)
beffri	0.000	-	-0.048*	-
	(0.033)		(0.027)	
wkdlen	0.129**	-	0.175***	-
	(0.063)		(0.051)	
beffor	0.000	-	0.000	-
	(0.000)		(0.000)	
act/adv	-0.173**	-	-0.052	-
	(0.078)		(0.063)	
ani	-0.316**	-	-0.145	-
	(0.128)		(0.105)	
com	0.032	-	0.027	-
	(0.064)		(0.052)	
doc	0.212	-	0.267	-
	(0.362)		(0.052)	
fant/sci	-0.175*	-	0.039	-
	(0.103)		(0.084)	
susp/hor	0.012	-	0.040	-
	(0.078)		(0.064)	
year	-0.083***	-0.082***	-0.050***	-0.050***
	(0.011)	(0.011)	(0.009)	(0.009)
pg	-0.182	-	-0.034	-
	(0.130)		(0.106)	
pg13	-0.179	-	0.100	-
	(0.136)		(0.111)	
r	-0.225	-	0.113	-
	(0.140)		(0.114)	
const	0.506**	0.433***	-0.493***	-0.188***
	(0.209)	(0.134)	(0.171)	(0.110)
R-squared	0.677	0.675	0.718	0.707
N	856	856	856	856
degrees of freedom	21	8	21	8
*p<0.1, **p<0.05, ***p<0.01				

Table 3.2: Regressions of log box office revenues (in millions)

Genre	# Movies	Cold Opens	Percent	Avg Weekend Log(BO)	Avg Metacritic	Avg imdb user rating	Imdb-Meta Correlation	Cold Dummy for Weekend Log(BO)
Act/Adv	143	5	0.03	3.827	48.238	5.911	0.7928	0.165
Animated	51	1	0.02	3.949	56.627	6.071	0.8456	-0.547
Comedy	322	21	0.07	3.352	42.096	5.413	0.7321	0.173
Doc	3	0	0.00	2.351	58.000	5.633	0.8663	-
Drama	145	3	0.02	3.240	50.510	6.275	0.6866	0.229
Fant/Sci	55	4	0.07	4.120	50.764	6.115	0.8771	0.121
Susp/Horr	137	25	0.18	3.334	41.401	5.698	0.7368	0.077
Overall	856	59	0.07	3.443	45.793	5.762	0.7641	0.155

Table 3.3: Data separated by genre

using the GDP deflator (`www.bea.gov`). Table 3.2 shows the regression results.

The point of this initial regression is not to estimate a model that captures the cold opening premium or the decision to cold open; that will be done in Sections 3.3 and 3.4. Instead, it is to determine whether there is a difference in the revenue between cold opened and screened movies. (Notice that if taken as a literal model for the industry equation 3.1 suggests that all or no movies should be cold opened depending on the sign of d.) Under the standard equilibrium assumption that all quality information of cold opened movies is revealed, we should see no difference in revenues, and the cold coefficient should be zero.[14]

The "cold" coefficient in the first row of Table 3.2 shows that cold opening a movie is positively correlated with the logs of opening weekend and total US box office (see Supplemental Section B.2, Table B.1 for a similar result with opening day data). These coefficients suggest that cold opening a movie increases revenue from 14–17%.[15,16] These effects persist when the "lean" regressions are run with the most significant variables (i.e., only including the highly significant variables cold, metacritic, theaters, budget, competition, star ranking, sequel or adaptation dummy, and year of release). The lean regressions show a more significant effect for opening weekend box office as conventional wisdom suggests, because critic reviews of cold opened movies are normally available after the Monday after the first weekend.[17] The coefficients also suggest that cold opening increases

[14] Alternatively, a switching regression model (similar to Borjas, 1987) for the choice to cold opened could be used to capture the cold opening premium and characterize the decision to cold open. We have instead chosen to describe the industry through a quantal response model (see Section 3.4).

[15] For the average gross of a cold opened movie, \$20 million, this is roughly \$3 million of box office revenue.

[16] Although we would consider it a regression with endogeneity, these results do not change using imdb.com user ratings instead of metacritic ratings.

[17] It is somewhat surprising that the effect of a cold opening continues after the first weekend when critical reviews are available. Intuitively, the cold opening effect should occur during the first weekend and then dissipate as moviegoers learn the true quality of a cold opened movie. However, a likely alternative explanation is that moviegoers infer quality from the first weekend's revenue (see De Vany and Walls (1996) for a model with such dynamics). Then the perceived

movie revenue by roughly the same amount as the previous regressions (14–17%). The regression results in Table 3.2 are generally sensible. Higher quality leads to higher box office — an increase in one metacritic point increases revenues by 2.1%. An extra ten million dollars in production budget is correlated with a 3% increase in revenues. The number of theaters opened, which often indicate expectations about movie revenues, have a very large effect.[18] For an increase of 1000 theaters movie revenue increases 86%. The averaged logged star power rankings have a negative correlation (higher numbers indicate lower rankings and less revenue). Adaptations and sequels increase box office by roughly 13%, a result which may foreshadow the recent explosion in the fraction of movies in this category.

3.2.1 Alternative Explanations of the Cold Opening Premium

It is possible that cold opened movies have some other characteristic omitted from the Table 3.2 regressions that causes these movies to generate apparently greater box office (an omitted variable bias). In this case, our regressions are not capturing the effect of cold opening, but are capturing the effect of an omitted variable that is correlated with cold opening. Since many controls are included, the most likely omitted variable that could be correlated with the decision to cold open is spending on publicity and advertising.[19] Omitting this variable would explain the cold opening premium if revenues increase with spending on advertising, and if advance screening and advertising are substitutes (i.e., distributors spend more on ads to compensate for cold opening). However, a senior executive at Fox Studios we interviewed contradicted this notion, suggesting that if anything distributors are *tighter* with their spending on advertising once the decision to cold-open is made (which happens late in the process, after the number of screens and most other variables have been determined). The executive's view was that distributors know cold-opened movies are not very good, and see high levels of ad spending on such movies as throwing good money after a bad movie.[20] Further, the industry appears to typically set advertising budgets as a fixed proportion

"effect" of a cold opening on post-first-weekend box office is a secondary result from cold opening affecting the first weekend's box office. The data agree with this assessment; if we run a regression on logged box office revenues after the first weekend (see Table 3.3), including logged first weekend with our other independent variables, then we find cold has an insignificantly negative effect ($-3\%, p \approx 0.5$) after first weekend revenue. The lean regression finds a significantly negative relationship between cold and post-first weekend box office, ($-10\%, p < 0.1$). Still, the increase in first weekend revenue and its correlation with post-opening weekend box office more than offset any negative effect of cold opening on post-opening-weekend box office.

[18]Theaters may be a proxy for ad budget as well, which may magnify their effect.

[19]Unfortunately, we found advertising budgets for only 445 of the 856 movies in our sample, and only 12 of the 59 cold openings.

[20]Our regression on log advertising budget suggests a cold opening is associated with a 10% drop in advertising budget, but recall this result is only based off 12 cold openings and is not significant (p ≈ 0.3; lean regression p ≈ 0.23).

dependent variable	cold opening (dummy) coefficient	std error	t-statistic	one-tailed significance
log total US box office	0.154	0.090	1.710	0.044
log total US rentals	-0.007	0.101	-0.067	-
log UK box office	-0.021	0.231	-0.090	-
log Mexico box office	-0.001	0.150	-0.010	-

Table 3.4: The cold opening coefficient in non-US box office markets

of production budgets (Vogel (2007) suggests $\frac{1}{2}$, an executive at Village Roadshow told us $\frac{2}{3}$). If these rules of thumb are true, then our production budget variable will pick up the omitted effect of advertising.[21]

Another potential problem is that annoyed critics might give cold-opened movies lower critical ratings than they would have if the movies were screened in advance (perhaps as a way of punishing the studios for making the movie unavailable).[22] This explanation seems unlikely since critics pride themselves on objectivity (for example, they rarely mention in late reviews of cold opened movies that the movie was unavailable in advance).

One way to test an omitted-variable bias is to look at the log total box office of the U.K. and Mexico, and log of US video rental data. In these markets, the possible deception of cold opening should be less effective because movies are almost always released in the U.K. and Mexico after the initial U.S. release, and home video rentals are always later. If information about the movie's quality is widely disseminated by these later releases, the cold opening effect should disappear. Table 3.4 reports the cold-opening coefficients (from regression including all variables as in 3.2). There is apparently no cold opening premium in these two foreign markets and the rental markets, which works against the hypothesis that the premium is due to an omitted-variable bias and is consistent with the hypothesis that some moviegoers are fooled by cold openings.[23]

Another way to check whether cold opened movies have any inherent differences in sensitivity

[21] A regression of production budget on marketing budget has $R^2 = 0.496$, indicating much of adverting budget is dependent on production budget.

[22] Litwak (1986) mentions this idea when describing a cold opening.

[23] Another explanation is that moviegoers of cold-opened movies are less sensitive to critic reviews. Then the high turnouts for cold-opened movies have nothing to do with the opening, but just the fact that given identically low critic reviews, cold-opened movies turn out more viewers. This explanation may appear appealing as the correlation of critic reviews and user reviews for cold-opened movies while high (0.5122) is much lower than the correlation of critic reviews and user reviews of non-cold opened movies (0.7649). However, this relationship likely results from the fact that cold-opened movies are on a smaller range of critic ratings ($\bar{x} \approx 25, s^2 \approx 11$). If we restrict non-cold opened movies to those with critic ratings under 40 ($\bar{x} \approx 29, s^2 \approx 8$) or above 60 ($\bar{x} \approx 70, s^2 \approx 7$), we find similar values for correlation (0.5128 and 0.5256, respectively).

to critic ratings is to examine the movies by genre. Comedies and Suspense/Horror movies account for 80% of cold openings and only 54% of all movies (see Table 3.3). If fans of these genres have less sensitivity to bad reviews (suggested by Reinstein and Snyder, 2005) and are more likely to go to a movie that has low critic ratings than fans of other genres then the cold opening premium could be a result of the selection of cold-opened movies to this genre.[24] Table 3.3 shows that this is not the case. Throughout genres moviegoers' correlation between critic reviews and self-reported reviews are all around 0.75. Further, the cold open premium is positive for all genres (6–21%) with the exception of animated, but that result is based off one movie (Doogal). However, given the few number of data points in each genre and low number of cold openings in most genres, the level of significance is low. At the very least, we can say it does not seem like the cold opening premium is a genre-specific effect.

Finally, our hypothesis is that limited iterated strategic thinking causes moviegoers to be "tricked," incorrectly overestimating the ex-ante quality of cold opened movies. Since moviegoers presumably go to these movies based on, among other factors, their perception of quality, a greater number of cold opening moviegoers will have negative impressions of their movie. Then if user reviews are representative of the audience that views a film, we should find lower user reviews for cold-opened movies holding everything else constant (like critic ratings). Using imdb.com user data and the usual independent variables, we find cold opened movies have a rating 0.4 points (out of 10) lower than non-cold opened movies. The result is highly significant ($p < 0.001$).

In the next section we will develop three structural models of strategic thinking by moviegoers and distributors and estimate behavioral parameters which measure the degree of limited strategic thinking for both groups. If some of these models can successfully explain the cold opening premium with similiar values to what has been observed before, that success is another piece of evidence that the premium is not due to an omitted variable.

3.3 The General Model

In designing a model of movie viewing and distributor choice, the aim is to create a model that can be analyzed with box office data, but allow us to estimate behavioral parameters of individual thinking.

[24]This explanation would not explain why distributors would be more likely withhold bad news in genres where the intended audience is the least receptive to bad news.

Every movie j has specific characteristics X_j and (integer) quality $q_j \in [0, 100]$. We assume that the distributor of movie j and moviegoers both know X_j. The game form is simple: Distributors observe q_j and then choose whether to open cold ($c_j = 1$) or to screen for critics in advance ($c_j = 0$). Moviegoers form a belief $E_m(q_j|c_j, X_j)$ about a movie that depends on its characteristics X_j and whether it was cold opened c_j.[25] Below we consider three models of belief formation. One is a standard equilibrium concept and two incorporate forms of limited strategic thinking.

The first assumption is that if a movie is screened to critics, its quality is then known to moviegoers:

Assumption 3.1. $E_m[q_j|0, X_j] = q_j$.

To model moviegoing and distributor decisions jointly, we use a quantal response approach in which moviegoers and distributors choose stochastically according to either utilities and expected profits. Since we have no data on individual choices or demographic market-segment data, we use a representative-agent approach to model moviegoers. Assumption 3.2 is that moviegoer utility is linear in movie characteristics and expected quality, subtracting the ticket price.

Assumption 3.2. $U(X_j, E_m(q_j|c_j, X_j)) = \alpha E_m(q_j|c_j, X_j) + \beta X_j - \hat{t} + \epsilon_j$

where α and β give the corresponding predictive utility associated with expected quality and other known characteristics of movies. The opportunity utility of not going to the movies is defined as zero.[26] In the quantal response approach, probabilities of making choices depend on their relative utilities. We use a logit specification (e.g., McFadden, 1974). The probability that the representative moviegoer will go to movie j with characteristics X_j and expected quality $E_m(q_j|c_j, X_j)$, at ticket price $\hat{t}$[27] is

$$p(X_j, E_m(q_j|c_j, X_j)) = \frac{1}{1 + e^{-\lambda_m\left(\alpha E_m(q_j|c_j, X_j) + \beta X_j - \hat{t} + \epsilon_j\right)}} \tag{3.2}$$

where λ_m is the sensitivity of responses to utility. Higher values of λ_m imply that the higher-utility choice is made more often. At $\lambda_m = 0$, choices are random.[28] As $\lambda_m \to \infty$, the probability of

[25] It is not crucial that moviegoers literally know whether a movie has been cold-opened or not (e.g., surveys are likely to show that many moviegoers do not know). The essential assumption for analysis is that beliefs are approximately accurate for pre-reviewed movies and formed based on some different behavioral assumption for cold-opened movies.

[26] This is without loss of generality because a constant term is included in the revenue regression, which in this model is equivalent to the estimated utility of not going to the movie.

[27] The term $\hat{t}$ is the average US ticket price in midyear 2003 (recall box office revenues are in 2003 dollars). For an explanation on why movie ticket prices do not differ by movies see Orbach and Einav (2007) or for a more general explanation, Barro and Romer (1987).

[28] This model implies that if $\lambda_m = 0$, the representative moviegoer will attend each movie in its first weekend with .5 probability. While that result may be unappealing, note that a multinomial specification (i.e, if $\lambda_m = 0$, the repre-

choosing the option with the highest utility converges to one (best-response).[29]

Expected box-office revenues are assumed to equal the probability of attendance by a representative moviegoer, times the population size N and ticket price $\hat{t}$, which is $R(X_j, E_m(q_j|c_j, X_j)) = N\hat{t}p(X_j, E_m(q_j|c_j, X_j))$. Note that the distributor's choice of c_j is assumed to enter the revenue equation solely through its effect on moviegoer expectations of quality $E_m(q_j|c_j, X_j)$.

The distributor's decision to screen the movie ($c_j = 0$) or open it cold ($c_j = 1$) is also modelled by a stochastic choice function based on a comparison of expected profits from the two decisions. Given assumption 3.1, the revenue from screening is $R(X_j, q_j)$ and the revenue from cold opening is $R(X_j, E_m(q_j|1, X_j))$. Given the same logit choice specification as for moviegoers, the probability of a distributor opening the movie cold is therefore given by assumption 3.3,

Assumption 3.3. $\pi(X_j, q_j) = 1/\left(1 + \exp\left[-\lambda_d\left[R(X_j, E_m(q_j|1, X_j)) - R(X_j, q_j)\right]\right]\right)$

where λ_d is the sensitivity of distributor responses to expected revenue.[30]

The logic of the model and our data (see Section 3.2 and Table 3.2) suggest that cold opening most strongly affects the first weekend's revenue (which may then affect cumulative revenue). Therefore, we use the first weekend's revenue to calibrate the models' revenue equations and distributor decisions in the next section. Then our probability and utility functions given in assumption 3.2 and equation 3.2 are based on the moviegoers' behavior in the first weekend. Results are similar (see Supplemental Section B.2) when total box office is used.

3.4 Models of Strategic Thinking

The crucial behavioral question is what moviegoers believe about the quality of a movie that is cold-opened—i.e., what is $E_m(q_j|1, X_j)$?—and how those expected beliefs influence the distributor's cold opening choice probability, $\pi(X_j, q_j)$. This section compares three models of beliefs: Quantal response equilibrium, cursed equilibrium, and a cognitive hierarchy.

sentative moviegoer will go to the movies with .5 probability and which movie he goes to will depend on its underlying characteristics) would be much more complicated to calculate and also has unappealing results. For instance, movies that open alone each weekend should have much higher box office than those that open with three other movies (generally not true). For these reasons and because it was the most apparent logit model, this model was chosen.

Additionally, this point is moot. The later λ_m estimates will be far from 0 (see Table 3.5). As it turns out when one looks at equation 3.5, it is apparent that it would require on average movies to make roughly \$800 million in their first weekend to push $\hat{\lambda}_m$ to 0. Instead this value can be thought of as an upper bound on movie revenue and a lower bound on rationality.

[29] See Luce and Raiffa (1957), Chen et al. (1997), McKelvey and Palfrey (1995, 1998).

[30] In many previous applications of these games to experimental datasets the response sensitivity parameters λ are the same since game payoffs are on similar payoff scales. We use two separate parameters here, λ_m and λ_d, because the payoffs are on the order of dollar-scale utilities for moviegoers and millions of dollars for distributors.

Quantal response equilibrium combines the stochastic choice functions described above with the standard equilibrium assumption that agents' beliefs about the behavior of other agents are statistically correct—in this case, moviegoers' beliefs reflect an understanding of the distributors' decisions, and vice versa the "QR" part of QRE reflects the fact that players do not choose best economic responses all the time. The "E" part suggests their expectations about other players' behavior are still correct (i.e., they are still in equilibrium).

The cursed equilibrium and cognitive hierarchy approaches both allow limits on strategic thinking which are parsimoniously parameterized by a single behavioral parameter.

In cursed equilibrium, moviegoers' beliefs about the quality of a cold-opened movie are a χ-weighted average of unconditional overall average quality (with weight χ) and the rationally-expected quality that fully anticipates distributors' decisions (with weight $1 - \chi$). The parameter χ is a measure of the degree of naïveté in the moviegoers' strategic thinking (i.e., to what extent beliefs about cold-opened movies are biased toward average unconditional quality). An alternative interpretation is that all χ is the fraction of moviegoers who do not find out whether a movie was reviewed or cold-opened (e.g., they did not bother to check whether there are reviews available or not). This study is not well-designed to distinguish the cursed hypothesis (that they know there is no review but draw the wrong inference) from the lack-of-information hypothesis. Below we will call these two parametrically equivalent hypotheses into "limited rationality" but will remind the reader in the conclusion that they are fundamentally different.[31]

In the cognitive hierarchy approach, there is a hierarchy of levels of strategic thinking. The lowest-level thinkers do not think strategically at all, and higher-level thinkers best-respond to correctly anticipated choices of lower-level thinkers. For parsimony, the frequencies of players at different levels in the cognitive hierarchy are characterized by a Poisson distribution with mean level parameter τ. Importantly, both models allow full rationality as a limiting case of their behavioral parameters. Full rationality corresponds to $\chi = 0$ and $\tau \to \infty$. Therefore, the data, will indicate the *degree* of moviegoer rationality and distribution expectation of moviegoer rationality.

[31]In some applications χ is more naturally interpreted as a fraction of people who are uninformed or not thinking strategically, which might be measured directly in surveys or methods to classify people into types. However, in our specific structural framework, box office revenues are not linear in expected beliefs (through assumption 3.2). So a model in which there are a fraction χ of people who use average quality for cold-opened movies, and a fraction $1 - \chi$ who form rational expectations is not exactly equivalent. (The difference is that between a nonlinear probability function of a weighted average and a weighted average of nonlinear probabiliites.)

3.4.1 Logistic Quantal Response Equilibrium (QRE) Model

In QRE, the moviegoers use Bayes' rule and rational expectations to infer the expected quality of movies that are cold-opened from the distributors actual choice probabilities. That is, $E_m^{qre}(q_j|1,X_j)$, the QRE expectation of moviegoers about the quality of unscreened movies is

$$\begin{aligned}
E_m^{qre}(q_j|1,X_j) &= \textstyle\sum_{q=0}^{100} qP\left(q|X_j,1\right) \\
&= \frac{\sum_{q=0}^{100} qP(1,X_j,q)}{P(1,X_j)} && \text{(Bayes' rule)} \\
&= \frac{\sum_{q=0}^{100} qP(1,X_j,q)}{\sum_{q=0}^{100} P(1,X_j,q)} && \text{(laws of probability)} \\
&= \frac{\sum_{q=0}^{100} qP(1|X_j,q)P(X_j,q)}{\sum_{q=0}^{100} P(1|X_j,q)P(X_j,q)} && \text{(laws of probability)} \\
&= \frac{\sum_{q=0}^{100} qP(1|X_j,q)P(X_j)P(q)}{\sum_{q=0}^{100} P(1|X_j,q)P(X_j)P(q)} && \text{(independence assumption)} \\
&= \frac{\sum_{q=0}^{100} q\pi(X_j,q)P(q)}{\sum_{q=0}^{100} \pi(X_j,q)P(q)} && \text{(definition in (A3)).}
\end{aligned} \tag{3.3}$$

Intuitively, for agents to form an expectation about the quality of a cold opened movie $E_m^{qre}(q_j|1,X_j)$, they must consider all possible levels of quality that a movie *could have* (hence the summations over all integers in [0,100]), and the conditional probability that it would be of that quality given its characteristics and the fact that a distributor decided to cold open it with probability $P\left(q|1,X_j\right)$ (which is equal to $\pi(q|1,X_j)$ under QRE). Using laws of probability, and the assumption that the probability of any movie's quality level $P(q)$ is independent from the probability of it having any other characteristics $P(X_j)$[32] then a cold opened movie's expected quality $E_m^{qre}(q_j|1,X_j)$ only depends on the probability that a distributor would cold open a movie with such characteristics for any quality $\pi\left(X_j,q\right)$, and the frequency of quality ratings $P(q)$. From this transformation we are able to calculate $E_m^{qre}(q_j|1,X_j)$ if $\pi\left(X_j,q\right)$ is known.

The cold opening probabilities $\pi(X_j,q)$ depend on estimated revenues from opening the movie cold or screening it (and revealing its quality, assuming (A1)). We use a transformation, then regression, to estimate the revenue as a function of X_j and q. The revenue equation is

$$\begin{aligned}
R(X_j,E_m(q_j|c_j,X_j)) &= N\hat{t}p(X_j,E_m(q_j|c_j,X_j)) \\
&= N\hat{t}/\left[1+e^{-\lambda_m\left(\alpha E_m(q_j|c_j,X_j)+\beta X_j-\hat{t}+\epsilon_j\right)}\right].
\end{aligned} \tag{3.4}$$

Rearranging terms and taking the logarithm, yields a specification which is easy to estimate because

[32] Supplemental Section B.2, Table B.3 shows the intercorrelation matrix. There is only one variable which has a correlation with quality higher than .20—namely, the budget (ρ=.28).

it is linear in characteristics X_j and expected quality $E_m(q_j|c_j, X_j)$,

$$\log\left(\frac{R(X_j, E_m(q_j|c_j, X_j))}{N\hat{t} - R(X_j, E_m(q_j|c_j, X_j))}\right) = -\lambda_m\left(\alpha E_m(q_j|c_j, X_j) + \beta X_j - \hat{t} + \epsilon_j\right). \tag{3.5}$$

The QRE is recursive: Moviegoers' beliefs about the quality of cold opened movies depend on which movies the distributors choose to open cold (through equation 3.3). The distributors' choice to open cold depends on moviegoers' beliefs about the quality of cold opened movies (through assumption 3.3).

Because of this recursive structure, we estimate the model using an iterative procedure (see B.3 for details). The procedure first uses the large number of screened movies (where quality is assumed to be known to moviegoers by 3.1) to estimate regression parameters that forecast revenues conditional on quality in 3.5. Then specific expected qualities for all cold opened movies are imputed using a maximum-likelihood procedure that chooses a distributor response sensitivity λ_d which explains actual decisions best and satisfies the rational expectations property. These inferred expected qualities are then added to qualities of screened movies to re-estimate equation 3.5 and the process iterates until parameters converge. Convergence means that parameters have been found such that both the representative moviegoer and the distributors best-respond (stochastically) and the moviegoer rational-expectations constraint on cold-opened movies equation 3.3 is satisfied.

Table 3.5 shows the regression results from six iterations from this process (which stopped according to the step 6 convergence definition in Supplemental Section B.3). The r-squared value, 0.682, shows our model has a reasonable fit with the data. The final log likelihood value, -205.7 implies that the (geometric) mean predicted probability of actual decisions for all movies is 0.79, much better than chance guessing and a little better than simply guessing that all movies have a cold opening probability equal to the 7% (59/856) base rate (which yields a value of -211.62). Standard error estimates, determined by 100 bootstraps of this process, are shown in Table 3.10 and will be discussed later.

3.4.2 Cursed Equilibrium

Eyster and Rabin (2005) created a model of "cursed equilibrium" to explain stylized facts like the winner's curse in auctions, and other situations in which some agents do not seem to infer the private information of other players from those players' actions. Their idea is that such an incomplete inference is consistent with agents not appreciating the degree to which other players' actions are

iteration	final
λ_m	1.288
crit	0.016
thtr	0.001
bud	0.003
comp	0.002
star	-0.037
sum	0.036
sq/adpt	0.102
beffri	0.007
wkdlen	0.098
beffor	0.000
act/adv	-0.149
ani	-0.233
com	0.018
doc	0.196
fant/sci	-0.120
susp/hor	0.037
year	-0.063
pg	-0.142
pg13	-0.123
r	-0.169
R-squared	0.682
N	856
degrees of freedom	20
λ_d	1.345
log likelihood	-205.712
Mean $E_m(q\|Xj,1)$ for λ_d	15.106

Table 3.5: The iterative estimation process for the QRE model after 6 iterations

conditioned on information.

In our context, for every cold opened movie, all moviegoers believe that the movie has quality equal to some weighted average of the true expected movie quality (given distributor decisions) and the average of all movies. That is,

$$E_m^{ce}(q_j|1) \equiv (1-\chi_m)\, E_m^{re}\,(q|X_j,1) + \chi_m \bar{q} \tag{3.6}$$

where $E_m^{re}\,(q|X_j,1)$ reflects rational expectations about distributor decisions. We use an iterative procedure nearly identical to the one used above to find a best-fitting value of χ_m. The difference in the procedure is that $E_m^{ce}(q_j|1)$ and χ_m are used in forecasting revenue rather than $E_m^{qre}(q_j|1)$, and hence are also used in predicting distributor decisions.

If $\chi_m = 0$ this model is equivalent to QRE. The best-fitting value, based on first weekend cold box office data, is $\widehat{\chi_m} = .922$ however, which indicates a high degree of curse.

Since the estimated correct expectation $E_m^{re}\,(q|X_j,1)$ for cold opened movies is low ($\overline{E_m\,(q|1)}$= 25), and average overall quality is much higher ($\bar{q} = 48$), cursed moviegoers vastly overestimate the quality of movies that are opened cold. Since box-office revenues are increasing in quality, the fact that cursed moviegoers overestimate the quality of cold opened movies is consistent with the box office premium found in the basic regressions in Section 3.2. Indeed, the best-fitting cursed parameter estimate given the expectations found in the previous model is $\hat{\chi}_{m1} = .922$. This parameter predicts an average log box office premium on weekend box office of 0.33 (an increase in revenue of 33%). This value is considerably higher than 15% estimate determined from our initial regression — i.e., it appears that the model implies too *little* rationality of moviegoers, compared to the revenue effect from regression (see Supplemental Section B.2 for more detail).

Eyster and Rabin applied their model to experimental data from Forsythe et al. (1989) on agents "blind bidding" for objects of unknown value, after the producers of the objects have decided whether to reveal their values. In their estimation, all values of $\chi \in (0,1]$ fit better than the no-curse value $\chi = 0$, and the best-fitting $\chi = .8$. This number is similar to our estimate. Both estimates indicate high degrees of curse, with agents inferring quality of unknown goods at levels much greater than the fully correct expectation.

However, ignoring the distributor decision function, the result above only describes what best fits first weekend box office revenues. If cursed equilibrium requires the moviegoer curse parameter χ_m to be the same as the distributor's estimate of the curse (as the model intends), then iterating the procedures leads to $\chi_m = E_d(\chi_m) = 0$, which is equivalent to the QRE restriction in which there is no curse.[33]

The intuition is simple: Given an apparent curse of $\chi_m = .922$, distributors should be cold-opening a lot more movies of low quality than they actually are. Within the simple structure of this model, the only way to explain their anomalous behavior is that they do not believe moviegoers are as cursed as the box office revenue data suggest they are.

3.4.3 A Cognitive Hierarchy Model

Cognitive hierarchy models assume the population is composed of individuals that do different numbers of steps of iterative strategic thinking. The lowest level (0-level) thinkers behave heuristically (perhaps randomly) and k level thinkers optimize against $k-1$ type thinkers.[34] Zero-level thinkers, as moviegoers, do not think about the distributor's actions of cold opening a movie. For any cold-opened movie they infer the movie's quality $E^0_m(q_j|X_j, 1)$ at random[35] by selecting any integer on [0,100] with equal probability. They will go to any movie with probability defined as an analogue of equation (2)

$$p_0\left(X_j, E^0_m(q_j|c_j, X_j)\right) = \sum_{q=0}^{100}(1/101)\frac{1}{1+e^{-\lambda_m\left(\beta X_j+\alpha q-\hat{t}+\epsilon_j\right)}} \tag{3.7}$$

where $E^0_m(q_j|c_j, X_j) \sim U[0, 100]$. Similarly, a 0-level distributor will cold open movies at random, that is,

$$\pi_0(q_j, X_j) = 1/2. \tag{3.8}$$

A 1-level moviegoer knows 0-level distributors cold open movies at random, and assumes all distributors behave in this manner. For each movie he calculates the expected quality given it has been

[33] That is, the maximum likelihood parameter for distributor's decisions is $\chi_d = 0$ (i.e., distributors act as if moviegoers have no curse). If we only try to fit first weekend box office revenue, we have $\chi_m = 0.922$. If we must use the same value for both, the process will converge in the same way as QRE with $\chi_m = E_d(\chi_m) = 0$.

[34] This classification differs from some other version of the cognitive hierarchy model (Camerer et al., 2004) which suggests k level thinkers optimizes against a distribution of 0,1,...$k-1$ level thinkers.

[35] In many games, assuming that 0-level players choose randomly across possible strategies is a natural starting point. However, the more general interpretation is that 0-level players are simple, or heuristic, rather than random. For example, in "hide-and-seek" games a natural starting point is to choose a "focal" strategy (see Crawford and Iriberri (2007a)). In our game, random choice by moviegoers would mean random attendance at movies. That specification of 0-level play doesn't work well because it generates far too much box office revenue. Another candidate for 0-level moviegoer play is to assume a cold-opened movie has sample-mean quality $\bar{q}$. For technical reasons, that does not work well either. It is admittedly not ideal to have special *ad hoc* assumptions for different games. Eventually we hope there is some theory of 0-level play that maps the game structure and a concept of simplicity or heuristic behavior into 0-level specifications in a parsimonious way.

cold opened as

$$
\begin{aligned}
E_m^1\left(q|X_j,1\right) &= \frac{\sum_{q=0}^{100} qP\left(q\right)\pi_0\left(q,X_j\right)}{\sum_{q=0}^{100} P\left(q\right)\pi_0\left(q,X_j\right)} \\
&= \frac{\sum_{q=0}^{100} qP\left(q\right)\frac{1}{2}}{\sum_{q=0}^{100} P\left(q\right)\frac{1}{2}} \\
&= \bar{q}. \qquad (3.9)
\end{aligned}
$$

A 1-level distibutor expects all moviegoers to behave like 0-level moviegoers; they will assign quality ratings to cold-opened movies at random from the uniform $U\left[0,100\right]$ distribution. The 1-level distributor will therefore cold-open movie j with probability

$$
\pi_1(q_j,X_j) = 1/\left(1+\exp\left[\lambda_d(\sum_{q=0}^{100}(1/101)R(X_j,q)-R(X_j,q_j))\right]\right). \qquad (3.10)
$$

Proceeding inductively, for any strategic level k, the values $E_m^{k-1}(q|1,X_j)$ and $\pi_{k-1}(q_j,X_j)$ are computed from response to k-1 level type beliefs and actions. The k-level distributor and moviegoer have beliefs and

$$
\pi_k(q_j,X_j) = 1/\left(1+\exp\left[\lambda_m(R(X_j,E_{k-1}(q|X_j,1))-R(X_j,q_j))\right]\right) \qquad (3.11)
$$

and

$$
E_k(q|X_j,1) = \frac{\sum_{q=0}^{100} qP\left(q\right)\pi_{k-1}\left(q,X_j\right)}{\sum_{q=0}^{100} P\left(q\right)\pi_{k-1}\left(q,X_j\right)} \qquad (3.12)
$$

which leads to moviegoing probability

$$
p_k\left(X_j,E_m^k(q_j|c_j,X_j)\right) = \frac{1}{1+e^{-\lambda_m\left(\beta X_j+\alpha E_m^k(q_j|c_j,X_j)-\hat{t}+\epsilon_j\right)}} \qquad (3.13)
$$

where every level-k distributor and moviegoer is playing a quantal response to the level-k-1 moviegoer and distributor respectively.

The cognitive hierarchy model of Camerer et al. (2004), based on dozens of structurally different experimental games, suggests that the proportion of thinkers in the population is often well

k	$E_k(q\|Xj,1)$	$\pi_k(q_j,X_j)$
0	U[0,100]	0.50
1	48.12	1.00
2	40.79	1.00
3	34.28	1.00
4	29.40	1.00
5	24.66	0.94
6	20.97	0.10
7	17.40	0.01
8	14.73	0.00
9	12.20	0.00
10	10.26	0.00

Table 3.6: Expected quality of *When a Stranger Calls* $(q = 27)$ given it is cold opened by level-k moviegoer and probability it is cold opened by level-k distributor in CH with QR model $(\lambda_d = 7.085)$

approximated by a one-parameter Poisson distribution with mean τ,

$$P\,(x = n|\tau) = \tau^n e^{-\tau}/n!, \tag{3.14}$$

where τ is the average number of steps of strategic thinking (and also the variance).

As an example, Table 3.6 shows moviegoer-inferred quality and distributor probability of cold opening for the movie *When a Stranger Calls*, for various levels of thinking and their proportions within the population with $\lambda_d = 2.755$, (a figure estimated from the data, as estimated below) if the distribution of levels is Poisson distributed.

To determine QRE parameters $\{\lambda_d, \lambda_m\}$ and additional CH parameters $\{\tau_d, \tau_m\}$, we use an iterative procedure for estimating values similar to the QRE procedure. The procedure is much easier, however, because level-k player behavior is determined by level-k-1 behavior. The iteration is a "do loop" for specific on λ_m, λ_d values, naturally truncated when the percentage of high level-k players is very small (which depends on τ). Looping through for various λ_m, λ_d makes it easy to then grid-search over the λ values and find best-fitting values, values of both τ and λ.

Table 3.7 shows the results of the iterative process for the CH model with QR. The process stopped after six iterations with a log likelihood value of -166.232, which is a significant improvement over the QRE model. The value for λ_d (7.085) is also much greater than for the QRE (λ=1.345), which indicates less noise in the estimated decision process and suggests a better fit.

Note that the estimated value of $\hat{\tau}^*_m = 1.12$ is lower but is in the ballpark[36] of estimates from

[36]The objective function (sum of squared residuals) is rather flat in the vicinity of the best-fitting τ_m, so values from 2–4 give comparable fits to $\hat{\tau}^*_m = 1.12$. So an ex ante prediction based on $\tau = 1.5$ from lab data would forecast reasonably well in this field setting.

iteration	final
λ_m	1.302
crit	0.017
thtr	0.001
bud	0.003
comp	0.002
star	-0.039
sum	0.037
sq/adpt	0.108
beffri	0.000
wkdlen	0.095
beffor	0.000
act/adv	-0.155
ani	-0.254
com	0.016
doc	0.228
fant/sci	-0.154
susp/hor	-0.018
year	-0.066
pg	-0.139
pg13	-0.149
r	-0.177
R-squared	0.682
N	856
degrees of freedom	20
λ_d	7.085
τ_d	8.554
log likelihood	-166.232
τ_m	1.12
Mean $E_m(q\|Xj,1)$ for τ_m, λ_d	45.412

Table 3.7: The iterative estimation process for the QRE Model with CH after 5 iterations

model	parameter estimates	mean predicted cold-opened weekend box office (N=59)	predicted weekend cold opening premium	(Average SSR in \$)$^{1/2}$
QRE	λ_m= 1.288	6.73	-5.47%	6.00
		(0.04)	(0.31)	(0.09)
Cursed Equilibrium	χ_m=0.922, λ_m=1.288	9.46	33.32%	3.86
		(0.08)	(0.80)	(0.05)
Cognitive Hierarchy k-1	τ_m=1.12, λ_m=1.302	8.96	35.70%	5.64
		(0.12)	(1.20)	(0.06)
Random	E(q)~U[0,100]	11.44	63.56%	5.74
		(0.09)	(0.83)	(0.06)

Table 3.8: Comparison of the three models for moviegoer predictions with bootstrapped standard errors ($N = 100$)

experimental games (around 1.5) and for the initial week of Swedish LUPI lotteries (2.98, Östling et al., 2007) and managerial IT decisions (2.67, Goldfarb and Yang (2007)). The parameter estimates lead to an average cold opening box office premium of 35.7% (see Table 3.8, which like the estimate for the cursed model is much higher than the regression estimate.)

3.4.4 Comparing Distributor Estimation across Models

Table 3.10 provides standard error estimates from 100 random bootstraps of the data set for each parameter and each model. These bootstrapped samples are then used to give standard error estimates for comparative statistics between the three models and other benchmarks in Tables 3.8 and 3.9. Among other things, Table 3.10 indicates the cognitive hierarchy model with quantal response fits distributor decisions (in terms of log likelihood) far better than any other model.

One other thing to note is that few of the cursed model bootstraps did not result with a best fitting χ_d value of 0. This endpoint value explains that randomly sampling the dataset does not take away distributors' reluctance to cold open movies. However, random sampling may not greatly change the best-fitting moviegoer's curse parameter (χ_m). Standard errors indicate that the moviegoer's value is far more variable.

Table 3.8 compares best-fitting parameter values in sums of squared residuals (for moviegoer decisions). The non-equilibrium cursed model best predicts the box office revenues of cold opened movies in terms of deviations from actual data, and the cognitive hierarchy model is second best. This is not surprising since both models predict a box office premium. Even a prediction that

model	parameter estimates	log likelihood	mean correct (of 856)	standard deviation	no. predicted to open cold
QRE	λ_d= 1.345	-205.76	737.28	7.37	113.11
		(2.16)	(1.70)	(0.07)	(2.10)
Cursed Equilibrium	λ_d=1.345, χ=0	-205.76	737.28	7.37	113.11
		(2.16)	(1.70)	(0.07)	(2.10)
Cognitive Hierarchy k-1	l_d=2.755, τ_d=8.550	-166.23	772.22	6.10	61.73
		(2.37)	(0.83)	(0.03)	(1.09)
Base Rate	p=59/856	-214.73	740.55	7.41	59.00
		(2.21)	(1.47)	(0.05)	(0.85)

Table 3.9: Predictions of cold opening choices of distributors with bootstrapped standard errors $(N = 100)$

	QRE	Cursed Eq	CH with QR
	1.288	1.288	1.302
λ_m	(0.005)	(0.005)	(0.004)
	1.345	1.345	7.085
λ_d	(0.018)	(0.018)	(30.299)
	-	0.922	-
χ_m		(0.008)	
	-	0.000	-
χ		(0.001)	
	-	-	1.260
τ_m			(0.121)
	-	-	8.567
τ_d			(0.045)
	15.106	15.106	44.666
$E_m(q\|Xj,1)$	(0.315)	(0.315)	(0.587)
	-205.712	-205.712	-166.232
Log MLE	(2.134)	(2.134)	(2.126)

Table 3.10: Predictions of cold opening choices of distributors with bootstrapped standard errors $(N = 100)$

moviegoers assume uniformly random quality (all 0-level thinkers) to cold opened movies fits the data better than the QRE model which assumes correct expectations for cold opened quality.

For distributor decisions (Table 3.10) the cursed model and QRE models perform identically since the best-fitting cursed parameter is zero, so the rational expectations part of the QRE explains behavior pretty well if we must assume the system is in equilibrium. All models are also an improvement over the baseline case which predicts all movies to be cold-opened with the same probability.

The CH model improves on the predictions of the other two models. The key to its relative success is that the model estimates a low τ for moviegoers (close to experimental estimates of τ around 1.5–2.5) but the distributor τ_d is much higher. These parameters express the intuition that some moviegoers are easily fooled—they think cold openings are close to random—but distributors do not think moviegoers are so easily fooled, which is why so few movies are cold-opened. The CH model also predicts the most number of opening decisions correctly because its high τ_d predicts very few movies will be cold opened but the higher λ_d predicts some movies will be cold opened because of noise. The bootstrapped standard errors show these results are reasonably robust and do not depend on only a few data points. Importantly, the bootstrapped standard error around the mean bootstrapped estimate $\tau_m = 1.260$ is .121.

3.5 Conclusion

In games in which information can be considered good or bad news, and may be strategically disclosed or withheld at no cost, the only equilibrium involves the information receiver believing all withheld information conveys the worst possible signal, and the information sender choosing to reveal all information (except the worst). However, these equilibria require many steps of iterated strategic thinking.

Numerous laboratory experiments have shown in a variety of games that a small number of steps of strategic thinking tends to explain data well, as parameterized by quantal response equilibrium (QRE), cursed equilibrium, and cognitive hierarchy (CH) approaches. These models explain both experimental results that are far from equilibrium and other results that are surprisingly close to equilibrium, even in one-shot games. This paper is the first to apply all three parametrized models to a naturally occurring field phenomenon, an example of "structural behavioral economics."

We study a market in which information senders (movie distributors) are strategically withholding information (the quality of their movie) from information receivers (moviegoers). We find evidence that a Bayesian-Nash equilibrium has not been reached in our data since cold-opened movies do not have the worst possible quality values ($q_j = 0$) *and* there is a "box office premium" movies that have been cold opened earn more than other pre-screened movies of similar characteristics. This premium does not exist in foreign or video rental markets, where movies are released after the initial US release (so reviews are widely available), which suggests that the premium is due to some moviegoers failing to realize that no review is a bad signal about quality.

The QRE and cursed equilibrium models do not do a great job of explaining these facts. The QRE model performs poorly because moviegoers should correctly anticipate that cold opened movies are of low quality, which is inconsistent with the cold opening box office premium. The CH model with a low τ_m to represent moviegoer naïveté and a high τ_d to represent distributor over-sophistication can represent the mismatch of moviegoer perceptions and the reluctance (given moviegoer perceptions) of distributors to cold open.

The mismatch of parameter values for moviegoers and distributors suggest that either moviegoers should learn that cold opened movies are bad, or distributors should learn to cold open more movies. It *does* appear that distributors are learning. Figure 3.1 shows cold opening decisions by year. Near the end of this paper's dataset (January 1, 2006), distributors began to cold open movies with much higher frequency.[37] The models in this paper would suggest distributors have cold opened more movies as a best-response to moviegoer behavior.

One should wonder if consumers will eventually learn. In fact, experiments on lemons (e.g. Lynch et al., 2001) suggest consumers will ultimately infer the quality low-quality goods (when information is withheld). However, primary moviegoers are of age 14–18. By the time these moviegoers would realize cold openings are of lower quality than the average movie, they no longer are primary moviegoers, and another generation is already entering the prime moviegoing years. In this way, the model can be thought of as having overlapping generations of consumers.

There are hints of a difference between consumer and producer strategies in previous studies. With the benefit of this study's results, the somewhat contrary results of Mathios (2000) and Jin and Leslie (2003) make sense. Mathios found that mandatory disclosure was necessary for consumers to buy low-fat salad dressing (i.e., with voluntary disclosure they did not infer what the absence of a "low-fat" label meant). Jin and Leslie found suggestive evidence that voluntary disclosure of health

[37] Through 2000–2005 distributors cold opened around 5–8% of widely released movies. In 2006 and 2007 distributors cold opened 19% (30/160) and 23% (30/131).

ratings improved health quality nearly as much as mandatory disclosure, but the data on revenue for non-disclosing restaurants was inconclusive. Given that one choice is a consumer decision (buying salad dressing) and one is a producer decision (improving health quality in a restaurant), we may have more evidence of a different level of iterative thinking by diverse consumers and expert producers. Of course, this is only speculation and more investigation is needed before any conclusions can be drawn from these types of investigations. As noted earlier, there are many markets with asymmetric information in which the failure to reveal information that is often revealed should be informative—if the receiver makes the proper strategic inference. Our approach and some of its technical details could be applied to these markets and to other markets with this property.

Chapter 4

Endogenous Time Preference and Personal Rules

4.1 Introduction

Although Adam Smith (1759/1892) introduced the concept of "self-command" into political economy nearly 250 years ago, economists did not treat internal conflict systematically until Strotz (1955–1956) analyzed Odysseus's Sirenic dilemma. All economists who recognized intrapersonal conflict after Strotz, agreed with Homer's intuition on two important points: (i) internal conflict is between a side representing present drives and another representing long-term interests, (ii) if given the opportunity, an individual may choose to limit his future actions, like Odysseus tying himself to the mast of his ship, to prevent future impulses from hampering long-term goals[1] (c.f., Thaler and Shefrin, 1981; Laibson, 1997; Loewenstein and O'Donoghue, 2004; Gul and Pesendorfer, 2001; Fudenberg and Levine, 2006).

Accepting the idea of internal conflict, economists are left with another problem: given no external commitment device and a repeated succession of *identical* choices between myopic and responsible alternatives, an individual may take the tempting alternative in each period, even if, in any period, he would choose to take all responsible alternatives over all tempting alternatives given the *aggregate* choice, forgoing that period's temptation.[2] For example, suppose an individual has a daily vice (e.g., spend an additional 0.1% of yearly income, consume an additional serving of a junk food, smoke an extra cigarette) and faces immediate temptation to do so everyday, and over the course of ten years, the daily vice leads to ruin (e.g. debt, diabetes/heart conditions, respiratory

[1] See chapter 2 for an example of experimental results consistent with intrapersonal conflict.

[2] Psychological literature uses a non-economic model, called "melioration," to model the choices between two alternatives under many short-term decisions. See Herrnstein and Prelec (1991) for a survey.

problems, respectively). However, if each day he prefers the vice, and his choice has no impact on future behavior, then each day he will choose the vice regardless of his feelings about the aggregate choice.[3] But if the responsible aggregate choice is preferred on each day while the daily choice is taken, then a pareto-efficient (to each period's self) outcome is not being realized in intrapersonal conflict.

"Personal rules" may remedy this predicament. Ainslie (1992) suggested that personal rules may serve as an internal commitment device to enforce an aggregate decision. Today's self chooses the responsible alternative because if he does not, tomorrow's self (who was going to also choose the responsible alternative) will retaliate and choose the tempting alternative. Much the way agents playing the repeated prisoner's dilemma may make rules (to cooperate for a limited period of time), Ainslie (1992) believed a present self would make an agreement with a future self. This has lead many psychologists to model the problem of self-control with the repeated prisoner's dilemma. (e.g., Yi and Rachlin, 2004)

While experimental evidence agrees with the end result—Kirby and Guastello (2001) showed individuals are less likely to take the myopic reward when choices are repeated compared to a single decision—even Ainslie admitted the repeated prisoner's dilemma is an imperfect comparison for making a multitude of repeated decisions under temptation. In intrapersonal conflict, if any future self $(t + k)$ defects and chooses the tempting alternative, the present self cannot retaliate (t), unlike in the repeated prisoner's dilemma. Thus, free from retaliation, future selves will likely choose temptation regardless of previous history. Knowing that his actions have no effect on future selves, the present self chooses temptation as well. Thus, the grim-trigger rule developed by Ainslie is not compatible with backward induction.[4]

This chapter models personal rules in a manner compatible with game theory. By making one more assumption, that committing (forgoing) an action makes it more (less) tempting in the future, this chapter shows that the "precedent effect" of Kirby and Guastello (2001) can be obtained with backward induction. In this manner personal rules are sequentially rational (see Kreps and Wilson, 1982) and the subgame perfect equilibrium to an extensive-form game.

The findings and theories of Baumeister and colleagues (Baumeister et al., 1994; Baumeister and Vohs, 2003; Muraven et al., 1999; Muraven and Baumeister, 2000) provide the main evidence

[3] In a way, the situation is similar to each period's self failing to contribute to a public good.

[4] The results of the repeated prisoner's dilemma are not compatible with backward induction either. In that game two players may make an agreement to cooperate with a finite end, but often they break that agreement shortly before the last period. (Selten and Stoecker, 1986; Camerer and Weigelt, 1988)

that temptation or immediate utility is endogenous to personal choice. Their findings suggest that willpower can be built up with frequent exercise and then rest, like a muscle. As it relates to this chapter, an individual who uses willpower to avoid a tempting alternative will be stronger, after rest, to all temptations (and thus that specific alternative). However, since this chapter involves temptation from only one choice, one could think of this as an assumption of habit formation[5] where the habit affects immediacy preference. However one thinks about the assumption,[6] the simple notion that people can develop "good" and "bad" habits and find them difficult to break is quite intuitive. This intuition and the belief that the habit acts on immediate impulses is roughly all it takes to accept the main assumption of this chapter.

Like other works involving personal rules (e.g., Ainslie, 1992; Benabou and Tirole, 2004) this chapter concerns a binary choice between two alternatives: one that is tempting and yields an immediate benefit, and one that is better long-term but yields no immediate benefit. Those two models differed on the description of personal rule, Ainslie believed rules were "perscriptive," they could effect agent's behavior, while Benabou and Tirole only believed they were "descriptive" and did not affect behavior.

The model shows as a result of backward induction, that at certain levels of immediacy-preference (or self-control) agents will forgo the myopic alternative if the choice is repeated, but will consume it if the choice will only be made once, a result that is consistent with the hypthesis of Ainslie (1992) and experimental evidence (Kirby and Guastello, 2001). However, an agent may defect in any given period, exhibiting the "procrastination effect" knowing that in the future he will exercise self-control. One remedy for this procrastination is for an agent to institute a personal rule, if we think of beliefs about future actions as personal rules, then there always exists a consistent personal rule that can eliminate this procrastination effect. If we restrict are rules to the descriptive type and thus only subgame perfect beliefs, we only find that agents may persevere under high costs. Regardless, when we expand our model to an infinite horizon we find that pure-strategy decision may not exist and no personal rule can be consistent.

[5]Economists have a long history with habit formation (see Duesenberry, 1949; Becker and Murphy, 1988). Of course, these models do not involve dynamic inconsistency.

[6]There are other psychological explanations for this habit component: it is possible that by avoiding a choice an individual makes himself less susceptible to cues concerning that choice and no longer feels the same temptation. Conversely giving into those cues make the choice stronger (see Laibson, 2001, for a cue-based model of addiction). It also could be by repeatedly avoiding (or succumbing) to a temptation an individual makes the action of avoiding (or succumbing) more of an automated process (see Schneider and Shiffrin, 1977; Shiffrin and Schneider, 1977, for an example of frequent practice making a difficult action automatic). Any of these three reasons would be enough to justify the main assumption of the chapter.

This chapter contributes to the literature on time preference, willpower, and personal rules by unifying all three. It incorporates psychological ideas from Ainslie (1992) and Baumeister et al. (1994), and the respective economic models of Benabou and Tirole (2004) and Ozdenoren et al. (2006). While similar, both sides have rarely if ever been united in an academic work. The chapter also provides an alternative to self signaling—endogenous immediacy preferece—to explain the reasons behind personal rules and precedent effects.

The chapter proceeds as follows: Section 4.2 presents an example, the model, and finite period solutions for both certain and stochastic cases. Section 4.3 provides an extensive form characterization of the model, and proves the existence of a solution for all periods. Section 4.4 develops a characterization for personal rules, and shows their effectiveness. Section 4.5 studies the problem in the infinite case. Section 4.6 and 4.7 provide testable implications of the model and conclude, respectively.

4.1.1 Literature Review

To model intrapersonal conflict, this chapter assumes that individuals have inconsistent preferences over time. Economic models have represented this dynamic discord literally, with time-discounting (Ainslie, 1975, 1992; Laibson, 1997), or metaphorically, with long-term and short-term selves (Thaler and Shefrin, 1981; Fudenberg and Levine, 2006) and multiple preferences (Gul and Pesendorfer, 2001). This chapter will use the quasi-hyperbolic model (Laibson, 1997) to represent an agent's time preferences in this chapter, because it is widely used, allows making a comparison to Benabou and Tirole (2004), and has parameters that have been estimated in experiments. That model features two discount terms, β and δ, to represent dynamic inconsistency. Agents discount all periods but the present once with the factor β and again with an factor, δ, which increases exponentially over future periods. The model expands on Laibson's framework by allowing β to dynamically change each period as individuals make decisions.

This is not the first model to suggest that time preference dynamically changes with agents' decisions. Becker and Mulligan (1997) created a model where time preference could be built up by increasing wealth. But that model only featured an exponential discount term—it did not represent intrapersonal conflict. Because nearly all intrapersonal conflicts manifest themselves in time preference, we justify the dynamic β term, an immediate preference, by using it as a representation of willpower. The idea of endogenous choice is also not new to economics, and is most commonly seen in papers on habit formation. (e.g., Duesenberry, 1949; Becker and Murphy, 1988).

Most of the considerable research done on willpower has treated it as a resource that behaves much like a muscle. Experiments showed that willpower is like a muscle because it can (i) be worn down with constant use over a period of time, (ii) be recharged after rest, and (iii) be built up after periodic use and rest. Baumeister and Vohs (2003) surveyed recent work on willpower and concluded that willpower can be depleted when individuals perform several tasks requiring willpower on the same day. They also found that individuals will ration willpower knowing that it will be depleted when they are presented with consecutive tasks that require willpower. Other studies suggest repeated use of willpower over a series of days will increase its capacity (Muraven et al., 1999; Muraven and Baumeister, 2000). That is, if an individual repeatedly performs a task requiring willpower, and his willpower is given sufficient time to recharge, he will become better at his conservation of willpower over time. Ozdenoren et al. (2006) provided an economic model of these results, a model which we use as a basis for changes in β. Unlike our model, their model had agents who made decisions in continuous time and did not exhibit dynamic inconsistencies. Perhaps because its agents did have dynamic conflict, their model did not feature any personal rules.

Ainslie (1975), drawing upon the ideas of Becker (1960) and Schelling (1960), first suggested that personal rules could solve dynamic conflicts. The first papers dedicated to personal rules (Schelling, 1985; Prelec and Herrnstein, 1991) did not provide a formal or single explanation of why personal rules were effective. Later, Ainslie (1992) argued that personal rules were an agreement that an individual would make with his future selves. Normally, the agreement would be to forsake present rewards in favor of long-term goals. If an individual violated the agreement, he would lose all benefits from it, because his future selves would not follow it again. Unfortunately, it is not always the case that a future self will credibly follow or break any rule. To make a personal rule "self-sustaining," Benabou and Tirole (2004) defined personal rules as solutions to a perfect Bayesian equilibrium; if an agent defects, he will not be maximizing utility. Expanding on Bodner and Prelec's (2003) concept of self-signaling, Benabou and Tirole argued that agents, who either have high or low willpower, signal information about their type to their future selves who have imperfect memory. Personal rules are a description of what each type does in states with high and low costs of perseverance.

This chapter uses the same binary decision of Benabou and Tirole's model. And both papers view β as a representation of willpower, and feature rules that describe an action for a given type and state. The similarities end there. This chapter features dynamic immediacy preference, while their paper assumes that an individual's willpower type is constant. In this chapter, low willpower types persevere against immediate impulses to become high willpower types in the next period. In their paper, low willpower types persevere against immediate impulses to deceive their future selves into believing they have high willpower. Additionally, this chapter models the type of "prescriptive" rules that Ainslie (1975, 1992) suggested in his writings.

4.2 The Model

Each morning an agent chooses what to eat for breakfast. He can choose the healthy option (e.g., a grapefruit) or the sweet, tempting option (e.g., a donut). The donut and grapefruit both give the agent utility (d and g respectively), but they are of different types.[7] The sweetness of the unhealthy donut gives the agent immediate utility, but because it is unhealthier it has less total utility $(g > d)$. The agent has quasi-hyperbolic preferences (Laibson, 1997)—he discounts the future with two seperate terms, β and δ. He discounts all non-immediate benefits by a factor $0 < \beta < 1$ and discounts all future periods by a standard exponential term $\delta < 0$. That is, an agent has a traditional utility function,

$$u(a) = \begin{cases} g & \text{if } a = \text{grapefruit (G)}, \\ d & \text{if } a = \text{donut (D)}, \end{cases} \tag{4.1}$$

and present-preferring utility function,

$$v(a, \beta) = \begin{cases} g & \text{if } a = \text{grapefruit (G)}, \\ \frac{d}{\beta} & \text{if } a = \text{donut (D)}, \end{cases} \tag{4.2}$$

where it is assumed $d, g > 0$ (food is better than no food). Therefore the agent will chose to eat the donut if $\beta < d/g$ and choose the grapefruit if $\beta \geq d/g$. However, from the point of view of the past periods, the agent would always prefer to eat the grapefruit.[8] With static β, agents can be divided into two groups, those with high willpower $(\beta_i \geq d/g)$ who will always choose the grapefruit, and those with low willpower $(\beta_i < d/g)$ who will always choose the donut.[9]

The situation in this chapter is more complicated. When facing repeated decisions, the choice

[7] Grapefruit and donut are just examples. If an agent were given *only* two choices for breakfast, he might make decisions for the sake of variety. In a more generalized form this is a decision between a healthy and tempting, but unhealthy breakfast.

[8] In other words, in any period our agent would seek external commitment (if available) to prevent himself from eating donuts in the future. See Section A.3 for an in-depth description of the quasi-hyperbolic model.

[9] This is the characterization of Benabou and Tirole (2004).

an agent makes in the present period will alter their future ability to make that decision. If the agent chooses to eat the donut in the current period, it becomes harder for him to resist impulses to eat the donut in the future. Conversely, avoiding the donut and eating the grapefruit in the present makes it easier to avoid the donut and eat the grapefruit in the future. This change will be modeled as an increase or decrease in β, the term that represents immediate preference.[10] This interpretation draws on psychological literature which suggests that willpower, usually defined as control over impulses, can be built up like muscle through repeated use[11] (Baumeister and Vohs, 2003; Muraven et al., 1999; Muraven and Baumeister, 2000).

Therefore, in this example, the agent has a more complicated decision to make with multiple periods and a dynamically changing β. He could choose to eat the donut, but that may make his future self more likely to eat the donut,[12] something he prefers not to have happen in the future. The model of the chapter will capture this decision.

4.2.1 The Finite Horizon Case

To formalize our previous example, consider the choice of eating the donut to be choosing D, and eating the grapefruit to be choosing G. That decision will be made over $T \in \mathbb{N}$, $T \geq 1$ periods. Equation 4.1 gives the utility function in the last period $(t = T)$ from the point of view of previous periods.

In period t the agent has β_t, a variable which represents his immediate preference. Because an agent has quasi-hyperbolic preferences, he will discount future utility by a standard discount factor δ and the immediate preference factor. In period t he will discount the utility of period $t+k$ by $\beta_t \delta^k$. As the example indicates, only d is experienced by the agent immediately, so all utility other than the experience of the donut (d) will be discounted once by β.[13] The function $v : A \times (0, 1) \to \mathbb{R}$ represents the agent's preferences under this preference. Equation 4.2 shows this function in the last period $(t = T)$.

Since no external commitment devices are available to the agent in period $T - 1$, the present-

[10]This interpretation differs from a traditional habit formation model (Becker and Murphy, 1988), as the parameters in the utility function (i.e., d and g do not change).

[11]Remember, because we are concerned only with the temptation increasing or decreasing for *one* activity, it is not necessarily the case that willpower needs to be reduced or increased for this assumption to hold. It could be instead that the action becomes endogenously more (or less) tempting with consumption (avoidance) because of cue-based (Laibson, 2001) or automatic processes (Schneider and Shiffrin, 1977; Shiffrin and Schneider, 1977).

[12]It is not always the case that eating the donut in period t makes one eat cake in $t + 1$. See Example 4.2 for the "procrastination effect."

[13]Equivalently, we often inflate d by β.

preferring utility function, $v(a, \beta)$ given in equation 4.2 will determine the period's action. That is, with only one period, the agent will choose the donut if $\beta_T < \frac{d}{g}$ and the grapefruit if $\beta_T \geq \frac{d}{g}$ (we are giving ties in preference to the grapefruit). We will define one more important term, $\hat{\beta}_1$, the threshold for choosing grapefruit over donut in the last period

$$\hat{\beta}_1 \equiv \frac{d}{g}. \tag{4.3}$$

The new addition to this model is that immediate preference and its numerical representation, β_t, will change with each decision.[14] To put it another way, the past matters, or the choices are "path dependent." The preference will increase in the next period when an agent chooses the tempting alternative (i.e., donut). It will decrease when one chooses the wholesome alternative (i.e., grapefruit). The value will represent the current "state." If β_{t+1} represents the next period's state given the initial state, β_t, and current action a_t then

$$\beta_{t+1}(\beta_t, a_t) = \begin{cases} \Gamma(\beta_t) & \text{if } a_t = \text{G} \\ \Delta(\beta_t) & \text{if } a_t = \text{D}, \end{cases} \tag{4.4}$$

where $\Delta(\beta) < \beta < \Gamma(\beta)$.

For notational ease, we will denote $z \in \mathbb{Z}$ iterations of Δ on β as $\Delta_z(\beta)$. The function Γ will follow the identical convention. For example, $\Gamma_6(\beta) \equiv \Gamma(\Gamma(\Gamma(\Gamma(\Gamma(\Gamma(\beta))))))$, and $\Delta_{-2} = \Delta^{-1}(\Delta^{-1}(\beta))$. We will put further restrictions on $\Delta(\beta)$ and $\Gamma(\beta)$.

Assumption 4.1 (convergence)**.** *For all $\beta \in \left(\underline{\beta}, \overline{\beta}\right)$, the transformation functions will converge, $\lim_{n\to\infty} \Delta_n(\beta) = \underline{\beta}$ and $\lim_{n\to\infty} \Delta_n(\beta) = \overline{\beta}$.*

Assumption 4.2 (monotonicity)**.** *Functions Γ and Δ are monotonic. Specifically, $\forall \beta'$, β'', $\beta' > \beta''$ implies $\Gamma(\beta') > \Gamma(\beta'')$ and $\Delta(\beta') > \Delta(\beta'')$. That is, Γ and Δ are increasing.*

Assumption 4.1 ensures that $\beta_t \in (\underline{\beta}, \overline{\beta})$ for any transformations undertaken in the model. Assumption 4.2 ensures that states do not jump one another—if an individual has a higher immediacy preference than another to begin a period, and they both take the same action, then that individual will still have the higher immediacy preference.

To solve this model when $T \geq 2$, an important distinction needs to be made. An agent who has quasi-hyperbolic preferences may be either be naïve or sophisticated. A "naïve" agent does

[14] Note that as immediacy preference increases, β decreases, and vice versa.

not see his future shift in preferences and believes he will make decisions in the future with his present period preferences. A "sophisticated" agent foresees his future shift in preferences and acts accordingly. Often the solutions between the naïve and sophisticated agents are very different. (O'Donoghue and Rabin, 1999) Both models will be considered in this chapter, but as lemma 4.1 will show, the naïve model has a very simple solution.

It is helpful to define the concept of solution to this dynamic programming problem in the "naïve" and "sophisticated" cases. These solutions can be thought as either dynamic optimization (see Section A.3 and/or Harris and Laibson, 2001) for T periods with inconsistent preferences or a subgame perfect equilibrium to an extensive game with T players and periods. Both will yield the same solution (see Section 4.3 for extensive form representations); the difference is in terminology. We will use terms consistent with dynamic programming, with an action function a_t^* and value function V_t^*. Alternatively, a_t^* could be thought of as player t's best response to $t+1,\ldots,T$.

In the $t = T$ case we have already shown that for either a naïve or sophisticated agent, we will have

$$a_T^*(\beta) = \begin{cases} G & \text{if } \beta \geq \hat{\beta}_1 \\ D & \text{if } \beta < \hat{\beta}_1 \end{cases} \tag{4.5}$$

and value functions V_T^* can be defined.

However at time $T-1$ a naïve hyperbolic does not believe he will have hyperbolic preferences in the next period. Since $g > d$ by assumption, he believes in the future he will always take the healthy option. His future value function can be thought of as

$$\breve{V}_T^* = g. \tag{4.6}$$

We can classify the naïve hyperbolics' decisions for all t.

Lemma 4.1 (naïve solution)**.** *For all T, for every $t \leq T$. The naïve hyperbolic will choose the healthy option if $\beta \geq \hat{\beta}_t$ and will choose temptation if $\beta < \hat{\beta}_t$. If $\delta < 1$, this will occur at $T = \infty$ as well.*

Proof. At T the naive hyperbolic prefers D if $d/\beta \geq g \rightarrow \beta \geq \frac{d}{g} = \hat{\beta}_1$. Then at $T-1$, the naïve hyperbolic believes $\breve{V}_T = g$. Since at any β, $V_T(\Gamma(\beta)) = V_T(\Delta(\beta))$, we have $\breve{a}_{T-1}(\beta) = G$ if $\beta \geq \hat{\beta}_1$, and $\breve{a}_{T-1}(\beta) = D$ if $\beta < \hat{\beta}_1$.

Suppose now, for some $1 \leq n < T$, $\breve{V}_{T-n} = \sum_{i=T-n}^{T} \delta^{i-T+n} g$ and $\breve{a}_{T-n}(\beta) = G$ if $\beta \geq \hat{\beta}_1$, and $\breve{a}_{T-n}(\beta) = D$ if $\beta < \hat{\beta}_1$. At $T-n-1$, the naïve agent believes a β term will not influence him at $T-n$, since $g > d$, this is a belief that at $T-n$ he will always choose G. Then $\breve{V}_{T-n-1}(\beta) = g + \delta\breve{V}_{T-n}(\beta) = \sum_{i=T-n-1}^{T} \delta^{i-T+n+1} g$, $\forall\beta$. Since $\forall\beta$, $V_{T-n}(\Gamma(\beta)) = V_{T-n}(\Delta(\beta))$, we have $\breve{a}_{T-n-1}(\beta) = G$ if $\beta \geq \hat{\beta}_1$ and $\breve{a}_{T-n-1}(\beta) = D$ if $\beta < \hat{\beta}_1$. This completes the inductive proof.

It follows that as $n < T$, $n \to \infty$, $V_{T-n} = \frac{g}{1-\delta}$, $\forall\beta$. Then the condition $V_{T-n}(\Gamma(\beta)) = V_{T-n}(\Delta(\beta))$ still holds. Thus $n \to \infty$, $\breve{a}_{T-n}(\beta) = G$ if $\beta \geq \hat{\beta}_1$ and $\breve{a}_{T-n}(\beta) = D$ if $\beta < \hat{\beta}_1$. $\square$

The rest of the chapter will be concerned with the sophisticated hyperbolics. Their solution is much more interesting. We begin by defining the sophisticated solution for any t periods.

Definition 4.1 (sophisticated solution)**.** For T periods $\{a_t^*, V_t^*\}$ is a sequence of actions and value functions $(a_t^*, V_t^*), \ldots, (a_T^*, V_T^*)$ where $a_t^* : (\underline{\beta}, \overline{\beta}) \to \{G, D\}$ and $V_t^* : (\underline{\beta}, \overline{\beta}) \to \mathbb{R}$. For $a_T^*(\beta_T)$ defined in equation 4.5, $V_T^* = u(a_T^*(\beta_T))$, and $\forall t < T$,

$$a_t^*(\beta) = \begin{cases} G & \text{if } v(G,\beta) + \delta V_{t+1}^*(\Gamma(\beta)) \geq v(D,\beta) + \delta V_{t+1}^*(\Delta(\beta)) \\ D & \text{if } v(G,\beta) + \delta V_{t+1}^*(\Gamma(\beta)) < v(D,\beta_t) + \delta V_{t+1}^*(\Delta(\beta)) \end{cases} \tag{4.7}$$

$$V_t^*(\beta) = \begin{cases} u(G) + \delta V_{t+1}^*(\Gamma(\beta)) & \text{if } a_t^*(\beta) = G \\ u(D) + \delta V_{t+1}^*(\Delta(\beta)) & \text{if } a_t^*(\beta) = D. \end{cases} \tag{4.8}$$

The logic of backward induction is involved in this definition. After (a_T^*, V_T^*) is defined, we cannot define (a_t^*, V_t^*) unless we know (a_{t+1}^*, V_{t+1}^*). The solution to these problems, like the cases without quasi-hyperbolic preferences, is solved by starting at the last period first and working back to the current period. In the sophisticated case, it follows from definition 4.1 and equation 4.5 that the solution for the Tth period is

$$V_T^*(\beta) = \begin{cases} g & \text{if } \beta \geq \hat{\beta}_1, \\ d & \text{if } \beta < \hat{\beta}_1. \end{cases} \tag{4.9}$$

Then for all $1 \leq t \leq T-1$, we have a solution if and only if $\forall\beta \in (\underline{\beta}, \overline{\beta})$

$$a_t^*(\beta) = \begin{cases} G & \text{if } \beta \geq d \big/ \left(g + \delta[V_{t+1}^*(\Gamma(\beta)) - V_{t+1}^*(\Delta(\beta))]\right), \\ D & \text{if } \beta < d \big/ \left(g + \delta[V_{t+1}^*(\Gamma(\beta)) - V_{t+1}^*(\Delta(\beta))]\right). \end{cases} \tag{4.10}$$

To explain the solutions in the finite period case, we need to define a few more terms. The term $\hat{\beta}_{T-t+1}$ is the lowest value of β such that an agent would prefer to choose the healthy alternative for all remaining periods $(T-t+1)$ rather than choose the unhealthy alternative for all the remaining periods, and the decision is feasible to do so. Alternatively, $\beta < \hat{\beta}_{T-t+1}$ implies the tempting alternative is preferred, because either β is so low the immediate temptation outweighs all future benefits, or choosing the wholesome alternative will still lead to temptation chosen in the next period (because $\Gamma(\beta) < \hat{\beta}_{T-t}$).

Definition 4.2 (threshold β)**.** For all $s \in \mathbb{N}$, $s > 1$, let

$$\hat{\beta}_s = \max\left\{\frac{d}{g + \sum_{i=1}^{s-1} \delta^i (g-d)}, \Gamma^{-1}\left(\hat{\beta}_{s-1}\right)\right\} \tag{4.11}$$

and

$$\hat{\beta}_\infty = \lim_{s\to\infty} = \frac{d}{g + \frac{\delta}{1-\delta}(g-d)}. \tag{4.12}$$

It follows that $\forall s > 1$, $\hat{\beta}_s \leq \frac{d}{g} = \hat{\beta}_1 < \overline{\beta}$. It turns out these threshold values $(\hat{\beta}_1, \hat{\beta}_{T-t+1})$ will be essential in describing any $t < T$ period solution. Proposition 4.1 shows that for any period t the actions and agent takes in with state β are clearly defined above and below these values.

Lemma 4.2. *For all T, t, $2 \leq t < T$.*

(i) If the solution for the value function for t is $V_t^*(\beta) = \sum_{i=t}^T \delta^{i-t} d$, $\forall \beta < \hat{\beta}_{T-t+1}$, then $a_{t-1}^*(\beta) = D$, $\forall \beta < \hat{\beta}_{T-t+2}$.

(ii) If the solution for the value function for t is $V_t^*(\beta) = \sum_{i=t}^T \delta^{i-t} g$, $\forall \beta \geq \hat{\beta}_1$, then $a_{t-1}^*(\beta) = G$, $\forall \beta \geq \hat{\beta}_1$.

Proof. For any $2 \leq t < T$, either $\hat{\beta}_{T-t+2} = \Gamma^{-1}(\hat{\beta}_{T-t+1})$ or $\hat{\beta}_{T-t+2} = \frac{d}{g+\sum_{i=1}^{T-t+1}\delta^i(g-d)}$. Suppose $\hat{\beta}_{T-t+2} = \Gamma^{-1}(\hat{\beta}_{T-t+1})$, then $\forall \beta$, $\beta < \hat{\beta}_{T-t+2}$, $\Delta(\beta) < \Gamma(\beta) < \Gamma(\hat{\beta}_{T-t+2}) \leq \hat{\beta}_{T-t+1}$. So $V_t^*(\Delta(\beta)) = \sum_{i=t}^T \delta^{i-t} d = V_t^*(\Gamma(\beta))$, and $a_{t-1}^*(\beta) = G$. If instead $\hat{\beta}_{T-t+2} = \frac{d}{g+\sum_{i=1}^{T-t+1}\delta^i(g-d)} = \frac{d}{g+\sum_{i=t-1}^{T}\delta^{i-t+1}(g-d)}$, then $\beta < \hat{\beta}_{T-t+2}$ implies $\frac{d}{\beta} + \sum_{i=t-1}^T \delta^i d > g + \sum_{i=t-1}^T \delta^i g \rightarrow \frac{d}{\beta} + \delta V_t(\Delta(\beta)) > g + \delta V_t(\Gamma(\beta))$, thus $a_{t-1}^*(\beta) = D$.

(ii) If $\beta \geq \hat{\beta}_1$, then $v(G,\beta) > v(D,\beta)$. Since $V_t^*(\Gamma(\beta)) \geq V_t^*(\Delta(\beta))$, by equation 4.7 we must have $a_{t-1}^*(\beta) = G$. $\square$

Proposition 4.1. *For all t, if $\beta \geq \hat{\beta}_1$ then $a_t^*(\beta) = G$. If $\beta < \hat{\beta}_{T-t+1}$, then $a_t^*(\beta) = D$.*

Proof. At $T-1$, equation 4.9 and lemma 4.2 shows us that $a^*_{T-1}(\beta) = G$, $\forall \beta \geq \hat{\beta}_1$, and $a^*_{T-1}(\beta) = D$, $\forall \beta < \hat{\beta}_1$. It follows that $V^*_{T-1}(\beta) = g + \delta g$, $\forall \beta \geq \hat{\beta}_1$ and $V^*(\beta)_{T-1} = d + \delta d$, $\forall \beta < \hat{\beta}_1$. Suppose for some n, $V^*_{T-n+1}(\beta) = \sum_{i=T-n+1}^{T} \delta^{i-T+n-1} d$, $\forall \beta < \hat{\beta}_{T-t+1}$ and $V^*_{T-n+1}(\beta) = \sum_{i=T-n+1}^{T} \delta^{i-T+n-1} g$, $\forall \beta > \hat{\beta}_1$, and $a^*_{T-n+1}(\beta) = D$, $\forall \beta < \hat{\beta}_{T-t+1}$ and $a^*_{T-n+1}(\beta) = G$, $\forall \beta > \hat{\beta}_1$. Then $V^*_{T-n}(\beta) = \sum_{i=T-n+1}^{T} \delta^{i-T+n-1} d$, $\forall \beta < \hat{\beta}_{T-t+1}$ and $V^*_{T-n}(\beta) = \sum_{i=T-n+1}^{T} \delta^{i-T+n-1} g$, $\forall \beta > \hat{\beta}_1$. It follows from lemma 4.2 that $a^*_{T-n}(\beta) = D$, $\forall \beta < \hat{\beta}_{T-t+1}$ and $a^*_{T-n}(\beta) = G, \forall \beta > \hat{\beta}_1$. □

Example 4.1 shows how over multiple periods, for states $\hat{\beta} < \hat{\beta}_1$ we still can have the agent choose the healthy alternative. The future reward for developing a positive habit and promise of G in the future incentivizes the non-myopic choice.

Example 4.1 (precedent effect)**.** *Suppose $T \geq 2$. The $t = 2$ solutions are given by equations 4.5 and 4.9. Then for $t = T-1$, if $\hat{\beta}_2 > \underline{\beta}$ we have*

$$a^*_{T-1}(\beta) = \begin{cases} G & \text{if } \beta \geq \hat{\beta}_2 \\ D & \text{if } \beta < \hat{\beta}_2 \end{cases} \tag{4.13}$$

$$V^*_{T-1}(\beta) = \begin{cases} g + \delta g & \text{if } \beta \geq \hat{\beta}_2 \\ d + \delta d & \text{if } \beta < \hat{\beta}_2 \end{cases} \tag{4.14}$$

if $\hat{\beta}_2 < \underline{\beta}$, $a^*_{T-1}(\beta) = G$ and $V^*_{T-1}(\beta) = g + \delta g$.

Proof. Consider $\hat{\beta}_2 \geq \underline{\beta}$. Then proposition 4.1 defines $a^*_{T-1}(\beta)$ on $(\underline{\beta}, \hat{\beta}_2) \cup [\hat{\beta}_1, \overline{\beta})$. If $\beta \in (\hat{\beta}_2, \hat{\beta}_1)$, $\Gamma(\beta) > \Gamma(\hat{\beta}_2) > \hat{\beta}_1$ and $\Delta(\beta) < \hat{\beta}_1$, then $V_T(\Gamma(\beta)) = g > d = V_T(\Delta(\beta))$. But $\beta > \hat{\beta}_2$ implies $\frac{d}{\beta} + \delta V_T(\Delta(\beta)) < g + \delta V_T(\Gamma(\beta))$, so we must have $a^*_{T-1}(\beta) = G$. The value function, V^*_{T-1}, is defined from this process.

If $\hat{\beta}_2 < \underline{\beta}$, proposition 4.1 defines $a^*_{T-1}(\beta)$ on $[\hat{\beta}_1, \overline{\beta})$. For $\hat{\beta}_1 > \beta > \underline{\beta} > \hat{\beta}_2$, a similar argument to above shows $a^*_{T-1}(\beta) = G$. The value function, V^*_{T-1}, is defined from this process. □

In experiments, Kirby and Guastello (2001) were able to show that individuals were more likely to move away from immediacy preference (choose the non immediate, but greater quantity award) more often in instances when they were told the *identical* choice was going to be repeated in the future. They called this observation the “precedent effect.” Our model is consistent with this observation. Example 4.1 at shows this may happen even if the choice is going to be repeated only one more time. Notice that $\hat{\beta}_2 < \hat{\beta}_1$, so agents at state levels $\hat{\beta}_2 < \beta < \hat{\beta}_1$ would choose the healthy alternative if their choice was going to be made once more, but would choose the tempting alternative in a single period. In future periods, the precedent effect will be mitigated by a "procrastination effect."

Example 4.2 (procrastination effect)**.** *Suppose $T \geq 3$, and we have $\Delta(\hat{\beta}_1) > \hat{\beta}_2$. The $t = T-1$ solutions are given by equations 4.13 and 4.14. Then for $t = T-2$, if $\hat{\beta}_3 > \underline{\beta}$ we have*

$$a^*_{T-2}(\beta) = \begin{cases} G & \text{if } \beta \in [\hat{\beta}_1, \overline{\beta}) \\ D & \text{if } \beta \in \left[\Delta^{-1}(\hat{\beta}_2), \hat{\beta}_1\right) \\ G & \text{if } \beta \in \left[\hat{\beta}_3, \Delta^{-1}(\hat{\beta}_2)\right) \\ D & \text{if } \beta \in (\underline{\beta}, \hat{\beta}_3). \end{cases} \tag{4.15}$$

Proof. Proposition 4.1 defines $a^*_{T-1}(\beta)$ on $(\underline{\beta}, \hat{\beta}_2) \cup [\hat{\beta}_1, \overline{\beta})$. For $\beta \in (\Delta^{-1}(\hat{\beta}_2), \hat{\beta}_1)$ we know that $V^*_{T-1}(\Gamma(\beta)) = g + \delta g = V^*_{T-1}(\Delta(\beta))$. Since $\beta < \hat{\beta}_1$ we have $a^*_{T-2}(\beta) = D$. For $\beta \in (\hat{\beta}_3, \Delta^{-1}(\hat{\beta}_2))$, $\Delta(\beta) < \hat{\beta}_2$, so $V_{T-1}(\Delta(\beta)) = d + \delta d$, and $\Gamma(\beta) > \Gamma(\hat{\beta}_3) > \hat{\beta}_2$, so $V_{T-1}(\Gamma(\beta)) = g + \delta g$. Since $\beta > \hat{\beta}_3$, we have $a^*_{T-2}(\beta) = G$. The value functions V^*_{T-2} are determined from a^*_{T-2}. □

So even though the agent will ultimately choose G in later periods, currently he chooses D on the interval $\left[\Delta^{-1}(\hat{\beta}_1), \hat{\beta}_1\right)$, exhibiting the procrastination effect. The effect is rather intutive: there is a long tradition of self-indulgence before periods of restraint, whether as a bachelor party, New Year's Eve, or Mardi Gras.

4.2.2 Stochastic Utility and Descriptive Rules

At this point, for the sake of consistency with Benabou and Tirole (2004) we will expand the model further, to include stochastic costs. Each period an agent encounters a high- or low-quality grapefruit $g_t \in \{g_L, g_H\}$ where $g_t = g_H$ with probability π, and $g_t = g_L$ with probability $1 - \pi$. Ignoring temptation, the low-quality grapefruits will be worse to eat than the donuts, but the high-quality grapefruits will still be better, $g_L < d < g_H$. This modification changes the solutions in the cases of 1 and 2 total periods. With only one period, agents will always choose the donut rather than endure bad grapefruit. With two periods and certain restrictions, agents may still choose the bad grapefruit.

We will rewrite utility functions, defined in equations 4.1 and 4.2 as

$$u(a,g) = \begin{cases} g & \text{if } a = G \\ d & \text{if } a = D. \end{cases} \tag{4.16}$$

$$v(a,g,\beta) = \begin{cases} g & \text{if } a = G \\ d/\beta & \text{if } a = D. \end{cases} \tag{4.17}$$

An agent who is a sophisticated quasi-hyperbolic discounter will have a pure-strategy solution to his problem defined below.

Definition 4.3. For T periods, (a, V) is a stochastic solution to the sophisticated model with finite periods if $\forall t, 1 \le t \le T$, we have $a_t : (\underline{\beta}, \overline{\beta}) \times \{g_h, g_l\} \to \{G, D\}$ and $V_t : (\underline{\beta}, \overline{\beta}) \times \{g_h, g_l\} \to \mathbb{R}$ s.t.

$$a^*_T(\beta, g) = \begin{cases} G & \text{if } v(G,\beta,g) \ge v(D,\beta,g) \\ D & \text{if } v(G,\beta,g) < v(D,\beta,g) \end{cases} \tag{4.18}$$

and

$$EV^*_T(\beta) = \pi u(a^*_T(\beta, g_H), g_H) + (1-\pi) u(a^*_T(\beta, g_L), g_L) \tag{4.19}$$

and $\forall 1 \le t < T$

$$a^*_t(\beta, g) = \begin{cases} G & \text{if } v(G,\beta,g) + \delta EV^*_{t+1}(\Gamma(\beta)) \ge v(D,\beta,g) + \delta EV^*_{t+1}(\Delta(\beta)) \\ D & \text{if } v(G,\beta,g) + \delta EV^*_{t+1}(\Gamma(\beta)) < v(D,\beta,g) + \delta EV^*_{t+1}(\Delta(\beta)) \end{cases} \tag{4.20}$$

$$EV^*_t(\beta) = \pi V^*_t(\beta, g_H) + (1-\pi) V^*_t(\beta, g_L) \tag{4.21}$$

$$V^*_t(\beta, g) = \begin{cases} u(G,g) + \delta EV^*_{t+1}(\Gamma(\beta)) & \text{if } a_t(\beta,g) = G \\ u(D,g) + \delta EV^*_{t+1}(\Delta(\beta)) & \text{if } a_t(\beta,g) = D. \end{cases} \tag{4.22}$$

Example 4.3. *If $T = 2$ we can define solutions as before with thresholds. Let $\hat{\beta}^H_L = \frac{d}{g_H}$. And let*

$$\hat{\beta}^L_2 = \max\left\{ \frac{d}{g_L + \delta\pi(g_H - d)}, \Gamma^{-1}(\hat{\beta}^H_1) \right\} \tag{4.23}$$

$$\hat{\beta}^H_2 = \max\left\{ \frac{d}{g_H + \delta\pi(g_H - d)}, \Gamma^{-1}(\hat{\beta}^H_1) \right\}. \tag{4.24}$$

Then if $\hat{\beta}_2^L \leq \overline{\beta}$, *and* $\Delta(\overline{\beta}) > \hat{\beta}_1^H > \Delta(\hat{\beta}_2^L)$, *Figure 4.1,* R_0 *represents the solution to* $t = T - 1$ *where* $\beta_1 \in \left[\underline{\beta}, \hat{\beta}_2^H\right)$, $\beta_2 \in \left[\hat{\beta}_2^H, \hat{\beta}_2^L\right)$, *and* $\beta_3 \in \left[\hat{\beta}_2^H, \Delta^{-1}\left(\hat{\beta}_1^H\right)\right]$ $\beta_4 \in \left(\Delta^{-1}\left(\hat{\beta}_1^H\right), \overline{\beta}\right]$. *If instead,* $\hat{\beta}_2^L \leq \overline{\beta}$, *and* $\Delta(\overline{\beta}) < \hat{\beta}_2^L$, *define* $\beta_3' \in \left[\hat{\beta}_2^L, \overline{\beta}\right)$, *then* R_1 *represents the solution to* $t = T - 1$. *See Supplemental Section C.1 for more detail.*

Example 4.3 reveals a solution that varies by type and state and provides us with an example of a "descriptive" rule. We refer to this decision as a descriptive rule, because no agent has altered his state or perceptions because of the rule, rather he is following a solution given his type. An observer might describe his actions as a rule, but whether he describes what he is doing as rule does not affect the outcome. Notice that Benabou and Tirole (2004) have a similar definition of a personal rule.

Definition 4.4 (Benabou and Tirole)**.** A sustainable behavioral rule is a perfect Bayesian equilibrium of the dynamic game between an individual's incarnations (possible types).

In Benabou and Tirole's model, two types with fixed self-control reached a perfect Bayesian equilibrium. The types that myopically preferred to "persevere" (G) under low costs (equivalent to $\beta \geq \hat{\beta}_1^H$) were called high types. The types that myopically preferred to give up under low costs were called low types. The solution described in example 4.3 contains four types, two low and two high, by that definition. We will define two low types $\beta_1 \in \left[\underline{\beta}, \hat{\beta}_2^H\right)$, $\beta_2 \in \left[\hat{\beta}_2^H, \hat{\beta}_2^L\right)$, and two high types $\beta_3 \in \left[\hat{\beta}_2^H, \Delta^{-1}\left(\hat{\beta}_1^L\right)\right]$, $\beta_4 \in \left(\Delta^{-1}\left(\hat{\beta}_1^H\right), \overline{\beta}\right)$. Under different conditions in 4.3 we have only three types, two low and one high, as before β_1, β_2, and additionally $\beta_3' \in \left[\hat{\beta}_2^L, \overline{\beta}\right)$. The types in this model do not interact; the actions of one type do not influence the other. The solution is a Bayesian optimal and we can interpret it as a personal rule.

Following the language of Benabou and Tirole, rule R_0 for β_3 and β_3' would be referred to as a "bright line" rule, because the type perseveres against both low and high quality; that type's rule is always to eat the grapefruit regardless of state. Those authors suggested that the types persevered to avoid *signaling* to themselves that they had little willpower. In this model, they persevered to avoid *having* low willpower. Rule R_0 for β_2 and β_4 and R_1 for β_2 resemble "flexible" rules because the types have flexibility under which state to eat the grapefruit.

4.3 An Extensive-Form Representation

Section 4.2.1 concerned viewing the model as a dynamic optimization problem with β_t as a state variable and dynamically inconsistent preferences. Harris and Laibson (2001) have shown that such

R_0	β_1	β_2	β_3	β_4
g_H	D	G	G	G
g_L	D	D	G	D

R_1	β_1	β_2	β_3'
g_H	D	G	G
g_L	D	D	G

Figure 4.1: Benabou and Tirole style descriptive rules from 4.3. Types β_3 under (R_0) and β_3' under (R_1) follow a "bright line" rule. Types β_2, β_2 under (R_0), and β_2 under (R_1) follow a "flexible" rule.

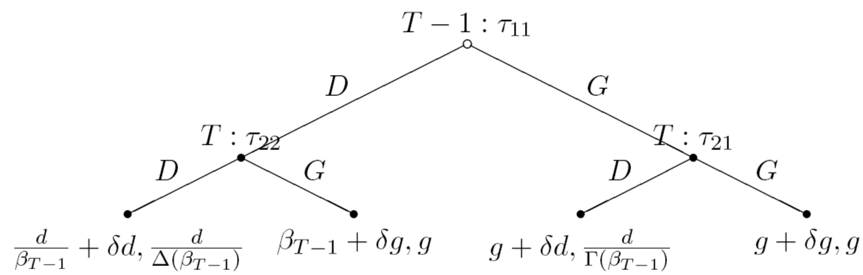

Figure 4.2: Extensive form equivalent for $t = T-1$ in model. Players are named T and $T-1$. Initial value is β_{T-1}. Decisions nodes are $X = \{\tau_{11}, \tau_{21}, \tau_{22}\}$.

models can describe a dynamically inconsistent consumer under certain assumptions. But Gul and Pesendorfer (2005) have shown that under wider assumptions such models may not have a solution. This section approaches the sophisticated solution in a game-theoretical manner to show that a solution to the problem exists for all finite periods.

Consider Example 4.1 in the previous section. Figure 4.2 expresses this game as a sequential game between two players. Notice at each decision node (of the set $X = \mathbb{T} \backslash Z$)[15] agent T will only choose G if $\beta_T \in \{\Delta(\beta_{T-1}), \Gamma(\beta_{T-1})\}$ is greater than or equal to $\frac{d}{g}$. That is, $\beta_T \geq \hat{\beta}_1$. Figure 4.3 shows the game with three players starting at $T = t-2$. Notice that the decision nodes $\tau_{ij} \in X$ are labeled in order.

Using the definition of Kreps and Wilson (1982) (see Section C.2 for the full version), it is possible to show that for any $n \in \mathbb{N}$, $t = T-n$ the model can be expressed as an extensive form game. In terms of terminology each "player" will be the representation of the agent at time t. In other words, there will be n players, but only one agent.

Lemma 4.3. *For any $n \in \mathbb{N}$, and initial state β_0 the model can be expressed as a extensive game with perfect information, nodes $\tau_{ij} \in \mathbb{T}$, players $T-n+1 \leq i \leq T$, actions $a_i \in \{G, D\}$, and*

[15] This is the set characterization of Kreps and Wilson (1982): Z is the set of terminal nodes and $\mathbb{T}$ is the set of all nodes. We use $\mathbb{T}$ instead of T to avoid confusion with the last period.

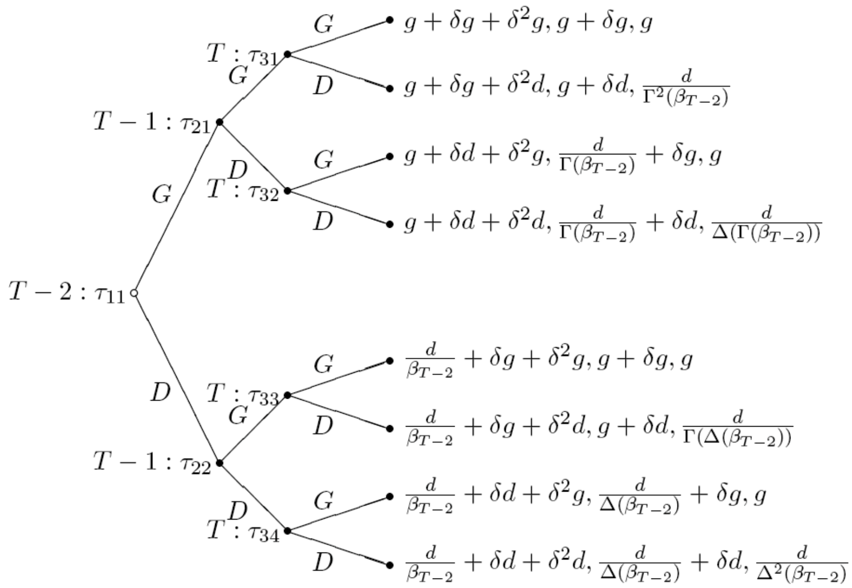

Figure 4.3: Extensive form equivalent for $t = T-2$ in model. Players are named $T, T-1, T-2$. Initial value is β_{T-2}. Decisions nodes are $X = \{\tau_{11}, \tau_{21}, \tau_{22}, \tau_{31}, \tau_{32}, \tau_{33}, \tau_{34}\}$.

utility functions

$$U_i = v(a_i, \beta|\tau_{ij}) + \sum_{k=i+1}^{T} \delta^k u(a_k) \tag{4.25}$$

where $\beta|\tau_{ij}$ is determined by the previous actions leading back to β_0 as defined by equation 4.4.

Proof. See Supplemental Section C.2. □

Theorem 4.1 (Myerson (1991), 4.7)**.** *A finite extensive form game with perfect information has at least one sequential equilibrium in pure strategies. Furthermore, for all generic games with perfect information, there is exactly one sequential equilibrium.*

Proposition 4.2. *For every $t \in \mathbb{N}$, there exists a sophisticated solution (defined in definition 4.1). The model has an equilibrium where each player $T - n \leq i \leq T$ plays $a_t \in \{D, G\}$ that is sequentially rational.*

Proof. From lemma C.1 and Myerson's theorem, we know a pure-strategy sequential equilibrium exists in the extensive form game defined in lemma C.1. Since this game has perfect information it is also a subgame perfect equilibrium. That means for any player i the function U_i (see equation 4.25) is maximized based on future player behavior (in other words, backward induction). This is identical to the definition of the model in definition 4.1 □

From Myerson's theorem we see that our model has at least one solution for any finite time horizon. The solution is also sequentially rational; since the game has full information this implies the solution is from backward induction (i.e., subgame perfect). Now that we know that we always have a solution our results (i.e., proposition 4.1) apply more generally. Generally the solution is unique, but the game could be considered a non-generic game if at a node a player were indifferent between G and D. However, we have assumed at this case each player will only choose G, so when multiple equilibria do exist we are only looking at one of them.

Now that we know there will always be a solution for this model, we turn our attention to personal rules.

4.4 Personal Rules

We will begin with a very broad definition of a personal rule: it is a mapping from decision nodes to actions. To begin, we will put no additional restrictions on the definition.

Definition 4.5. For any $n \in N$. For each player i, $T - n + 1 \leq i \leq T$ define $\mathbb{T}_i$ the set of future nodes as $\mathbb{T}_i = \{\tau_{kj} : k > i - (T - n + 1) + 1\}$ and future decision nodes as $X_i = \mathbb{T} \backslash Z$. Then for every i, $T - n + 1 \leq i \leq T$, each μ_i is a mapping of future decision nodes to future actions, that is $\mu_i : X_i \rightarrow A_i$ where $A_i = \{G, D\}^{T-i-1}$. An agent's personal rule is the set $\mu = \{\mu_{T-n+1}, \ldots, \mu_T\}$.

From this characterization and proposition 4.2, the sophisticated solution for any β_0 and n can be expressed as a personal rule, provided an action is added at the initial node. If we only consider rules consistent with subgame perfection, then it is a "descriptive" personal rule—it *describes* a type of equilibrium but does not influence it. Benabou and Tirole (2004) also used descriptive personal rules in their model (see Section 4.2.2), but their results were the result of a perfect Bayesian equilibrium.

4.4.1 Prescriptive Rules

Ainslie (1992) defines personal rules as agreements an agent makes with his future self. Schelling (1985) suggests that personal rules are principles people "impose on themselves" to "govern their future behavior." These other definitions imply that making a personal rule changes an agent's perception of the problem he will solve. This type of rule, will be defined as "prescriptive," because it is determined before an agent chooses a solution to a problem, and it influences that solution.

The rules in the previous solution were descriptive because they satisfied neither criterion. It should be noted that with definition 4.5, because the definition is not limited to sequentially rational rules, there is great freedom to make a personal rule that alter decisions by the agent. However, in the next section the rules will follow the conventions of the chapter before and will consist of beliefs that are sequentially rational.

4.4.1.1 Ainslie's Example

Ainslie (1992) provides an example that can distinguish descriptive personal rules from prescriptive. We will modify the example slightly to fit the model of this chapter, but the spirit of Ainslie will be preserved. In the example, the agent faces sequential choices ($T = 10$). In the first period, the agent myopically prefers to give up rather than persevere ($\beta_1 < \hat{\beta}_1$), and would prefer to persevere for all periods rather than give up for all periods ($\beta_1 > \hat{\beta}_{10}$). Ainslie suggests that the agent forces himself to persevere in each period by playing a "grim-trigger strategy," a strategy which is similar to strategies used in repeated prisoner's dilemma games with certain endings. The agent tells himself that if he gives up once, he will never persevere again. He and his future selves have made an agreement: persevere in the past and they will persevere. His concept of a rule is prescriptive to decision making. Using the model with certain costs, we can find values that match Ainslie description of the situation, even without defining a personal rule, in this way it can still be thought of as a descriptive rule.

Example 4.4. *Let $T = 10$. Given g, d, Γ, Δ, and $\beta_0 \in \left(\underline{\beta}, \overline{\beta}\right)$, suppose that $\forall t$, $1 \leq t \leq 9$, $\Gamma_{t-1}(\beta_0) < \Delta^{-1}(\beta_{10-t+1})$, and $\beta_0 > \hat{\beta}_{10}$. An agent in state β_0 in period 1 will choose the grapefruit every period. If he gave up once, he would choose the donut and would continue to do so for the remainder of the periods.*

Proof. First, $\beta_0 > \hat{\beta}_{10}$ implies $\Gamma_{t-1}(\beta_0) > \hat{\beta}_{10-t}$, for all $1 < t \leq 9$, by definition 4.2. Then at $t = 10$, $\Gamma_9(\beta_0) > \hat{\beta}_1$, and $\Delta(\Gamma_9(\beta_0)) < \hat{\beta}_1$ the agent will choose (G) after nine previous Gs and, trivially, choosing D leads to never choosing G again (since there are no more choices). Let us proceed inductively: consider any $t-1$ where at $a_t^*(\Gamma_{t-1}(\beta_0)) = G$ and our agent will choose G indefinitely until T. For $t-2$, since $\Gamma_{t-2}(\beta_0) > \hat{\beta}_{10-t}$, our agent prefers the aggregate choice of G to D until T, and $\Gamma_{t-2}(\beta_0) < \Delta^{-1}(\beta_{10-t})$ implies giving up once will once will lead to giving up for all future periods. Thus the agent chooses G. This completes the inductive argument. $\square$

As the example shows, Ainslie's example can be duplicated in the model by having each period's state such that if an agent chooses donut, he will be have dropped below the lower threshold of self-control and will always choose D. Nonetheless this situation only applies to a small set of $\{\beta_t, \Gamma, \Delta\}$ and requires strict assumptions on the transfer functions.

Definition 4.6. For any $n \in N$ let μ be a all-or-nothing rule defined as, $\forall T-n+1 \leq i < T$

$$\mu_i(\tau_{kj}) = \begin{cases} G & \text{if } j = 1 \\ D & \text{if } j > 1 \end{cases} \tag{4.26}$$

where $k > i-(T-n+1)+1$.

In this type of rule, the initial player believes that G will only be chosen at the highest node of the tree, if he defects he believes no one will choose G again, all other future nodes have the same condition; one defection leads to permanent defection. While this description may sound outlandish, consider that several real life policies for abstinence from alcohol and strict diets fit this description. (Marlatt, 2005; Baumeister et al., 1994)

If we allow rules that are not sequentially rational, they are very powerful in the finite period case. We define a looser restriction of consistency on personal rules. It is modeled from the concept of Kreps and Wilson (1982). Basically, a consistent rule is defined in such a way that no player will ever witness it proved false.

Definition 4.7. For any μ define a set of best responses $\hat{A}_i = \operatorname{argmax}_{a_i} U_i(a_i, \mu_{-i})$ where U_i is defined as in equation 4.25. Let $\hat{X}$ be all possible nodes that could be reached from action set $\hat{A}$. If for every $\hat{x} \in \hat{X}$, $a \in \hat{A}$, $\mu(\hat{x}) = a(\hat{x})$ then we say that our rule μ is consistent.[16]

Proposition 4.3. *Suppose at t and $\beta_0 > \hat{\beta}_{T-t+1}$, agent t defines an all-or-nothing rule. Then an agent will choose G for all periods. Furthermore, the beliefs are consistent.*

Proof. Consider the decision node $\tau_{(T-t)1}$, with state $\Gamma_{T-t}(\beta_0)$ and player $T-1$. Player $T-1$'s beliefs μ_{T-1} suggest $\mu_{T-1}(\tau_{(T-t)1}) = G$ and $\mu_{T-1}(\tau_{(T-t)2}) = D$. Since $\beta_0 > \hat{\beta}_{T-t-1}$, $\Gamma_{T-t}(\beta_0) > \hat{\beta}_2$, so $T-1$ prefers G based on beliefs, and $\Gamma_{T-t+1}(\beta_0) > \hat{\beta}_1$ so T prefers G. Thus $T-1$ will choose G, believing future T's will choose G, and that belief is consistent because T chooses G.

Now consider any n where $T-n+1$ chooses $\tau_{(T-t+n)1}$. Then $\mu_{T-1}(\tau_{(T-t+n+1)1}) = G$ and $\mu_{T-1}(\tau_{(T-t+n+1)2}) = D$. Since $\beta_0 > \hat{\beta}_{T-t+1}$ we must have $\Gamma_{T-t+n-1}(\beta_0) > \hat{\beta}_n$ so by beliefs

[16] This concept is similar to the idea of "self-confirming" from citetfude93, although our model has perfect information.

$T-n+1$ chooses G, and since $T-n+1$ chooses G at $\tau_{(T-t+n+1)1}$, this belief is consistent. This completes our inductive argument. $\square$

In the finite period case, without any stochastic parameters, the all-or-nothing rule is very effective. However, it should be noted that given these stochastic parameters the rule might not work as well as more agents would given in and choose D given the g_L state. In the infinite period the all-or-nothing rule does not work as well.

4.5 The Infinite Horizon Case

In the infinite period case there is a continuum of β and an infinite number of periods, so it is not possible to represent the game as a finite extensive-form game. By defining a very specific type of solution, we can find another equilibrium. A stationary[17] solution with infinite periods is defined as:

Definition 4.8. The twople, $(\sigma^*_\infty, V^*_\infty) : \left(\underline{\beta}, \overline{\beta}\right)^2 \to ([0,1], \mathbb{R})$ is a stationary infinite period solution if $\forall \beta \in \left(\underline{\beta}, \overline{\beta}\right)$,

$$0 \leq \sigma^*_\infty(\beta) \leq 1 \tag{4.27}$$

$$V^*_\infty(\beta) = \sigma^*_\infty(\beta)\left[u(G) + \delta V^*_\infty(\Gamma(\beta))\right] + (1 - \sigma^*_\infty(\beta))\left[u(G) + \delta V^*_\infty(\Delta(\beta))\right] \tag{4.28}$$

$$\sigma^*_\infty(\beta) = \begin{cases} 1 & \text{if } v(\beta, G) + \delta V^*_\infty(\Gamma(\beta)) > v(\beta, G) + \delta V^*_\infty(\Gamma(\beta)) \\ 0 & \text{if } v(\beta, G) + \delta V^*_\infty(\Gamma(\beta)) < v(\beta, G) + \delta V^*_\infty(\Gamma(\beta)). \end{cases} \tag{4.29}$$

Notice that we are suggesting that there is a solution where every node, β, plays the same solution over time.[18] We make one more assumption to ensure that for values above $\hat{\beta}_1$, G will always be chosen.

Assumption 4.3. *There exists a $\beta' < \overline{\beta}$ s.t. $\beta' > \frac{d(1-\delta)}{g-\delta d}$.*

Let $\breve{\beta} \equiv \frac{d(1-\delta)}{g-\delta d}$.

Corollary 4.1. *For any function $\frac{d}{1-\delta} < v < \frac{g}{1-\delta}$,*

$$v(D, \beta) + v(\Delta(\beta)) > v(G, \beta) + v(\Gamma(\beta)), \forall \beta < \hat{\beta}_\infty \tag{4.30}$$

$$v(D, \beta) + v(\Delta(\beta)) < v(G, \beta) + v(\Gamma(\beta)), \forall \beta > \breve{\beta}. \tag{4.31}$$

[17] The solution is called stationary because it does not depend on history, only the current β.

[18] Some may find this concept reminiscent of forward induction (Kohlberg and Mertens, 1986).

Proof. Follows directly from definitions. □

Corollary 4.2. *If a solution exists satisfying 4.27, 4.28, 4.29 it must have*

$$v(D,\beta) + v(\Delta(\beta)) < v(G,\beta) + v(\Gamma(\beta)), \forall\beta \geq \hat{\beta}_1. \tag{4.32}$$

Proof. From Corollary 4.1, we know if a solution existed $v(\beta) = \frac{g}{1-\delta}, \forall\beta \geq \breve{\beta}$ It follows that $\forall\beta > \Gamma^{-1}(\breve{\beta})$, we will see action (G) because their value function is at an upper bound, and they prefer G anyway. Then this argument will hold for $\forall\beta > \Gamma^{-1}(\breve{\beta})$ and this process can continue indefinitely $\Gamma^{-t}(\breve{\beta}) < \hat{\beta}$ at which point we will have shown all $\beta \geq \hat{\beta}$ have (G). □

Much like proposition 4.1 did for the finite period, corollaries 4.1 and 4.2 provide a set of boundaries on states outside of which we always know what actions will take place. We can think of the interval $[\hat{\beta}_\infty, \hat{\beta})$ as the *interior* of the state space. We are now ready to make a fixed-point argument. It is actually very similar to Nash (1950) if we think about having each β between $\left(\hat{\beta}_\infty, \hat{\beta}_1\right)$ as representing a player with its own utility function.

Theorem 4.2. *Suppose*

$$V^*_\infty(\beta) = \begin{cases} \frac{g}{1-\delta} & \text{if } \beta \geq \hat{\beta}_1 \\ \frac{d}{1-\delta} & \text{if } \beta < \hat{\beta}_\infty. \end{cases} \tag{4.33}$$

For β_0 let there only be finite n number of states created from all transformations of Γ and Δ where $\hat{\beta}_\infty \leq \beta_t < \hat{\beta}_1$. Then there is a vector $\tilde{\omega}$ that satisfies 4.27, 4.28, 4.29.

Proof. See supplemental section C.3. □

Corollary 4.3. *Provided for β_0 there are only a finite n number of states created from all transformations of Γ and Δ where $\hat{\beta}_\infty \leq \beta_t < \hat{\beta}_1$. There exits a vector $\tilde{\sigma}$ that satisfies 4.27, 4.28, 4.29.*

Proof. The results of Theorem 4.2, corollaries 4.1 and 4.2 are all compatible with each other and 4.27, 4.28, 4.29. Since a fixed point exists (Theorem 4.2), it can exist with the conditions of corollaries 4.1 and 4.2. □

Thus in the infinite period case as long as there are only finitely many possible transformations that can be done to β_0 inside the interior, an equilibria always exists. The next section will investigate a special case where the Γ and Δ functions are inverses. Note that in this special case, there are always finitely many states in the interior, and thus a solution always exists.

4.5.1 One-Step Cases

The class of infinite solutions that can exist is quite broad. This section concerns a specific class of solutions, those in the case where Γ and Δ are inverses. In this case each move is one step to a higher or lower state.

Assumption 4.4. *The functions Γ and Δ are inverse of each other $\Gamma = \Delta^{-1}$ and $\Gamma^{-1} = \Delta$.*

Definition 4.9 (interior nodes)**.** For any $\hat{\beta}_\infty \leq \beta < \hat{\beta}_1$ let n denote the number of reachable nodes between $\hat{\beta}_1$ and $\hat{\beta}_\infty$. That is

$$n = 1 + \max_{\{i:\Gamma_i(\beta)<\hat{\beta}_1\}} i + \max_{\{j:\Delta_i(\beta)\geq\hat{\beta}_\infty\}} j. \tag{4.34}$$

Also for any n, β we will number the interior nodes in increasing order $N = \{\beta_n, \ldots, \beta_1\}$.

For the sake of notational simplicity, define $u \equiv g - d$, $d_i \equiv 1 - \frac{d}{\beta_i} - d$, $\bar{u} = \frac{u}{1-\delta}$. It follows that $d_1 < \ldots < d_n < \bar{u}$.

Example 4.5. *If $n = 2$ and $\delta \geq \sqrt{1 - \frac{u}{d_1}}$ we will have $a^*_\infty(\beta_i) = G, \forall i \in \{1, 2\}$. Otherwise we will have*

$$p_1 = 1 - \frac{\bar{u} - d_2}{\delta d_1} \qquad p_2 = \frac{\bar{u} - d_1}{\delta d_2}. \tag{4.35}$$

Proof. See Appendix C.4. □

Proposition 4.4. *For any even*[19] *$n \in 2\mathbb{N}$, there exists a stochastic solution where*

$$\begin{aligned} p_i &= 1 - \frac{\bar{u} + \frac{n-i-1}{2}u + \frac{i-1}{2}\delta u - \sum_{m=i+1}^{n/2} d_{2m} + \delta \sum_{m=i}^{\frac{m-1}{2}} d_{2m+1}}{\delta d_i}, \quad \forall i < n, i \in 2\mathbb{N}+1 \\ p_j &= \frac{\bar{u} + \frac{j-2}{2}u\frac{n-i}{2}\delta u - \sum_{m=1}^{m/2} d_{2m-1} - \delta \sum_{m=i}^{\frac{m-2}{2}} d_{2m}}{\delta d_j} \quad \forall j < n, j \in 2\mathbb{N} \end{aligned} \tag{4.36}$$

provided p_i, $p_j \in (0,1)$ $\forall\, j, i < n$, $j \in 2\mathbb{N}$ $i \in 2\mathbb{N}+1$.

[19]Notice that if n is odd we can have the stochastic solution over k-1 nodes, these conditions depend on δ, and to a lesser extent the other parameters.

Proof. See Appendix C.4. □

Given the existence of these mixed strategy equilibria in the infinite period case, and possible non-existence of pure-strategy equilibria, it becomes easy to see that there may not be a consistent all-or-nothing rule for every infinite period game. In the finite period setting without variable costs the all-or-nothing rules were consistent and very effective at causing and agent to persevere indefinitely. With infinite periods this does not happen: consider any β_0 where $\beta_0 < \Gamma(\beta_0) < \hat{\beta}_1$ if $\sigma(\Gamma(\beta_0)) \in (0,1)$ we do not have a consistent rule. Example 1 suggests if $n \geq 2$ and $\delta < \sqrt{1 - \frac{u}{d_1}}$ we will likely have no consistent rules.

4.6 Testable Implications

There are some general patterns of this model that could be examined by empirical methods.

1. *Precedent and Procrastination effect.* With repeated tasks over finite periods we should find that more people avoid temptation when the choice is repeated rather than as an a single decision (e.g., Kirby and Guastello, 2001, also see example 4.1). But the procrastination effect may interfere with this result, and lead to more choices of temptation with multiple periods (see example 4.2). However, the procrastination effect only works in a few distinct ways. First, we will never see a pattern like {G,G,D,D}. For any precedent effect to occur, the last period must be the non-tempting choice. Further a procrastination effect also only occurs at periods greater than the last two. So, generally, if a sequence of actions contains any decisions away from temptation, this model predicts it must include the two decisions.

2. *Descriptive Rules.* The results of this study under stochastic costs are nearly identical to Benabou and Tirole (2004). Both models can explain low and high types persevering under high and low conditions. However they occur for different reasons. In an experiment with perfect monitoring (suppose subjects record the difficulty of their task) their model predicts low willpower types should not persevere because they cannot signal future selves. Monitoring will have no effect on this type of model. Also under their model, any signal of great strength should lead to more signs of strength, while the signal without the action would have no effect on the predictions of this model.

3. *Prescriptive Rules.* Under this model prescriptive all-or-nothing rules are very effective in finite periods, but not in infinite periods. If we find can find an instance in an experiment

where individuals appear to be using an all-or-nothing strategy, an interesting test would be to see if that would continue in an infinite horizon environment.

4. *Habit Formation.* The main assumption of the study is that immediacy peference for a given item increases (or decreases) with use (or disuse). Charness and Gneezy (2007) have shown evidence of habit with excercise, that after offering a treatment group incentives to go to the gym, that group was more likely to continue to go to the gym than a control even after all incentives had been removed.

   However, it is not clear whether this is a habit being built up on a standard utility function (i.e., Becker and Murphy, 1988) or immediacy preference parameter. One way our model could be tested is through an experiment for four weekly lab sessions. In the session subjects have the opportunity to partake in some tempting task, or menial. Each session they use a second-price mechanism to evaluate their true cost of doing the test. In one treatment the mechanism could be used instantaneosly to ascertain β on the other it could be done the session before. Standard habit formation models would suggest both values would be equal across treatments, but the value to do the task should decrease when it is performed. Alternatively, this model suggests the group giving immediate preferences should have higher valuations than the control group, but that immediate group's preferences should change dependent on whether they do the task. The test could evaluate this model against standard habit formation and also may lead to some explanation of the functions Γ and Δ that underly β's changes.

4.7 Conclusion

This chapter attempts to solve the dilemna brought about by repeated choices for a dynamically inconsistent agent. In every period the agent may wish to take the sensible alternative always over the tempting one, but will not make that choice on an individual level if does not affect future outcomes. Several explanations have attempted to explain how an individual can remedy this predicament, most notably Ainslie (1992) suggests an individual can make an agreement with future selves and Benabou and Tirole (2004) suggest an individual perseveres to give positive signals of willpower to future selves.

This chapter gives a different explanation, that individuals persevere to build up their resistance to the current temptation in later periods. This result is based off the findings of Baumeister et al.

(1994), but it can be thought of as adding a habit formation component to immediacy preference. With this assumption the model can create a precedent effect—the tendency of individuals to forgo temptation when making a repeated choice—in the form of a subgame perfect (and sequential) equilibrium among selves.

It is this assumption that allows for so strong of a descriptive rule. Without altering one's preferences, in the last period an individual would always want to seek temptation, and there would not be a backward induction solution that suggests agents will persevere. The unique assumption of this model is very helpful in giving agents rational incentives to persevere: if they do, they will become better in the future.

However, this incentive has its drawbacks, if an agent knows he will become better in future anyway (because the incentive will also affect his future self) he will exhibit a procrastination effect. He will seek temptation now, knowing his future is unaffected, because his future self will pick up the slack. The effect is very similiar to those who attend Mardi Gras, bachelor parties, or New Year's Eve parties before making strong commitments (i.e., Lent, marriage, New Year's resolutions).

By representing the model as an extensive form game, it is shown that a backward induction solution exists for all periods. The extensive form also allows for a more general form of personal rules. The subgame perfect form of the rules is classified as a descriptive rule and is similiar to other studies (Benabou and Tirole, 2004). A set of beliefs that moves an individual off the subgame perfect path is classified as a descriptive rule. Those rules are characterized on whether they are consistent, meaning each player never learns his beliefs are wrong.

The idea of a descriptive rule may be criticized because it is an incorrect belief, but by allowing a concept of consistency, we allow rules that are irrational, but never tested. There is evidence that these type of consistent rules exist. Individuals who do not use drugs admit to believing the first use of cocaine or heroin would lead to unavoidable addiction (Miller and Brown, 1991), a falsehood (Peele, 1989; Robins et al., 1975; Baumeister et al., 1994).

When the infinite horizon problem is examined it is simplified so that the solution is classified without history, the state value β is all that determines an action. Under this form there always exists an equilibrium, albeit sometimes a stochastic one. When the results are stochastic, personal rules cannot be consistent, thus implying personal rules are favored in finite period horizons.

Section 4.6 provides guidelines for this theory to be tested against others, especially models of self-signaling and standard habit formation.

Bibliography

Ainslie, George, "Specious Reward—Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin*, 1975, *82* (4), 463–496.

—, *Picoeconomics: The Interaction of Successive Motivational States within the Person*, Cambridge: Cambridge University Press, 1992.

Allen, Todd W. and Christopher D. Carroll, "Individual Learning About Consumption," *Macroeconomic Dynamics*, 2001, *5* (2), 255–271.

Anderson, Lisa R. and Charles A. Holt, "Information Cascades in the Laboratory," *The American Economic Review*, 1997, *87* (5), 847–862.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 2001, *15* (3), 47–68.

Ballinger, T. Parker, Eric Hudson, Leonie Karkoviata, and Nathaniel T. Wilcox, "Saving Performance and Cognitive Abilities," Stephen F. Austin State University, 2006.

—, Michael G. Palumbo, and Nathaniel T. Wilcox, "Precautionary Saving and Social Learning across Generations: An Experiment," *Economic Journal*, 2003, *113* (490), 920–947.

Barro, Robert J. and Paul M. Romer, "Ski-Lift Pricing, with Applications to Labor and Other Markets," *The American Economic Review*, 1987, *77* (5), 875–890.

Baumeister, Roy F. and Kathleen D. Vohs, "Willpower, Choice and Self-Control," in George F. Loewenstein, D. Read, and Roy F. Baumeister, eds., *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*, New York: Russell Sage, 2003.

—, Todd F. Heatherton, and Dianne M. Tice, *Losing Control: How and Why People Fail at Self-Regulation*, San Diego: Academic Press, 1994.

Becker, Gary S. and Casey B. Mulligan, "The Endogenous Determination of Time Preference," *The Quarterly Journal of Economics*, 1997, *112* (3), 729–758.

— **and Kevin M. Murphy**, "A Theory of Rational Addiction," *The Journal of Political Economy*, 1988, *96* (4), 675–700.

Becker, Howard S., "Notes on the Concept of Commitment," *The American Journal of Sociology*, 1960, *66* (1), 32–40.

Benabou, Roland and Jean Tirole, "Willpower and Personal Rules," *Journal of Political Economy*, 2004, *112* (4), 848–886.

Bernasconi, Michele and Oliver Kirchkamp, "Why Do Monetary Policies Matter? An Experimental Study of Saving and Inflation in an Overlapping Generations Model," *Journal of Monetary Economics*, 2000, *46* (2), 315–343.

Bernheim, B. Douglas and Antonio Rangel, "Addiction and Cue-Triggered Decision Processes," *American Economic Review*, 2004, *94* (5), 1558–1590.

Bodner, Ronit and Dražen Prelec, "The Diagnostic Value of Actions in a Self-Signaling Model," in I. Brocas and J.D. Carillo, eds., *The Psychology of Economic Decisions*, Vol. 1, USA: Oxford University Press, 2003.

Borjas, George J., "Self-Selection and the Earnings of Immigrants," *The American Economic Review*, 1987, *77* (4), 531–553.

Browning, Martin and Annamaria Lusardi, "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature*, 1996, *34* (4), 1797–1855.

Camerer, Colin F., "Three Cheers—Psychological, Theoretical, Empirical—for Loss Aversion," *Journal of Marketing Research*, 2005, *42* (2), 129–133.

— , "Behavioral Economics," in Richard Blundell, Whitney Newey, and Torsten Persson, eds., *Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress of the Econometric Society*, Vol. 2, Cambridge: Cambridge University Press, 2006.

— **and Keith Weigelt**, "Experimental Tests of a Sequential Equilibrium Reputation Model," *Econometrica*, 1988, *56* (1), 1–36.

__ **, Teck-Hua Ho, and Juin-Kuan Chong**, "A Cognitive Hierarchy Model of Games," *The Quarterly Journal of Economics*, 2004, *119* (3), 861–898.

Carbone, Enrica, "Demographics and Behaviour," *Experimental Economics*, 2005, *8* (3), 217–232.

__ **and John D. Hey**, "The Effect of Unemployment on Consumption: An Experimental Analysis," *The Economic Journal*, 2004, *114* (497), 660–683.

Carroll, Christopher D., "The Buffer-Stock Theory of Saving—Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 1992, *1992* (2), 61–135.

__ , "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 1997, *112* (1), 1–55.

__ , "A Theory of the Consumption Function, with and without Liquidity Constraints," *Journal of Economic Perspectives*, 2001, *15* (3), 23–45.

__ **, Changyong Rhee, and Byungkun K. Rhee**, "Does Cultural Origin Affect Saving Behavior? Evidence from Immigrants," *Economic Development and Cultural Change*, 1999, *48* (1), 33–50.

__ **, Judy Overland, and David N. Weil**, "Saving and Growth with Habit Formation," *American Economic Review*, 2000, *90* (3), 341–355.

Charness, Gary and Uri Gneezy, *Incentives to Exercise*, SSRN, 2007.

Chen, Hsiao-Chi, James W. Friedman, and Jacques-Francois Thisse, "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach," *Games and Economic Behavior*, 1997, *18* (1), 32–54.

Choi, James J., David Laibson, and Brigitte C. Madrian, "Reducing the Complexity Costs of 401(K) Participation through Quick Enrollmenttm," in David A. Wise, ed., *Developments in the Economics of Aging*, Chicago: University of Chicago Press, (forthcoming).

__ **,** __ **, and** __ , "$100 Bills on the Sidewalk: Suboptimal Saving in 401(K) Plans,", Harvard University, 2005.

__ **,** __ **,** __ **, and Andrew Metrick**, "Optimal Defaults," *American Economic Review Papers and Proceedings*, 2003, *93* (2), 180–185.

Chua, Zhikang and Colin F. Camerer, "Experiments on Intertemporal Consumption with Habit Formation and Social Learning,", California Institute of Technology, 2004.

Cochrane, John H., "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *The American Economic Review*, 1989, *79* (3), 319–337.

Crawford, Vincent P., "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions," *American Economic Review*, 2003, *93* (1), 133–149.

_ and Nagore Iriberri, "Fatal Attraction: Salience, Naivete, and Sophistication in Experimental 'Hide-and-Seek' Games," *American Economic Review*, 2007, *97* (5).

_ and _ , "Level-K Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica*, 2007, *75* (6).

Darlin, Damon, "A Contrarian View: Save Less and Still Retire with Enough,", *The New York Times*, January 27, 2007.

De Vany, Arthur and W. David Walls, "Bose-Einstein Dynamics and Adaptive Contracting in the Motion Picture Industry," *The Economic Journal*, 1996, *106* (439), 1493–1514.

DeFrank, Thomas M., "Bush Zips Lip Again Refuses to Reply to Drug-Use Questions,", *New York Daily News*, August 21, 1999.

DeGeorge, Francois, Jayendu Patel, and Richard Zeckhauser, "Earnings Management to Exceed Thresholds," *Journal of Business*, 1999, *72* (1), 1–33.

DellaVigna, Stefano, "Psychology and Economics: Evidence from the Field," University of California, 2007.

_ and M. Daniele Paserman, "Job Search and Impatience," *Journal of Labor Economics*, 2005, *23* (3), 527–588.

Duesenberry, James Stemble, *Income, Saving, and the Theory of Consumer Behavior*, Cambridge, MA: Harvard University Press, 1949.

Dye, Ronald A., "Disclosure of Nonproprietary Information," *Journal of Accounting Research*, 1985, *23* (1), 123–145.

— **and Sri S. Sridhar**, "Industry-Wide Disclosure Dynamics," *Journal of Accounting Research*, 1995, *33* (1), 157–174.

Eliashberg, Jehoshua and Steven M. Shugan, "Film Critics: Influencers or Predictors," *Journal of Marketing*, 1997, *61* (1), 68–78.

Ellison, Glenn, "Bounded Rationality in Industrial Organization," in Richard Blundell, Whitney K. Newey, and Torsten Persson, eds., *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Cambidge: Cambridge University Press, 2006, pp. 142–174.

Eyster, Erik and Matthew Rabin, "Cursed Equilibrium," *Econometrica*, 2005, *73* (5), 1623–1672.

Farrell, Joseph, "Voluntary Disclosure: Robustness of the Unraveling Result, and Comments on Its Importance," in Ronald E. Grieson, ed., *Antitrust and Regulation*, Lexington, MA: Lexington Books, 1986.

Fehr, Ernst and Peter K. Zych, "Do Addicts Behave Rationally?" *Scandinavian Journal of Economics*, 1998, *100* (3), 643–662.

— **and** — , "Intertemporal Choice under Habit Formation," in Charles R. Plott and Vernon L. Smith, eds., *Handbook of Results in Experimental Economics*, (forthcoming).

Fishman, Michael J. and Kathleen M. Hagerty, "Mandatory Versus Voluntary Disclosure in Markets with Informed and Uninformed Customers," *Journal of Law, Economics and Organization*, 2003, *19* (1), 45–63.

Forsythe, Robert, R. Mark Isaac, and Thomas R. Palfrey, "Theories and Tests of "Blind Bidding" in Sealed-Bid Auctions," *The RAND Journal of Economics*, 1989, *20* (2), 214–238.

Friedman, Milton, *Essays in Positive Economics*, Chicago: University of Chicago Press, 1953.

— **and National Bureau of Economic Research.**, *A Theory of the Consumption Function*, Princeton: Princeton University Press, 1957.

Fuchs-Schundeln, Nicola and Matthias Schundeln, "Precautionary Savings and Self-Selection: Evidence from the German Reunification 'Experiment,"' *The Quarterly Journal of Economics*, 2005, *120*, 1085–1120.

Fudenberg, Drew and David Levine, "A Dual Self Model of Impulse Control," *American Economic Review*, 2006, *95* (5), 1449–1476.

Gale, William G., John Sabelhaus, and Robert E. Hall, "Perspectives on the Household Saving Rate," *Brookings Papers on Economic Activity*, 1999, *1999* (1), 181–224.

Goldfarb, Avi and Botao Yang, "Are All Managers Created Equal?," SSRN, 2007.

Goldstein, Daniel G., Eric J. Johnson, and William F. Sharpe, "Measuring Consumer Risk-Return Tradeoffs," London Business School, 2006.

Grossman, Sanford Jay, "The Informational Role of Warranties and Private Disclosure About Product Quality," *Journal of Law and Economics*, 1981, *24*, 461–483.

Gul, Faruk and Wolfgang Pesendorfer, "Temptation and Self-Control," *Econometrica*, 2001, *69* (6), 1403–1435.

_ and _ , "The Revealed Preference Theory of Changing Tastes," *Review of Economic Studies*, 2005, *72* (2), 429–448.

Harris, Christopher and David Laibson, "Dynamic Choices of Hyperbolic Consumers," *Econometrica*, 2001, *69* (4), 935–957.

Haveman, Robert, Karen Holden, Barbara Wolfe, and Shane Sherlund, "Do Newly Retired Workers in the United States Have Sufficient Resources to Maintain Well-Being?" *Economic Inquiry*, 2006, *44* (2), 249–264.

Herrnstein, Richard J. and Dražen Prelec, "Melioration: A Theory of Distributed Choice," *The Journal of Economic Perspectives*, 1991, *5* (3), 137–156.

Hey, John D., "A Pilot Experimental Investigation into Optimal Consumption under Uncertainty," in Shlomo Maital, ed., *Applied Behavioural Economics*, Vol. 2, New York: New York University Press, 1988, pp. 653–667.

_ and Valentino Dardanoni, "Optimal Consumption under Uncertainty: An Experimental Investigation," *The Economic Journal*, 1988, *98* (390), 105–116.

Holland, John H., "Properties of the Bucket Brigade Algorithm," in John J. Grefenstette, ed., *Proceedings of the First International Conference on Genetic Algorithms and Their Applications*, Hillsdale, NJ: Lawrence Erlbaum, 1985, pp. 1–7.

Houser, Daniel, Michael Keane, and Kevin McCabe, "Behavior in a Dynamic Decision Problem: An Analysis of Experimental Evidence Using a Bayesian Type Classification Algorithm," *Econometrica*, 2004, *72* (3).

Jin, Ginger Zhe, "Competition and Disclosure Incentives: An Empirical Study of HMOs," *The RAND Journal of Economics*, 2005, *36* (1), 93–112.

_ **and Phillip Leslie**, "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards," *Quarterly Journal of Economics*, 2003, *118* (2), 409–451.

Johnson, Eric J., Simon Gaechter, and Andreas Herrmann, "Exploring the Nature of Loss Aversion," SSRN, 2006.

Jovanovic, Boyan, "Truthful Disclosure of Information," *The Bell Journal of Economics*, 1982, *13* (1), 36–44.

Jung, Woon-Oh and Young K. Kwon, "Disclosure When the Market Is Unsure of Information Endowment of Managers," *Journal of Accounting Research*, 1988, *26* (1), 146–153.

Kahneman, Daniel and Amos Tversky, "Prospect Theory—Analysis of Decision under Risk," *Econometrica*, 1979, *47* (2), 263–291.

Kirby, Kris. N. and Barbarose Guastello, "Making Choices in Anticipation of Similar Future Choices Can Increase Self-Control," *Journal of Experimental Psychology—Applied*, 2001, *7* (2), 154–164.

Kohlberg, Elon and Jean-Francois Mertens, "On the Strategic Stability of Equilibria," *Econometrica*, 1986, *54* (5), 1003–1037.

Kotlikoff, Laurence J., Stephen Johnson, and William Samuelson, "Can People Compute? An Experimental Test of the Life-Cycle Consumption Model," in Laurence J. Kotlikoff, ed., *Essays on Savings, Bequests, Altruism and Life-Cycle Planning*, Cambridge, MA: MIT Press, 2001.

Kreps, David M. and Robert Wilson, "Sequential Equilibria," *Econometrica*, 1982, *50* (4), 863–894.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, 1997, *112* (2), 443–477.

— , "A Cue-Theory of Consumption," *Quarterly Journal of Economics*, 2001, *116* (1), 81–119.

— , Andrea Repetto, and Jeremy Tobacman, "A Debt Puzzle," in Edmund S. Phelps and Philippe Aghion, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Princeton: Princeton University Press, 2003, pp. 228–266.

— , — , and — , "Estimating Discount Functions with Consumption Choices over the Lifecycle," Harvard University, 2007.

Levitt, Stephen D. and Stephen J. Dubner, *Freakonomics—a Rogue Economist Explores the Hidden Side of Everything*, New York: William Morrow, 2005.

Litwak, Mark, *Reel Power: The Struggle for Influence and Success in the New Hollywood*, New York: William Morrow and Co., Inc., 1986.

Loewenstein, George and Ted O'Donoghue, "Animal Spirits: Affective and Deliberative Influences on Economic Behavior," 2004.

Luce, R. Duncan and Howard Raiffa, *Games and Decisions*, New York: Wiley, 1957.

Lusardi, Annamaria, Jonathan Skinner, and Steven Venti, "Savings Puzzles and Savings Policies in the United States," *Oxford Review of Economic Policy*, 2001, *17* (1), 95–115.

Lynch, Michael, Ross Miller, Charles R. Plott, and Russell Porter, "Product Quality, Informational Efficiency, and Regulations in Experimental Markets," in Charles R. Plott, ed., *Collected Papers on the Experimental Foundations of Economics and Political Science*, Economists of the Twentieth Century, Cheltenham, UK ; Northampton, MA USA: Edward Elgar, 2001.

Malkiel, Burton G. and Atanu Saha, "Hedge Funds: Risk and Return," *Financial Analysts Journal*, 2005, *61* (6), 80–88.

Marlatt, G. Alan, "Chapter 1: Relapse Prevention: Theoretical Rationale and Overview of the Model," in G. Alan Marlatt and Dennis M. Donovan, eds., *Relapse Prevention: Maintenance Strategies in the Treatment of Addictive Behaviors*, 2nd ed., New York: Guilford Press, 2005, pp. 1–44.

Mathios, Alan D., "The Impact of Mandatory Disclosure Laws on Product Choices: An Analysis of the Salad Dressing Market," *The Journal of Law and Economics*, 2000, *43* (2), 651–677.

Matthews, Steven and Andrew Postlewaite, "Quality Testing and Disclosure," *The RAND Journal of Economics*, 1985, *16* (3), 328–340.

McClure, Samuel. M., Keith. M. Ericson, David I. Laibson, George Loewenstein, and Johnathan D. Cohen, "Time Discounting for Primary Rewards," *Journal of Neuroscience*, 2007, *27* (21), 5796–5804.

McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," in Paul Zarembka, ed., *Frontiers and Econometrics*, New York: Academic Press, 1974.

McKelvey, Richard D. and R. Palfrey Thomas, "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 1995, *10* (1), 6–38.

_ and _ , "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1998, *1* (1), 9–41.

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 1981, *12* (2), 380–391.

Miller, William R. and Janice M. Brown, "Self-Regulation as a Conceptual Basis for the Prevention and Treatment of Addictive Behaviors," in Nick Heather, William R. Miller, and Janet Greely, eds., *Self-Control and the Addictive Behaviors*, Sydney: Maxwell Macmillian Publishing, 1991.

Muraven, Mark and Roy. F. Baumeister, "Self-Regulation and Depletion of Limited Resources: Does Self-Control Resemble a Muscle?," *Psychological Bulletin*, 2000, *126* (2), 247–259.

_ , _ , and Diane M. Tice, "Longitudinal Improvement of Self-Regulation through Practice: Building Self-Control Strength through Repeated Exercise," *Journal of Social Psychology*, 1999, *139* (4), 446–457.

Myerson, Roger B., *Game Theory: Analysis of Conflict*, Cambridge, MA: Harvard University Press, 1991.

Nagel, Rosemarie, "Unraveling in Guessing Games: An Experimental Study," *The American Economic Review*, 1995, *85* (5), 1313–1326.

Nash, John F., "Equilibrium Points in N-Person Games," *Proceedings of the National Academy of Sciences of the United States of America*, 1950, *36* (1), 48–49.

O'Donoghue, Ted and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 1999, *89* (1), 103–124.

Orbach, Barak Y. and Liran Einav, "Uniform Prices for Differentiated Goods: The Case of the Movie-Theater Industry," *International Review of Law and Economics*, 2007, *27* (2), 129–153.

Östling, Robert, Joseph Tao yi Wang, Eileen Chou, and Colin F. Camerer, "Field and Lab Convergence in Poisson Lupi Games," California Institute of Technology, 2007.

Ozdenoren, Emre, Stephen Salant, and Dan Silverman, "Willpower and the Optimal Control of Visceral Urges," University of Michigan, 2006.

Patton, Jim H., Matthew S. Stanford, and Ernest S. Barratt, "Factor Structure of the Barratt Impulsiveness Scale," *Journal of Clinical Psychology*, 1995, *51* (6), 768–774.

Peele, Stanton, *Diseasing of America: Addiction Treatment out of Control*, Lexington, MA: Lexington Books, 1989.

Plott, Charles R. and Shyam Sunder, "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models," *The Journal of Political Economy*, 1982, *90* (4), 663–698.

_ and _ , "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica*, 1988, *56* (5), 1085–1118.

Prelec, Dražen and George Loewenstein, "The Red and the Black: Mental Accounting of Savings and Debt," *Marketing Science*, 1998, *17* (1), 4–28.

_ and R. J. Herrnstein, "Preferences and Principles, Alternative Guidelines for Choice," in Richard Zeckhauser, ed., *Strategic Reflections on Human Behavior*, Cambridge, MA: MIT Press, 1991.

Read, Daniel, George Loewenstein, and Matthew Rabin, "Choice Bracketing," *Journal of Risk and Uncertainty*, 1999, *19* (1-3), 171–197.

Reinstein, David A. and Christopher M. Snyder, "The Influence of Expert Reviews on Consumer Demand for Experience Goods: A Case Study of Movie Critics," *Journal of Industrial Economics*, 2005, *53* (1), 27–51.

Rick, Scott I., Cynthia E. Cryder, and George Loewenstein, "Tightwads and Spendthrifts," *Journal of Consumer Research*, 2008, *34* (6), 767–782.

Robins, Lee N., John E. Helzer, and Darlene H. Davis, "Narcotic Use in Southeast—Asia and Afterward - Interview Study of 898 Vietnam Returnees," *Archives of General Psychiatry*, 1975, *32* (8), 955–961.

Schelling, Thomas C., *The Strategy of Conflict*, Cambridge, MA: Harvard University Press, 1960.

— , "Enforcing Rules on Oneself," *Journal of Law, Economics and Organization*, 1985, *1* (2), 357–374.

Schneider, Walter and Richard M. Shiffrin, "Controlled and Automatic Human Information Processing: I. Detection, Search, and Attention," *Psychological Review*, 1977, *84* (1), 1–66.

Scholtz, John Karl, Ananth Seshadri, and Surachai Khitatrakun, "Are Americans Saving Optimally for Retirement?" *Journal of Political Economy*, 2006, *114* (4), 607–643.

Selten, Reinhard and Rolf Stoecker, "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames: A Learning Theory Approach," *Journal of Economic Behavior & Organization*, 1986, *7* (1), 47–70.

Shavell, Steven, "Acquisition and Disclosure of Information Prior to Sale," *The RAND Journal of Economics*, 1994, *25* (1), 20–36.

Shiffrin, Richard M. and Walter Schneider, "Controlled and Automatic Human Information Processing: II. Perceptual Learning, Automatic Attending and a General Theory," *Psychological Review*, 1977, *84* (2), 127–190.

Shiv, Baba and Alexander Fedorikhin, "Spontaneous Versus Controlled Influences of Stimulus-Based Affect on Choice Behavior," *Organizational Behavior and Human Decision Processes*, 2002, *87* (2), 342–370.

Silverman, B. W., *Density Estimation for Statistics and Data Analysis* Monographs on Statistics and Applied Probability, New York: Chapman and Hall, 1986.

Simmons, Jacqueline, "A 'Thumbs up' Pulls in the Audience," *Wall Street Journal*, March 25, 1994.

Smith, Adam, *The Theory of Moral Sentiments*, New York: Prometheus Books, 1759/1892.

Stahl, Dale O. and Paul W. Wilson, "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 1995, *10* (1), 218–254.

Strahan, Erin J., Steven J. Spencer, and Mark P. Zanna, "Subliminal Priming and Persuasion: Striking While the Iron Is Hot," *Journal of Experimental Social Psychology*, 2002, *38* (6), 556–568.

Strotz, Robert H., "Myopia and Inconsistency in Dynamic Utility Maximization," *The Review of Economic Studies*, 1955–1956, *23* (3), 165–180.

Tanaka, Tamomi, Colin F. Camerer, and Quang Nguyen, "Poverty, Politics, and Preferences: Experimental and Survey Data from Vietnam," California Institute of Technology, 2006.

Thaler, Richard H., "Mental Accounting Matters," *Journal of Behavioral Decision Making*, 1999, *12* (3), 183–206.

_ **and Hersh M. Shefrin**, "An Economic Theory of Self-Control," *The Journal of Political Economy*, 1981, *89* (2), 392–406.

_ **and Shlomo Benartzi**, "Save More Tomorrow (Tm): Using Behavioral Economics to Increase Employee Saving," *Journal of Political Economy*, 2004, *112* (s1), S164–S187.

Venti, Steven, "Choice, Behavior and Retirement Saving," in G. Clark, A. Munnell, and M. Orzsag, eds., *Oxford Handbook of Pensions and Retirement Income*, Oxford: Oxford University Press, 2006.

Verrecchia, Robert E., "Essays on Disclosure," *Journal of Accounting and Economics*, 2001, *32* (1-3), 97–180.

Viscusi, W. Kip, "A Note on "Lemons" Markets with Quality Certification," *The Bell Journal of Economics*, 1978, *9* (1), 277–279.

Vogel, Howard L., *Entertainment Industry Economics*, Cambridge: Cambridge University Press, 2007.

Wang, Joseph T., "Is Last Minute Bidding Bad?" California Institute of Technology, 2006.

Yi, Richard and Howard Rachlin, "Contingencies of Reinforcement in a Five-Person Prisoner's Dilemma," *Journal of the Experimental Analysis of Behavior*, 2004, *82* (2), 161–76.

Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *The Quarterly Journal of Economics*, 1989, *104* (2), 275–298.

Appendix A

Supplemental Materials for Learning and Visceral Temptation

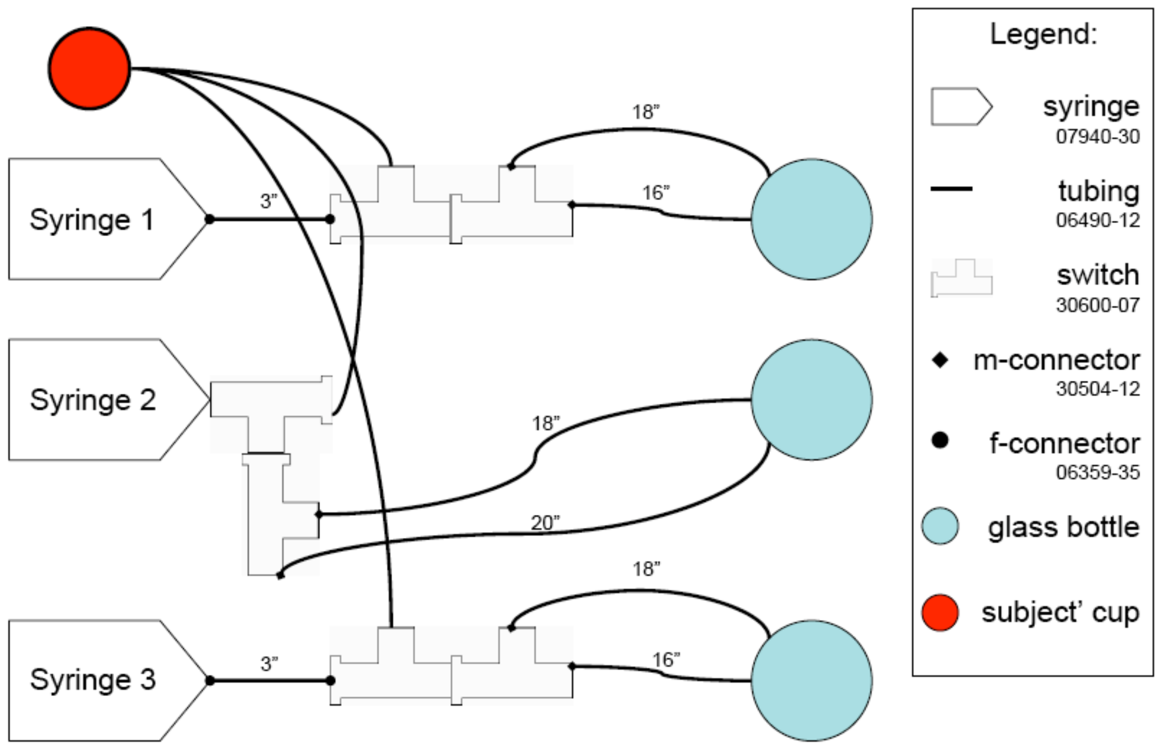

Figure A.1: Diagram of 3-syringe beverage delivery apparatus used in Study 2

A.1 Diagram of the Apparatus

Diagram A.1 shows the beverage delivery apparatus. Tubing from Cole-Parmer was used to deliver an exact amount of beverage into a cup that a subject would drink each period.

A.2 Experimental Instructions

INSTRUCTIONS (Version 1):

Before we begin, there are some rules in this experiment that are necessary for its validity. First we have asked you not to drink liquids for four hours before this experiment. We also will ask you not to drink during this experiment, except for the liquid rewards you receive. Additionally, you will be asked to consume a salty snack before this experiment begins. To ensure that you do not break these rules, if you need to leave the experimental room for a bathroom break, the experimenter or female assistant of the experimenter will monitor that you do not drink liquids during this time. If you feel at any time your health is at risk in this experiment, please tell the experimenter and the session will be stopped. You will receive all earnings up to that point and your show up payment. At the end of the experiment you will be asked to take a small questionnaire. At that time the experimenter will provide you with any beverages you may require.

What you need to know about this experiment.

In this experiment, we are interested in how you make your spending and saving decisions over a 30 period 'lifetime'. You will make these decisions for money and for a liquid reward. The instructions will explain how the computer interface works. It will also explain how the decisions you make determine the amount of money you will earn. The money for the experiment has been provided by a research foundation. If you follow the instructions, and think carefully before making your decisions, you can earn a considerable amount of money. This will be paid to you in cash at the end of the experiment.

You will play 5 sequences of a 30 period spending/saving game.

There will be 5 sequences of a 30 period game. In each sequence you will make 30 decisions in a row. The third sequence will be different from the other four. It will be for a liquid reward rather than cash. This round will be explained later. You will receive a fixed amount of money for participating in this round. The point totals from the other four sequences will be calculated to determine the total amount you will earn at the end of the experiment. The game will be played on a Microsoft Excel workbook. Table A below shows an example of what the **first** period of **one non-liquid reward** sequence might look like. An explanation of the liquid reward sequence will be given at the end of these instructions.

Period	Expected Salary	Adjustment Factor	Actual Salary	Available Cash	Lifestyle Index	Spending Choice	Total Savings	Points Obtained	
1	**100.00**	**1.321**	**132.10**	**132.10**	**10.00**			**nil**	**Next Period**
2	105.00								
3	110.25								
4	115.76								
5	121.55								

Table A (First period)

The x-axis of the graph shows the possible adjustment factors (where 0 is the lowest possible factor). The y-axis shows how likely it is that an adjustment factor on the x-axis will actually randomly occur. Notice that the **most common** adjustment factors are **less than one** (because the curve is very tall for values on the x-axis between 0 and 1); but some of the adjustment factors are very large. This means that **most** of the time, your actual salary will be below your expected salary, but sometimes your actual salary will be much larger than your expected salary. In fact, half the time the adjustment factor will be below 0.607 and half the time the adjustment factor will be above 0.607. About 10% of the time the adjustment factor will be very low, 0.168 or less, and about 10% of the time it will be very high, 2.185 or above. (In case you are curious, we can tell you that the statistical distribution of the adjustment factor is generated by taking a "normal" or "bell curve" distribution, then taking the mathematical constant "e" (which is roughly 2.718) raised to a power equal to the number drawn from the bell curve distribution.)

Note also that each adjustment factor is statistically independent of the factors in early periods. This means that whether the factor is particularly high or low does not depend on whether it was high or low in the previous periods.

Table B below shows three <u>example</u> sequences of 30 adjustment factors (Sequence A, B and C), randomly drawn using the above-described distribution.

	Sequence A		Sequence B		Sequence C
Period	Adjustment Factors		Adjustment Factors		Adjustment Factors
1	1.364		0.845		0.624
2	0.461		2.464		2.660
3	0.498		0.403		2.643
4	0.223		0.199		1.298
5	0.323		0.413		0.840
6	0.108		0.296		0.389
7	0.283		0.199		0.530
8	0.588		0.926		2.592
9	4.793		1.989		0.599
10	0.780		1.601		1.246
11	2.721		0.230		0.674
12	0.334		1.270		0.159
13	2.203		0.715		1.586
14	1.363		0.404		0.129
15	0.289		0.100		0.471
16	0.194		0.170		0.309
17	0.369		0.426		0.364
18	1.296		0.604		0.703
19	0.256		0.248		1.120
20	0.308		1.033		0.219
21	0.767		1.441		0.780
22	0.671		0.910		0.049
23	0.578		0.198		0.486

24	0.956		1.665		0.446
25	2.000		1.636		0.265
26	1.782		0.174		0.549
27	0.140		0.482		0.276
28	0.384		0.342		0.406
29	0.087		0.929		0.457
30	1.692		1.625		0.367

Table B

Please note that these sequences of adjustment factors are only examples of what a sequence of 30 adjustment factors might look like; the actual sequences of adjustment factors you will get in your experiment will be different, even though the underlying probability distribution from which it was drawn is the same.

Period	Expected Salary	Adjustment Factor	Actual Salary	Available Cash	Lifestyle Index	Spending Choice	Total Savings	Points Obtained	
1	100.00	1.321	132.10	132.10	10.00	60.00	72.10	38.49	
2	**105.00**	**0.345**	**36.23**	**108.33**	**67.00**			**nil**	**Next Period**
3	110.25								
4	115.76								
5	121.55								

Table C (Second Period)

<u>Actual Salary</u> = Expected Salary multiplied by Adjustment Factor.

Each period, the actual salary is equal to the expected salary in that period times the adjustment factor. For example, in table C, the actual salary in period 1 is given by: 100 x 1.321 = 132.10. A low adjustment factor in period 2 (0.345) means that the actual salary in that period is only 36.23, which is much lower than the expected salary of 105.00. Keep in mind that these adjustment factors are just examples. When you participate in the experiment and make your own decisions, the adjustment factors will probably be different.

<u>Available Cash</u> = Last Period's Savings + Current Period's Actual Salary.

Remember that the one decision you must make in each period is how much of your available cash to spend. In table C, suppose you decide to spend 60.00 in Period 1. The total savings for period 1 is then your available cash (equal to actual salary because there was no past savings before period 1) minus your spending choice, which is 132.10 – 60.00 = 72.10. Please note that you <u>do not</u> earn interest on savings.
In period 2 of the table above, your actual salary is 36.23. Therefore, your available cash for period 2 is your savings left over from period 1, which was 72.10, plus your actual salary in period 2, which is 36.23. The total is 72.10 + 36.23 = 108.33, which will be automatically calculated for you and shown in the available cash column.

Spending earns you points. Make your <u>Spending Choice</u> in the yellow box.

Enter your spending choice each period in the yellow box. For your spending decision, the corresponding **Points Obtained** will be shown in the green box.

In table D below, entering a spending decision of 60.00 in period 2 will show that you can get 25.18 points for that period. You can try out different levels of spending before you make your final decision, by entering different values in the yellow spending choice box. Every time you input a value and press "enter" the computer will calculate how many points obtained you would get from that spending choice.

Period	Expected Salary	Adjustment Factor	Actual Salary	Available Cash	Lifestyle Index	Spending Choice	Total Savings	Points Obtained	
1	100.00	1.321	132.10	132.10	10.00	60.00	72.10	38.49	
2	**105.00**	**0.345**	**36.23**	**108.33**	**67.00**	**60.00**	**48.33**	**25.18**	**Next Period**
3	110.25				106.90				
4	115.76								
5	121.55								

Table D

The number of points obtained also depends on your Lifestyle Index.

For each level of spending, the number of points you earn is dependent on your lifestyle index. At a higher lifestyle index, you will earn a smaller amount of points for a given level of spending than when you are at a lower lifestyle index. In table D for example, the lifestyle index is at a higher level in period 2 (67.00) than in period 1 (10.00), therefore, the same level of spending of 60.00 yields a lower level of points in period 2.

A point transformation table is placed on your desk. Table E below shows part of this point transformation table. At each level of lifestyle index, it displays the number of points you can get at different levels of spending.

		Lifestyle Index							
		10	20	50	100	150	200	250	300
	5	-60.2	-190.3	-651.5	-1548.7	-2544.4	-3609.9	-4730.7	-5897.4
	10	3.2	-44.7	-214.2	-544.0	-910.0	-1301.7	-1713.7	-2142.6
	20	28.5	13.5	-39.6	-142.8	-257.4	-380.0	-508.9	-643.2
	40	36.7	32.5	17.5	-11.7	-44.0	-78.7	-115.1	-153.1
	60	**38.5**	36.5	**29.6**	16.0	1.0	-15.0	-31.9	-49.5
	80	39.1	38.0	34.0	26.2	17.6	8.4	-1.4	-11.5
Spending	100	39.4	38.7	36.1	31.1	25.5	19.5	13.2	6.6
Choice	120	39.6	39.1	37.3	33.7	29.8	25.6	21.2	16.6
	140	39.7	39.3	38.0	35.4	32.5	29.4	26.1	22.7
	160	39.8	39.5	38.5	36.4	34.2	31.8	29.3	26.7
	180	39.8	39.6	38.8	37.2	35.4	33.5	31.5	29.5
	200	39.9	39.7	39.0	37.7	36.3	34.7	33.1	31.4
	220	39.9	39.7	39.2	38.1	36.9	35.6	34.3	32.9
	240	39.9	39.8	39.3	38.4	37.4	36.3	35.2	34.0
	260	39.9	39.8	39.4	38.6	37.8	36.9	35.9	34.9

Table E (Point Transformation Table)

As you can see, a lifestyle index of 10 and a spending choice of 60 gives you 38.5 points. However, if the lifestyle index is 50 and you spend the same level of 60, the points obtained will be at a lower level of 29.6.

You are advised to look up this table before you make your spending choice. Alternatively, you can find out how many points you can earn by inputting different values of spending in the yellow box.
The lifestyle index grows with spending.

In general, the lifestyle index for a period is calculated by taking the value of the index from the previous period times .70, and adding in the previous period's spending. For example, in table D, the lifestyle index for period 2 is calculated as shown:
0.7 * 10.00 *(1st Period Lifestyle Index)* + 60.00 *(1st Period Spending)* = 67.00.
Likewise, if spending is again 60.00 in the second period, the lifestyle index for period 3 is: 0.7 * 67.00 *(2nd Period Lifestyle Index)* + 60.00 *(2nd Period Spending)* = 106.90
When you enter a spending level each period, the lifestyle index for the **next** period will be automatically calculated and shown.

A lifestyle conversion table is also provided on your desk. It shows you how your lifestyle index in the next period is dependent on how much you spend in the current period. Table F below shows <u>part</u> of this lifestyle conversion table.

		Lifestyle Index, Current Period						
		10	**20**	**50**	**100**	**150**	**200**	**250**
	10	17	24	45	80	115	150	185
	20	27	34	55	90	125	160	195
	40	47	54	75	110	145	180	215
	60	**67**	74	95	130	165	200	235
	80	87	94	115	150	185	220	255
	100	107	114	135	170	205	240	275
Spending Level, Current Period	**120**	127	134	155	190	225	260	295
	140	**147**	154	175	210	245	280	315
	160	167	174	195	230	265	300	335
	180	187	194	215	250	285	320	355
	200	207	214	235	270	305	340	375
	220	227	234	255	290	325	360	395
	240	247	254	275	310	345	380	415
	260	267	274	295	330	365	400	435
	280	287	294	315	350	385	420	455
	300	307	314	335	370	405	440	475
	320	327	334	355	390	425	460	495

Table F (Lifestyle Conversion Table)

As you can see, choosing a spending level of 60, when lifestyle index is 10, will result in a lifestyle index in the **next period** of 67. If you decide to spend more, for example 140, then your lifestyle index for the **next period** will be at a higher level of 147.

Note that the more you spend in the current period the higher your lifestyle index will be in future periods. The **point transformation table** (table E) shows that for any particular level of spending, you earn fewer points if the lifestyle index is higher. So if you spend a lot in early periods, you will receive many points in those periods, but you also increase the lifestyle index for future periods, which will then reduce the points you obtain in future periods.

You cannot spend more than your available cash.

Each period, you are not able to spend more than the available cash you have. If you choose a spending level greater than the cash you have, the program will tell you to lower your spending.

Proceed to the next period when you have made your spending choice.

Once you have thought carefully about how much to spend each period, proceed to the next period by using your mouse to **click once** on the pink box labelled 'Next Period'. **Please note that the program prevents you from returning to earlier periods to change your spending choice.** Therefore, please be careful not to click the 'Next Period' box before you enter your spending decision, because you will not be able to return to change it.

Once you have completed each 30 period sequence, proceed to the next sequence of 30 periods by clicking the 'Continue' link, which will appear at the bottom right of your screen.

Please note that the sequence of adjustment factors will be **different** in each of the 5 sequences, but the overall statistical distribution of possible adjustment factors will be the same. Once you have completed all 5 sequences, a screen will appear to tell you your overall points obtained from all 5 sequences.

The computer will automatically spend <u>all</u> available cash in the <u>last</u> period of each sequence.

Available cash from one sequence will <u>not</u> be carried over to the next sequence. This means that the computer will be automatically spend <u>all</u> remaining available cash in period 30 of each sequence.

How your earnings are determined:

After you make your spending choice each period, the points you obtain that period, in addition to all points you obtain in previous periods will be tallied at the bottom of the screen. Some of the point outcomes each period will be negative but your total points from each sequence should be positive.

The total points you obtained from your four non liquid reward rounds will be calculated and will be converted to cash at a rate:
14 points = $0.03
466.67 points = $1
Your earnings from the money rounds, in addition to the **$5.00** show-up payment and a **$25.00** fixed payment for the liquid round, will be paid to you in cash when you leave the laboratory. They will be rounded up to the nearest dollar.

For the third sequence you will make decisions for a liquid reward.

The third sequence will be identical to the other sequences except your point totals will be converted to mL of liquid and dispensed in a cup to your right. Each mL of liquid will be equal to two points. Table G shows the first period of the liquid reward round:

Period	Expected Salary	Adjustment Factor	Actual Salary	Available Cash	Lifestyle Index	Spending Choice	Total Savings	Liquid Max (mL)	Liquid (mL)	
1	100.00	1.321	132.10	132.10	10.00			19.84	Nil	Next Period
2	105.00									
3	110.25									
4	115.76									
5	121.55									

The first eight columns are identical to those in the monetary reward rounds, and were explained in the instructions previously. 'Liquid Max (mL)' is the maximum amount of liquid that can be delivered at the end of the period. To achieve this amount you must spend all your available cash. Liquid (mL) is the actual amount of liquid that will be delivered to you at the end of the period. It is equal to half of 'Total Points' from previous rounds.

After you have made your spending choice and clicked 'Next Period' the program will ask you to close Microsoft Excel. The liquid reward will be delivered to you at the amount specified under the liquid column. You then will have a 60 second break, before the program opens. You must consume all of the liquid reward during this break. You cannot save it for future periods.

If you sustain a negative result for any round, you will not receive liquid until you have produced enough positive periods to offset that result. For example, if you sustain a negative liquid mL total of -20 in period 8, 10 in period 9, and 15 in period 10, you will receive no liquid reward in period 8 or 9 but 5 mL of liquid reward in period 10.

Please note that if the liquid reward is a carbonated beverage, the volume of liquid may be slightly different than the value in the excel spreadsheet. This is due to the carbonation gas being measured as liquid in the syringe pump. It is unavoidable when using carbonated beverages.

After you have made decisions for thirty periods and received your liquid rewards, there will be an additional ten dummy periods. In these periods you will not make a decision or receive any liquid reward. They are a necessary part of this experiment and we appreciate your patience in this matter. These ten dummy periods should last about ten minutes.

Here is a brief summary of what you need to know.

You will be making decisions in 5 sequences of 30 periods. In each period you will have some available cash and will choose a level of spending. In the third sequence your decisions will determine the liquid rewards you receive over that round. All of the other sequences are important in determining your overall cash earnings, because your earnings will depend on the point total of sequences 1, 2, 4 and 5.

Expected salary grows at 5% each period. The actual salary that you get depends on a random adjustment factor that occurs during each period. These factors are randomly determined and the adjustment factor in one period does not depend on whether the previous period's adjustment factor was high or low. The available cash you have during each period is the actual salary you get in the current period plus the level of savings that was left over from the previous period.

The level of points (mL of liquid reward) you can get during each period depends on the level of spending you make, as well as your lifestyle index. More spending this period increases the lifestyle index for next period. A higher level of lifestyle requires a higher level of spending than before to obtain the same level of points. The point transformation table on your desk will give you a better idea on how this works.

Take as much time as you like to make your spending decision in each period. Please note that your spending level in each period cannot exceed the available cash you have. Remember that you cannot go back to earlier periods to change your spending level once you have clicked on the 'Next Period' box. Therefore, please make sure that you have correctly entered your final spending decision in the yellow spending choice box before proceeding to the next period.

The total points you have obtained for all four non-liquid reward sequences will be calculated and converted to cash. In the liquid reward round your decisions will determine how many mL of liquid you will receive between periods. Remember that the liquid reward round is very similar to the other rounds. The mL of liquid you receive at the end of each period is equivalent to half your period's points total in any other round.

If these instructions were not clear to you, or you have a question of any sort, please tell the experimenter now.

If you don't have any questions, please attempt the short quiz on the following page before you start the experiment. These questions will test whether you have fully understood the instructions. Once you are done with the questions, the experimenter will come by to check your answers. **If your answers are not right, the experimenter will give the correct answer and help you understand how the tables and instructions should enable you to give the correct answers.**

You can only start the experiment when all your answers are correct.

Quiz

Answer questions 1-7 for non-liquid reward sequences. (Sequences 1,2,4,5).

1) If you spend 60.00 this period, and your lifestyle index is 50.00, how many points will you obtain?

Ans: ________________

2) If you spend 80.00 this period, and your lifestyle index is 250, how many points will you obtain?

Ans: ________________

3) If you spend 450.00 this period, and your lifestyle index is 700, how many points will you obtain?

Ans: _________________

4) If you increase your spending level from 60.00 to 100.00, and your lifestyle index is 100.00, how many additional points will you get?

Ans: ________________

5) Your expected salary in period 2 is 150.00. The adjustment factor is 0.500 in the same period. Total savings from period 1 was 40.00. How much available cash do you have in period 2?

Ans: ________________

6) Your lifestyle index is 50 in period 1. If you decide to spend 60.00 in the same period, what would be the level of lifestyle index in period 2?

Ans: ________________

7) In period 20, your lifestyle index is 200.00. You decide to spend 120.00.
   a) How many points will you get?
   b) What will your lifestyle index be in period 21?

Ans: ____________________________________

8) Answer questions 1-4 again except assume they have been asked for the liquid reward sequence. That is substitute 'mL liquid reward' for 'points'.

Ans: ________________
Ans: ________________
Ans: ________________
Ans: ________________

Δ	Time preference factor (assumed constant)
X_S	Total cash/resources available in period *s* ("cash-on-hand")
S_S	Savings in period *s* (portion of X_s not consumed)
C_S	Spending in period *s*
R	Gross interest rate each period
H_{S-1}	Habit stock from period *s-1*
$u(C_S, H_{S-1})$	Utility
Y_S	Actual income in period *s*
P_S	Permanent labor income in period *s*
$P_{s+1} = G P_s$	*G* is the growth rate of permanent income each period (assumed constant)
η_s	Income shock in period, a random variable, drawn from a distribution

Table A.1: Notation for the buffer stock model

A.3 The Experimental Problem in Formalized Detail

This section repeats the description of the experimental problem, and theoretical predictions of the exponential and quasi-hyperbolic models in much greater detail. The text should provide enough description for the reader to understand the predictions of the models, but this section adds greater depth to that understanding and includes similar predictions of the dual-self model.

The experimental design implements the assumptions of the buffer stock savings model of Carroll et al. (2000). Agents earn income each period, subject to stochastic independent shocks from a distribution they know. In each period they choose how to divide their available cash the previous buffer stock, plus new income—between spending and savings. Utility in each period depends upon a ratio of current consumption to a habit index. The habit index is a depreciated sum of previous consumption. Agents should maximize the discounted utility from consumption over the remainder of their lifetimes, which is a dynamic programming problem. The variables in this dynamic program are listed in Table A.1.

A.3.1 Exponential Discounting

Assuming exponential discounting of future rewards, the consumer's maximization problem is

$$\max E_t \left[\sum_{s=t}^{T} \Delta^{s-t} u(\tilde{C}_s, \tilde{H}_s \right]. \tag{A.1}$$

The utility function incorporates constant relative risk-aversion (CRRA) and habit formation as follows:

$$u(C_t, H_{t-1}) + k + \frac{\theta}{1-\rho} \left(\frac{C_t + \hat{\epsilon}}{H_{t-1}^{\gamma}} \right)^{1-\rho}. \tag{A.2}$$

The parameter ρ is the coefficient of relative risk aversion, and γ determines how strongly previous habitual consumption affects current utility (e.g., if $\gamma = 0$ there is no effect of habit).[1]

The habit stock is given by $H_t = \lambda H_{t_1} + C_t$ where λ is a depreciation rate (as in Fehr and Zych, 1998). Actual income each period is equal to permanent income multiplied by an income shock, $Y_t = P_t \eta_t$, where η_t is a random variable drawn from a distribution each period. The value function for cumulative future utility, at period t, depends on three state variables: permanent income P_t, habit H_{t-1}, and available savings X_t. The optimal value function is

$$V_t(P_t, X_t, H_{t-1}) = \max_{c_t} \left\{ u(C_t, H_{t-1}) + \Delta E_t \left[V_{t+1}(P_{t+1}, X_{t+1}, H_t) \right] \right\} \tag{A.3}$$

or, writing out the state variables,

$$V_t(P_t, X_t, H_{t-1}) = \max_{c_t} \left\{ u(C_t, H_{t-1}) + \Delta E_t \left[V_{t+1}(GP_t, R(X_t - C_t) + \eta_{n+t}, GP_t, \lambda H_{t-1} + C_t) \right] \right\} \tag{A.4}$$

subject to constraints

$$S_t = X_t - C_t, S_t \geq 0 \tag{A.5}$$

$$X_{t+1} = S_t + \eta_{n+1} P_{t+1} \tag{A.6}$$

$$H_t = \lambda H_{t-1} + C_t. \tag{A.7}$$

To make the problem easier to solve computationally, the state variable P_t can be eliminated by

[1] Since $\rho = 3$, the term k is the upper asymptote of utility, θ is a scaling parameter, and $\hat{\epsilon}$ bounds the utility function from below. In the experiments, $\hat{\epsilon} = 2.7$, similar to Ballinger et al. (2003).

normalizing each variable by permanent income.[2]

In the experiment, most parameter values were roughly calibrated to those measured in actual savings data from the U.S. Carroll (1992) found income shocks η_t to be lognormally distributed with a mean value of one and a standard deviation of 0.2. We use η_t drawn from a lognormal distribution. We pick $\sigma = 1$ rather than .2, to create more income variation. This makes the problem more challenging for subjects, and therefore gives a better chance of observing a range of conditions under which performance is very bad or surprisingly good. Permanent income grows according to $P_{t+1} = (1.05)P_t$ with $P_t = 100$. The discount factor and gross interest rate are both set equal to one $(\Delta = 1, R = 1)$. The risk-aversion coefficient is $\rho = 3$, an estimate often used in consumption studies which seems to fit many types of aggregate data. For habit formation, $\gamma = 0.6$, depreciation $\lambda = 0.7$, and the initial habit is $H_0 = 10$. There are thirty periods in this experiment so $T = 30$.

Exponential discounting is dynamically consistent—the current tradeoff between two future points which is reflected in current decisions is preserved when those future points eventually arrive. We discuss two approaches—β-δ quasi-hyperbolic discounting, and a dual-self model in which a foresightful "planner" tries to restrain a myopic "doer" from spending too much.

A.3.2 Quasi-hyperbolic Discounting

In this section we show how these models work and contrast them with the exponential model. These models are also of special interest in study 2, which uses beverage rewards rather than money (as in study 1). In study 2, in one lifecycle thirsty subjects earn the same number of utility points as in the study with money, but x points are converted into $x/2$ milliliters of cola. In the immediate condition they drink the cola right away. In the delayed condition they "order in advance," so that spending decisions in period t determine the amount of cola that can be drunk in period $t + 10$. Dynamically consistent subjects should make the same decisions in these immediate and delayed conditions. However, under β-δ discounting or dual-self models, subjects may "spend" more (i.e., earn more points which are converted to beverage) in the immediate condition than in the delayed condition. For notational simplicity define ω,

$$\omega(\tilde{C_t}, \tilde{H_{t-1}}) = v\left(\frac{\tilde{C_t}, \tilde{H_{t-1}}}{2}\right) \tag{A.8}$$

[2]That is, $x_t = X_t/P_t, h_t = H_t/P_{t-1}, and \epsilon_t = eps\hat{i}lon/P_t$

where $\omega(\tilde{C}_t, \tilde{H_{t-1}})$ is the beverage reward (in terms of a beverage utility function v) associated with consumption decision $\tilde{C}_t$ and habit $\tilde{H}_{t-1}$.

In the quasi-hyperbolic discounting (or present-bias) model, current utilities have a weight of one, and utilities t periods in the future $(t > 1)$ have a weight of $\tilde{\beta}\tilde{\delta}^t$.

In the β-δ, the implicit tradeoff between future periods is *not* necessarily the same as the tradeoff which is made when the future arrives. Therefore, the model requires a behavioral assumption about whether current agents are "naïve" or "sophisticated" about their own future behavior.

Naïve agents believe—incorrectly—that the weights they currently apply to future periods are the same as the (relative) weights they will apply when the future arrives. Intuitively, even though all future periods are discounted by a present bias β, the naïve discounter believes that in future evaluations there will no such present bias. This model corresponds to chronic procrastination which is justified by the hope that starting tomorrow, the activity that has been put off for so long will finally get done.

Sophisticated discounters have a present bias, but correctly realize that they will have a present bias in the future, too. A crucial difference is that sophisticated discounters will seek external commitment devices (to restrain the present bias they know they will have) while naïve discounters do not. Our view is that it is too early in the empirical literature to consider only one model, when both can be considered and compared, so we develop both here and calibrate them on the experimental data in Section 2.5.

A.3.2.1 Sophistication

Similiar to the design of Harris and Laibson (2001), optimal consumption can be determined by backward induction because we have a finite number of periods T. In the last period the subject will solve

$$\ddot{V}_T(X_T, H_{T_1}, P_T) = \omega(C'_T, \tilde{H}_{T-1}) \tag{A.9}$$

where

$$\ddot{V}_T(X_T, H_{T_1}, P_T) = \omega(C'_T, \tilde{H}_{T-1}). \tag{A.10}$$

Assuming sophistication, the optimization problem can then be solved recursively using equations A.11 and A.12.

$$C'_t = \operatorname*{argmax}_{c_t} \omega(C_t, \tilde{H}_{t-1}) + \tilde{\beta}\tilde{\delta}E_t\left[\ddot{V}_{t+1}(GP_t, X_t - C_t + \eta_{t+1}GP_t, \lambda H_{t-1} + c_t)\right] \tag{A.11}$$

$$\ddot{V}(P_t, X_t, H_{t-1}) = \omega(C_t, \tilde{H}_{t-1}) + \tilde{\beta}\tilde{\delta}E_t\left[\ddot{V}_{t+1}(GP_t, X_t - C_t + \eta_{t+1}GP_t, \lambda H_{t-1} + c_t)\right] \quad \text{(A.12)}$$

Notice that the function $\ddot{V}_t$ is different than a typical dynamic programming value function as in equation A.3. In that equation the value function is the maximum of the current consumption utility plus the (discounted) continuation value function conditioned on that level of consumption. But the possibility of dynamic inconsistency requires us to create a pseudo-value function $\ddot{V}_t$ instead. Equation A.11 dictates that a sophisticated agent will maximize utility consistent with her present preferences. However, the sophisticated hyperbolic knows that in the future she will not apply the same weights and make the same tradeoffs, so she needs a way to keep track of consumption utilities in future periods without aggregating them into a conventional value function. Here, $\ddot{V}_t$ is a pseudo-value function which is simply a sum of future utilities from consumption, discounted at the exponential rate $\tilde{\delta}$.

A.3.2.2 Naïveté

A naïve agent believes that her future decisions will be made as if she is an exponential discounter.[3] A naïve agent therefore creates a value function $\breve{V}_t(P_t, X_t, Ht_1)$ which exhibits present bias but uses the exponential value function $\breve{V}_t(P_t, X_t, Ht_1)$ (a modified version of equation A.3) with $\Delta = \tilde{\delta}$ and $u = \omega$ in forecasting future utilities.

$$C_t^*(P_t, X_t, H_{t-1}) = \underset{C_t}{\operatorname{argmax}}\, \omega(C_t, H_{t-1}) + \tilde{\beta}\tilde{\delta}E_t\left[\breve{V}_{t+1}(GP_t, X_t - C + t + \eta_{t+1}GP_t, \lambda Ht_1 + C_t)\right] \quad \text{(A.13)}$$

where

$$\breve{V}_t(P_t, X_t, Ht_1) = \max_{C_t} \omega(C_t, H_{t-1}) + \tilde{\delta}\left[\breve{V}_{t+1}(GP_t, X_t - C + t + \eta_{t+1}GP_t, \lambda Ht_1 + C_t)\right]. \quad \text{(A.14)}$$

Figure A.2 shows an example consumption path which compares sophisticated and naïve hyperbolic consumption paths for $\delta = 0.9$ and $\beta = 0.8$, compared to the optimal path (with $\delta = \beta = 1$) from Figure 2.2.

As O'Donoghue and Rabin (1999) have stressed, present bias and sophistication can interact in interesting ways. Generally, a naïve person exhibits more present bias than a sophisticated one. However, a sophisticated person who is sufficiently present-biased can succumb to temptation im-

[3] The reason is that her current weights on all future periods (for $t > 1$), $\tilde{\beta}\tilde{\delta}^t$, imply relative tradeoffs in future periods in which the $\tilde{\beta}$ terms divide out for optimization.

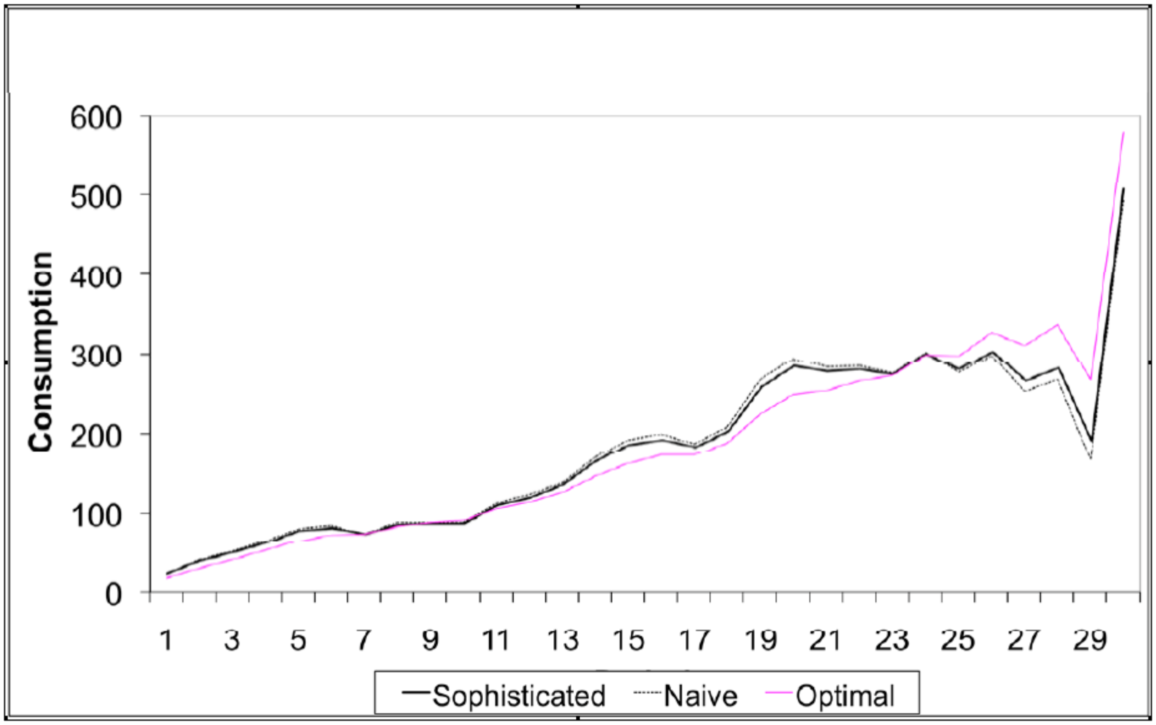

Figure A.2: Quasi-hyperbolic consumption path for naïve and sophisticated cases ($\beta = 0.8, \delta = 0.9$)

mediately because she knows her future self will too, while a naïve person might postpone temptation because she thinks in the future she will be more patient than she currently is. In the model, as A.2 indicates both sophisticated and naïve hyperbolic discounting cause an individual to overconsume relative to optimal in this experiment. The difference between the two paths is very small [as in Angeletos et al, 2002], but the naïve consumer does consume a little more than the sophisticated one in early periods.

A.3.2.3 Immediate and Delayed Beverage Rewards

To compare the immediate and delayed beverage reward conditions, we first assume that the utility of beverage is linear in volume. Then $u(C_T, H_{t-1} = \tilde{v}\left(\frac{u(C_t,H_{t-1})}{2}\right) = \tilde{\omega}(C_t, H_{t-1})$. For simplicity, we also assume that subjects do not satiate in beverage, and utilities are additively separable across periods. Even if these assumptions do not hold, there is no reason to think that they are violated more or less in the two conditions (immediate and delayed). In the delayed condition subjects do not receive the beverage amount they decided upon in period t until period $t + 10$. Since there is no immediate reward, all future consumption has a weight of $\tilde{\beta}$ (along with discount factors $\tilde{\delta}$) and the $\tilde{\beta}$ terms divide out in optimization. Then each subject will solve:

$$V_t'(P_t, X_t, H_{t-1}) = \max_{C_t} \{u(C_t, H_{t-1}) + \Delta E_t \left[V_{t+1}(GP_t, X_t - C + t + \eta_{t+1} GP_t, \lambda H t_1 + C_t)\right]\} \tag{A.15}$$

subject to constraints A.5, A.6, and A.7. Notice that $V'(P_t, X_t, H_{t-1}) = \frac{1}{2} V_t(P_t, X_t, H_{t_1})$ (from equation A.3) if $\Delta = \tilde{\delta}$. In this case both value functions will have the same optimal consumption path. If we assume $\tilde{\beta} < 1$ (present bias), subjects in the immediate condition will weigh payoffs in the earlier periods more heavily than subjects in the delayed condition weigh them, and will choose to consume more beverage in early periods; this spending over the total beverage-maximizing level will be greater in the earlier periods of the immediate condition. As a result, $\tilde{\beta} < 1$ predicts that subjects in the immediate condition will receive a lower beverage total because they are deviating from the optimal beverage maximizing total.

A.3.3 A Dual-Self Planner-Doer Model

Another way to model dynamic inconsistency is by positing two systems, or a "dual self," which interact to create behavior (see Thaler and Shefrin, 1981; Bernheim and Rangel, 2004; Loewenstein and O'Donoghue, 2004). For brevity, we focus on just one of these models, the Fudenberg and Levine (2006) approach. They assume long-run and short-run players, much as in Thaler and Shefrin's earlier "planner-doer" model. For consistency with the hyperbolic discounting model, assume that the long-run player (L) has a linear discount factor for beverage, $\tilde{\delta}$. L also knows that the short-run player (S) will spend all resources in a given period if L does not exercise self-control. In equilibrium L will choose a strategy from histories $m \in M$ and states to actions $\sigma_{SC} : M \times Y \to A$ to maximize the following reduced form objective function:

$$U_t = \sum_{s=t}^{T} \tilde{\delta}^{t-s} \int \left[\omega(y, 0, a) - q(y, a)\right] d\pi_t(y(m)). \tag{A.16}$$

In our application the only actions are spending decisions, so a is replaced by $\tilde{C}_t$, and the current state is defined as $y = (P_t, X_t, H_{t-1})$. The function $q(y, a)$ is the self-control cost of the long-run player enforcing spending $\tilde{C}_t$ in state y. The function π_t is the measure associated with histories of previous short-run actions and a given state. Histories are irrelevant to the long-run player (except as summarized by the state variables); only the probabilities of other states are relevant. Thus equation

A.16 can be rewritten as (A.17),

$$U_t = \sum_{s=t}^{T} \tilde{\delta}^{t-s} \int \left[\omega(P_t, X_t, H_{t-1}, 0, \tilde{C}_t) - q(y, a)\right] d\mu_((P_t, X_t, H_{t-1}) \tag{A.17}$$

where $d\mu_((P_t, X_t, H_{t-1})$ is the probability measures of the states. If the subject is an expected utility maximizer, than her problem becomes very similar to the hyperbolic case. She will solve

$$\max_{\tilde{C}_t} \sum_{s=t}^{T} \tilde{\delta}^{t-s} \left[\omega(P_t, X_t, H_{t-1}, 0, \tilde{C}_t) - q(y, a)\right]. \tag{A.18}$$

In the delayed condition the short-run player has no control over how much utility she receives because the decision determining her current utility was made ten periods ago. According to assumption 4 of Fudenberg and Levine (2006) this means L can make all choices without exerting a self-control cost. So her optimization procedure will be identical to the traditional exponential-discounting case. L will maximize

$$\max E_t \left[\sum_{s=t}^{T} \tilde{\delta}^{s-t} \omega(\tilde{C}_s, \tilde{C}_{s-1})\right]. \tag{A.19}$$

The optimal value function will be

$$V_t'(P_t, X_t, H_{t-1}) = \max_{C_t} \{u(C_t, H_{t-1}) + \Delta E_t \left[V_{t+1}(GP_t, X_t - C + t + \eta_{t+1} GP_t, \lambda H t_1 + C_t)\right]\} \tag{A.20}$$

subject to constraints A.5, A.6, and A.7 as before. In the immediate condition, if there are positive self-control costs in restraining S's spending in each period, the subject will consume more than is optimal (since S's myopic ideal is to consume everything). The implication is that subjects in the immediate condition will spend more than is optimal in early periods, and will therefore spend more than the delayed-condition subjects (who, by assumption, optimize). Hence the planner-doer and hyperbolic discount models both predict more early consumption, and less overall consumption, in the immediate condition compared to the delayed condition. To make a more precise prediction (and comparison between theories) requires a detailed specification of the utility costs of self-control, which is an important topic that lies beyond the scope of this paper.

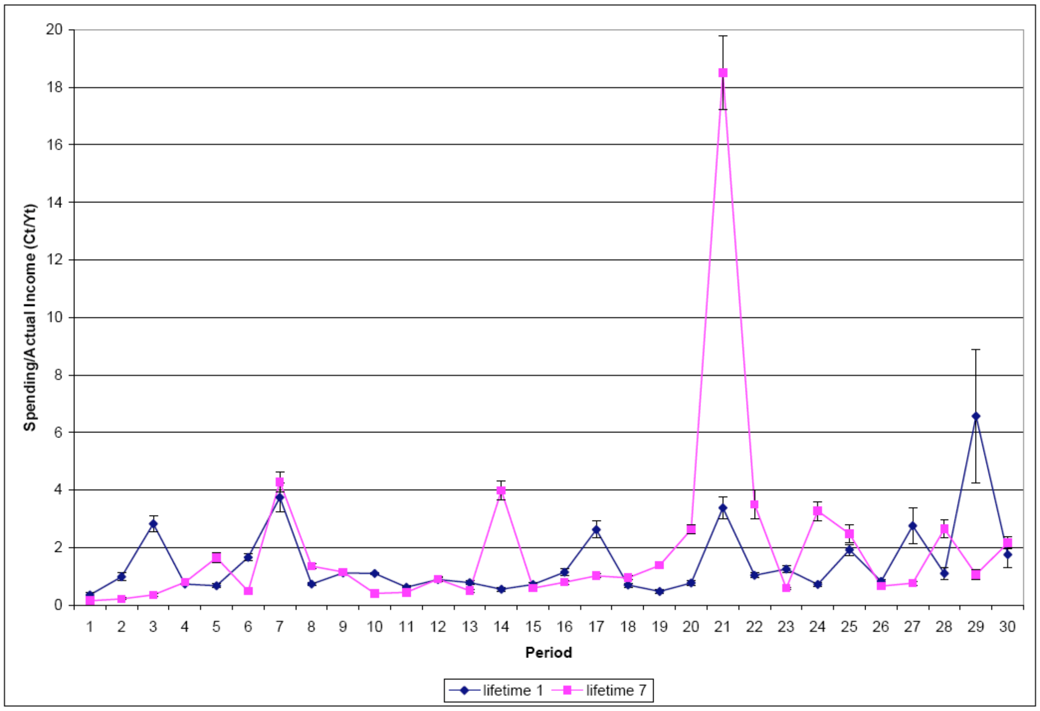

Figure A.3: Spending as a proportion of actual income each period, lifetimes 1 & 7, private social learning condition ($N = 36$)

A.4 Rules of Thumb: Spending as a Proportion of Cash-on-Hand

These figures show that simple rule-of-thumb models of constant consumption across periods are badly rejected by actual behavior. Figures A.3 and A.4 show that consumption as a proportion of current income fluctuates a lot across periods. Figures A.5 and A.6 show that consumption as a proportion of cash-on-hand fluctuates a lot, as well. However, those figures show that learning could conceivably be modeled as a shift in the consumption function from one that exhibits little trend across periods (in lifetime 1) to one which is clearly increasing across periods and jumps up sharply in the last couple of periods. (Keep in mind that the software automatically spends all available cash-on-hand in the last period, so the consumption/cash-on-hand ratio is necessarily equal to 1 in the the last period 30.)

A.5 Myopic Loss Aversion

A widely used concept in behavioral economics which might apply here is myopic loss-aversion. Loss-aversion is the idea that people are disproportionately averse to making decisions that create

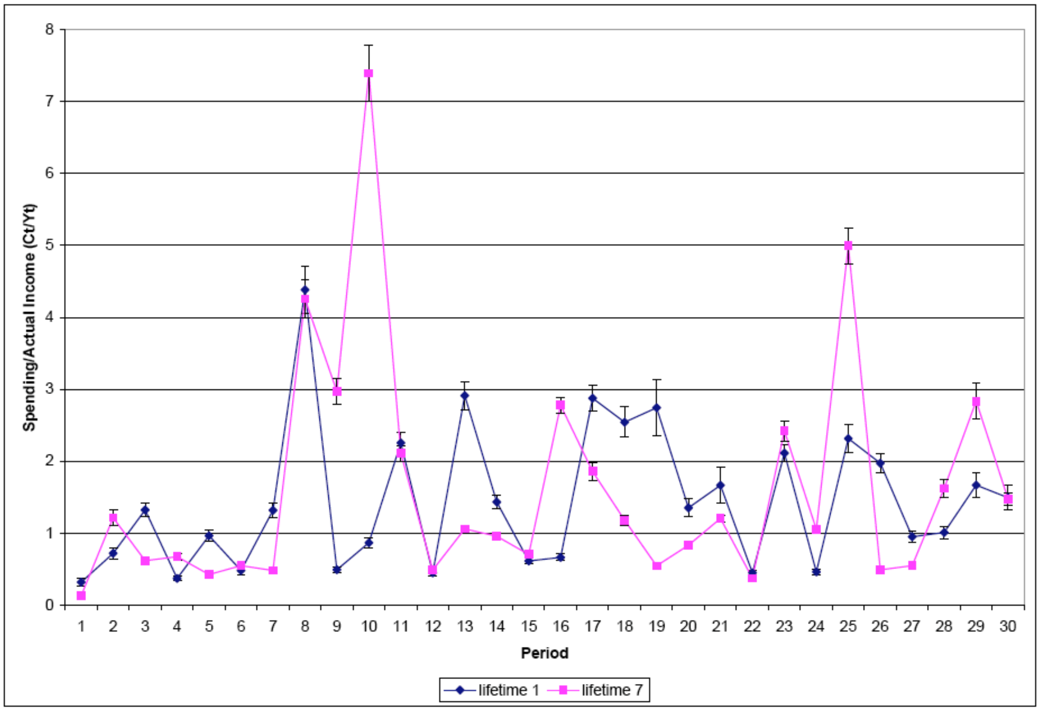

Figure A.4: Spending as a proportion of actual income each period, lifetimes 1 & 7, social learning condition $(N = 36)$

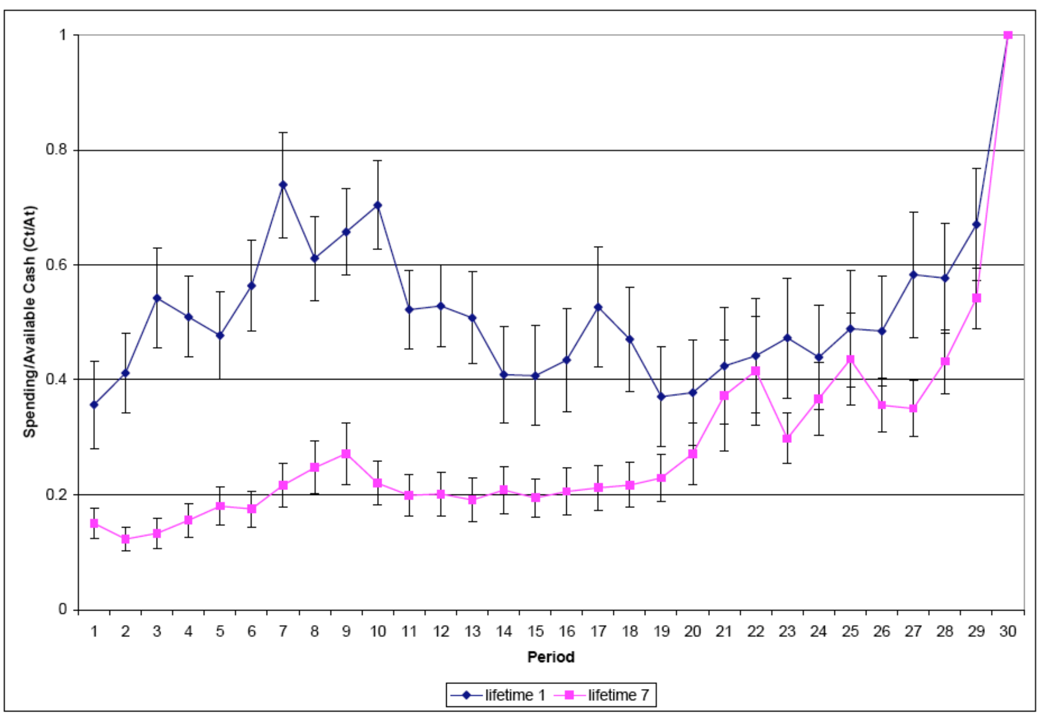

Figure A.5: Proportion of available cash spent each period, lifetimes 1 & 7, private learning condition $(N = 36)$

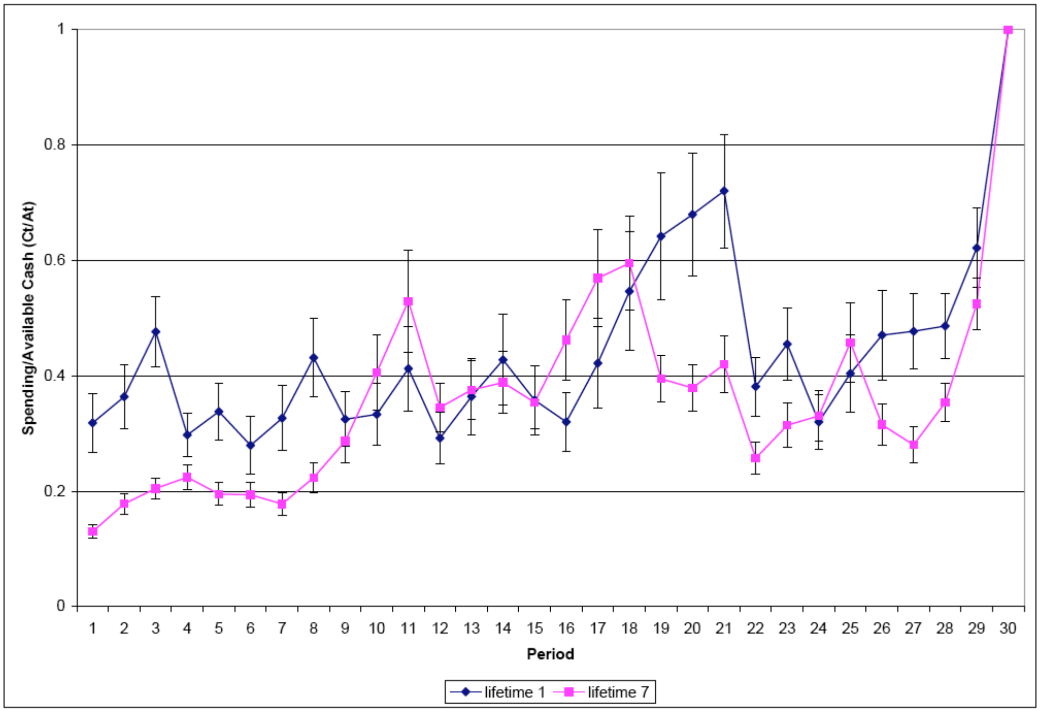

Figure A.6: Proportion of available cash spent each period, lifetimes 1 & 7, private learning condition $(N = 36)$

nominal losses relative to a point of reference (see Kahneman and Tversky, 1979; Camerer, 2005, 2006; Goldstein et al., 2006). Myopic loss-aversion means that people focus on losses only in a small segment of time or a part of a portfolio, neglecting the benefits of decision rules which aggregate losses and gains across choice sets. In our setting, one hypothesis from myopic loss-aversion is that subjects will be unusually reluctant to choose a consumption level that generates a period-specific utility which is negative (assuming zero is a reference point).[4] Figure A.7 tests this hypothesis using study 1 data, by plotting nominal utility losses in each period from actual consumption on the y-axis, and corresponding losses that would have resulted from conditionally-optimal consumption (for utilities between -50 and +50, $n = 14,228$). There is a sharp visible drop-off in points between the bottom and top halves of the Figure A.7 scatter plot. It appears that subjects hate to lose a small amount of nominal utility, even when they should take a small loss to build up savings (as shown in Fehr and Zych, forthcoming).[5] A piecewise-linear jackknife regression through the origin gives coefficients in the domain of positive and negative conditionally

[4]Subjects sometimes input a series of consumption levels, trying to find the value that would give a positive utility. Unfortunately, the software did not capture these attempts; data like these would be useful to understand the nature of loss-aversion and its persistence.

[5]The result is reminiscent of DeGeorge et al. (1999) finding that small negative earnings announcements, and small year-to-year drops, are relatively rare for corporations.

optimal utilities of 0.79 and 0.15, respectively. The ratio of these two slopes is 5.2. Another way to see the effect is to plot histograms of small actual and optimal utilities (between -10 and +10), across all subjects and periods (see Figure A.8). In the actual period-by-period utilities there is a huge spike on the slight positive side. This spike is not at all evident in the corresponding distribution of optimal utilities, so the preference for a small positive utility and aversion to loss is not normative.

Figures A.9 and A.10 show actual-optimal utility scatter plots like FigureA.7, for the money periods and beverage periods of study 2. For money (Figure A.9), the jackknife regression gives positive and negative slopes of 0.88 and 0.10 (a ratio of 8.8). For beverage (Figure A.9), the slopes are 0.92 and 0.62 (a ratio of 1.49). The difference between money and beverage is consistent with the idea that in the domain of beverage, subjects know that a large loss creates a debt that they must pay off before getting more sips, so they are more willing to accept small losses rather than run up large debts. Johnson et al. (2006) also show variations in loss-aversion across domains.

The myopia underlying Figures A.7–A.10 is surprising. The subjects make 210 separate money decisions in study 1, and 120 decisions in study 2. They know that the utilities in each of those periods will be added up at the end to determine their total money earnings. (The software even updates the total points for each lifecycle every period and shows the total at the bottom of the screen.) There is no good normative reason to avoid a small loss in any single period (as Figure A.8 showed). These data are a reminder that a complete theory of theory of loss-aversion and its interaction with a myopic focus needs to account for how broadly decisions are bracketed or lumped together (Thaler, 1999; Read et al., 1999).

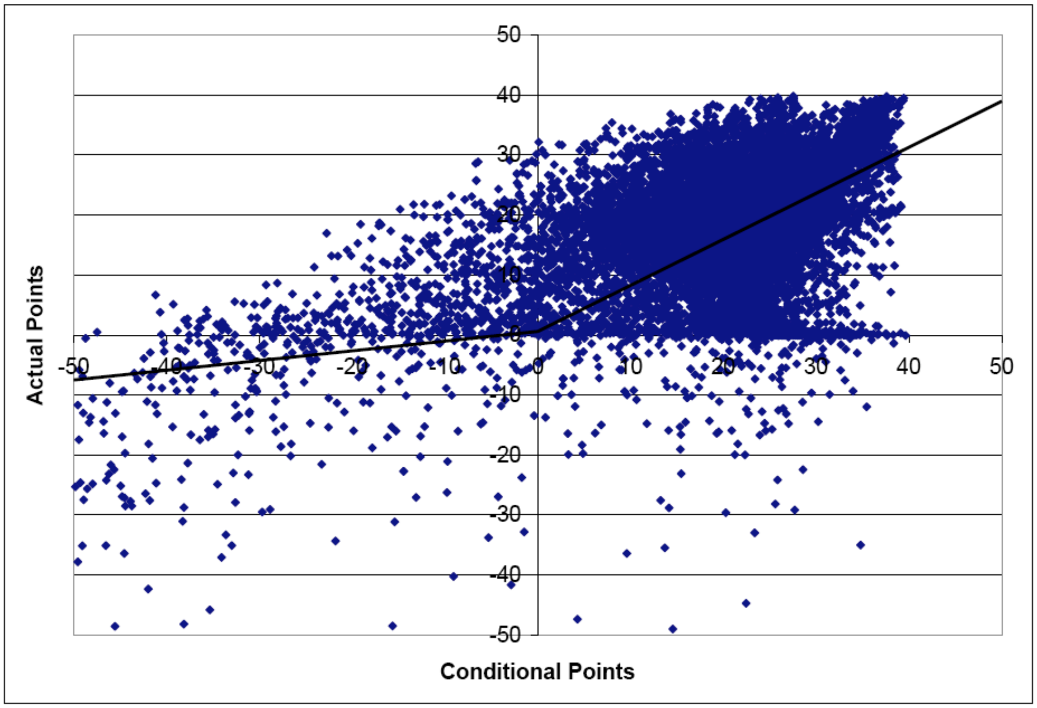

Figure A.7: Actual (y) and conditionally optimal (x) utilities, study 1, observations between -50 and +50 $(N = 14,228)$

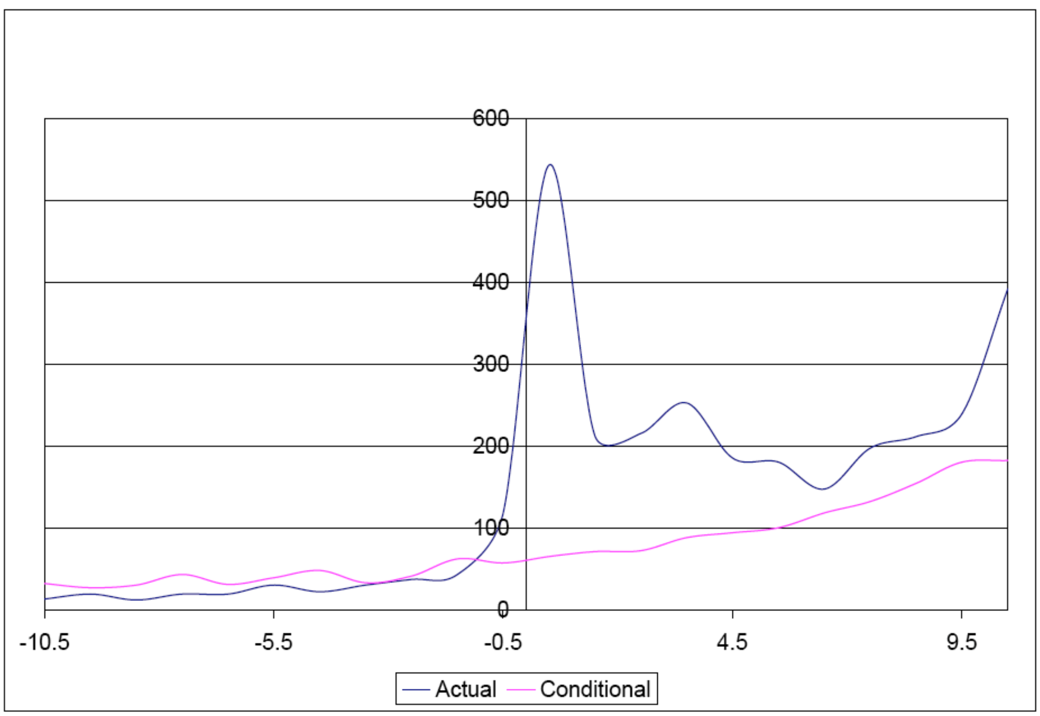

Figure A.8: Frequency of actual and conditionally optimal utilities, study 1, observations between -10 and +10 $(N = 14,228)$

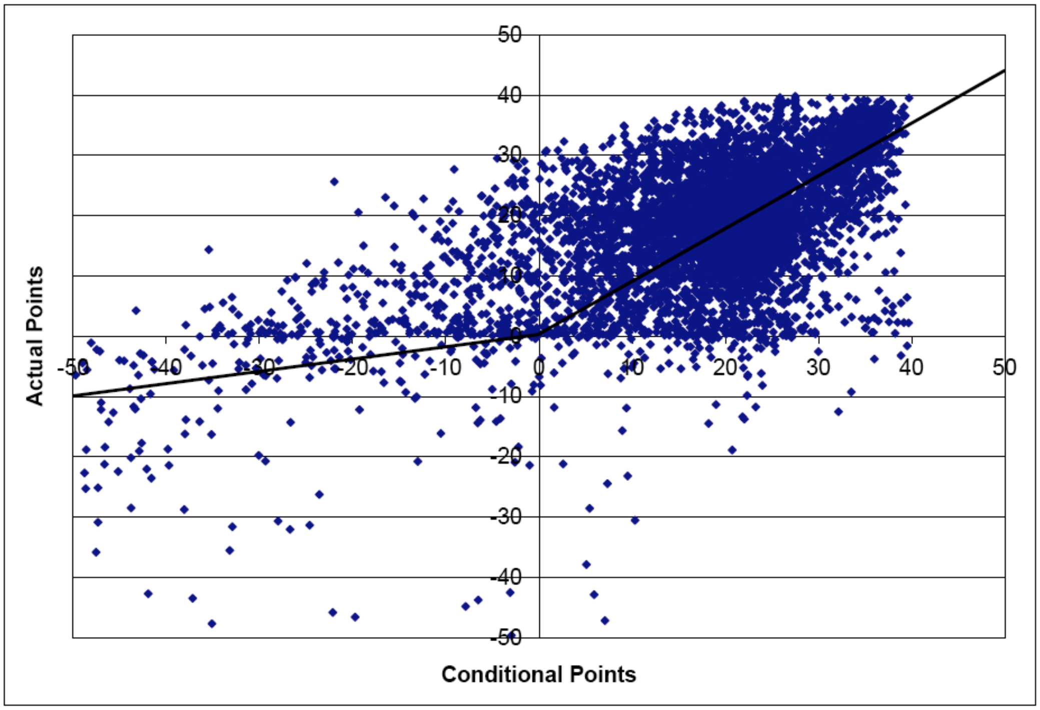

Figure A.9: Actual (y) and conditionally optimal (x) utilities, study 2, observations between -50 and +50 $(N = 5,840)$

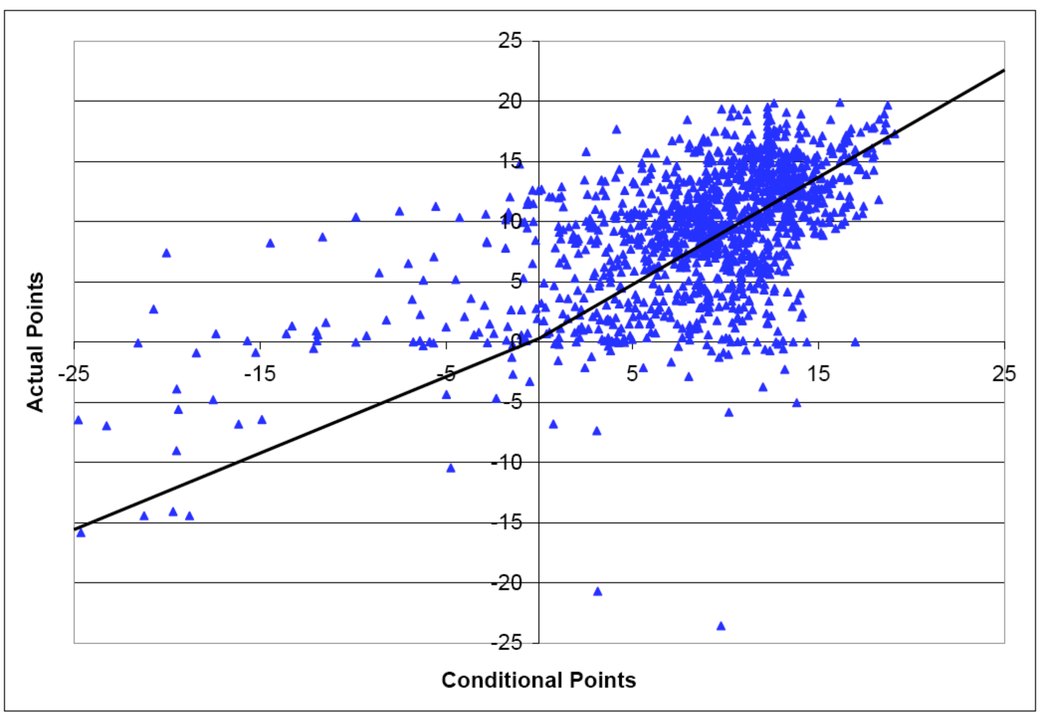

Figure A.10: Actual (y) and conditionally optimal (x) ml of beverage, observations between -25 and +25 $(n = 1346)$

Appendix B

Supplemental Materials for To Review or Not Review?

B.1 Description of Variables

To determine if a movie was cold opened $(c_j = 1)$ we examined the dates on three or four major news publications (the Los Angeles Times, New York Times, San Francisco Chronicle, and New York Post). If the dates of reviews in any of these publications were later than the release date we examined the reasoning behind the late reviews. A movie was classified a "cold open" if at least one source stated the movie was not screened for critics before release (in most cases, all of the available sources did not have advance reviews).

Weekend and total US box office data were obtained from a *FilmSource* database (Nielsen EDI, `www.filmsource.com`). The *FilmSource* database also included the number of theaters that showed a movie during its first weekend, the number of days in the opening weekend, and if the movie was released before Friday (generally only for anticipated blockbusters). *FilmSource* also gave a description of the genre of the movie, its MPAA rating (G, PG, PG-13, R), and whether the movie was adapted from previous source material.

Production budget information came from imdb.com for most movies, and from boxoffice-mojo.com or the-numbers.com for those missing from imdb.com. Budget data were available for 856 of the 890 movies, including 59 or the 62 cold openings (95%). Of this set, 832 movies also had the first day's box office data available on imdb.com including 59 of the 62 cold openings.

The imdb.com database was used to determine the star power rating of each movie's stars. Each week imdb.com determined this value by ranking the number of searches done on the imdb.com site for every person affiliated with movies. The most searched star would have value 1. Since there are over one million stars on imdb.com, we took the natural logarithm of the star ranking to reduce effect of unknown stars with very high numbers. We averaged the logged star ranking for the top two stars for each movie during its opening week.

Three other variables, competition (the average production budget of other movies released on the same opening weekend), the summer dummy variable (whether the movie was released in June, July and August), and the year of release variable (2000=-3, 2001=-2, 2002=-1, 2003=0, 2004=1, 2005=2, 2006=3) were calculated from the previous data.[1]

B.2 Supplemental Tables and Figures

Table B.1 shows the regressions done on logged opening day box office in the full and lean regressions. Table B.2 shows regressions done on logged post-opening-weekend box office in the full and lean regressions including logged opening weekend box office as a regressor. Table B.3 shows the correlations between the regressors in the full regressions done on box office. Table B.4 shows values after each iteration in estimating the QRE parameters. Table B.5 shows values after each iteration in estimating the CH parameters with QR.

Figure B.1 shows the predicted probabilities of cold opening by critic quality and actual decision. About half the movies cold-opened (squares) have high predicted probabilities and half have low probabilities. There is a clear relation between quality and the predicted probability of cold opening. The model is on the right track but the correlation is far from perfect. Figure B.2 shows the expected quality of each movie given it was cold opened vs. the actual quality. Since moviegoers correctly infer expectations in the QRE, about half the cold opened movies have quality less than their actual critic ratings. The low value of λ_d causes a wide variety of expectation values for cold opened movies between 0 and 40.

Figure B.3 which represents the estimated probability that each movie will be cold opened, with actual cold-openings plotted in red. The implied line from the scatter plot is much clearer in Figure B.3 than in Figure B.1 for the QRE model. The cold opened movies on average have higher probabilities of being cold opened, so the model fits better. Figure B.2 shows the estimated value of the expected quality belief moviegoers would have if each movie had been cold opened. This is almost a constant because of the relatively high λ_d and lower τ_m. Moviegoers expect few cold opening decisions to be the result of quantal response, but they mostly expect distributors to be best

[1] The regressions had similar results when dummy variables for year were used instead of one year variable.

variable\regression	log 1st day	log 1st day
cold	0.119*	0.153**
	(0.083)	(0.086)
crit	0.012***	0.010***
	(0.001)	(0.001)
thtr	0.875***	0.822***
	(0.037)	(0.035)
bud	0.002***	0.001*
	(0.001)	(0.001)
comp	0.015**	0.001
	(0.001)	(0.001)
star	-0.026*	-0.066***
	(0.015)	(0.014)
sum	0.101**	-
	(0.046)	
sq/adpt	0.139***	0.097**
	(0.043)	(0.044)
beffri	0.000	-
	(0.000)	
wkdlen	-0.001	-
	(0.059)	
beffor	0.000	-
	(0.000)	
act/adv	0.082	-
	(0.073)	
ani	-0.264**	-
	(0.120)	
com	0.021	-
	(0.061)	
doc	0.493	-
	(0.334)	
fant/sci	0.131	-
	(0.096)	
susp/hor	0.035	-
	(0.074)	
year	-0.041***	-0.037***
	0.011	(0.011)
pg	-0.105	-
	(0.130)	
pg13	0.376***	-
	(0.127)	
r	0.384***	-
	(0.130)	
const	-1.843***	-1.10***
	(0.196)	(0.134)
R-squared	0.670	0.622
N	833	833
degrees of freedom	21	8
*p<0.1, **p<0.05, ***p<0.01		

Table B.1: Regressions of log opening day revenues (in millions)

variable\regression	log post-opening weekend	log post-opening weekend
log rwkd	1.210*** (0.024)	1.177*** (0.026)
cold	-0.035 (0.051)	-0.091* (0.054)
crit	0.009*** (0.001)	0.009*** (0.001)
thtr	-0.134*** (0.030)	-0.076** (0.030)
bud	1.21*** (0.001)	0.000 (0.001)
comp	0.001*** (0.000)	0.002*** (0.000)
star	-0.016* (0.009)	-0.001 (0.009)
sum	0.035 (0.028)	-
sq/adpt	-0.020 (0.026)	-0.013 (0.028)
beffri	0.087*** (0.019)	-
wkdlen	-0.121*** (0.036)	-
beffor	0.000 (0.000)	-
act/adv	-0.166*** (0.044)	-
ani	-0.238*** (0.073)	-
com	0.001 (0.036)	-
doc	-0.197 (0.206)	-
fant/sci	-0.329*** (0.059)	-
susp/hor	-0.048 (0.044)	-
year	-0.042*** (0.007)	-0.041*** (0.007)
pg	-0.201*** (0.074)	-
pg13	-0.427*** (0.077)	-
r	-0.52*** (0.079)	-
const	0.556*** (0.119)	-0.060 (0.026)
R-squared	0.916	0.899
N	856	833
degrees of freedom	21	8
*p<0.1, **p<0.05, ***p<0.01		

Table B.2: Regressions of log box office revenues after first weekend (in millions)

	Cold	crit	theater s	rBud	Actor 1	Actor 2	Hol	adapt or seq	BefFri	Wkd Len	Before Foreign	Act/ Adv	Ani- mated	Com- edy	Docu- mentary	Fant/ Sci	Susp/ Horr	YEAR	PG	PG-13	R
Cold	1.00	-0.33	-0.15	-0.17	-0.06	0.12	-0.05	-0.08	-0.03	0.01	0.01	-0.06	-0.05	-0.01	-0.02	0.00	0.20	0.07	-0.11	0.06	0.03
crit	-0.33	1.00	0.16	0.28	0.10	-0.11	0.04	0.13	0.12	0.00	0.04	0.06	0.16	-0.18	0.04	0.08	-0.12	0.12	0.04	-0.06	-0.01
theaters	-0.15	0.16	1.00	0.58	0.07	-0.38	0.11	0.26	0.07	0.02	-0.12	0.22	0.16	-0.11	-0.10	0.19	-0.03	0.15	0.17	0.10	-0.28
rBud	-0.17	0.28	0.58	1.00	0.12	-0.40	0.14	0.24	0.19	0.04	-0.08	0.32	0.09	-0.26	-0.07	0.28	-0.16	0.01	0.05	0.11	-0.16
Actor 1	-0.06	0.10	0.07	0.12	1.00	-0.05	0.15	0.05	0.11	0.05	0.01	0.01	0.06	0.06	-0.01	0.00	-0.08	-0.01	0.10	0.02	-0.11
Actor 2	0.12	-0.11	-0.38	-0.40	-0.05	1.00	-0.06	-0.09	-0.05	0.01	0.10	-0.15	0.25	0.03	0.17	-0.11	0.00	0.02	0.17	-0.17	-0.03
Hol	-0.05	0.04	0.11	0.14	0.15	-0.06	1.00	0.04	0.08	-0.11	0.01	0.06	0.00	0.02	0.06	0.06	-0.03	-0.02	0.04	0.06	-0.10
adapt or seq	-0.08	0.13	0.26	0.24	0.05	-0.09	0.04	1.00	0.13	0.02	-0.05	0.11	0.05	-0.16	-0.05	0.14	0.02	0.10	0.09	-0.04	-0.05
BefFri	-0.03	0.12	0.07	0.19	0.11	-0.05	0.08	0.13	1.00	0.09	-0.05	0.01	0.05	-0.02	0.04	0.07	-0.09	-0.03	0.02	-0.02	-0.02
WkdLen	0.01	0.00	0.02	0.04	0.05	0.01	-0.11	0.02	0.09	1.00	-0.03	-0.01	0.02	-0.02	-0.02	0.03	-0.05	-0.03	0.01	0.04	-0.06
ForBef	0.01	0.04	-0.12	-0.08	0.01	0.10	0.01	-0.05	-0.05	-0.03	1.00	0.09	-0.01	-0.07	-0.01	-0.01	0.03	-0.02	-0.02	-0.04	0.05
Act/Adv	-0.06	0.06	0.22	0.32	0.01	-0.15	0.06	0.11	0.01	-0.01	0.09	1.00	-0.11	-0.35	-0.03	-0.12	-0.20	0.04	-0.07	0.09	-0.01
Animated	-0.05	0.16	0.16	0.09	0.06	0.25	0.00	0.05	0.05	0.02	-0.01	-0.11	1.00	-0.20	-0.01	-0.07	-0.11	0.00	0.24	-0.23	-0.18
Comedy	-0.01	-0.18	-0.11	-0.26	0.06	0.03	0.02	-0.16	-0.02	-0.02	-0.07	-0.35	-0.20	1.00	-0.05	-0.20	-0.34	-0.02	0.05	0.14	-0.16
Documentary	-0.02	0.04	-0.10	-0.07	-0.01	0.17	0.06	-0.05	0.04	-0.02	-0.01	-0.03	-0.01	-0.05	1.00	-0.02	-0.03	0.04	-0.03	-0.02	0.04
Fant/Sci	0.00	0.08	0.19	0.28	0.00	-0.11	0.06	0.14	0.07	0.03	-0.01	-0.12	-0.07	-0.20	-0.02	1.00	-0.11	0.00	0.11	0.04	-0.11
Susp/Horr	0.20	-0.12	-0.03	-0.16	-0.08	0.00	-0.03	0.02	-0.09	-0.05	0.03	-0.20	-0.11	-0.34	-0.03	-0.11	1.00	0.04	-0.19	-0.10	0.28
YEAR	0.07	0.12	0.15	0.01	-0.01	0.02	-0.02	0.10	-0.03	-0.03	-0.02	0.04	0.00	-0.02	0.04	0.00	0.04	1.00	0.09	0.02	-0.10
PG	-0.11	0.04	0.17	0.05	0.10	0.17	0.04	0.09	0.02	0.01	-0.02	-0.07	0.24	0.05	-0.03	0.11	-0.19	0.09	1.00	-0.41	-0.30
PG-13	0.06	-0.06	0.10	0.11	0.02	-0.17	0.06	-0.04	-0.02	0.04	-0.04	0.09	-0.23	0.14	-0.02	0.04	-0.10	0.02	-0.41	1.00	-0.66
R	0.03	-0.01	-0.28	-0.16	-0.11	-0.03	-0.10	-0.05	-0.02	-0.06	0.05	-0.01	-0.18	-0.16	0.04	-0.11	0.28	-0.10	-0.30	-0.66	1.00

Table B.3: Correlation between variables

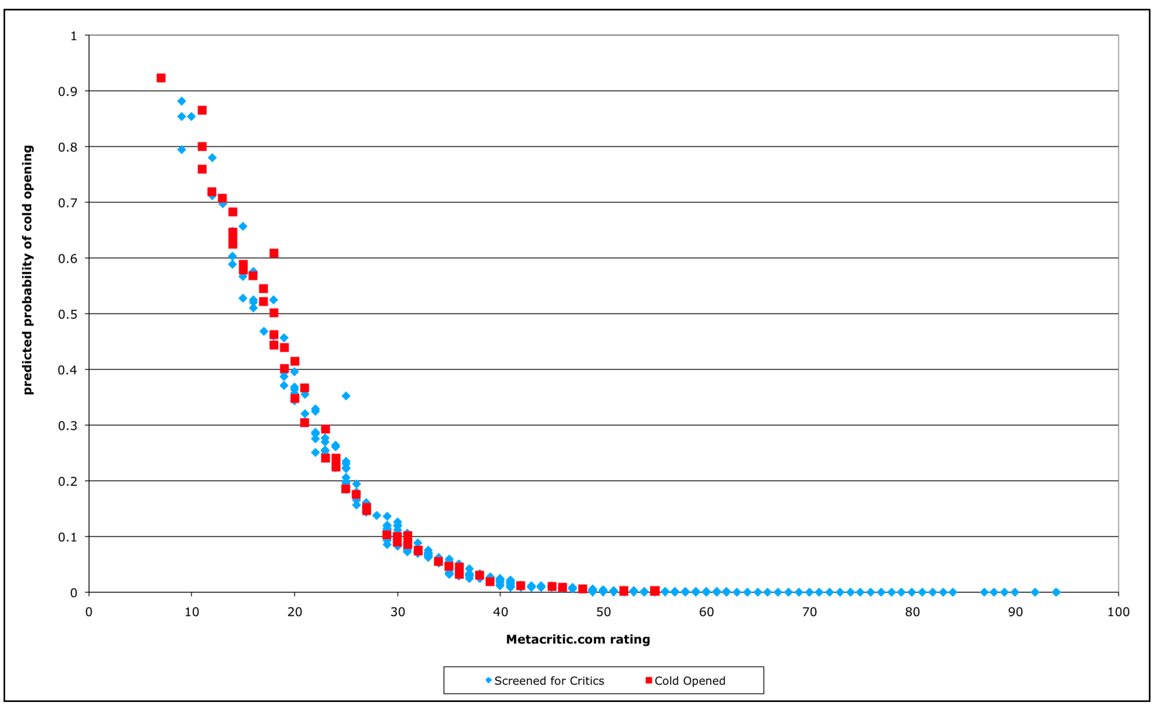

Figure B.1: Probability of movie being cold opened in QRE model by critic rating $(\lambda_d = 1.345)$

variable\iteration	1st	2nd	3rd	4th	5th	6th
λ_m	1.285	1.290	1.282	1.291	1.288	1.288
crit	0.017	0.016	0.015	0.016	0.016	0.016
thtr	0.001	0.001	0.001	0.001	0.001	0.001
bud	0.003	0.003	0.003	0.003	0.003	0.003
comp	0.002	0.002	0.002	0.002	0.002	0.002
star	-0.041	-0.036	-0.037	-0.037	-0.037	-0.037
sum	0.057	0.040	0.035	0.037	0.036	0.036
sq/adpt	0.101	0.105	0.101	0.105	0.102	0.102
beffri	0.001	0.001	0.007	0.004	0.007	0.007
wkdlen	0.101	0.100	0.097	0.099	0.098	0.098
beffor	0.000	0.000	0.000	0.000	0.000	0.000
act/adv	-0.140	-0.151	-0.153	-0.148	-0.149	-0.149
ani	-0.247	-0.236	-0.232	-0.234	-0.233	-0.233
com	0.023	0.022	0.014	0.022	0.018	0.018
doc	0.211	0.193	0.197	0.196	0.196	0.196
fant/sci	-0.140	-0.120	-0.121	-0.119	-0.120	-0.120
susp/hor	0.001	0.025	0.042	0.028	0.037	0.037
year	-0.073	-0.062	-0.061	-0.064	-0.063	-0.063
pg	-0.151	-0.141	-0.142	-0.141	-0.142	-0.142
pg13	-0.165	-0.124	-0.122	-0.126	-0.123	-0.123
r	-0.205	-0.164	-0.170	-0.166	-0.169	-0.169
R-squared	0.683	0.682	0.678	0.682	0.681	0.682
N	797	856	856	856	856	856
degrees of freedom	20	20	20	20	20	20
λ_d	1.478	1.356	1.367	1.350	1.345	-
log likelihood	-191.308	-198.569	-196.965	-203.406	-205.712	-
Mean $E_m(q\|Xj,1)$ for λ_d	14.159	13.264	14.894	15.319	15.106	-

Table B.4: The iterative estimation process for the QRE model

variable\iteration	1st	2nd	3rd	4th	5th
λ_m	1.285	1.302	1.302	1.302	1.302
crit	0.017	0.017	0.017	0.017	0.017
thtr	0.001	0.001	0.001	0.001	0.001
bud	0.003	0.003	0.003	0.003	0.003
comp	0.002	0.002	0.002	0.002	0.002
star	-0.041	-0.039	-0.039	-0.039	-0.039
sum	0.057	0.036	0.037	0.037	0.037
sq/adpt	0.101	0.107	0.108	0.108	0.108
beffri	0.001	0.000	0.000	0.000	0.000
wkdlen	0.101	0.095	0.095	0.095	0.095
beffor	0.000	0.000	0.000	0.000	0.000
act/adv	-0.140	-0.154	-0.155	-0.155	-0.155
ani	-0.247	-0.253	-0.254	-0.254	-0.254
com	0.023	0.016	0.016	0.016	0.016
doc	0.211	0.226	0.228	0.228	0.228
fant/sci	-0.140	-0.152	-0.153	-0.153	-0.154
susp/hor	0.001	-0.016	-0.018	-0.018	-0.018
year	-0.073	-0.066	-0.066	-0.066	-0.066
pg	-0.151	-0.139	-0.139	-0.139	-0.139
pg13	-0.165	-0.148	-0.149	-0.149	-0.149
r	-0.205	-0.177	-0.177	-0.177	-0.177
R-squared	0.683	0.682	0.674	0.682	0.682
N	797	856	856	856	856
degrees of freedom	20	20	20	20	20
λ_d	6.816	7.136	7.090	7.085	-
τ_d	8.567	8.550	8.554	8.554	-
log likelihood	-166.424	-166.226	-166.231	-166.232	-
τ_m	1.26	1.15	1.13	1.12	-
Mean $E_m(q\|Xj,1)$ for τ_m, λ_d	44.666	45.254	45.359	45.412	-

Table B.5: The iterative estimation process for the QRE model with CH

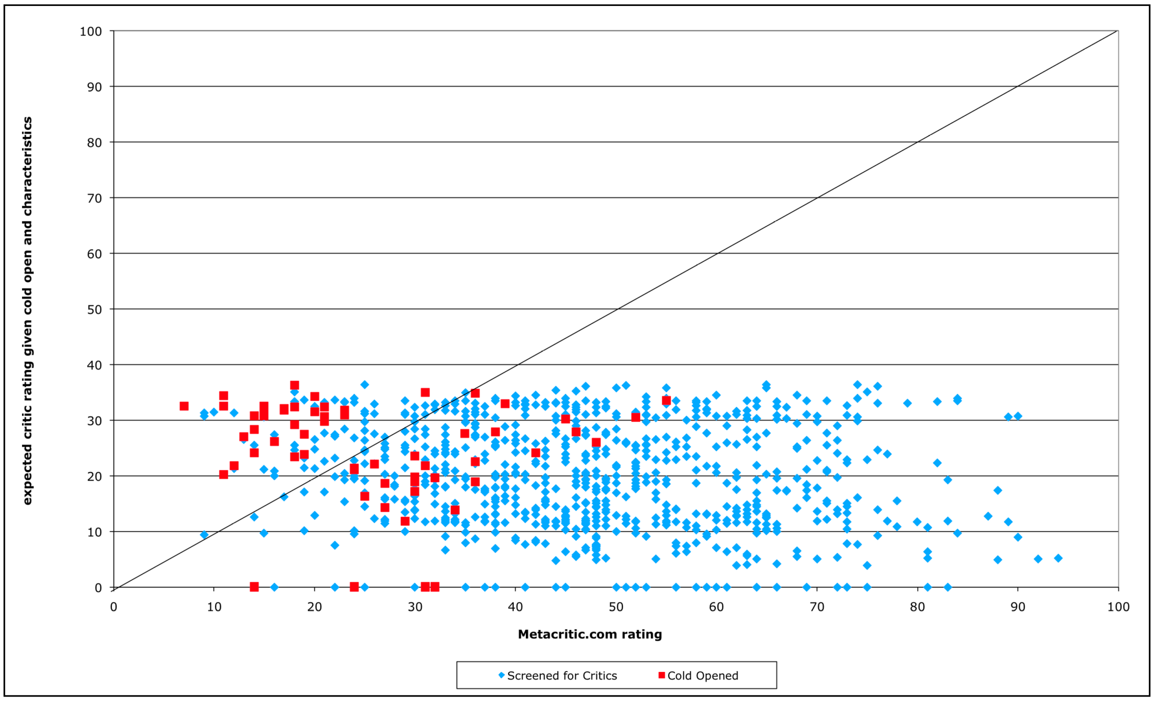

Figure B.2: Expected movie quality given it is cold opened in QRE model by critic rating ($\lambda_d = 1.345$)

responding to 0 or 1 level moviegoers who do not associate quality with movies or the decision to cold open. Note that cold opened movies (red squares) tend to have expected quality above actual quality and screened movies have expected quality below actual quality.

Tables B.6 and B.7 show the sum of squares and log likelihood for the various values of estimates of τ_d and $\{\tau_d, \lambda_d\}$, respectively.

τ_m	(Average SSR)$^{1/2}$
0	5.14
1.12*	5.08
2	5.08
4	5.14
6	5.35
8	5.59
10	5.81

Table B.6: Average squared difference between predicted and actual weekend revenues (in $) for all cold openings ($N = 59$) by moviegoer sophistication in CH model with QR ($\lambda_d = 7.085$)

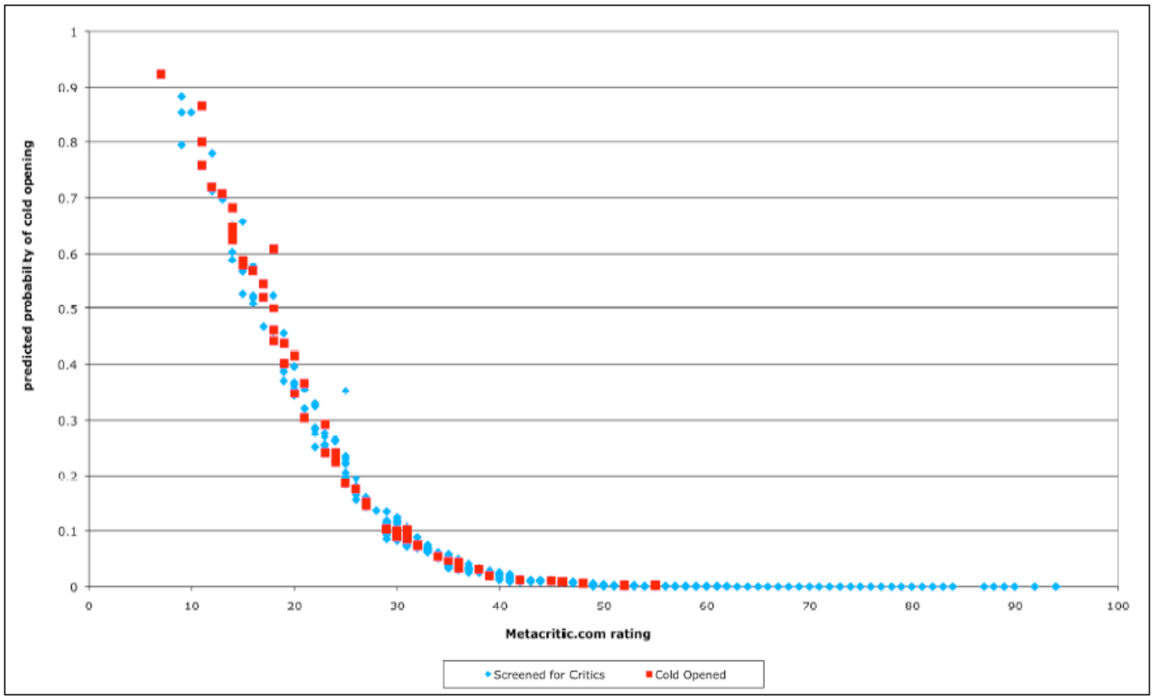

Figure B.3: Probability of movie being cold opened in CH model with QR by critic rating ($\lambda = 7.085, \tau_d = 8.567$)

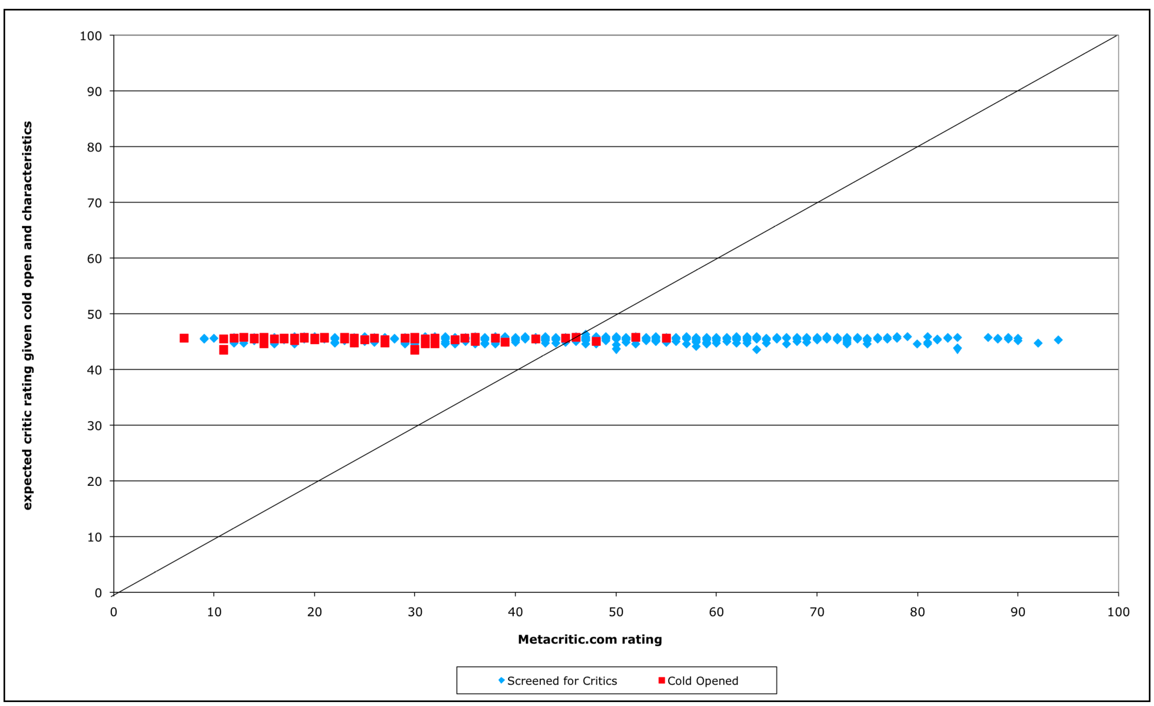

Figure B.4: Expected movie quality given it is cold opened in CH model with QR by critic rating ($\lambda = 7.085, \tau_d = 8.567$)

$\tau_d \backslash \lambda_d$	0	3	6	7.085*	9
0	-593.33	-593.33	-593.33	-593.33	-593.33
2	-593.33	-783.20	-798.29	-800.62	-803.45
4	-593.33	-441.44	-438.66	-438.31	-438.17
6	-593.33	-248.56	-230.67	-228.38	-226.32
8	-593.33	-183.70	-169.26	-168.39	-167.89
8.567*	-593.33	-177.43	-166.43	**-166.23**	-166.44
10	-593.33	-173.94	-175.62	-177.55	-179.92

Table B.7: Log likelihood for all distributor release decisions ($N = 856$) in CH model with QR by distributor sophistication

B.3 Details of Iterative Estimation Procedures (QRE, Cursed, CH)

This section provides the general iterative procedure for obtaining estimates for the relevant parameters of the QRE, cursed and CH models.

1. The iteration counter begins at $i = 1$.

2. The coefficients in equation 3.5 are estimated using a linear regression,

$$\log\left(\frac{y_j}{N\hat{t} - y_j}\right) = (-\lambda_m\alpha)E_m(q_j|c_j, X_j) - (\lambda_m\beta)X_j - (\lambda_m)\hat{t} - (\lambda_m)\epsilon_j \tag{B.1}$$

   assuming $N = 300 \times 10^6$ and $\hat{t} = 5.34$.[2] In iteration $i = 1$ only the 797 movies which are screened to critics ($c_j = 0$) are used. Using assumption 3.1, the observed q_j is substituted for the unobserved expectation $E_m(q_j|0, X_j)$ for these movies. Then all the independent and dependent variables are measured and we can estimate the regression easily.[3] In later iterations, expected quality values $\big(E^{qre}_{m_i}(q_j|c_j, X_j)$, $E^{ce}_{m_i}(q_j|c_j, X_j)$, or $E^{ch}_{m_i}[E_k(q_j|c_j, X_j)|\tau_d]\big)$ after iteration i will have been computed, and a regression on the full sample can be run.

3. Since simply using $\hat{R}(X_j, E_m(q_j|c_j, X_j)) = N\hat{t}\left(1 + \exp[\widehat{\lambda_m\alpha}E_m(q_j|c_j, X_j) + \widehat{\lambda_m\beta}X_j + \widehat{\lambda_m}\hat{t}]\right)$ to estimate $R(X_j, E_m(q_j|c_j, X_j)$ would produce biased estimates, non-parametric kernel regression techniques are used. A consistent Gaussian kernel regression is used to estimate revenue from the parameter estimates from equation B.1.

$$\begin{aligned}\hat{R}(X_j, E_m(q_j|c_j, X_j)) &= \hat{m}\left(-\lambda_m(\alpha E_m(q_j|c_j, X_j) + \beta X_j - \hat{t})\right) \\ &= \frac{\sum_{l\in J_i} K_h(g(j) - g(l))y_l}{\sum_{l\in J_i} K_h(g(j) - g(l))}\end{aligned} \tag{B.2}$$

[2]Results are highly similar for $N = 100 \times 10^6$, 200×10^6, and $\hat{t} = 5.34$.

[3]A crucial maintained assumption below is that the coefficient on expected quality, α, in determining moviegoer attendance, and hence revenue, is the same for known-quality (screened) and unknown-quality (cold opened) movies.

where $g(j) = -\widehat{\lambda_m}\left(\hat{\alpha}E_m(q_j|c_j, X_j) + \hat{\beta}X_j - \hat{t}\right)$, and K is the Gaussian kernel, $K_h(x) = h\frac{1}{\sqrt{2\pi}}e^{\frac{x}{2h}}$, with bandwidth, $h = 0.9w\|J_i\|^{-1/5}$ where $w = \min(s_y, IQR_y/1.34)$ (from Silverman, 1986) and J_i is the current iteration's set of movies (with length 797 for iteration 1, 856 thereafter).

4. The regression results from step 2 give iteration-i coefficients $\hat{\alpha}_i$ and $\hat{\beta}_i$ and a response sensitivity $\hat{\lambda}_{m,i}$. Step 3 gives estimated revenue equation $\hat{R}_i$ from these parameters for different values of X_j and q_j. From equation (3) we have

$$\begin{aligned}
&E_m^{qre}(q_j|X_j,1) = \frac{\sum_{q=0}^{100} q\pi(X_j,q)P(q)}{\sum_{q=0}^{100}\pi(X_j,q)P(q)} \\
&\Rightarrow E_m^{qre}(q_j|X_j,1)\sum_{q=0}^{100}\pi\left(X_j,q\right)P(q) = \sum_{q=0}^{100} q\pi\left(X_j,q\right)P(q) \\
&\Rightarrow \sum_{q=0}^{100}\pi\left(X_j,q\right)P(q)\left[E_m^{qre}(q_j|X_j,1) - q\right] = 0 \\
&\Rightarrow \frac{\sum_{q=0}^{100} P(q)[E_m^{qre}(q_j|X_j,1)-q]}{1+\exp\left(\lambda_d\left(\hat{R}\left(X_j,E_m^{qre}(q_j|X_j,1)\right)-\hat{R}(X_j,q_j)\right)\right)} = 0
\end{aligned} \tag{B.3}$$

where the last step follows from the definition of $\pi(X_j, q)$ (assumption 3.3). All the terms in B.3 can be estimated from regression coefficients ($\hat{\alpha}_i$, $\hat{\beta}_i$, $\hat{\lambda}_{m,i}$ from step 2), determined from the revenue equation $\hat{R}_i$ (from step 3), fit from the quality distribution $P(q)$, or fixed by assumption ($\hat{t}$, N), except for λ_d and $E_{m_i}^{qre}(q_j|X_j,1)$. To create an iteration of estimates of $E_m^{qre}(q_j|X_j,1)\ \forall j$ we fix a value of λ_d and solve B.3 for each movie j. Next, using fixed λ_d, and newly calculated estimates of $E_{m_i}^{qre}(q_j|X_j,1)$ for each movie, along with the estimated parameters in step 2 and revenue equations in step 3, the predicted iteration-i probability $(\hat{\pi}_i(X_j, q_j, \lambda_d))$ that each movie j will be cold opened can be computed from assumption 3.3.

Additionally, in the cursed procedure: For any fixed χ_d, the newly calculated estimates of $E_m^{qre}(q_j|X_j,1)$ (see above equation B.3) for each movie can be converted to $E_m^{ce}(q_j|X_j,1)$, by equation 3.6. With those values, along with the estimated parameters in step 2 and revenue equations in step 3, the predicted iteration-i probability $(\hat{\pi}_i(X_j, q_j, \lambda_d)$ that each movie j will be cold opened can be computed from assumption 3.3 *for each value of* χ_d.

The ch procedure obtains the probabilities that each movie is cold opened differently:

For a given λ_d and τ_d, we use our estimated values $\hat{\alpha}_i$, $\hat{\beta}_i$, $\hat{\lambda}_{m,i}$, and estimated revenue equation $\hat{R}_i$ to estimate $\pi_{ki}(q_j, X_j)$, $E_{ki}(q|X_j,1)$, and $\hat{R}_i(E_k(q|X_j,1))$ for $k = 0\ldots\bar{k}$ using equations 3.9–3.12.[4] Since the probability of a given distributor being level k is $P\left(x=n|_d\right) = \tau_d^n e^{-\tau}/n!$ and the probability of that distributor cold opening given he is level k is $\pi_k(q_j, X_j)$, the total probability that a movie is cold opened is

$$\hat{\pi}_i(X_j, q_j, \lambda_d, \tau_d) = \sum_{k=0}^{\bar{k}} \pi_k(q_j, X_j) \times \tau_d^n e^{-\tau}/n! \tag{B.4}$$

5. Step 4 is performed repeatedly for a grid search over sets of values of $\lambda_d \in A_i$ (or $(\lambda_d, \chi_d) \in \{A_i, B\}$, $(\lambda_d, \tau_d) \in \{A_i, B_i\}$), where the grid search becomes progressively finer across iterations i.[5]

   The maximum likelihood estimate ω ($\lambda^*_{d,i}$, (λ^*_d, χ^*_d), or (λ^*_d, τ_d)) is chosen from the set Ω_i (A_i, $\{A_i, B\}$, or $\{A_i, B_i\}$).[6] That value satisfies

$$\begin{aligned} \omega_i^* &= \underset{\omega\in\Omega_i}{\operatorname{argmax}}\, L(\omega) \\ &= \underset{\omega_i\in\Omega_i}{\operatorname{argmax}} \prod_j \left[\hat{\pi}_i\left(X_j, q_j, \omega\right) c_j \times \left(1 - \hat{\pi}_i\left(X_j, q_j, \omega\right)\right)\left(1 - c_j\right)\right] \end{aligned} \tag{B.5}$$

   where $L(\omega)$ is the joint probability that distributors would choose to screen and cold open each of the 856 movies in the exact manner they did under the QRE (or cursed, CH) model with parameter(s) ω.[7]

6. The value for the maximum likelihood parameter $\lambda^*_{d,i}$ determined from the last step 3.5 is then used in equation B.3 to solve for iteration-i values of $E^{qre}_{m_i}(q_j|c_j, X_j)$ for each of the 59 cold opened movies.

   For the ch procedure: The maximum likelihood value $\lambda^*_{d,i}$ is used to compute the population-

---

[4]We used $\bar{k} = 40$, because given regular τ values the probability of $k > 40$ is nearly zero.

[5]The initial $\lambda_{d,i}$ grid is $A_1 = \{1, 1.25, ..., 2\}$. The second grid A_2 takes an interval of values in increments of .1 around the maximum likelihood estimate $\lambda^*_{d,1}$. The next grids A_i take on values of values of width .05, 0.01, 0.005, and 0.001 around the maximum likelihood estimate $\lambda^*_{d,i-1}$.

For cursed: The initial $\lambda_{d,i}$ grid is $A_1 = \{1, 1.25, ...2\}$. The second grid A_2 takes an interval of values in increments of .1 around the maximum likelihood estimate $\lambda^*_{d,1}$. The next grids A_i take on values of values of width .05, 0.01, 0.005 and 0.001, around the maximum likelihood estimate $\lambda^*_{d,i-1}$. The grid for χ_d, B is always {0,0.005,...1}.

For ch: The initial $\lambda_{d,i}$ grid is $A_1 = \{1, 2, ..., 10\}$ and $\tau_{d,i}$ grid is $B_1 = \{0.05, 0.1, ..., 10\}$. The second grid A_2 takes an interval of values in increments of .1 around the maximum likelihood estimate $\lambda^*_{d,1}$. The next grids A_i take on values of values of width .05, 0.01, 0.005 and 0.001, around the maximum likelihood estimate $\lambda^*_{d,i-1}$. For grids $i \geq 2, B_i = 8.001, ...9$.

[6]In the early steps of iteration (i.e., steps 1–3) this value is determined by interpolating inside the grid to achieve more decimal precision.

[7]This process takes roughly 15 minutes (8 minutes for CH) for each λ_d on a single PC running Mathematica 5.2.

averaged expectation for each of the 59 cold opened movies in the sample with

$$E[E_k(q|X_j,1)|\tau_m] = \sum_{k=0}^{\bar{m}} \pi_k(q_j, X_j) E_k(q|X_j,1). \tag{B.6}$$

The value of τ_m that minimizes the squared residuals in equation 3.5 is considered the best estimator for this step, that is

$$\tau^*_{m,i} = \underset{\tau_m}{\operatorname{argmin}} \sum_{j:c_j=1} \left(\hat{R}_i\left(X_j, E\left[(E_k(q|X_j,1)|\tau_m]\right) - y_j\right)^2 \tag{B.7}$$

where $\hat{R}_i(...)$ is estimated from the kernel estimation B.1 in step 2.

Now we have a full set of quality measures q_j and expected qualities for every movie.

7. The process is stopped when the regression values and parameter estimates (λ^*_m, ω^*_m) from the current iteration i are all within .001 of those from iteration $i-1$. Otherwise, the process is repeated with the iteration counter increased by one, starting with the regression step 2. *For the cursed procedure:* When the process converges, a new value χ_m is calculated to minimize the sum of squares between predicted and actual values over all cold openings. That value is determined by

$$\chi^*_m = \underset{\chi_m}{\operatorname{argmin}} \sum_{j:c_j=1} \left(\hat{R}\left(X_j, (1-\chi_m)\, E^{re*}_m\left(q|X_j,1\right) + \chi_m \bar{q}\right) - y_j\right)^2 \tag{B.8}$$

where $\hat{R}(...)$ is the last estimate done in step 3.

8. The process is repeated 100 more times with different bootstrapped data sets. A bootstrapped data set is created by randomly sampling with replacement from the 856 movies in the original data set. Parameter estimates are obtained by repeating steps 1–7. Standard errors (see Table 3.10) are calculated by taking the standard deviation of these 100 parameter estimates.[8]

[8]Depending on the bootstrap and number of iterations, the process for a single bootstrap takes 2–6 hours (2–6 for cursed, 2–12 for ch) on a single PC running Mathematica 5.2

Appendix C

Supplemental Materials for Endogenous Time Preference and Personal Rules

C.1 Stochastic Grapefruit Quality Example for T=2

Definition C.1. For T periods, $\{(a, V)\}$ is a stochastic solution to the sophisticated model with finite periods if $\forall t, 1 \leq t \leq T$, we have $a_t : (\underline{\beta}, \overline{\beta}) \times \{g_H, g_L\} \to \{G, D\}$ and $V_t : (\underline{\beta}, \overline{\beta}) \times \{g_H, g_L\} \to \mathbb{R}$ s.t.

$$a_T^*(\beta, g) = \begin{cases} G & \text{if } v(G, \beta, g) \geq v(D, \beta, g) \\ D & \text{if } v(G, \beta, g) < v(D, \beta, g) \end{cases} \tag{C.1}$$

and

$$EV_T^*(\beta) = \pi u(a_T^*(\beta, g_H), g_H) + (1 - \pi) u(a_T^*(\beta, g_L), g_L) \tag{C.2}$$

and $\forall 1 \leq t < T$

$$a_t^*(\beta, g) = \begin{cases} G & \text{if } v(G, \beta, g) + \delta EV_{t+1}^*(\Gamma(\beta)) \geq v(D, \beta, g) + \delta EV_{t+1}^*(\Delta(\beta)) \\ D & \text{if } v(G, \beta, g) + \delta EV_{t+1}^*(\Gamma(\beta)) < v(D, \beta, g) + \delta EV_{t+1}^*(\Delta(\beta)) \end{cases} \tag{C.3}$$

$$EV_t^*(\beta) = \pi V_t^*(\beta, g_H) + (1 - \pi) V_t^*(\beta, g_L) \tag{C.4}$$

$$V_t^*(\beta, g) = \begin{cases} u(G, g) + \delta EV_{t+1}^*(\Gamma(\beta)) & \text{if } a_t(\beta, g) = G \\ u(D, g_t) + \delta EV_{t+1}^*(\Delta(\beta)) & \text{if } a_t(\beta, g_t) = D. \end{cases} \tag{C.5}$$

Example C.1. *If $T = 2$ we can define solutions as before with thresholds. Let $\hat{\beta}_L^H = \frac{d}{g_H}$. And let*

$$\hat{\beta}_2^L = \max\left\{\frac{d}{g_L + \delta\pi(g_H - d)}, \Gamma^{-1}(\hat{\beta}_1^H)\right\} \tag{C.6}$$

$$\hat{\beta}_2^H = \max\left\{\frac{d}{g_H + \delta\pi(g_H - d)}, \Gamma^{-1}(\hat{\beta}_1^H)\right\}. \tag{C.7}$$

Then if $\hat{\beta}_2^L \leq \overline{\beta}$, and $\Delta(\hat{\beta}_2^L) > \hat{\beta}_1^H$ the solution to $t = T - 1$ is

$$a_{T-1}^*(\beta, g) = \begin{cases} D & \text{if } \beta \in \left(\Delta^{-1}(\hat{\beta}_1^H), \overline{\beta}\right) \text{ and } g = g_L \\ G & \text{if } \beta \in \left[\hat{\beta}_2^H, \overline{\beta}\right] \text{ and } g = g_H \\ G & \text{if } \beta \in \left[\hat{\beta}_2^L, \Delta^{-1}\left(\hat{\beta}_1^H\right)\right] \text{ and } g = g_L \\ D & \text{if } \beta \in \left[\hat{\beta}_2^H, \hat{\beta}_2^L\right) \text{ and } g = g_L \\ D & \text{if } \beta \in \left(\underline{\beta}, \hat{\beta}_2^H\right) \end{cases} \tag{C.8}$$

if $\hat{\beta}_2^L \leq \overline{\beta}$, and $\Delta(\overline{\beta}) < \hat{\beta}_1^H$ the solution to $t = T - 1$ is

$$a_{T-1}^*(\beta, g) = \begin{cases} G & \text{if } \beta \in \left[\hat{\beta}_2^L, \overline{\beta}\right) \text{ and } g = g_L \\ G & \text{if } \beta \in \left[\hat{\beta}_2^H, \overline{\beta}\right) \text{ and } g = g_H \\ D & \text{if } \beta \in \left(\hat{\beta}_2^H, \hat{\beta}_2^L\right) \text{ and } g = g_L \\ D & \text{if } \beta \in \left(\underline{\beta}, \hat{\beta}_2^H\right). \end{cases} \tag{C.9}$$

Proof. For $\beta < \hat{\beta}_2^H < \hat{\beta}_1^H$, $v(G, \beta, g) + \delta EV_T^*(\Gamma(\beta)) < v(D, \beta, g) + \delta EV_T^*(\Delta(\beta))$, so $a_{T-1}^*(\beta, g) = D, \forall g$. If $\beta \in [\hat{\beta}_2^H, \hat{\beta}_2^L)$ and $\hat{\beta}_1^H > \Delta(\hat{\beta}_2^L)$, then $EV_T^*(\Gamma(\beta)) = \pi g_H + (1-\pi)d$ and $EV_T^*(\Gamma(\beta)) = d$. Then $a_T^*(\beta, g_H) = G$ and $a_T^*(\beta, g_L) = D$.

If $\hat{\beta}_1^H < \Delta(\hat{\beta}_2^L)$, then for $\Delta^{-1}(\hat{\beta}_1^H) < \beta < \overline{\beta}$, $EV_T(\Gamma(\beta)) = EV_T(\Delta(\beta))$ and $a_{T-1}^*(\beta, g_L) = D$, $a_{T-1}^*(\beta, g_H) = G$. If instead, $\Delta(\overline{\beta}) < \hat{\beta}_1^H$, for $\hat{\beta}_2^L < \beta < \overline{\beta}$, $EV_T(\Delta(\beta)) = d$ and $EV_T(\Gamma(\beta)) = \pi g_H + (1-\pi)d$, so $a_{T-1}^*(\beta, g) = G$. A similar argument shows that if $\Delta(\overline{\beta}) > \hat{\beta}_1^H > \Delta(\hat{\beta}_2^L)$ for $\beta \in [\hat{\beta}_2^L, \Delta^{-1}(\hat{\beta}_1^H))$, $a_{T-1}^*(\beta, g) = G$. □

C.2 Characterization of the Extensive Form Game

We will use the Kreps and Wilson (1982) definition of an extensive form game.

Definition C.2 (Kreps and Wilson (1982))**.** In a finite extensive form game the following are specified:

1. *The physical order of play.* The physical order of play is given by a finite set of nodes $\mathbb{T}^1$ with a binary relation on that represents precedence.

2. *The choices available to each player when it is his turn to move.* A finite set A of actions and a function, $\alpha : \mathbb{T}\backslash W \to A$ that labels each non-initial node with the last action taken to reach it. A is required to be one-to-one.

3. *Rules for determining whose move it is at any point.* A finite set of players I and a function $\iota(x) : X \to I$ that assigns to each decision node the player whose turn it is.

4. *The information a player has whenever it is his turn to move.* Information possessed by players is represented by a partition H of X that divides the decision nodes into information sets.

5. *The payoffs to the players as functions of the moves they select.* For each player i, the payoff function $u^i : Z \to R$ assigns a real-valued von Neumann-Morgenstern utility to each outcome.

6. *The initial conditions that begin the game (that is the actions of nature).* Player i's intial assesment ρ^i is a probability measure on the set W of states or initial nodes.

Lemma C.1. *For any $n \in \mathbb{N}$, and initial state β_0 the model can be expressed as an extensive game.*

Proof. For any n we can define a finite set of nodes $\mathbb{T}$ where $|\mathbb{T}| = \sum_{k=0}^{n} 2^k$. Index the nodes so that for any $1 \leq i \leq n$ there are 2^{i-1} nodes of the form τ_{ik}. Starting with τ_{11}, each node τ_{ij} is the parent of nodes $\left\{\tau_{(i+1)(2j)}, \tau_{(i+1)(2j-1)}\right\}$ which gives a binary relation satisfying (1). Then there are 2^n terminal nodes, denoted $\tau_{(n+1)k}$. Decision nodes are $x_{ij} = \tau_{ij}\ \forall i, T - n \leq i \leq T$.

The game always begins at node τ_{11} and this is known to all players $T - n + 1, \ldots, T$, satisfying (6).

We have $T - n + 1$ players. The function $\iota(x_i j) = i + T - n$ assigns movement at every decision node to a player. Satisfying (3) Player $T - n + 1$ moves first.

All players have perfect knowledge so H=X. Satisfying (4).

[1] Used instead of Kreps and Wilson's T to avoid confusion.

At every node x_{ij}, player $T-n+i$ may make actions from the set $\{D, G\}$. Also define

$$\alpha(\tau_{ij}) \rightarrow= \begin{cases} G & \text{if } j \in 2\mathbb{N}, \\ D & \text{if } j \in 2\mathbb{N}+1 \end{cases} \tag{C.10}$$

satisfying (2).

Given state β_0 at τ_{11} the state β_{ij} at node τ_{ij} $(\beta|\tau_{ij})$ is defined recursively as

$$\beta|\tau_{ij} = \begin{cases} \Gamma\left(\beta|\tau_{(i-1)(\frac{j}{2})}\right) & \text{if } j \in 2\mathbb{N} \\ \Delta\left(\beta|\tau_{(i-1)(\frac{j+1}{2})}\right) & \text{if } j \in 2\mathbb{N}+1. \end{cases} \tag{C.11}$$

Then given a sequence of actions $a_{T-n+1}\ldots a_T$ and player i, $T-n+1 \leq i \leq T$. Condition (2) gives us a unique node that was reached by that sequence for any i. Let τ_{ij} be that node. Then let

$$U_i = v(a_i, \beta|\tau_{ij}) + \sum_{k=i+1}^{T} \delta^k u(a_k) \tag{C.12}$$

then u is a real-valued von Neumann-Morgenstern utility function for each outcome, satisfying (5). □

C.3 Proof of Theorem 4.2, the Existence of an Infinite Horizon Solution

Theorem C.1. *Suppose*

$$V^*_\infty(\beta) = \begin{cases} \frac{g}{1-\delta} & \text{if } \beta \geq \hat{\beta}_1 \\ \frac{d}{1-\delta} & \text{if } \beta < \hat{\beta}_\infty. \end{cases} \tag{C.13}$$

For β_0 let there only be finite n number of states created from all transformations of Γ and Δ where $\hat{\beta}_\infty \leq \beta_t < \hat{\beta}_1$. Then there is a vector $o\tilde{mega}$ that satisfies 4.27, 4.28, 4.29.

Proof. Let $B = \{\beta_0, \ldots, \beta_n\}$ be the ordered set of all possible transformations of β_0 by Γ, Δ where $\hat{\beta}_\infty < \beta < \hat{\beta}_1$. We first must show that there exists a V^*_∞ satisfying 4.28 that is defined over B.

Let I be the set of functions defined on $[0,1]^n$ that map to $\mathbb{R}^n$, that are continuous and concave.

Define operator $T : I \rightarrow I$ s.t.

$$(Tv)(\sigma_i) = \sigma_i \left(g + \delta v_{\Gamma(i)}(\sigma_i)\right) + (1-\sigma_i)\left(d + \delta v_{\Delta(i)}(\sigma_i)\right) \tag{C.14}$$

where

$$\Gamma(i) = j \text{ s.t. } \beta_j = \Gamma(\beta_i) \text{ and } \Gamma(\beta_i) \leq \beta_n \tag{C.15}$$

$$\Gamma(i) = j \text{ s.t. } \beta_j = \Delta(\beta_i) \text{ and } \Gamma(\beta_i) \geq \beta_1 \tag{C.16}$$

$$V_{\Gamma(i)} = \frac{g}{1-\delta}, \text{ if } \Gamma(\beta_i) > \beta_n \tag{C.17}$$

$$V_{\Delta(i)} = \frac{d}{1-\delta}, \text{ if } \Gamma(\beta_i) < \beta_1 \tag{C.18}$$

Now for any $u, v \in I$, $\sigma \in [0,1]^n$

$$\begin{aligned}
|(Tv)(\sigma_i) - (Tu)(\sigma_i)| &= |\sigma_i \left(g + \delta v_{\Gamma(i)}(\sigma_i)\right) + (1-\sigma_i)\left(d + \delta v_{\Delta(i)}(\sigma_i)\right) - \\
&\quad \sigma_i \left(g + \delta u_{\Gamma(i)}(\sigma_i)\right) + (1-\sigma_i)\left(d + \delta u_{\Delta(i)}(\sigma_i)\right)| \\
&\leq \delta \left|\sigma_i(v_{\Gamma(i)}(\sigma_i) - u_{\Gamma(i)}(\sigma_i)) + (1-\sigma_i)(v_{\Delta(i)}(\sigma_i) - u_{\Delta(i)}(\sigma_i))\right| \\
&\leq \delta \left|\sigma_i \sup_{\sigma_i}[v(\sigma_i) - u(\sigma_i)] + (1-\sigma_i)\sup_{\sigma_i}[v(\sigma_i) - u(\sigma_i)]\right| \\
&\leq \delta|u_i - v_i|
\end{aligned} \tag{C.19}$$

Now we have a contraction mapping, so there exists a unique $Tv^* = v^*$. Since T maps continuous concave functions to continuous concave functions (because convex combinations of continuous and concave functions are continuous and concave), and v^* is the limit of any repeated mapping (including concave and continuous functions) v^* (henceforth V^*) must be continuous and concave.

To proceed further, let us define

$$U_i(\sigma) = \sigma_i(g + \delta V^*_{\Gamma(i)}(\sigma)) + (1-\sigma_i)\left(\frac{d}{\beta_i} + \delta V^*_{\Delta(i)}(\sigma)\right) \tag{C.20}$$

Now let us defined a best response function, $r(\sigma) : [0,1]^n \rightarrow [0,1]^n$ s.t. $r_i(\sigma) \in \text{argmax}_{\sigma_i} U_i(\sigma_i, \sigma_{-i})$ $\forall i$. We now must show a fixed point exists ($\sigma \in r(\sigma)$), to show our solution holds.

First, note that $[0,1]^n$ is a compact, convex, nonempty subset of Euclidean space.

Second, since $U_i(\sigma)$ is linear with respect to σ_i and $V(\sigma)$ and $V(\sigma)$ is continuous with respect to σ, U is continuous with respect to σ. Therefore $r(\sigma)$ is non-empty and by Berge's maximum theorem $r(\sigma)$ is upper-hemi-continuous.

Finally, since $V(\sigma)$ is concave, and U is a convex combination of that function plus constants, U is also concave. Then $r(\sigma)$ is convex.

By Kakatani's fixed point theorem, there exists a $\sigma \in r(\sigma)$. Therefore at every node β_i there is a σ that satisfies the solution to the system. $\square$

C.4 One-Step Cases

For notational simplicity, V_i will denote $V^*_\infty(\beta_i)$ in this section.

Example C.2. *If $n = 2$ and $\delta < \sqrt{1 - \frac{u}{d_1}}$ we will have $a^*_\infty(\beta_i) = G, \forall i \in \{1, 2\}$. Otherwise we will have*

$$p_1 = 1 - \frac{\bar{u} - d_2}{\delta d_1} \qquad p_2 = \frac{\bar{u} - d_1}{\delta d_2}. \tag{C.21}$$

Proof. Let us use notation $V_1 = V^*(\beta_1)$ and $V_2 = V^*_t(\beta_2)$. By definition, $V_1 = p_1\hat{u} + (1 - p_1)\delta V_2)$, and $V_2 = p_2(u + \delta V_1))$. Solving we have $p = \frac{V_1 - \delta V_2}{\hat{u} - V_2\delta}$ and $p_2 = \frac{V_2}{u + \delta d_2}$. If $\hat{u} = d_1 + \delta V_2$ and $d_2 = u + \delta V_1$, then we have $p_1 = 1 - \frac{\hat{u} - d_2}{\delta d_1}$ $p_2 = \frac{\bar{u} - d_1}{\delta d_2}$ Now we must show $p_2, p_1 \in (0, 1)$. Since $\bar{u} - d_i > 0$, $p_1 < 1$, $p_2 > 0$. Our given indicates $\delta < \sqrt{1 - \frac{u}{d_1}} < \sqrt{1 - \frac{u}{d_2}}$ because $d_2 > d_1$. Then $\forall\ i \in \{1, 2\}$, $\delta < \sqrt{1 - \frac{u}{d_i}}$. We know $\delta^2 < 1 - \frac{u}{/}d_i$ and this implies $u < (1 - \delta^2)d_i$. Let $j = \{1, 2\}\backslash\{i\}$, then $u < (1 - \delta^2)d_i + (1 - \delta)d_j$ which implies $\bar{u} < \delta d_i + d_j$. Therefore, $p_2, p_1 \in (0, 1)$ $\square$

Proposition C.1. *For any even $n \in 2\mathbb{N}$, there exists a stochastic solution where*

$$\begin{aligned}
p_i &= 1 - \frac{\bar{u} + \frac{n-i-1}{2}u + \frac{i-1}{2}\delta u - \sum_{m=i+1}^{n/2} d_{2m} + \delta \sum_{m=i}^{\frac{m-1}{2}} d_{2m+1}}{\delta d_i}, \quad \forall i < n, i \in 2\mathbb{N} + 1 \\
p_j &= \frac{\bar{u} + \frac{j-2}{2}u\frac{n-i}{2}\delta u - \sum_{m=1}^{m/2} d_{2m-1} - \delta \sum_{m=i}^{\frac{m-2}{2}} d_{2m}}{\delta d_j} \forall j < n, \quad j \in 2\mathbb{N}
\end{aligned} \tag{C.22}$$

provided p_i, $p_j \in (0, 1)$.

Proof. For even n, we may write

$$\begin{aligned} V_1 &= p_1\hat{u} + (1-p_1)\delta V_2, \\ V_i &= p_i(u + \delta V_{i-1}) + (1-p_i)\delta V_{i+1}, \quad \forall i,\ i \in \mathbb{N},\ 1 < i < n \\ V_n &= p_n(u + \delta V_{n-1}). \end{aligned} \tag{C.23}$$

Then solving for p_i we have

$$\begin{aligned} p_1 &= \frac{V_1 - \delta V_2}{\bar{u} - \delta V_2}, \\ p_i &= \frac{V_i - \delta V_{i-1}}{u + \delta V_{i-1} - \delta V_{i+1}}, \quad \forall i,\ i \in \mathbb{N},\ 1 < i < n \\ p_n &= \frac{v_n}{u + \delta V_{n-1}}. \end{aligned} \tag{C.24}$$

If we have a stochastic solution we also have indifference between the two joining nodes. That is,

$$\begin{aligned} d_1 &= \bar{u} - \delta V_2 \\ d_i &= u + \delta V_{i-1} - \delta V_{i+1} \quad \forall i,\ i \in \mathbb{N},\ 1 < i < n \\ d_n &= u + \delta V_{n-1}. \end{aligned} \tag{C.25}$$

Solving for V_i we have

$$\begin{aligned} V_i &= \frac{1}{\delta}\left(\sum_{m=i}^{n/2} d_{2m} - \frac{n-i-1}{2}u\right) \quad \forall i < n, i \in 2\mathbb{N}-1 \\ V_i &= \frac{1}{\delta}\left(\bar{u} + \frac{i-2}{2}u - \sum_{m=1}^{n/2-1} d_{2m-1}\right). \end{aligned} \tag{C.26}$$

Substituting equations C.26 into C.24 yields equations C.22. □