

**Modeling the Psychology of Consumer and Firm Behavior
with Behavioral Economics***

Teck H. Ho
University of California, Berkeley
Berkeley, CA 94720
Email: hoteck@haas.berkeley.edu

Noah Lim
University of Houston
Houston, TX 77204
Email: noahlim@uh.edu

Colin F. Camerer
California Institute of Technology
Pasadena, CA 91125
Email: camerer@hss.caltech.edu

* Direct correspondence to the first author. This research is partially supported by NSF Grant SBR 9730187. We thank Wilfred Amaldoss, Botond Koszegi, George Loewenstein, John Lynch, Robert Meyer, Drazen Prelec, and Matt Rabin for their helpful comments. We are especially grateful to the late journal editor, Dick Wittink, for inviting and encouraging us to undertake this review. Dick was a great supporter of inter-disciplinary research. We hope this review can honor his influence and enthusiasm by spurring research that spans both marketing and behavioral economics.

ABSTRACT

Marketing is an applied science that tries to explain and influence how firms and consumers actually behave in markets. Marketing models are usually applications of economic theories. These theories are general and produce precise predictions, but they rely on strong assumptions of rationality of consumers and firms. Theories based on rationality limits could prove similarly general and precise, while grounding theories in psychological plausibility and explaining facts which are puzzles for the standard approach.

Behavioral economics explores the implications of limits of rationality. The goal is to make economic theories more plausible while maintaining formal power and accurate prediction of field data. This review focuses selectively on six types of models used in behavioral economics that can be applied to marketing.

Three of the models generalize consumer preference to allow (1) sensitivity to reference points (and loss-aversion); (2) social preferences toward outcomes of others; and (3) preference for instant gratification (quasi-hyperbolic discounting). The three models are applied to industrial channel bargaining, salesforce compensation, and pricing of virtuous goods such as gym memberships. The other three models generalize the concept of game-theoretic equilibrium, allowing decision makers to make mistakes (quantal response equilibrium), encounter limits on the depth of strategic thinking (cognitive hierarchy), and equilibrate by learning from feedback (self-tuning EWA). These are applied to marketing strategy problems involving differentiated products, competitive entry into large and small markets, and low-price guarantees.

The main goal of this selected review is to encourage marketing researchers of all kinds to apply these tools to marketing. Understanding the models and applying them is a technical challenge for marketing modelers, which also requires thoughtful input from psychologists studying details of consumer behavior. As a result, models like these could create a common language for modelers who prize formality and psychologists who prize realism.

1. INTRODUCTION

Economics and psychology are the two most influential disciplines that underlie marketing. Both disciplines are used to develop models and establish facts,¹ in order to better understand how firms and customers *actually* behave in markets, and to give advice to managers.² While both disciplines have the common goal of understanding human behavior, relatively few marketing studies have integrated ideas from the two disciplines. This paper reviews some of the recent research developments in “behavioral economics”, an approach which integrates psychological insights into formal economic models. Behavioral economics has been applied fruitfully in business disciplines such as finance (Barberis and Thaler 2003) and organizational behavior (Camerer and Malmendier forthcoming). This review shows how ideas from behavioral economics can be used in marketing applications, to link the psychological approach of consumer behavior to the economic models of consumer choice and market activity.

Because behavioral economics is growing too rapidly to survey thoroughly in an article of this sort, we concentrate on six topics. Three of the topics are extensions of the classical utility function, and three of the topics are alternative methods of game-theoretic analysis to the standard Nash-Equilibrium analysis.³ A specific marketing application is described for each idea.

It is important to emphasize that the behavioral economics approach *extends* rational-choice and equilibrium models; it does not advocate abandoning those models entirely. All of the new preference structures and utility functions described here generalize the standard approach by adding one or two parameters, and the behavioral game theories generalize standard equilibrium concepts in many cases as well. Adding parameters allows us to detect when the standard models work well and when they fail, and to measure empirically the importance of extending the standard models. When the standard methods fail, these new tools can then be used as default alternatives to describe and influence markets. Furthermore,

¹ The group which uses psychology as its foundational discipline is called “behavioral researchers” and the group which uses economics is called “modelers”. Unlike economics and psychology where groups are divided based on problem domain areas, the marketing field divides itself mainly along methodological lines.

² Marketing is inherently an applied field. We are always interested in both the descriptive question of how actual behavior occur and the prescriptive question of how one can influence behavior in order to meet a certain business objective.

³ There are several reviews of the behavioral economics area aiming at the economics audience (Camerer 1999, McFadden 1999, Rabin 1998; 2002). Camerer et al (2003b) compiles a list of key readings in behavioral economics and Camerer et al (2003a) discusses the policy implications of bounded rationality. Our review reads more like a tutorial and is different in that we show how these new tools can be used and we focus on how they apply to typical problem domains in marketing.

there are usually many delicate and challenging theoretical questions about model specifications and implications which will engage modelers and lead to progress in this growing research area.

1.1 Desirable Properties of Models

Our view is that models should be judged according to whether they have four desirable properties—generality, precision, empirical accuracy, and psychological plausibility. The first two properties, generality and precision, are prized in formal economic models. The game-theoretical concept of Nash equilibrium, for example, applies to any game with finitely-many strategies (it is general), and gives exact numerical predictions about behavior with zero free parameters (it is precise). Because the theory is sharply defined mathematically, little scientific energy is spent debating what its terms mean. A theory of this sort can be taught around the world, and used in different disciplines (ranging from biology to political science), so that scientific understanding and cross-fertilization accumulates rapidly.

The third and fourth desirable properties that models can have—empirical accuracy and psychological plausibility— have generally been given more weight in psychology than in economics, until behavioral economics came along.⁴ For example, in building up a theory of price dispersion in markets from an assumption about consumer search, whether the consumer search assumption accurately describes experimental data (for example) is often considered irrelevant in judging whether the theory of market prices built on that assumption might be accurate. (As Milton Friedman influentially argued, a theory’s conclusions might be reasonably accurate even if its assumptions are not.) Similarly, whether an assumption is psychologically plausible— consistent with how brains works, and with data from psychology experiments—was not considered a good reason to reject an economic theory.

The goal in behavioral economics modeling is to have *all* four properties, insisting that models both have the generality and precision of formal economic models (using mathematics), *and* be consistent with psychological intuition and experimental regularity. Many psychologists believe that behavior is context-specific so it is impossible to have a common theory that applies to all contexts. Our view is that we don’t know whether general theories fail until general theories are compared to a set of separate customized models of different domains. In principle, a general theory could include context-sensitivity as part of the theory and would be very valuable.

⁴ We are ignoring some important methodological exceptions for the sake of brevity. For example, mathematical psychology theories of learning which were popular in the 1950s and 1960s, before the “cognitive revolution” in psychology, resembled modern economic theories like the EWA theory of learning in games described below, in their precision and generality.

The complaint that economic theories are unrealistic and poorly-grounded in psychological facts is not new. Early in their seminal book on game theory, Von Neumann and Morgenstern (1944) stressed the importance of empirical facts:

“...it would have been absurd in physics to expect Kepler and Newton without Tycho Brahe, and there is no reason to hope for an easier development in economics.”

Fifty years later, the game theorist Eric Van Damme (1999), a part-time experimenter, thought the same:

“Without having a broad set of facts on which to theorize, there is a certain danger of spending too much time on models that are mathematically elegant, yet have little connection to actual behavior. At present our empirical knowledge is inadequate and it is an interesting question why game theorists have not turned more frequently to psychologists for information about the learning and information processes used by humans.”

Marketing researchers have also created lists of properties that good theories should have, which are similar to those listed above. For example, Little (1970) advised that

“A model that is to be used by a manager should be simple, robust, easy to control, adaptive, as complete as possible, and easy to communicate with.”

Our criteria closely parallel Little’s. We both stress the importance of simplicity. Our emphasis on precision relates to Little’s emphasis on control and communication. Our generality and his adaptive criterion suggest that a model should be flexible enough so that it can be used in multiple settings. We both want a model to be as complete as possible so that it is both robust and empirically grounded.⁵

1.2 Six Behavioral Economics Models and their Applications to Marketing

Table 1 shows the three generalized utility functions and three alternative methods of game-theoretic analysis which are the focus of this paper. Under the generalized preference structures, decision makers care about both the final outcomes as well as changes in outcomes with respect to a reference point and they are loss averse. They are not purely self-interested and care about others’ payoffs. They exhibit a taste for instant gratification and are not exponential discounters as is commonly assumed. The new methods of game-theoretic analysis allow decision makers to make mistakes, encounter surprises, and learn in response to feedback over time. We shall also suggest how these new tools can increase the validity of marketing models with specific marketing applications.

⁵ See also Leeflang et al (2000) for a detailed discussion on the importance of these criteria in building models for marketing applications.

Table 1: Behavioral Economics Models

Behavioral Regularities	Standard Assumptions	New Specification (Reference Example)	New parameters (Behavioral Interpretation)	Marketing Application
I. Generalized Utility Functions				
Reference-Dependence and Loss Aversion	Expected Utility Hypothesis	Reference-Dependent Preferences Kahneman and Tversky (1979)	ω (weight on transaction utility) μ (loss-aversion coefficient)	Business-to-Business Pricing Contracts
Fairness and Social Preferences	Pure Self-Interest	Inequality Aversion Fehr and Schmidt (1999)	γ (envy when others earn more) η (guilt when others earn more)	Salesforce Compensation
Impatience and Taste for Instant Gratification	Exponential Discounting	Hyperbolic Discounting Laibson (1997)	β (preference for immediacy, “present bias”)	Price Plans for Gym Memberships
II. New Methods of Game-Theoretic Analysis				
Noisy Best-Response	Best-Response Property	Quantal Response Equilibrium McKelvey and Palfrey (1995)	λ (“better response” sensitivity)	Price Competition with Differentiated Products
Thinking Steps	Rational Expectations Hypothesis	Cognitive Hierarchy Camerer et al (2004)	τ (average number of thinking steps)	Market Entry
Adaptation and Learning	Instant Equilibration	Self-Tuning EWA Ho et al (2004)	λ (“better response” sensitivity)*	Lowest-Price Guarantees

*There are two additional behavioral parameters ϕ (change detection, history decay) and ζ (attention to foregone payoffs, regret) in the self-tuning EWA model. These parameters need not be estimated; they are calculated based on feedback.

This paper makes three contributions:

1. Describe some important generalizations of the standard utility function and robust alternative methods of game-theoretic analysis. These examples show that it is possible to simultaneously achieve generality, precision, empirical accuracy and psychological plausibility with behavioral economics models.
2. Demonstrate how each generalization and new method of game-theoretic analysis works with a concrete marketing application example. In addition, we show how these new tools can influence how a firm goes about making its pricing, product, promotion, and distribution decisions with examples of further potential applications.
3. Discuss potential research implications for behavioral and modeling researchers in marketing. We believe this new approach is one sensible way to integrate research between consumer behavior and economic modeling.

The rest of the paper is organized as follows. In each of sections 2-7, we discuss one of the utility function generalizations or alternative methods of game-theoretic analysis listed in Table 1 and describe an application example in marketing using that generalization or method. Section 8 describes potential applications in marketing using these new tools. Section 9 discusses research implications for behavioral researchers and modelers and how they can integrate their research to make their models more predictive of market behavior. The paper is designed to be appreciated by two audiences. We hope that psychologists, who are uncomfortable with broad mathematical models, and suspicious of how much rationality is ordinarily assumed in those models, will appreciate how relatively simple models can capture psychological insight. We also hope that mathematical modelers will appreciate the technical challenges in testing these models and in extending them to use the power of deeper mathematics to generate surprising insights about marketing.

2. REFERENCE DEPENDENCE

2.1 Behavioral Regularities

In most applications of utility theory, the attractiveness of a choice alternative depends on only the final outcome that results from that choice. For gambles over money outcomes, utilities are usually defined over final states of wealth (as if different sources of income which are fungible are combined in a single “mental account”). Most psychological judgments of sensations, however, are sensitive to points of reference. This reference-dependence suggests decision makers may care about changes in outcomes as well as the final outcomes themselves. Reference-dependence, in turn, suggests that when the point of reference against which

outcomes is compared is changed (due to “framing”), the choices people make are sensitive to the change in frame. A well-known and dramatic example of this is the “Asian disease” experiment in Tversky and Kahneman (1981):

Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:

“Gains” Frame

If Program A is adopted, 200 people will be saved. (72%)

If Program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved. (28%)

“Loss” Frame

If Program C is adopted, 400 people will die. (22%)

If Program D is adopted, there is one-third probability that nobody will die and a two-thirds probability that 600 people will die. (78%)

In this empirical example, one group of subjects ($n=152$) were asked to choose between Programs A and B. Another group ($n=155$) choose between Programs C and D. The percentages of program choice are indicated in parentheses above. Note that Programs A and C yield the same final outcomes in terms of the actual number of people who will live and die Programs B and D have the same final outcomes too. If decision makers care only about the final outcomes, the proportion of decision makers choosing A (or B) in the first group should be similar to that choosing C (or D) in the second group. However, the actual choices depend dramatically on whether the programs are framed as gains or losses. When the problem is framed in terms of gains, the reference point is the state where no lives are saved, whereas when framed as losses, the reference point becomes the state where no lives are lost. In the “Gains” frame, most decision makers choose the less risky option (A) while they choose the more risky option (D) in the “Loss” frame. In other words, decision makers are sensitive to the manipulation of reference point and are risk-averse in gain domains but risk-seeking in loss domains. Framing effects like these have been replicated in many studies (see Camerer 1995 for a review), including gambles for real money (Camerer 1988), although the results sometimes depend on features of the problem.

The concept of reference-dependence preference has also been extended to the analysis of choice without risk (Tversky and Kahneman 1991). In a classic experiment that has been replicated many times, one group of subjects is endowed with a simple consumer good, such as a coffee mug or expensive pen. The subjects who are endowed with the good are asked the least

amount of money they would accept to sell the good. Subjects who are not endowed with the good are asked how much they would pay to buy one. Most studies find a striking “instant endowment effect”: Subjects who are endowed with the good name selling prices which are about twice as large as the buying prices. This endowment effect (Thaler 1980) is thought to be due to a disproportionate aversion to giving up or losing from one’s endowment, compared to the value of gaining, an asymmetry called “loss aversion”. Endowing an individual with an object shifts one’s reference point to a state of ownership and the difference in valuations demonstrates that the disutility of losing a mug is greater than the utility of gaining it.

There is an emerging neuroscientific basis for reference-dependence and loss aversion. Using fMRI analysis, Knutson and Peterson (2005) finds different regions of activity for monetary gain and loss. Recordings of activity in single neurons of monkeys show that neural firing rates respond to relative rather than the absolute levels of stimuli (Schultz and Dickinson 2000).⁶

Like other concepts in economic theory, loss-aversion appears to be general in that it spans domains of data (field and experimental) and many types of choices (see Camerer 2001, 2005). Table 2 below summarizes some economic domains where loss-aversion has been found. The domain of most interest to marketers is the asymmetry of price elasticities (sensitivity of purchases to price changes) for price increases and decreases. Elasticities are larger for price increases than for decreases, which means that demand falls more when prices go up than it increases when prices go down. Loss-aversion is also a component of models of context-dependence in consumer purchase, such as the compromise effect (Simonson 1989, Simonson and Tversky 1992, Tversky and Simonson 1993, Kivetz et al 2004). Loss-aversion has been suggested by finance studies of the large premium in returns to equity (stocks) relative to bonds and the surprisingly few number of announcements of negative corporate earnings and negative year-to-year earnings changes. Cab drivers appear to be averse toward “losing” by falling short of a daily income target (reference point), so they supply labor until they hit that target. Disposition effects refer to the tendency to hold on to money-losing assets (stocks and housing) too long, rather than sell and recognize accounting losses. Loss-aversion also appears at industry levels, creating “anti-trade bias”, and in micro decisions of monkeys trading tokens for food rewards.⁷

⁶ That is, receiving a medium squirt of juice, when the possible squirts were small or medium, activates reward-encoding neuron more strongly than when the same medium squirt is received, and the foregone reward was a large squirt.

⁷ The “endowment effect” has been subject to many “stress” tests. Plott and Zeiler (2005) find that endowment effects may be sensitive to the experimental instructions used. Unlike Camerer et al (1997), Farber (2004, 2005) finds only limited evidence of income-target labor supply of cab drivers. Trading experience can also help to reduce the degree of endowment effects. For example List (2003) finds that

Table 2: Evidence of Loss Aversion

Economic Domain	Study	Type of Data	Estimated Loss Aversion Coefficient
Instant endowment effects for goods	Kahneman et al (1990)	Field data (survey), goods experiments	2.29
Choices over money gambles	Kahneman and Tversky (1992)	Choice experiments	2.25
Asymmetric price elasticities	Putler (1992) Hardie et al (1993)	Supermarket scanner data	2.40 1.63
Loss-aversion for goods relative to money	Bateman et al (forthcoming)	Choice experiments	1.30
Loss-aversion relative to initial seller "offer"	Chen et al (2005)	Capuchin monkeys trading tokens for stochastic food rewards	2.70
Aversion to losses from international trade	Tovar (2004)	Non-tariff trade barriers, US 1983	1.95-2.39
Reference-dependence in two-part distribution channel pricing	Ho-Zhang (2005)	Bargaining experiments	2.71
Surprisingly few announcements of negative EPS and negative year-to-year EPS changes	DeGeorge et al (1999)	Earnings per share (EPS) changes from year to year for US firms	n.r.*
Disposition effects in housing	Genesove & Mayer (2001)	Boston condo prices 1990-1997	n.r.
Disposition effects in stocks	Odean (1998)	Individual investor stock trades	n.r.
Disposition effects in stocks	Weber and Camerer (1998)	Stock trading experiments	n.r.
Daily income targeting by NYC cab drivers	Camerer et al(1997)	Daily hours-wages observations (three data sets)	n.r.
Equity premium puzzle	Benartzi and Thaler (1995)	US stock returns	n.r.
Consumption: Aversion to period utility loss	Chua and Camerer (2004)	Savings-consumption experiments	n.r.

*n.r. indicates that the studies did not estimate the loss aversion coefficient directly.

endowment effects disappear among experienced traders of sports collectibles. Genesove and Mayer (2001) find lower loss-aversion among owners who invest in housing, compared to owners who live in their condominiums, and Weber and Camerer (1998) find that stockholders do not buy back losing stocks if they are automatically sold, in experiments. Kahneman et al (1990:1328) anticipated this phenomenon, noting that "there are some cases in which no endowment effect would be expected, such as when goods are purchased for resale rather than for utilization."

2.2 The Generalized Model

The Asian Disease example, the endowment effect, and the other empirical evidence, suggests that a realistic model of preference should capture the following three empirical regularities:

1. Outcomes are evaluated as changes with respect to a reference point. Positive changes are framed as gains or negative changes as losses.
2. Decision makers are risk-averse in gain domains and risk-seeking in loss domains (the “reflection effect”).
3. Decision makers are loss-averse. That is, losses generate proportionally more disutility than equal-sized gains.

Prospect theory (Kahneman and Tversky 1979) is the first formal model of choice that captures these three empirical regularities. Extending their insight, Koszegi and Rabin (2004) model individual utility $u(x | r)$ so that it depends on both the final outcome (x) and a reference point (r). Specifically, $u(x | r)$ is defined as:

$$u(x | r) \equiv v(x) + t(x | r)$$

where $v(x)$ represents the intrinsic utility associated with the final outcome (independent of the reference point) and $t(x | r)$ is the transaction or change utility associated with gains and losses relative to the reference point r .⁸ This model generalizes the neoclassical utility function by incorporating a transaction component into the utility function. If $t(x | r) = 0$ the general function reduces to the standard one used in rational choice theory. An important question is how the reference point is determined. We will generally use the typical assumption that the reference point reflects the status quo before a transaction, but richer and more technically interesting approaches are worth studying.

We assume $v(x)$ is concave in x . For example, the intrinsic utility can be a power function⁹ given by $v(x) = x^k$. In Koszegi and Rabin’s formulation, $t(x | r)$ is assumed to have several simple properties. First assume $t(x | r) = t(x - r)$ and define $t(y) = t(x | r)$ to economize on notation. The crucial property of $t(y)$ is

⁸ This functional form makes psychological sense because it is unlikely that changes from a reference point are the *only* carrier of utility. If so, then a salesperson expecting a year-end bonus of \$100,000 and receiving only \$95,000 would be just as unhappy as one expecting \$10,000 and getting only \$5,000. The two-piece function also allows us to compare standard reference-independent preferences as a special case (when $t(x|r)=0$) of the more general form.

⁹ In the power form, $v(x) = -x^k$ if x is negative.

$\frac{t'_-(0)}{t'_+(0)} \equiv \mu > 1$, where $t'_+(0) \equiv \lim_{y \rightarrow 0} t'(|y|)$ and $t'_-(0) \equiv \lim_{y \rightarrow 0} t'(-|y|)$. The parameter μ is

the coefficient of loss-aversion: it measures the marginal utility of going from a small loss to zero, relative to the marginal utility of going from zero to a small gain. In a conventional (differentiable) utility function $\mu = 1$. If $\mu > 1$ then there is a “kink” at the reference point.

A simple $t(y)$ function that satisfies the Koszegi-Rabin properties is:

$$t(y) = \begin{cases} \omega \cdot v(y) & \text{if } y \geq 0, \\ -\mu \cdot \omega \cdot v(|y|) & \text{if } y < 0, \end{cases}$$

where $\omega > 0$ is the weight on the transaction utility relative to the intrinsic utility $v(x)$ and μ is the loss-aversion coefficient.

This reference-dependent utility function can be used to explain the endowment effect. Suppose a decision-maker has preferences over amounts of pens and dollars, denoted $x = (x_p, x_d)$. Because there are two goods, the reference point also will have two dimensions, $r = (r_p, r_d)$. Since the choice involves two dimensions, a simple model is to assume that the intrinsic utilities for pens and dollars can be evaluated separately and added up, so that $v(x) = v(x_p) + v(x_d)$. Make the same assumption for the transactional components of utility, $t(x|r) = t(y) = t(y_p) + t(y_d)$ where $y_p = x_p - r_p$ and $y_d = x_d - r_d$, as well. For simplicity, we let $k = 1$ so that $v(x_p) = bx_p$ and $v(x_d) = x_d$, where $b > 0$ represents the relative preference for pens over dollars. The decision maker's utility can now be expressed as $u(x_p, x_d; y_p, y_d) = bx_p + x_d + t(y_p) + t(y_d)$, where

$$t(y_p) = \begin{cases} \omega \cdot v(y_p) = \omega \cdot b \cdot y_p & \text{if } y_p \geq 0 \\ -\mu \cdot \omega \cdot v(|y_p|) = -\mu \cdot \omega \cdot b \cdot |y_p| & \text{if } y_p < 0 \end{cases}$$

$$t(y_d) = \begin{cases} \omega \cdot v(y_d) = \omega \cdot y_d & \text{if } y_d \geq 0 \\ -\mu \cdot \omega \cdot v(|y_d|) = -\mu \cdot \omega \cdot |y_d| & \text{if } y_d < 0 \end{cases}$$

In a typical endowment effect experiment, there are three treatment conditions—choosing, selling, and buying. In the first treatment, subjects are asked to state a dollar amount, their “choosing price” P_C (or cash value), such that they are indifferent between gaining a pen or gaining the amount P_C . Since they are not endowed with anything, the reference points are $r_p = r_d = 0$. The utility from gaining 1 pen is the pen's intrinsic utility, which is bx_p , or simply b (since one pen means $x_p = 1$). The transaction difference is $y_p = x_p - r_p = 1 - 0 = 1$.

Given the specification of $t(y_p)$ above (and the fact that the transaction is a gain), the transaction utility is $\omega \cdot b \cdot 1$. Therefore, the total utility from gaining 1 pen is

$$\text{Utility (gain 1 pen)} = b + \omega \cdot b$$

A similar calculation for the dollar gain P_C and its associated transaction utility gives

$$\text{Utility (gain } P_C) = P_C + \omega \cdot P_C$$

Since the choosing price P_C is fixed to make the subject indifferent between gaining the pen and gaining P_C , one solves for P_C by equating the two utilities, $P_C + \omega \cdot P_C = b + \omega \cdot b$, which yields $P_C = b$.

In the second treatment, subjects are asked to state a price P_S which makes them just willing to sell the pen they are endowed with. In this condition the reference points are $r_p = 1$ and $r_d = 0$. The intrinsic utilities from having no pen and gaining P_S are $0 + P_S$. The transaction differences are $y_p = -1$ and $y_p = P_S$. Plugging these into the $t(y)$ specification (keep in mind that $y_p < 0$ and $y_d > 0$) and adding up all the terms gives

$$\text{Utility (lose 1 pen, gain } P_S) = P_S - \mu \cdot \omega \cdot b \cdot 1 + \omega \cdot P_S$$

The utility of keeping the pen is $\text{Utility (keep 1 pen, gain 0)} = b$ (there are no transaction utility terms because the final outcome is the same as the reference point. Since P_S is the price which makes the subject indifferent between selling the pen at that price, the value of P_S must make the two utilities equal. Equating and solving gives $P_S = \frac{b(1 + \mu\omega)}{1 + \omega}$.

In the third treatment, subjects are asked to state a maximum buying price P_B for a pen. Now the reference points are $r_p = r_d = 0$. The intrinsic utilities are $b \cdot 1$ and $-P_B$ for pens and dollars respectively. Since the pen is gained, and dollars lost, the transaction differences are $y_p = 1$ and $y_d = -P_B$. Using the $t(y)$ specification on these differences and adding up terms gives a total utility of:

$$\text{Utility (gain 1 pen, lose } P_B) = b \cdot 1 - P_B + \omega \cdot b \cdot 1 - \mu \cdot \omega \cdot P_B$$

Since the buying price is the maximum, the net utility from the transaction must be zero. Setting the above equation to 0 and solving gives $P_B = \frac{b(1 + \omega)}{1 + \mu\omega}$. Summarizing results in the three

treatments, when $\omega > 0$ and $\mu > 1$ the prices are ranked $P_S > P_C > P_B$ because

$\frac{b(1 + \mu\omega)}{1 + \omega} > b > \frac{b(1 + \omega)}{1 + \mu\omega}$. That is, selling prices are higher than choosing prices, which are higher than buying prices. But note that if either $\omega = 0$ (transaction utility does not matter) or $\mu = 1$ (there is no loss-aversion), then all three prices are equal to the value of the pen b , so there is no endowment effect.¹⁰

2.3 Marketing Application: Business-to-Business Pricing Contracts

A classic problem in channel management (and in industrial organization more generally) is the “channel coordination” or “double marginalization” problem. Suppose an upstream firm (a manufacturer) offers a downstream firm (a retailer) a simple linear price contract, charging a fixed price per unit sold. This simple contract creates a subtle inefficiency: When the manufacturer and the retailer maximize their profits independently, the manufacturer does not account for the externality of its pricing decision on the retailer’s profits. If the two firms become vertically integrated and so that in the merged firm the manufacturing division sells to the retailing division using an internal transfer price, the profits of the merged firm would be higher than the total profits of the two separate firms, because the externality becomes internalized.

Moorthy (1987) had the important insight that even when manufacturer and retailer operates separately, the total channel profits can be equal to that attained by a vertically-integrated firm if the manufacturer offers the retailer a two-part tariff (TPT) contract that consists of a lump-sum fixed fee F and a marginal wholesale per-unit price w . In this simplest of nonlinear pricing contracts, the manufacturer should simply set w at its marginal cost. Marginal-cost pricing eliminates the externality and induces the retailer to buy the optimal quantity. However, marginal-cost pricing does not enable the manufacturer to make any profits, but setting a lump-sum fixed F does so. The retailer then earns the markup (retail price p minus w) on each of q units sold, less the fee F , for a total profit of $(p - w) \cdot q - F$.

While two-part tariffs are often observed in practice, it is difficult to evaluate whether they lead to efficiency as theory predicts. Furthermore, there are no experiments showing whether subjects set fees F and wholesale unit prices w at the levels predicted by the theory. A behavioral possibility is that two-part contracts might seem aversive to retailers, because they suffer an immediate loss from the fixed fee F , but perceive later gains from selling above the

¹⁰ In standard consumer theory, selling prices should be very slightly higher than buying prices because of a tiny “wealth effect” (prospective sellers start with more “wealth” in the form of the pen, than buyers do). Rational consumers can choose to effectively “spend” some of their pen-wealth on a pen, by asking a higher selling price. This effect disappears in our analysis because of the linear assumption of the utility function x^k which is assumed for simplicity.

wholesale price w they are charged. If retailers are loss-averse they may resist paying a high fee F even if it is theoretically efficiency-enhancing.

Ho and Zhang (2004) did the first experiments on the use of two-part tariffs in a channel and study their behavioral consequences. The results show that contrary to the theoretical prediction, channel efficiency (the total profits of the two separate firms relative to the theoretical 100% benchmark for the vertically integrated firm) is only 66.7%. The standard theoretical predictions and some experimental statistics are shown in Table 3. These data show that the fixed fees F are too low compared to the theoretical prediction (actual fees are around 5, when theory predicts 16). Since F is too low, to maintain profitability the manufacturers must charge a wholesale price w which is too high (charging around 4, rather than the marginal cost of 2). As a result, the two-part contracts are often rejected by retailers.

The reference-dependence model described in the previous section can explain the deviations of the experimental data from the theoretical benchmark. With a two-part tariff contract, the retailer's transaction utility occurs in two stages. In the first stage, it starts out with a reference profit of zero but is loss averse with respect to paying the fixed fee F . Its transaction utility if it accepts the contract is simply $-\omega \cdot \mu \cdot F$ where ω is the retailer's weight on the transaction component of utility and μ is the loss aversion coefficient as specified in the previous section. In the second stage, the retailer realizes a final profit of $(p - w) \cdot q - F$, which represents a gain of $(p - w) \cdot q$ relative to a reference point of $-F$ (its new reference point after the first stage). Hence, its transaction utility in the second stage is $\omega \cdot ((p - w) \cdot q)$. The retailer's overall utility is the intrinsic utility from net profits $(p - w) \cdot q - F$ in the entire game, plus the two components of transaction utility, $-\omega \cdot \mu \cdot F$ and $\omega \cdot ((p - w) \cdot q)$. Adding all three terms gives a retailer utility U_R of:

$$U_R = (1 + \omega) \left[(p - w) \cdot q - \left(\frac{1 + \omega \cdot \mu}{1 + \omega} \right) F \right]$$

Note that when $\omega = 0$, the reference-dependent model reduces to the standard economic model, and utility is just the profit of $(p - w) \cdot q - F$. When $\mu = 1$ (no loss aversion) the model just scales up retailer profit by a multiplier $(1 + \omega)$, which reflects the hedonic value of an above-reference-point transaction. When there is loss-aversion ($\mu > 1$), however the retailer's perceived loss after paying the fee F has a disproportionate influence on overall utility. Using the experimental data, the authors estimated the fixed fee multiplier $\frac{1 + \omega \cdot \mu}{1 + \omega}$ to be 1.57, much larger than the 1.0 predicted by standard theory with $\omega = 0$ or $\mu = 1$. Given this

estimate, Table 3 shows predictions of crucial empirical statistics, which generally match the wholesale and retail prices (w and p), the fees F , and the contract rejection rate, reasonably well. A value of $\omega = 0.5$ implies a loss aversion coefficient of 2.71.

Table 3: Two-Part Tariff Model Predictions and Experimental Results

Decisions	Standard Theory Prediction*	Experimental Data	Reference-Dependence Prediction
Wholesale cost w	2	4.05	4.13
Fixed fee F	16	4.61	4.65
Reject contract?	0%	28.80%	34.85%
Retail price p	6	6.82	7.06

*marginal cost of the manufacturer is 2, demand is $q=10-p$ in the experiment.

3. SOCIAL PREFERENCES

3.1 Behavioral Regularities

Standard economic models usually assume that individuals are purely self-interested, that is, they only care about earning the most money for themselves. Self-interest is a useful simplification but is clearly a poor assumption in many cases. Self-interest cannot explain why decision makers seem to care about fairness and equality, are willing to give up money to achieve more equal outcomes, or to punish others for actions which are perceived as selfish or unfair. This type of behavior points to the existence of social preferences, which defines a person's utility as a function of her own payoff *and* others' payoffs.

The existence of social preferences can be clearly demonstrated in an "ultimatum" price-posting game. In this game, a monopolist retailer sells a product to a customer by posting a price p . The retailer's marginal cost for the product is zero and the customer's willingness-to-pay for the product is \$1. The game proceeds as follows: the retailer posts a price $p \in [0,1]$ and the customer chooses whether or not to buy the product. If she buys, her consumer surplus is given by $1-p$, while the retailer's profit is p ; if she chooses not to buy, each party receives a payoff of zero. If both parties are purely self-interested and care only about their own payoffs, the unique subgame perfect equilibrium to this game would be for the retailer to charge $p = 0.99$ (assuming the smallest unit of money is a penny), anticipating that the customer accepts the price and earns a penny of surplus. Many experiments have been conducted to test the validity of this prediction in such "ultimatum" games (Camerer 2003, chapter 2). The results are markedly different from the prediction of the pure self-interested model and are characterized by three empirical regularities: (1) The average prices are in the region of \$0.60 to \$0.70, with the median and modal prices in the interval [\$0.50, \$0.60]; (2) There are hardly any prices above \$0.90 and very high prices often result in no purchases (rejections) – for example,

prices of \$0.80 and above yield no purchases about half the time; and (3) There are almost no prices in the range of $p < \$0.50$; that is, the retailer rarely gives more surplus to the consumer than to itself.

These results can be easily explained as follows: customers have social preferences which lead them to sacrifice part of their own payoffs to punish what they consider an unfair price, particularly when the retailer's resulting monetary loss is higher than that of the customer. The retailer's behavior can be attributed to both social preferences and strategic behavior: They either dislike creating unequal allocations, or they are selfish but rationally anticipate the customers' concerns for fairness and lower their prices to maximize profits.

3.2 The Generalized Model

One way to capture a concern for fairness mathematically is by applying models of inequality aversion. These models assume that decision makers are willing to sacrifice to achieve more equitable outcomes if they can. Fehr and Schmidt (1999) formalize a simple model of inequity-aversion¹¹ in terms of differences in players' payoffs. Their model puts different weight on the payoff difference depending on whether the other player earns more or less. For the two-player model (denoted 1 and 2), the utility of player 1 is given by:

$$U_1(x_1, x_2) = \begin{cases} x_1 - \eta \cdot (x_1 - x_2), & x_1 > x_2 \\ x_1 - \gamma \cdot (x_2 - x_1) & x_2 > x_1 \end{cases}$$

where $\gamma \geq \eta$ and $0 \leq \eta < 1$.¹² In this utility function, γ captures the loss from disadvantageous inequality (envy), while η represents the loss from advantageous inequality (guilt). For example, when $\gamma = 0.5$ and Player 1 is behind, she is willing to give up a dollar only if it reduces Player 2's payoffs by \$3 or more (since the loss of \$1 is less than reduction in envy of 2γ). Correspondingly, if $\eta = 0.5$ and Player 1 is ahead, then she is just barely willing to give away enough to Player 2 to make them even (since giving away \$ x reduces the disparity by \$ $2x$, and hence changes utility by $-x + \eta \cdot 2 \cdot x$). The assumption $\gamma \geq \eta$ captures the fact that envy is stronger than guilt. If $\gamma = \eta = 0$, then the above model reduces to the standard pure self-interest model.

To see how this model can explain the empirical regularities of the ultimatum price-posting game, suppose that both the retailer and the customer have inequity-averse preferences

¹¹ Bolton and Ockenfels (2000) have a closely related model which assumes that decision makers care about their own payoffs and their relative *share* of total payoffs.

¹² The model is easily generalized to n players, in which case envy and guilt terms are computed separately for each opponent player, divided by $n-1$, and added up.

that are characterized by the specific parameters (γ, η) .¹³ Recall that if both of them are purely self-interested, that is $\gamma = \eta = 0$, the retailer will charge the customer \$0.99, which the customer will accept. However, suppose we observe the customer reject a price of \$0.90. In this case, we know that γ must be greater than 0.125 if customers are rational (since rejecting earns 0, which is greater than $0.1 - \gamma(0.9 - 0.1)$ if and only if $\gamma > 0.125$). What is the equilibrium outcome predicted by this model? Customers with envy parameter γ are indifferent to rejecting a price offer of $p^* = \frac{1+\gamma}{1+2\gamma}$ (Rejecting gives $0 - \gamma(0 - 0)$ and accepting gives $1 - p - \gamma(p - (1 - p))$); setting these two expressions equal to be equal gives p^* . If we assume that retailers do not feel too much guilt, that is $\eta < 0.5$,¹⁴ then retailers will want to offer a price that customers will just accept. Their optimal price is therefore $p^* = \frac{1+\gamma}{1+2\gamma}$. If $\gamma = 0.5$ for example, the retailer's maximum price is \$0.75. This price is consistent with the empirical observations that p typically ranges from \$0.50 to \$0.70 and that very low offers are rejected. The model can also explain why that almost no retailer charge less than \$0.50 - because doing so results in less profit *and* in more envy.

Many other models of social preferences have been proposed. Charness and Rabin (2002) suggest a model in which players care about their own payoffs, the minimum payoff, and the total payoff. In a two-player game this model reduces to the Fehr-Schmidt form, but in multi-player games it can explain why one do-gooder player may sacrifice a small amount to create social efficiency.¹⁵

Inequality-aversion models are easy to use because a modeler can just substitute inequality-adjusted utilities for terminal payoffs in a game tree and use standard equilibrium concepts. Another class of models of social preferences is the "fairness equilibrium" approach of Rabin (1993) and Dufwenberg and Kirchsteiger (2004). In these models, players form beliefs about other players' kindness and other players' perceived kindness, and their utility function includes a term that multiplies a player's kindness (which can be positive or negative) with the

¹³ More insight can be derived from mixture models which specify distributions of (γ, η) across people or firms.

¹⁴ If retailers have $\eta = 0.5$ then they are indifferent between cutting the price by a small amount ε , sacrificing profit, to reduce guilt by $2\varepsilon\eta$, so any price in the interval $[\$0.50, p^*]$ is equally good. If $\eta > 0.5$ then they strictly prefer an equal-split price of \$0.50. Offering more creates too much guilt, and offering less creates envy.

¹⁵ The authors presented a more general 3-parameter model which captures the notion of reciprocity. We choose to ignore reciprocity and focus on a more parsimonious 2-parameter model of social preferences.

expected kindness of the other player. These models clearly capture the notion of reciprocity – players prefer to be positively kind to people who are positively kind to them, and to be hostile in response “negative kindness”. Assuming that beliefs are correct in equilibrium, one can derive a “fairness equilibrium” by maximizing the players’ utility functions. These models are harder to apply however, because branches in a game tree that were not chosen may affect the perceptions of kindness so backward induction cannot be applied in a simple way.

3.3 Marketing Application: Salesforce Compensation

The literature on salesforce management has mainly focused on how a manager should structure its compensation plans for a salesperson. If the effort level of the salesperson cannot be contracted upon or is not fully observable, then a self-interested salesperson will always want to shirk (provide the minimum level of effort) if effort is costly. Hence, the key objective for the manager (principal) revolves around designing incentive contracts that prevent moral hazard by the salesperson (agent). For example, an early paper by Basu et al (1985) shows that if a salesperson’s effort is not linked to output in a deterministic fashion, then the optimal compensation contract consists of a fixed salary and a commission component based on output.

Inequality-aversion and reciprocity complicate this simple view. If agents feel guilt or repay kindness with reciprocal kindness, then they will not shirk as often as models which assume self-interest predict (even in one-shot games where there are no reputational incentives). In fact, experimental evidence from Fehr et al (2004) suggests that incentive contracts that are designed to prevent moral hazard may not work as well as implicit bonus contracts if there is a proportion of managers and salespeople who care about fairness. Although the authors consider a slightly different principal-agent setting from that of the salesforce literature, their experimental findings are closely related and serve as a good potential application for marketing.

In their model, the manager can choose to offer the salesperson either of two contracts: a Bonus Contract (BC) or an Incentive Contract (IC). The salesperson’s effort e is observable, but any contract on effort must be verified by a monitoring technology which is costly. The costs of effort $c(e)$ are assumed to be convex (see Table 4 for experimental parameters).

Under the BC, the manager offers a contract (w, e^*, b^*) , where w is a prepaid wage, e^* is a requested effort level, and b^* is a promised bonus for the salesperson. However, both requested effort and the promised bonus are not binding, and there is no legal or reputational recourse. If the salesperson accepts the contract she earns the wage w immediately and chooses an effort e in the next stage. In the last stage, the manager observes effort e accurately and decides whether to pay an actual bonus $b \geq 0$. (which can be below, or even above, the

promised bonus b^*). The payoffs for the manager and salesperson with the BC are $\pi_M = 10 \cdot e - w - b$ and $\pi_S = w - c(e) + b$ respectively.

Under the IC, the manager can choose whether to invest $K=10$ in the monitoring technology. If she does, she offers the salesperson a contract (w, e^*, f) that consists of a wage w , a demanded effort e^* and a penalty f . The penalty f (which is capped at a maximum of 13 in this model) is automatically imposed if the manager verifies that the salesperson has shirked ($e < e^*$). While the monitoring technology is perfect when it works, it works only 1/3 of the time. With the IC, the expected payoffs for the manager and salesperson, if the manager invests in the monitoring technology are:¹⁶

$$\begin{aligned} \text{If } e \geq e^*, \\ \pi_M &= 10 \cdot e - w - K \\ \pi_S &= w - c(e) \\ \\ \text{If } e < e^*, \\ \pi_M &= 10 \cdot e - w - K + 0.33f \\ \pi_S &= w - c(e) - 0.33f \end{aligned}$$

Table 4: Effort Costs for Salesperson

E	1	2	3	4	5	6	7	8	9	10
$c(e)$	0	1	2	4	6	8	10	13	16	20

There is a large gain from exchange in this game if salespeople can be trusted to choose high effort. A marginal increase in one unit of effort earns the manager an incremental profit of 10, but costs the agent only 1 to 4 units. Therefore, the first-best outcome in this game is for the manager to forego investing in the monitoring technology and for the salesperson to choose $e=10$, giving a combined surplus of $10 \cdot e - c(e) = 80$. Under the IC, the optimal contract would be $(w=4, e^*=4, f=13)$ resulting in $\pi_M = 26$ and $\pi_S = 0$.¹⁷ With the BC, a self-interested manager will never pay any bonus in the last stage. Since the salesperson knows this, she will choose $e=1$. Therefore, the optimal contract will be $(w=0, e^*=1, b^*=0)$, yielding $\pi_M = 10$ and $\pi_S = 0$. Hence, if the manager has a choice between the two contracts, standard economic theory with self-interested preferences predicts that the manager will always choose the IC over the BC. Intuitively, if managers don't expect the salespeople to believe their bonus promises,

¹⁶ Alternatively, the manager can choose $K=0$ and offer only a fixed wage w .

¹⁷ The salesperson would choose the requested level of effort of 4 as deviating (by choosing an effort level of 0) leads to negative expected payoffs. Hence the manager earns $10 \cdot (4) - 4 - 10 = 26$ and the salesperson gets $4 - 4 = 0$.

and think salespeople will shirk, then they are better off asking for a modest enough effort ($e=4$), enforced by a probabilistic fine in the IC, so that the salespeople will put in some effort.

A group of subjects (acting as managers) were asked first to choose a contract form (either IC or BC) and then make offers using that contract form to another group of subjects (salespersons). Upon accepting a contract offer from a manager, a salesperson chose his effort level. Table 5 shows the theoretical predictions and the actual results of the data collected using standard experimental economics methodology (abstract instructions, no deception, repetition to allow learning and equilibration, and performance-based experimental payments).¹⁸

Table 5: Predicted and Actual Outcomes in the Salesforce Contract Experiment

	Incentive Contract (IC)		Bonus Contract (BC)	
	Prediction	Actual(Mean)	Prediction	Actual(Mean)
Manager's Decisions				
Choice (%)	100	11.6	0	88.4
Wage	4	24.0	0	15.2
Effort Requested	4	5.7	1	6.7
Fine	13	10.6	n.a.	n.a.
Bonus Offered	n.a.	n.a.	0	25.1
Bonus paid	n.a.	n.a.	0	10.4
Salesperson's Decisions				
Effort	4	2.0	1	5.0
Outcomes				
π_M	26	-9.0	10	27.0
π_S	0	14.4	0	17.8

Contrary to the predictions of standard economic theory, managers choose to offer the BC contract 88% of the time. Salespeople reciprocate by exerting a higher effort than necessary (an average of 5 out of 10) which is quite profitable for firms. In their paper, the authors also reported that actual ex-post bonus payments increase in actual effort, which implies that managers reward salespersons' efforts (like voluntary "tipping" in service professions). As a result of the higher effort levels, the payoffs for both the manager and the salesperson

¹⁸ We thank Klaus Schmidt for providing data that were not available in their paper.

(combined surplus) are higher in the BC than in the IC. Overall, these observed regularities cannot be reconciled with a model with purely self-interested preferences.

The authors show that the results of the experiment are consistent with the inequality-aversion model of Fehr and Schmidt (1999) when the proportion of fair-minded managers and salespersons (with $\gamma, \eta > 0.5$) in the market is assumed to be 40%. For the BC, there is a pooling equilibrium where both the self-interested and fair-minded managers offer $w=15$, with the fair-minded manager paying $b=25$ while the self-interested manager pays $b=0$ (giving an expected bonus of 10). The self-interested salesperson will choose $e=7$, while the fair-minded salesperson chooses $e=2$, giving an expected effort level of 5. The low effort exerted by the fair-minded salesperson is attributed to the fact she dislikes the inequality in payoffs whenever she encounters the self-interested manager with a probability of 0.6.

For the IC, the authors show that it is optimal for the self-interested manager to offer the contract ($w=4, e^*=4, f=13$). The fair-minded manager however will choose ($w=17, e^*=4, f=13$) that results in an equal division of surplus when $e=4$. A purely self-interested salesperson will accept and obey the contracts offered by both the self-interested and fair-minded managers. However, the fair-minded salesperson will only accept and obey the contracts of the fair-minded manager.

Comparing the BC and IC, the average level of effort is higher in the former (effort level of 5 versus 4), resulting in a higher expected combined surplus. Consequently, both the self-interested and fair-minded managers prefer the BC over the IC. This example illustrates how reciprocity can generate efficient outcomes in principal-agent relations when standard theory predicts rampant shirking.¹⁹

4. HYPERBOLIC DISCOUNTING

4.1 Behavioral Regularities

The Discounted-Utility (DU) framework is widely used to model intertemporal choice, in economics and other fields (including behavioral ecology in biology). The DU model assumes that decision makers make current choices which maximize the discounted sum of instantaneous utilities in future periods. The most common assumption is that decision makers discount the future utility at time t by an exponentially declining discount factor, $d(t)=\delta^t$

¹⁹ Other experiments show that the strength of reciprocal effort in similar “gift exchange” experiments is sensitive to framing effects (Hannan et al forthcoming) and to the gains from pure trust (Healy 2004). There probably are many other conditions which increase or decrease the strength of reciprocity. For example, self-serving bias in judgments of fairness (e.g., Babcock and Loewenstein 1997) will probably decrease it and communication will probably increase it.

(where $0 < \delta < 1$).²⁰ Formally, if u_τ is the agent's instantaneous utility at time τ , her intertemporal utility in period t , U^t , is given by:

$$U^t(u_t, u_{t+1}, \dots, u_T) \equiv u_t + \sum_{\tau=t+1}^T \delta^{\tau-t} u_\tau$$

The DU model was first introduced by Samuelson (1937) and has been widely adopted mainly due to the analytical convenience of “summarizing” agents’ future preferences by using a single constant parameter δ . The exponential function $d(t) = \delta^t$ is also the only form that satisfies *time-consistency* — when agents make plans based on anticipated future tradeoffs, they still make the same tradeoffs when the future arrives (provided there is no new information).

Despite its simplicity and normative appeal, many studies have shown that the DU model is problematic empirically.²¹ In economics, Thaler (1981) was the first to show that the per-period discount factor δ appears to decline over time (following Ainslie 1975 and others in psychology). Thaler asked subjects to state the amount of money they would require in 3 months, 1 year and 3 years later in exchange for receiving a sum of \$15 immediately. The respective median responses were \$30, \$60 and \$100, which imply average *annual* discount rates of 277% over 3 months, 139% over 1 year and 63% over 3 years. The finding that discount rates decline over time has been corroborated by many other studies (e.g., Benzion et al 1989, Holcomb and Nelson 1992, Pender 1996). Moreover, it has been shown that a hyperbolic discount function of the form $d(t) = 1/(1+mt)$ fits data on time preferences better than the exponential form does.

Hyperbolic discounting implies that agents are relatively farsighted when making tradeoffs between rewards at different times in the future, but pursue immediate gratification when it is available. Recent research in neuroeconomics (McClure et al 2004) suggests that hyperbolic discounting can be attributed to competition of neural activities between the affective and cognitive systems of the brain.²² A major consequence of hyperbolic discounting is that the behavior of decision makers will be *time-inconsistent*: decision makers might not make the same decision they expected they would (when they evaluated the decision in earlier

²⁰ The discount factor δ is also commonly written as $1/(1+r)$, where r is the discount rate.

²¹ In this section, we focus on issues relating to time discounting rather than other dimensions of intertemporal choice (Loewenstein 1987, Loewenstein and Thaler 1989, Loewenstein and Prelec 1992, 1993, Prelec and Loewenstein 1991). Frederick et al (2002) provides a comprehensive review of the literature on intertemporal choice. Zauberman and Lynch (2005) show that decision makers discount time resources more than money.

²² These findings also address to a certain extent the concerns of Rubinstein (2003) over the psychological validity of hyperbolic discounting. Rubinstein argues that an alternative model based on similarity comparisons is equally appealing, and presents some experimental evidence. Gul and Pesendorfer (2001) offer a different model which explains some of the same regularities as hyperbolic discounting, based on a preference for inflexibility (or a disutility from temptation).

periods) when the actual time arrives. Descriptively, this property is useful because it provides a way to model self-control problems and procrastination (e.g. O'Donoghue and Rabin 1999a).

4.2 The Generalized Model

A useful model to approximate hyperbolic discounting introduces one additional parameter into the standard DU framework. This generalized model is known as the $\beta - \delta$ “quasi-hyperbolic” or the “present-biased” model. It was first introduced by Phelps and Pollak (1968) to study transfers from parents to children, and then borrowed and popularized by Laibson (1997). With quasi-hyperbolic discounting, the decision maker’s weight on current (time t) utility is 1 while the weight on period τ 's utility ($\tau > t$) is $\beta\delta^{\tau-t}$. Hence, the decision maker’s intertemporal utility in period t , U^t , can be represented by:

$$U^t(u_t, u_{t+1}, \dots, u_T) \equiv u_t + \beta \sum_{\tau=t+1}^T \delta^{\tau-t} u_\tau$$

In the $\beta - \delta$ model, the parameter δ captures the decision maker’s “long-run” preferences, while β (which is between 0 to 1) measures the strength of the taste for immediate gratification or in other words, the degree of present bias. Lower values of β imply a stronger taste for immediacy. Notice that the discount factor placed on the next period after the present is $\beta\delta$, but the incremental discount factor between *any* two periods in the future is $\frac{\beta\delta^{t+1}}{\beta\delta^t} = \delta$. Decision

makers act today as if they will be more patient in the future (using the ratio δ), but when the future arrives the discount factor placed on the next period is $\beta\delta$. In the special case of $\beta = 1$, the model reduces to the standard DU framework. This special case is also important in that it is sometimes used as the benchmark by which the welfare effects of hyperbolic discounting are made. The (β, δ) model has been applied to study self-control problems such as procrastination and deadline-setting (O'Donoghue and Rabin 1999a, 1999b, 2001) and addiction (O'Donoghue and Rabin 1999c, 2002, Gruber and Koszegi 2001).²³

A natural question that arises is whether decision makers are aware that they are discounting hyperbolically. One way to capture agents’ self-awareness about their self-control is to introduce beliefs about their own future behavior (O'Donoghue and Rabin 2001, 2003). Let $\hat{\beta}$ denote the agent’s belief about β . Agents can be classified into two types. The first type is the naïf, who is totally unaware that he is a hyperbolic discounter and believes he discounts

²³ See O'Donoghue and Rabin (2000) for other interesting applications.

exponentially ($\beta < \hat{\beta} = 1$). The second type is the sophisticate ($\beta = \hat{\beta} < 1$), who is fully aware of his time-inconsistency and make decisions that rationally anticipate these problems.²⁴ The sophisticate will seek external self-control devices to commit himself to acting patiently in the future (Ariely and Wertenbroch 2002), but the naïf will not.

An example will illustrate how hyperbolic discounting and agents' beliefs about their preferences affect behavior. For simplicity, we assume $\delta = 1$. The decision maker faces two sequential decisions:

1. Purchase decision: In period 0, he must decide between buying a Small (containing 1 serving) or Large (containing 2 servings) pack of chips. The Large pack of chips comes with a quantity discount so it has a lower price per serving.
2. Consumption decision: In period 1, he must decide on the number of servings to consume. If he bought the Small pack, he can consume only 1 serving. However, if he bought the Large pack, he has to decide between eating 2 servings at once or eating 1 serving and conserving the second serving for future consumption.

The consumer receives an immediate consumption benefit as a function of the number of servings he eats minus the price per serving he paid. However, since chips are nutritionally unhealthy, there is a cost that is incurred in the period 2. This cost is a function of serving size consumed in period 1. Numerical benefits and costs for each of the purchase and consumption decision are given in Table 6:

Table 6: Benefits and Costs of Consumption by Purchase Decision

Purchase Decision Consumption Decision	Instantaneous Utility in Period 1	Instantaneous Utility in Period 2
Small 1 serving	2.5	-2
Large 1 serving	3	-2
2 servings	6	-7

Two assumptions are reflected in the numbers in Table 6. First, even though the consumer eats 1 serving, the consumption benefit is higher when she buys the Large pack because of the quantity discount (price per serving is relatively lower). Second, eating 2 servings at once is 3.5

²⁴ Of course, we can also have consumers who are aware that they are hyperbolic discounters but underestimate its true magnitude on their behavior ($\beta < \hat{\beta} < 1$).

times as bad as consuming 1 serving, reflecting the costs of exceeding one's daily "threshold" for unhealthy food.

Now we can figure out how the naïf ($\beta < \hat{\beta} = 1$) and the sophisticate ($\beta = \hat{\beta} < 1$) will behave, assuming that $\beta = 0.5$. We also contrast their behavior with that of the time-consistent rational consumer with $\beta = 1$. Using our generalized model, the intertemporal utility of the consumer who is faced with the purchase and consumption decisions in period 0 and period 1 are as follows:

Table 7: Utilities of the Consumer in Purchase and Consumption Decisions

	Rational	Naïf	Sophisticates
Purchase Decision			
(Period 0)			
Small	$2.5 - 2$	$\beta \cdot (2.5 - 2)$	$\beta \cdot (2.5 - 2)$
Large	$\text{Max}\{U_{1L}, U_{2L}\}$ $= \text{Max}\{3-2, 6-7\}$	$\beta \cdot \text{Max}\{U_{1L}, U_{2L}\}$ $= \beta \cdot \text{Max}\{3-2, 6-7\}$	$\beta \cdot U_{j^*L}$ where $j^* = \text{argmax}\{\text{Large} - j \text{ Serving in Period 1}\}$
Consumption Decision			
(Period 1)			
Small – 1 Serving	$2.5 - 2$	$2.5 - \beta \cdot 2$	$2.5 - \beta \cdot 2$
Large – 1 Serving	$3 - 2$	$3 - \beta \cdot 2$	$3 - \beta \cdot 2$
Large – 2 Servings	$6 - 7$	$6 - \beta \cdot 7$	$6 - \beta \cdot 7$

The term U_{jL} is the net flow of utility of consuming j servings, evaluated in Period 0, conditional on buying the Large pack. Consequently, the rational, naïf, and sophisticate separate themselves into the following purchase and consumption decisions:

Table 8: Decisions of the Rational, Naïf and Sophisticate

	Rational	Naïf	Sophisticate
Purchase Decision	Large	Large	Small
(Period 0)			
Small	0.5	0.25	0.25
Large	1	0.5	-0.5
Consumption	1 Serving	2 Servings	1 Serving
Decision (Period 1)			
Small - 1 Serving	<i>N.A.</i>	<i>N.A.</i>	1.5
Large – 1 serving	1	2	<i>N.A.</i>
Large - 2 Servings	-1	2.5	<i>N.A.</i>

Start with the rational consumer. In period 0, she buys the Large pack to take advantage of the quantity discount. When period 1 arrives, she has no self-control problem and eats only 1 serving and saving the other serving for the future.²⁵ Her forecasted utility is 1 and that is her actual utility (see Table 8).

The naïf also buys the Large pack in period 0, but for a different reason. In making his period 0 purchase decision, he mistakenly anticipates applying a discount factor of 1 when faced with the one versus two-serving choice in period 1 (see Table 7). As a result, he thinks he will consume only 1 serving in period 1. Given this plan, buying the Large pack appears to be superior in current discounted utility ($\beta \cdot 1$) to buying the Small pack ($\beta \cdot 0.5$). However, when period 1 arrives, eating 2 servings gives utility, *at that point in time*, of $6 - \beta \cdot 7$, which is greater than $3 - \beta \cdot 2$ for eating only 1 serving. The key point is that the naïf makes a forecasting error about his own future behavior: in period 0, he chooses as if he will be comparing in period 1 between utilities of 3-2 versus and 6-7, neglecting the β weight that will actually appear and discount the high future cost in period 1, making her eager to eat both servings at once. Notice that as a result, her actual utility, evaluated at period 0 is not 0.5 but $0.5(6 - 7) = -0.5$.

The sophisticate forecasts accurately what she will do if she buys the Large pack. That is, the Table 7 entries for utilities of consuming from the Large pack when period 1 arrives are

²⁵ To keep the example simple we do not include the benefits and costs of eating the leftover serving in future periods; the results remain unchanged even if we include them.

exactly the same for the naïf and the sophisticate. The difference is that the sophisticate anticipates this actual choice when planning which pack to buy in period 0. As a result, the sophisticate deliberately buys the Small pack, eats only one serving, and has both a forecasted and actual discounted utility of 0.25. The crucial point here is that the naïf does not plan to eat both servings, so she buys the Large pack. The sophisticated knows he can't resist so he buys the Small pack.²⁶

4.3 Marketing Application: Price Plans for Gym Memberships

Hyperbolic discounting is most likely to be found for products that involve either immediate costs with delayed benefits (visits to the gym, health screenings) or immediate benefits with delayed costs (smoking, using credit cards, eating), and temptation. Della Vigna and Malmendier (2004) examined the firm's optimal pricing contracts in the presence of consumers with hyperbolic preferences for gym memberships. Their three-stage model is set-up as follows:

At time $t=0$, the monopolist firm offers the consumer a two-part tariff with a membership fee F and a per-use fee p . The consumer either accepts or rejects the contract. If she rejects the contract, she earns a payoff of \bar{u} at $t=1$, the firm earns nothing, and the game ends. If she accepts the contract, the consumer pays F at $t=1$ and then decides between exercise (E) or non-exercise (N). If she chooses E, she incurs a cost c and pays the firm the usage fee p at $t=1$. She earns delayed health benefits $b > 0$ at $t=2$. If she chooses N, her cost is 0, and her payoffs at $t=2$ are 0 too. It is assumed that the consumer learns her cost c at the end of $t=0$, after she has made the decision to accept or reject the contract. However, before she makes that decision, she knows the cumulative distribution $G(c)$ from which c is drawn. ($G(c)$ is the percentage of consumers with a cost of c or less).²⁷ The firm incurs a set-up cost $K \geq 0$ whenever the consumer accepts the contract and a unit cost a if the customer chooses E. The consumer is a hyperbolic discounter with parameters $(\beta, \hat{\beta}, \delta)$. For simplicity, it is also assumed that the firm is time-consistent with a discount factor δ .

²⁶ However, it is not always the case that the sophisticated agents exhibit more self-control of this sort than naïves do. O'Donoghue and Rabin (1999a) present examples in which the sophisticated agents know they will succumb eventually and so they succumb sooner than the naïf agents. This is a serious problem in treating addictions to drugs and alcohol—addicts who are too experienced use the likelihood of relapsing in the future as an excuse to use immediately. Maintaining an illusion that sobriety will last a long time, even if naïve, can therefore be helpful. That is why programs like AA's 12 steps stress taking it "one day at a time", to prevent having too much foresight which leads to unraveling and immediate use.

²⁷ The unknown unit cost is just a modeling device to inject a probability of going to the gym or not into the analysis in a sensible way. It also captures the case where people are not genuinely sure how much they will dislike exercise, or like the health benefits that result, when they commit to a membership.

For the naïve hyperbolic consumer choosing to exercise, her decision process can be described as follows. At $t=0$, the utility from choosing E is $\beta\delta \cdot (\delta b - p - c)$, while the payoff from N is 0. Hence, she chooses E if $c \leq \delta b - p$. However, when $t=1$ actually arrives, choosing E yields only $\beta\delta b - p - c$, and so the consumer actually chooses E only if $c \leq \beta\delta b - p$. The naïve hyperbolic consumer mispredicts her own future discounting process, and hence overestimates the net utility of E when she buys the membership. So the actual probability that the consumer chooses to exercise is the percentage chance that her cost is below the cost threshold $\beta\delta b - p$, which is just $G(\beta\delta b - p)$. Hence, the consumer chooses to exercise less often than she plans to when she buys the membership. The difference between the expected and actual probability of exercise reflected by $G(\delta b - p) - G(\beta\delta b - p)$, which is always positive (since $\beta < 1$ and $G(c)$ is smaller when c is smaller). In addition, if one allows for an intermediate case where the consumer can be partially naïve ($\beta < \hat{\beta} < 1$) about her time-inconsistent behavior, the degree by which she overestimates her probability of choosing E is $G(\hat{\beta}\delta b - p) - G(\beta\delta b - p)$. Unlike the naïf or partially naïve consumers, the fully sophisticated consumer ($\beta = \hat{\beta} < 1$) displays no overconfidence about how often she will choose E. Overall, the consumer's expected net benefit, at $t=0$ when she accepts the contract, is $\beta\delta \left[-F + \int_{-\infty}^{\hat{\beta}\delta b - p} (\delta b - p - c) dG(c) \right]$.²⁸

The rational firm anticipates this, and its profit-maximization problem is given by:

$$\begin{aligned} & \max_{F,p} \delta \{F - K + G(\beta\delta b - p)(p - a)\} \\ & \text{such that } \beta\delta \left[-F + \int_{-\infty}^{\hat{\beta}\delta b - p} (\delta \cdot b - p - c) dG(c) \right] = \beta\delta \bar{u} \end{aligned}$$

The “such that” constraint reflects the fact that the firm, as a monopolist, can fix contract terms which make the consumer just indifferent between going and earning the expected benefit, or rejecting and earning the discounted rejection payoff $\beta\delta \bar{u}$. The firm maximizes its own discounted profits, which is the fixed fee F , minus its fixed costs K , times the percentage of time it collects user fees because the consumer chooses E (the term $G(\beta\delta b - p)$), the same

²⁸ The discount factor $\beta\delta$ reflects the weight on utilities coming in period $t=1$. F is the fixed membership fee, which is paid in $t=1$. The complex integral term $\left[\int_{-\infty}^{\hat{\beta}\delta b - p} (\delta \cdot b - p - c) dG(c) \right]$ is the expected net benefit, $\delta b - p - c$, integrated over the distribution of possible gym costs $G(c)$. The integral goes from the lowest possible cost to an upper bound of $\hat{\beta}\delta b - p$ because that is the highest cost value at which the consumer knows she will actually go to the gym when the decision point arrives. The forecasted weight $\hat{\beta}$ appears in this upper bound because that forecast determines the consumer's expected guess, at $t=0$, of how she will evaluate costs and benefits at $t=1$ when she decides whether to actually go to the gym.

probability of E which shows up in the consumer's expected utility calculation at $t=1$), times the net profit from the user fees, $p-a$.

Della Vigna and Malmendier (2004) start with the case in which consumers are time-consistent ($\beta = 1$). Then the firm simply sets p^* equal to marginal cost a and chooses F^* to satisfy the consumer's participation constraint (the "such that" constraint above). More interestingly, when $\beta < 1$ the firm's optimal contract involves setting the per-use fee below marginal cost ($p^* < a$) and the membership fee F above the optimal level F^* for time-consistent consumers. This result can be attributed to two reasons: first, the below-cost usage fee serves as a commitment device for the sophisticate to increase her probability of exercise. (Sophisticates like paying a higher membership fee coupled with a lower per-use fee, since they know they will be tempted to skip the gym unless the per-use fee is low.) Second, the firm uses the below-cost per-use fee coupled with an increase in F^* to exploit the naïf's overconfidence about future exercise: the naïf will accept the contract and pay F^* , but exercises (and pays $p^* < a$) less often than she thinks she would. To support these theoretical results, the authors presented empirical evidence that shows that the industry for health club memberships typically charges high membership fees and very low (and often zero) per-use fees. Furthermore, in their study the average membership fee is around \$300/year. For most gyms, consumers also have the option of paying no membership fee but a higher per-use fee (around \$15/visit). The average consumer who paid a typical \$300 fee goes to the gym so rarely that their effective per-use cost is \$19/visit; they would have been better off not buying the membership and just paying on a per-use basis. This type of forecasting mistake is precisely what the naïve hyperbolic consumer does.

It is interesting to contrast this model and empirical findings on gym memberships with a similar study of telephone calling plans by Miravete (2003). He finds that consumers choose calling plans, given the number of minutes they call locally and long-distance, which are very close to optimal, in contrast to the findings about health club under-use. There are two important differences between phone use and health club use. First, there is less temptation to use the phone too much or too little, in contrast to going to the gym. Second, because of telephone deregulation, there is intense competition among long-distance providers during the period of Miravete's data. (In contrast, since people prefer going to nearby gyms, a health club almost has a local monopoly.) To grab market share, providers were very aggressive about poaching the customers of other firms, by educating customers on how much they could save by switching to a better plan. This competition helped correct consumer mistakes in choosing the wrong plans. The general lesson here is that how behavioral mistakes and patterns affect

market equilibrium prices and quantities will depend only on consumer psychology, *and* on behavior of firms and on industrial organization more generally, including regulation²⁹ (see Ellison 2005). This is a crossroads at which combining studies of consumer psychology and careful economic modeling could be very useful.

NEW METHODS OF GAME-THEORETIC ANALYSES

Game theory is a mathematical system for analyzing and predicting how humans and firms will behave in strategic situations. It has been a productive tool in many marketing applications (Moorthy 1985). The field of game theory primarily uses the solution concept of Nash Equilibrium (hereafter NE) and various refinements of it (i.e., mathematical additions which restrict the set of NE and provide more precision). Equilibrium analysis makes three assumptions: (1) *strategic thinking*, i.e., players form beliefs based on an analysis of what others might do; (2) *optimization*, i.e., players choose the best action(s) given those beliefs to maximize their payoffs; and (3) *mutual consistency*, i.e., their best responses and others beliefs' of their actions are identical (or put more simply, players' beliefs about what other players will do are accurate). Taken together, these assumptions impose a high degree of rationality on the players in the game. Despite these strong assumptions, NE is an appealing tool because it does not require the specification of any free parameter (once the game is defined) in order to arrive at a prediction. Furthermore, the theory is general because in games with finitely-many strategies and players, there is always *some* Nash equilibrium (sometimes more than one). Thus, for any marketing application you can imagine, if the game is finite the theory can be used to derive a precise prediction. Furthermore, the theory tells you what to expect if one of the parameters describing the game changes. And if two or more parameters change at the same time, the theory tells you what *net* effect to expect.

The advent of laboratory techniques to study economic behavior involving strategic interaction has tested the predictive validity of NE in many classes of games, in hundreds of studies (see Camerer 2003 for a comprehensive review). The accumulated evidence so far suggests that there are many settings in which NE does not explain actual behavior well, although in many other settings it is remarkably accurate.³⁰ The fact that NE sometimes fits poorly has spurred researchers to look for alternative theories that are as precise as NE but have

²⁹ One can imagine regulations which would correct naïf consumer mistakes automatically. Apparently, in Germany there is a law that at the end of a month, phone users should be charged, regardless of how many calls they make, based on the offered plan which is cheapest for them. A similar regulation could force gyms, for example, to refund part of the membership F if a gym member used the gym so rarely that she would have been better off paying no membership and paying only a larger per-use fee.

³⁰ For example, the predictions of NE in games with unique mixed-strategy equilibrium are often close to the empirical results *in aggregate*.

more predictive power, typically at the cost of introducing one additional behavioral parameter. The next three sections discuss three alternative approaches to predicting what will happen when players actually play these games. These alternative theories relax one or more of the strong assumptions underlying NE, and try to make predictions by introducing only one additional free parameter which has a psychological interpretation.

In sections 5 and 6 we introduce two alternative solution concepts: Quantal-Response Equilibrium (QRE) (McKelvey and Palfrey 1995) which relaxes the assumption of optimization, and the Cognitive-Hierarchy (CH) model (Camerer et al 2004), which relaxes the assumption of mutual consistency. Both the QRE and CH models are one-parameter empirical alternatives to Nash equilibrium and have been shown to predict more accurately than Nash equilibrium in hundreds of experimental games.³¹ In section 7, we describe the self-tuning Experience-Weighted Attraction (EWA) learning model (Camerer and Ho 1998, 1999; Camerer et al 2002; Ho et al 2004). This model relaxes both the best-response and mutual consistency assumptions and describes precisely how players learn over time in response to feedback. The self-tuning EWA model nests the standard reinforcement and Bayesian learning as special cases and is a general approach to model adaptive learning behavior in settings where people play an identical game repeatedly.

5. QUANTAL-RESPONSE EQUILIBRIUM

5.1 Behavioral Regularities

Table 9 shows a game between two players. The Row Player's strategy space consists of actions $A1$ and $A2$, while the Column player's chooses between $B1$ and $B2$. The game is a very simple model of "hide-and-seek" (also known as "matching pennies") in which one player wants to match another player's numerical choice (e.g., $A1$ responding to $B1$), and another player wants to mismatch (e.g., $B1$ responding to $A2$). The row player earns either 9 or 1 from matching on $(A1, B1)$ or $(A2, B2)$ respectively. The column player earns 1 from mismatching on $(A1, B2)$ or $(A2, B1)$.³²

³¹ In an equilibrium, the beliefs about what other players do are accurate (condition 3 above). While beliefs may not be accurate initially, as players learn from experience their beliefs are likely to converge to equilibrium beliefs. Thus, NE can be thought of as a useful prediction of the limiting behavior after learning.

³² This game is the payoff matrix for Game 2 in Ochs (1995).

Table 9: Asymmetric Hide-and-Seek Game

	B1 <i>q</i>	B2 <i>1-q</i>	Empirical Frequency (<i>N</i> =128)	Nash Equilibrium	QRE Prediction
<i>A1</i> <i>p</i>	9, 0	0, 1	0.54	0.50	0.65
<i>A2</i> <i>1-p</i>	0, 1	1, 0	0.46	0.50	0.35
Empirical Frequency	0.33	0.67			
Nash Equilibrium	0.10	0.90			
QRE Prediction	0.35	0.65			

The empirical frequencies of each of the possible actions, averaged across many periods of an experiment conducted on this game, are also shown in Table 9.³³ What is the NE prediction for this game? We start by observing that there is no pure-strategy NE for this game³⁴ so we look for a mixed-strategy NE. Let us suppose that the Row player chooses *A1* with probability *p* and *A2* with probability *1-p*, and the Column player chooses *B1* with probability *q* and *B2* with probability *1-q*. In a mixed-strategy equilibrium, players actually play a probabilistic mixture of two or more strategies. If their valuation of outcomes is consistent with expected utility theory, they only prefer playing a mixture if they are indifferent between each of their pure strategies (i.e., if the expected utilities of the mixed strategies are the same). This property gives a way to compute the equilibrium mixture probabilities *p* and *q*.³⁵ The mixed-strategy NE for this game turns out to be [(0.5 *A1*, 0.5 *A2*), (0.1 *B1*, 0.9 *B2*)]. Comparing this with the empirical frequencies, we find that NE prediction is close to actual behavior by the Row players, whereas it under-predicts the choice of *B1* for the Column players.

If one player plays a strategy that deviates from the prescribed equilibrium strategy, then according to the optimization assumption in NE, the other player must best-respond and deviate from NE as well. In this case, even though the predicted NE and actual empirical

³³ The data is taken from the first row of Table IX in McKelvey and Palfrey (1995).

³⁴ In a pure-strategy equilibrium, players guess perfectly accurately what the other players will do (and because it is an equilibrium their guesses are correct). In a mixed-strategy equilibrium, players only predict the correct probability distributions (or mixtures) of actions by others. In games with mixed NE, players have a strategic incentive to “mix it up” or make different choices every time. If they are playing repeatedly, this strategic incentive extends over time: Players should randomize so their current choices are not predictable from previous choices.

³⁵ Assuming that the Row player is mixing between *A1* and *A2* with probabilities *p* and *1-p*, the expected payoff for the Column player from choosing *B1* is $p \cdot 0 + (1-p) \cdot 1$, and the expected payoff from choosing *B2* is $p \cdot 1 + (1-p) \cdot 0$. Equating the two expressions gives a solution $p=0.5$. Similarly, the Row player’s expected payoffs from *A1* and *A2* are $q \cdot 9 + (1-q) \cdot 0$ and $q \cdot 0 + (1-q) \cdot 1$. Equating the two expressions gives a solution $q=0.1$.

frequencies almost coincide for the Row player, the players are not playing a NE jointly, because the Row player should have played differently given that the Column player has deviated quite far from the mixed-strategy NE (playing *BI* 33% of the time rather than 10%).

5.2 The Generalized Model

Quantal-Response Equilibrium (QRE) relaxes the assumption that players always choose the best action(s) given their beliefs by incorporating “noisy” or “stochastic” best-response. However, the theory builds in a sensible principle that actions with higher expected payoffs are chosen more often - players “*better-respond*”, rather than “*best-respond*”. Mathematically, the QRE nests NE as a special case. QRE also has a mathematically useful property that all actions are chosen with strictly positive probability (“anything can happen”).³⁶ Behaviorally, this means that if there is a small chance that other players will do something irrational which has important consequences, then players should take this into account in a kind of robustness analysis.

The errors in the players’ QRE best-response functions are usually interpreted as decision errors in the face of complex situations or as unobserved latent disturbances to the players’ payoffs (i.e., the players are optimizing given their payoffs, but there is a component of their payoff that only they understand). In other words, the relationship between QRE and Nash equilibrium is analogous to the relationship between stochastic choice and deterministic choice models.

To describe the concept more formally, suppose that player i has J_i pure strategies indexed by j . Let π_{ij} be the probability that player i chooses strategy j in equilibrium. A QRE is a probability assignment π (a set of probabilities for each player and each strategy) such that for all i and j , $\pi_{ij} = \sigma_{ij}(\bar{u}_i(\pi))$ where $\bar{u}_i(\cdot)$ is i ’s expected payoff vector and $\sigma_{ij}(\cdot)$ is a function mapping i ’s expected payoff of strategy j onto the probability of strategy j (which depends on the form of the error distribution). A common functional form is the logistic quantal response

function with $\sigma_{ij}(\bar{u}_i) = \frac{e^{\lambda \cdot \bar{u}_{ij}}}{\sum_{k=1}^{J_i} e^{\lambda \cdot \bar{u}_{ik}}}$. In equilibrium, we have

$$\pi_{ij} = \frac{e^{\lambda \cdot \bar{u}_{ij}}}{\sum_{k=1}^{J_i} e^{\lambda \cdot \bar{u}_{ik}}}$$

³⁶ The fact that all strategies have strictly positive probabilities means there is no need to restrict attention to subgame perfect Nash equilibria which eliminate incredible threats, or to impose ad hoc assumptions of “trembling” in normal form games. The credibility of possible threats and the amount of trembling (in a statistical sense) are determined endogenously by the payoff loss from those threats or trembles.

where $\bar{u}_{ij}(\pi)$ is player i 's expected payoff from choosing j given that the other player(s) choose their strategy according to the equilibrium profile π . It is easy to show mathematically that π_{ij} is larger if the expected payoff $\bar{u}_{ij}(\cdot)$ is larger (i.e., the better responses are played more often). The parameter λ is the payoff sensitivity parameter. The extreme value of $\lambda = 0$ implies that player i chooses among the J_i strategies equally often (since there are J_i strategies, each strategy is played with probability $1/J_i$); that is, they do not respond to expected payoffs at all. At the other extreme, QRE approaches NE when $\lambda \rightarrow \infty$, so that the strategy with the highest expected payoff is played almost all the time. As λ increases, higher and higher probability is put on strategies with better payoffs.

We now illustrate how to derive the QRE for a given value of λ in the game in section 5.1. Again, suppose the Row player chooses $A1$ with probability p and $A2$ with probability $1-p$, while the Column player chooses $B1$ with probability q and $B2$ with probability $1-q$. Then the expected payoffs from playing $A1$ and $A2$ are $q*9+(1-q)*0 = 9q$ and $q*0+(1-q)*1 = 1-q$ respectively. Therefore

$$\pi_{RowA1} = p = \frac{e^{\lambda*9q}}{e^{\lambda*9q} + e^{\lambda*(1-q)}}$$

Similarly, the expected payoffs to $B1$ and $B2$ for the Column player are $1-p$ and p respectively, so

$$\pi_{ColumnB1} = q = \frac{e^{\lambda*(1-p)}}{e^{\lambda*(1-p)} + e^{\lambda*p}}$$

Notice that q is on the right hand side of the first equation, which determines p , and p is on the right hand side of the second equation, which determines q . To a psychologist, it might look like the logic is “circular” because p depends on q and q depends on p . But that mutual dependence is what delivers mathematical precision. For any value of λ , there is only one pair of (p, q) values that solves the simultaneous equations and yields a QRE. If $\lambda = 2$ for example, the QRE predictions are $p^*=0.646$ and $q^*=0.343$ which are closer to the empirical frequencies than the NE predictions are.

Using the actual data, a precise value of λ can be estimated using maximum-likelihood methods. The estimated λ for the QRE model for the asymmetric hide-and-seek game in section 5.1 is 1.95. The negative of the log likelihood of QRE (an overall measure of goodness of fit) is 1721, a substantial improvement over a random model benchmark ($p=q=0.5$) which has a fit of 1774. The NE prediction has a fit of 1938, which is worse than random (because of the extreme prediction of $q=0.1$).

5.3 Marketing Application: Price Competition with Differentiated Products

We now apply QRE to study price competition between firms and contrast its predictions with NE. This application is from Friedman et al (1995), in which the authors derived the QRE for a differentiated goods Bertrand duopoly. In this model, there are two firms who choose prices p_1 and p_2 . After the firms choose prices simultaneously, the quantity of products they sell is given by two demand functions:

$$\begin{aligned} q_1(p_1, p_2) &= a - bp_1 + cp_2 \\ q_2(p_1, p_2) &= a - bp_2 + cp_1 \end{aligned}$$

where a, b and $c > 0$. For simplicity, assume the marginal cost of production for each firm is zero and also that $c < 2b$. The fact that $c > 0$ implies that the two firms' products are imperfect substitutes, so that if p_2 is higher than consumers will demand more of product 1. We consider only the case when the payoffs are non-negative, so the firms' strategy space is the interval $[0, \frac{a}{b}]$. It can be easily shown that the symmetric pure-strategy NE in which both firms choose the same price is:

$$p^{NE} = \frac{a}{2b - c}$$

To begin deriving the QRE, let \bar{p} denote the firm's *expected* price in the symmetric QRE in which both firms choose the same distribution of prices. (Because the demand function is linear in both prices, the firm's demand depends only on its own price and on the expected price of other firm, as long as the resulting demand is not negative.) The authors assume a quantal-response function as follows:

$$f(p) = \frac{[p(a - bp + c\bar{p})]^\lambda}{\int_0^{\frac{a}{b}} [p(a - bp + c\bar{p})]^\lambda dp}$$

where $f(p)$ is the firm's density function of prices over the interval $[0, \frac{a}{b}]$ and $\lambda \geq 0$.

To simplify the analysis, the authors solve for a QRE assuming that $\lambda = 1$.

Letting $K = \int_0^{\frac{a}{b}} p(a - bp + c\bar{p}) dp$, the expected price \bar{p} in the symmetric QRE is then:

$$\begin{aligned}
\bar{p} &= \int_0^{\frac{a}{b}} p \cdot f(p) dp \\
&= \frac{1}{K} \int_0^{\frac{a}{b}} p^2 (a - bp + c\bar{p}) dp \\
&= \frac{1}{K} \left[\left(\frac{a}{b} \right)^3 \left(\frac{a + c\bar{p}}{3} \right) - \left(\frac{a}{b} \right)^4 \left(\frac{b}{4} \right) \right] \\
&= \frac{\left(\frac{a}{b} \right)^3}{K} \left[\frac{a + c\bar{p}}{3} - \frac{a}{4} \right]
\end{aligned}$$

Also, we have:

$$\begin{aligned}
K &= \int_0^{\frac{a}{b}} p(a - bp + c\bar{p}) dp \\
&= \left(\frac{a}{b} \right)^2 \left(\frac{a + c\bar{p}}{2} - \frac{a}{3} \right)
\end{aligned}$$

Hence, \bar{p} is given implicitly by the following expression

$$\bar{p} = \frac{a}{b} \left[\frac{a + 4c\bar{p}}{2a + 6c\bar{p}} \right]$$

Solving, we obtain

$$\bar{p} = \frac{a}{6bc} \left[\sqrt{(b-2c)^2 + 6bc} - (b-2c) \right]$$

In this case, we have $p^{NE} > \bar{p}$, since

$$\begin{aligned}
\frac{a}{2b-c} &> \frac{a}{6bc} \left[\sqrt{(b-2c)^2 + 6bc} - (b-2c) \right] \\
\Leftrightarrow 6bc + 2(2b-c)(b-2c) &> (2b-c)^2 \\
\Leftrightarrow (2b-c)3c &> 6bc
\end{aligned}$$

which is satisfied when $c > 0$. Hence, when $\lambda = 1$, QRE predicts a price dispersion and more competitive prices on average compared to the pure-strategy NE price. To see this, note that

when players are pricing purely randomly ($\lambda = 0$), the mean price is $\frac{a}{2b}$, which is less than the

Nash price $\frac{a}{2b-c}$ when $c > 0$. Since quantal response tends to drive prices downward, when

one firm charges a lower price, the other firm wants to price lower as well when products are

substitutes.³⁷ The authors also predict that \bar{p} is increasing in λ and approaches the NE solution when λ is large.

6. COGNITIVE-HIERARCHY MODEL

6.1 Behavioral Regularities

In a dominance solvable game called the p -beauty contest (Nagel 1995; Ho et al 1998), a group of players are each asked to pick a number from 0 to 100, and the player who chooses a number that is closest to $2/3$ the average of the numbers chosen by all players (including the player herself) wins a cash prize. If there is more than one winner (i.e., two or more players choose the same closest number), the prize is divided equally. The NE prediction for this game is that all players will choose the number 0 and the reasoning for every player goes like this: “even if everyone else picks 100, I should pick no more than 67 (i.e., $(2/3) \cdot 100$). However, all the other players apply the same reasoning, so I should pick no more than 45 (i.e., $((2/3)^2 \cdot 100)$) but again, all the other players know this as well...” and the choice of the player unravels to 0.

The p -beauty contest has been played across diverse subject groups and Table 10 gives a sample of the actual average choices from some of these groups:

Table 10: Average Choice in p -beauty Contests

Subject Pool	Group Size	Sample Size	Mean
Caltech Board	73	73	49.4
80 year olds	33	33	37.0
High School Students	20-32	52	32.5
Economics PhDs	16	16	27.4
Portfolio Managers	26	26	24.3
Caltech Students	3	24	21.5
Game Theorists	27-54	136	19.1

Source: Table II in Camerer et al (2004).

The data shows that the NE prediction of 0 explains first-period choices very poorly in this game. Furthermore, telling someone to play the equilibrium choice is bad advice. However, choices are heterogeneous across subject pools. The data raises an important question: is there an alternative to NE that generates more accurate predictions and captures the heterogeneity across different subjects? The Cognitive-Hierarchy (CH) model (Camerer et al 2004) is one such model.

³⁷ That is, differentiating firm 1’s profit function with respect to its price p_1 gives the “reaction function” $p_1^* = (a + cp_2)/2b$. If p_2 is lower then firm 1 wants to choose a lower p_1^* .

6.2 The Generalized Model

The CH model relaxes the assumption of mutual consistency by allowing players' beliefs of others' actions to be different from actual choices made by others. Unlike NE, the CH model assumes there are some players who have not carefully thought through what other players are likely to do, and so will be surprised. It uses an iterative process which formalizes Selten's (1998: 421) intuition that 'the natural way of looking at game situations is not based on circular concepts, but rather on a step-by-step reasoning procedure'. It captures heterogeneity in players' reasoning abilities and explicitly models the decision rule used by each type of players.

To begin, the CH model captures the possibility that some players use zero steps of thinking, that is, they do not reason strategically at all.³⁸ It is assumed that these 0-step players randomize equally among all available strategies. Players who are 1-step thinkers choose action(s) to maximize their payoffs believing all others are using zero steps (i.e., players are "over-confident" in that they think all others use less steps of thinking). Proceeding inductively, the model assumes that K -step players think all others use zero to $K-1$ steps. The model assumes that the frequencies of K step thinkers, $f(K)$ is given by:

$$f(K) = \frac{e^{-\tau} \cdot K^{\tau}}{K!}.$$

The Poisson characterization is appealing because it has only one free parameter τ (which is both its mean and variance). A large value of τ implies that players are sophisticated and undertake many steps of iterative reasoning.

Step- K players' beliefs of the proportions of the lower steps are obtained by dividing the actual proportions of lower-step types by the sum of their frequencies so that they add up to one as follows:

$$g_K(h) = \frac{f(h)}{\sum_{h=1}^{K-1} f(h)} \text{ for } h \leq K-1 \text{ and } g_K(h) = 0 \text{ for } h \geq K$$

As K gets larger, these normalized beliefs converge to the actual frequencies $f(h)$. That is, K is a measure of how well-calibrated a player's beliefs are, like "strategic intelligence". Step- K thinkers ($K > 0$) are assumed to compute expected payoffs given their beliefs, and choose the strategy that yields the highest expected payoff.

³⁸ These players can be thought of as being clueless, overwhelmed, or simply more willing to make a random guess.

We illustrate how to derive the CH prediction for a specific value of τ using the p -beauty contest as an example. The value of τ assumed in this case is 1.55, which gives us the frequencies of the step- K players shown in the first column of Table 11:

Table 11: Predictions of the CH model for the p -Beauty Contest ($\tau=1.55$)

Steps (K)	Frequency	Perceived Average	Choice
0	0.212	n.a.	50
1	0.329	50	33
2	0.255	39	26
3	0.132	35	23
4	0.051	33	22
5	0.016	32	21
6	0.004	32	21
Aggregate	1.000		32.67

In the CH model, the step-0 players do not think strategically - in this example, their choices are distributed uniformly from 0 to 100. The step-1 players think that all the other players are step-0 thinkers and best-respond: since their perceived average is 50, they will choose $\frac{2}{3} \cdot 50 = 33$.³⁹ The step-2 players think that the proportions of the step-0 and step-1 players are 0.39 (from $0.212/(0.212+0.329)$) and 0.61 respectively, and so they perceive the average to be $0.39 \cdot 50 + 0.61 \cdot 33 = 39.6$ (39 including their own choice), and hence choose 26 as a best-response. Following the same reasoning, we can compute the choices for players of higher steps. Finally, we compute the predicted choice of the CH model by weighting the K -step players' choices by the proportions of players doing each of K steps. For $\tau=1.55$, the CH model predicts that the average would be 32.67. The estimated τ s and the respective CH predictions for the subject pools are given in Table 12 below. In general, Camerer et al (2004) found that the value of τ which best explains a wide variety of games is around 1.5.

Table 12: Estimated τ and CH prediction for Different Subject Pools

Subject Pool	Mean	CH Prediction	Estimated τ
Caltech Board	49.4	43.1	0.50
80 year olds	37.0	36.9	1.10
High School Students	32.5	32.7	1.60
Economics PhDs	27.4	27.5	2.30
Portfolio Managers	24.3	24.4	2.80
Caltech Students	21.5	23.0	3.00
Game Theorists	19.1	19.1	3.70

³⁹ In this example we do not round the numbers up because the perceived average by the step- K thinker is lower than the average of the lower-step thinkers once we include the choice of the step- K thinker.

Like some of the other behavioral economics models described above, there is tentative evidence from neuroscience which is consistent with models of limited strategic thinking like CH. To a neuroscientist, if a player makes an accurate equilibrium guess about another player, this suggests that perhaps the brain circuitry a player uses to make her own choice is also being used to guess accurately what the other player will choose (e.g., by “putting yourself in the other player’s shoes” - or brain). Bhatt and Camerer (forthcoming) found that when players’ choices and beliefs were in equilibrium, there was a lot of overlap in brain areas when choosing and guessing, as if equilibrium is “a state of mind” (as well as a mathematical restriction on belief accuracy). When they were not in equilibrium, there was much more activity making one’s own choice, compared to guessing, as would be expected of a 1-step thinker.

6.3 Marketing Application: Market Entry

A central concept in marketing strategy is what markets a firm should enter. Entry should depend on whether a business leverages a firm’s brand equity (“umbrella branding”), and has cost or economies-of-scope advantages relative to its competitors. But entry also depends on a sensible forecast of how large the market will be, and how many other firms will enter. Even if a firm has competitive advantages in a particular market, if too many firms enter - perhaps because they underestimate the amount of competition or are optimistic about their relative advantage, then the firm should stay out until the “shake out” period when overly optimistic firms fail (e.g. Camerer and Lovo 1999).

The simplest way to study the effect of forecasts of competitors is with a simple model that strips away cost advantages. Suppose each of N firms simultaneously decide whether to enter a market or stay out (denoted 1 and 0, respectively). Denote the total market demand by $d < N$ (where d is expressed as a fraction of the number of firms, so $0 < d < 1$). If d or fewer firms enter (i.e. supply is equal to or less than demand), then the entrants each earn a profit of 1. If more than d firms enter (i.e. supply is greater than demand), then all the entrants earn 0. If a firm does not enter, it earns a profit of 0.5.

To keep the mathematics simple, we assume that firms are risk-neutral and there are infinitely many “atomistic” entrants. In this case, firms only care about whether the fraction of others entering is above d or not: if the fraction of others entering is below d , they should enter; but if the fraction of others entering is above d , the firm is better off staying out and earning 0.5. The NE in this case is a mixed-strategy equilibrium in which firms randomize and enter with

probability d . When the number of firms is large ($N \rightarrow \infty$) the law of large numbers implies that the fraction of entry is d .

To derive the prediction of the CH model, denote the entry function of step- K firms for a given demand d by $e(K, d)$. This function maps the demand d and thinking step K into a decision to enter (1) or stay out (0). Denote the interim total entry function for all steps up to and including K by $E(K, d)$. In the CH model, a step- K firm thinks it is competing against a normalized Poisson distribution of players who do 0, 1, 2, ..., $K-1$ steps of thinking. Thus, the step- K firm thinks it is facing an interim total entry rate of $E(K-1, d) = \sum_{h=0}^{K-1} g_K(h) \cdot e(h, d)$. Hence, if $E(K-1, d) < d$ the step- K firm should enter ($e(K, d) = 1$); if $E(K-1, d) > d$ the firm should stay out ($e(K, d) = 0$). The prediction of the CH model of how many firms will enter for each value of d is the limiting case $E(\infty, d)$, which depends on the value of τ .

Let's work through the first couple of steps of entry predicted by the CH model. Step-0 firms will randomize, so their entry function is $e(0, d) = 0.5$ for all values of d (think of these firms as ignoring competitors or market size). Step-1 firms believe that if d is below 0.5, there will be too much entry (because $e(0, d) = 0.5 > d$) so they stay out. However, if d is above 0.5, then they anticipate $e(0, d) = 0.5 < d$ so there will not be enough entry and they can profit by entering. Step-2 firms think they are facing a mixture of step-0 and step-1 firms. In the Poisson CH model, the relative proportions of these two types of firms are $g_2(0) = \frac{1}{1+\tau}$ and

$g_2(1) = \frac{\tau}{1+\tau}$.⁴⁰ For $d < 0.5$ the expected interim entry function $E(1, d)$ is

$g_2(0) \cdot 0.5 + g_2(1) \cdot 0 = \frac{0.5}{1+\tau}$. Thus, when $d < \frac{0.5}{1+\tau}$ there is too much entry, due to the entry by

the step-0 firms and so the step-2 firms stay out. But when $\frac{0.5}{1+\tau} < d < 0.5$, there is not enough

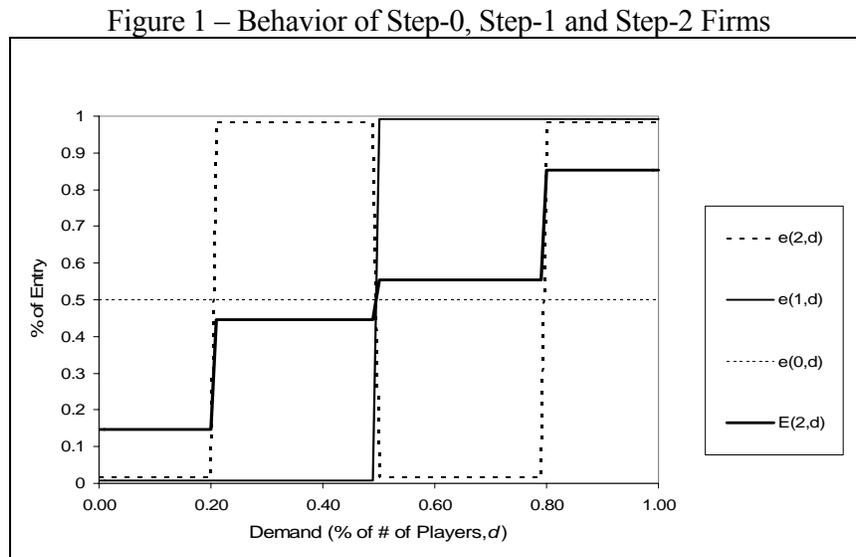
entry because the step-0 firms' entry scared away the step-1 firms and so the 2-step firms enter. Similar calculations (see Camerer et al 2004) show that for $d > 0.5$, the expected interim entry

function is $E(1, d) = g_2(0) \cdot 0.5 + g_2(1) \cdot 1 = \frac{0.5 + \tau}{1 + \tau}$. Step-2 firms stay out when

⁴⁰ The raw frequencies of step-0 and step-1 competitors are $f(0)=e^\tau$ and $f(1)=e^\tau \tau / 1!$. The conditional frequency $g_2(0)=f(0)/(f(0)+f(1))=1/(1+\tau)$ and similarly for $g_2(1)$.

$0.5 < d < \frac{0.5 + \tau}{1 + \tau}$ and enter when $d > \frac{0.5 + \tau}{1 + \tau}$. If $\tau=1.5$, then step-2 firms enter for $0.2 < d < 0.5$ and for $d > 0.8$.

Figure 1 below illustrates the entry functions $e(0,d)$, $e(1,d)$, $e(2,d)$ and $E(2,d)$ for the value of $\tau=1.5$. The interim entry function $E(2,d)$ (the thickest bold line) which mixes entry from the step-0 to step-2 firms is a step function which crudely approximates the NE prediction, which is the identity line (not shown) at which entry equals d . The key point is that the step-2 firm "smooths" the entry function $E(2,d)$, bringing it closer to the identity line than the interim function $E(1,d)$ (not shown). Adding in the step-3 and higher-step firms smooths the entry function even more.



In actual experiments conducted on market entry games like this, there are three empirical regularities (see Rapoport and Seale forthcoming and chapter 7 in Camerer 2003):

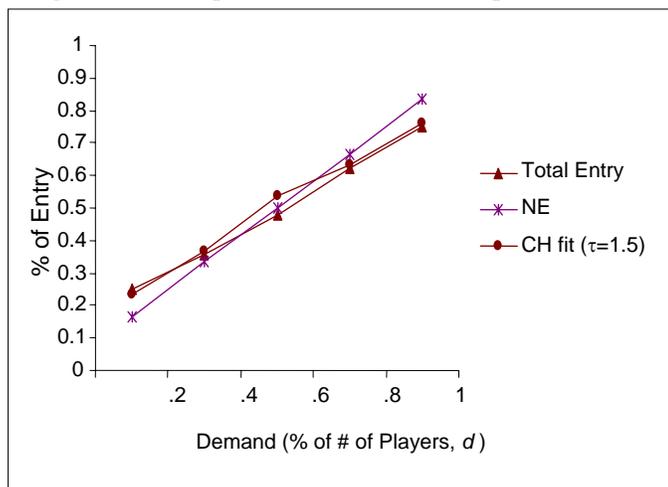
1. Entry tends to be monotonically increasing in demand d , even in one-shot games and early periods of repeated games;
2. There is too much entry for low values of d , and too little entry for high values of d ;
3. Individual profiles of entry rates for different d values in within-subject designs (where players choose whether to enter for different d values) tend to show some simple step functions (like the 1-step players) and more jagged out-in-out-in step functions (e.g. Rapoport et al 2002).

The first regularity is surprising. We usually expect equilibration to take learning, communication, evolution or some other dynamic process. But approximating the equilibrium of d entering happens even in one-shot games with no communication or dynamics. Remarking on

the surprising similarity between predicted entry rates across values of d and actual entry in pilot experiments he conducted, Kahneman (1988) wrote that “to a psychologist, it looks like magic”. The CH model can explain this magic that Kahneman noted.⁴¹ Suppose the market entry game was played sequentially with a smaller group of say, 20 firms. If the demand d was 40% (so that 8 out of 20 firms could profit), then after 8 firms have entered, the rest would stay out. While the game is not actually played sequentially, the higher-step firms *anticipate* what lower-step thinking firms will do, which smooths out the entry function in a “pseudo-sequential” way. It is as if the higher-step firms make entry decisions *after* the lower-step firms do, smoothing the entry function.

Moreover, the CH model also explains the second and third regularities. The second regularity results because of the lingering effect of the 0-step firms, which over-enter at low values of d and under-enter at high values of d . The third regularity is evident from Figure 1: As firms do more steps of thinking, their entry functions become more jagged, or nuanced - they are sensitive to subtle differences in which levels of demand d are attractive for entry. Figure 2 below shows that the CH model (for $\tau=1.5$) captures the three regularities for the experimental data in Camerer et al (2004).

Figure 2 – Comparison of CH with Empirical Data



The p -beauty contest discussed in 6.1 and 6.2 shows how a model like CH can explain limited equilibration, in which one-shot or first-period data are far from the NE. The market entry game requires exactly the opposite explanation - the puzzle is the “magic” of how players get so close to equilibrium instantly, not why they are so far away. The fact that the same CH model can

⁴¹ The entry rate $E(\infty, d)$ is monotonically increasing in d if and only if $1+2\tau < e^\tau$, or $\tau < 1.25$ (Camerer et al 2004). While this value is slightly lower than is typically estimated econometrically from experimental data, the only non-monotonicity is a small dip in which more firms may enter for d just below 0.5 than for d just above 0.5.

explain both rapid equilibration in the market entry games and the lack of equilibration in the p -beauty contest shows the generality of the model. While the behavioral economics approach is generally critical of the empirical limits of the rational-choice approach, good models should also be able to explain when rational-choice emerges in the limit, or surprisingly quickly;⁴² and the CH model can do so.

7. LEARNING

7.1 Behavioral Regularities

Equilibrium concepts like NE study behavior after equilibration, at the point where players have come to guess accurately what other players do. One force that produces equilibration is learning from feedback. That is, players' actions at a point in time reflect their past experiences with those actions, which are used to anticipate future actions of others. An example that illustrates learning is the "median-action" coordination game studied by Van Huyck et al (1991). This game involves a group of n players (where n is an odd number). The choice of each player i consists of choosing an integer action x_i from 1 to 7. Let M denote the median of all the players' actions. The payoff function for each player is given by $\pi_i(x_i, M) = 0.1M - 0.05(M - x_i)^2 + 0.6$. The possible payoffs are given in Table 13 below:

Table 13: Payoffs for the Median-Action Game

		Median Value of X Chosen						
		7	6	5	4	3	2	1
Your Choice X	7	1.30	1.15	0.90	0.55	0.10	-0.45	-1.10
	6	1.25	1.20	1.05	0.80	0.45	0.00	-0.55
	5	1.10	1.15	1.10	0.95	0.70	0.35	0.10
	4	0.85	1.00	1.05	1.00	0.85	0.60	0.25
	3	0.50	0.75	0.90	0.95	0.90	0.75	0.50
	2	0.05	0.40	0.65	0.80	0.85	0.80	0.65
	1	-0.50	-0.05	0.30	0.55	0.70	0.75	0.70

This game captures the key features of economic situations in which there is a motive for conformity, because players are penalized for choosing x_i which is different than the median M (through the penalty of $0.05(M - x_i)^2$). At the same time, players have a common

⁴² A remarkable study with ducks (Harper 1982) makes the same point. Two behavioral ecologists went to a pond in England where ducks gather. One experimenter tossed small bread balls to the ducks every five seconds, and another tossed bread balls nearby every 10 seconds. In equilibrium, twice as many ducks should gather where balls are being tossed every five seconds (since the rate of reward is twice as fast). In fact, the ducks sorted into two groups, approximating the equilibrium, in less than a minute, which suggests there may be some variation in duck rationality much as in the simple CH model (and as seen in humans playing for money rewards).

incentive to coordinate on a higher median choice, because of the premium $0.1M$ added to profits. So all the players want the median number to be high, which creates a motive to choose a large number, but nobody wants to choose a number larger than what they expect the median will be.

In this game there are many NE: For each of the values 1 to 7, if players think the median will equal that value, then the best response is $x_i = M$, so the median M will result. So all of the integers 1 to 7 are pure-strategy NE (there are many mixed equilibria as well). Hence, the seven cells on the diagonal of Table 13 represent the possible pure-strategy NE outcomes. In this example, the NE is not precise at all. This imprecision leaves room for an empirical theory to pin down what happens when players play for the first time (like the CH theory described in the previous section), and to specify a path of play over time as players learn. Before theories of limited strategizing and learning came along in behavioral game theory, however, theorists were inclined to ask which general deductive principles players might use to select among one of many equilibria (these are called “selection principles”).

Three deductive principles of equilibrium selection may be useful in predicting what will be chosen in the median-action game. The first selection principle is “payoff dominance” - players will choose the equilibrium that is best for everyone. The payoff dominant equilibrium is for everyone to choose 7, earning \$1.30 each. The second selection principle is based on the concept that people will maximize their “security” and choose the action whose smallest payoffs is the largest (maximizing the minimum payoff). Choosing 3 guarantees a player will make a minimum of at least \$0.50, and since this is the highest guaranteed payoff for any strategy, the maximin rule selects the strategy of 3. The third selection principle is the 1-step rule introduced previously, which chooses the strategy with the highest expected value, assuming all medians are equally likely. The 1-step rule selects the strategy of 4.

Van Huyck et al (1991) conducted six experimental sessions in which a group of nine players made decisions for ten periods using the above payoff table.⁴³ The median choice in each period for each session is given in Table 14 below.⁴⁴

⁴³ In the experiments each of the nine subjects chooses an action simultaneously and the median is then announced. Hence subjects only have the histories of their own actions and the medians.

⁴⁴ The data are taken from the first six rows (Treatments Gamma and Gammadm) of Table III in Van Huyck (1991).

Table 14: Median Choice in the Median-Action Game across Ten Periods

	Period									
	1	2	3	4	5	6	7	8	9	10
Session 1	4	4	4	4	4	4	4*	4	4*	4*
Session 2	5	5	5	5	5	5	5	5	5	5
Session 3	5	5	5	5	5	5	5	5	5	5*
Session 4	4	4	4	4	4	4*	4*	4*	4*	4*
Session 5	4	4	4	4*	4*	4*	4*	4*	4*	4*
Session 6	5	5	5	5	5	5	5	5*	5*	5*

*indicates a session in which all nine players choose the same action.

The results exhibit four regularities: first, the principles of payoff dominance and secure (maximin) actions are not supported: the medians were never 7 nor 3 in any period. Second, the median choices in each session are strongly influenced by the initial median (in all cases, the period 10 median is exactly the initial period 1 median). Third, not only does pure-strategy NE fail to provide a determinate prediction, but it also performs rather poorly as all the players choose the same action in only 19 out of 60 total periods (those periods are denoted by an “*” in Table 14). Last and most importantly for our purpose, equilibration seems to result from learning over time, because the NE (in which all subjects play the same strategy) are more likely to be played in the later periods. In the following section we present a model that is able to capture some of these dynamics.

7.2 The Generalized Model

Before we begin, some notation is necessary. For player i , there are m_i strategies, denoted s_{ij} (the j -th strategy for player i). Strategies actually chosen by i in period t , and by all other players (who are denoted $-i$) are $s_i(t)$ and $s_{-i}(t)$ respectively. Player i 's ex-post payoff of choosing strategy s_{ij} in time t is $\pi_i(s_{ij}, s_{-i}(t))$ and the actual payoff received is $\pi_i(s_i(t), s_{-i}(t)) \equiv \pi_i(t)$. For player i , strategy j has a numerical attraction $A_{ij}(t)$ at the end of period t after updating from the period's experience. $A_{ij}(0)$ is the initial attraction before the game starts. Attractions of period t determine choice probabilities in period $t+1$ through a logistic stochastic response function:

$$P_{ij}(t+1) = \frac{e^{\lambda \cdot A_{ij}(t)}}{\sum_{k=1}^{m_i} e^{\lambda \cdot A_{ik}(t)}}$$

where λ measures the degree of payoff sensitivity. Note that as in the QRE model discussed earlier, $\lambda = 0$ is a random response and $\lambda = \infty$ is best response. In the self-tuning EWA model, attractions are updated by

$$A_{ij}(t) = \frac{\phi_i(t)N(t-1)A_{ij}(t-1) + [\xi_{ij}(t) + (1 - \xi_{ij}(t))I(s_{ij}, s_i(t))]\pi_i(s_{ij}, s_{-i}(t))}{\phi_i(t)N(t-1) + 1}$$

where $I(x, y)$ is an indicator function (equal to one if $x=y$ and zero otherwise) (see Camerer and Ho, 1998; Camerer and Ho 1999, Camerer et al 2002, Ho et al 2004). That is, previous attractions are multiplied by an experience weight $N(t - 1)$, decayed by a weight $\phi_i(t)$, incremented by either the actual payoff received (when $I(s_{ij}, s_i(t)) = 1$) or by $\xi(t)$ times the payoff that could have been received (when $I(s_{ij}, s_i(t)) = 0$), and normalized by dividing by $\phi_i(t)N(t-1)$. Note that an equilibrium model is a special case of non-learning in which the initial attractions $A_{ij}(0)$ are derived from equilibrium calculations, and the initial experience weight $N(0)$ is very large.

EWA is a hybrid of reinforcement and fictitious play (belief-learning) models. Standard reinforcement models assume that only actual choices are reinforced (i.e., when $\xi_{ij}(t) = 0$) (Erev and Roth 1998). The weighted fictitious play belief learning model corresponds to a case in which all choices are reinforced with equal weights (i.e., all foregone payoffs are reinforced by a weight of $\xi(t) = 1$). We initialize $N(0)$ by setting it equal to one and update it according to $N(t) = \phi_i(t)N(t-1) + 1$. The initial attractions $A_{ij}(0)$ are computed numerically by using the C-H model with $\tau = 1.5$.⁴⁵

⁴⁵ The EWA model can also be written in a slightly different “temporal difference” form which links it to models in machine learning and neuroscience, like so:

$$A_{ij}(t) - A_{ij}(t-1) = \frac{[\xi_{ij}(t) + (1 - \xi_{ij}(t))I(s_{ij}, s_i(t))](\pi_i(s_{ij}, s_{-i}(t)) - A_{ij}(t-1))}{\phi_i(t)N(t-1)}$$

In this form, the updating implies that attractions are changed by a fraction of the “prediction error” $\pi_i(s_{ij}, s_{-i}(t)) - A_{ij}(t-1)$, the difference between the received payoff and the attraction. These models interpret attractions as predictions which are constantly revised in response to feedback. Dividing the weighted prediction error by $\phi_i(t)N(t-1)$ means that, since $N(t-1)$ is increasing over time, the weight on incremental prediction errors is steadily falling, so the rate of learning is slowing down.

The change-detector function $\phi_i(t)$

The decay rate $\phi_i(t)$ which weights lagged attractions is often called “forgetting”. While forgetting obviously does occur, the more interesting variation in $\phi_i(t)$ is a player’s perception of how quickly the learning environment is changing. The function $\phi_i(t)$ should therefore be capable of “detecting change”. When a player sense that other players are changing, a self-tuning $\phi_i(t)$ should dip down, putting less weight on old “stale” experience. However, this self-tuning process should also not over-react to small changes due to fluctuations.

The core of the $\phi_i(t)$ change-detector function is a “surprise index”, which is the difference between other players’ recent strategies and their strategies in previous periods. Define a history vector, across the other players’ strategies k , which records the historical frequencies (including the most recent period t) of the choices made by other players. A vector element is $h_{ik}(t) = \frac{\sum_{\tau=1}^t I(s_{-ik}, s_{-i}(\tau))}{t}$. The recent ‘history’ $r_{ik}(t)$ is a vector of 0’s and 1’s which has a one for strategy $s_{-ik} = s_{-i}(t)$ and 0’s for all other strategies s_{-ik} (i.e. $r_{ik}(t) = I(s_{-ik}, s_{-i}(t))$). The surprise index $S_i(t) = \sum_k^{m-i} (h_{ik}(t) - r_{ik}(t))^2$ sums up the squared deviations between the cumulative history vector $h_{ik}(t)$ and the immediate recent history vector $r_{ik}(t)$ across all the strategies.

This surprise index varies from zero to two. To map it into a sensible change detection value, we use

$$\phi_i(t) = 1 - 0.5S_i(t)$$

Because $S_i(t)$ is between zero and two, $\phi_i(t)$ is always (weakly) between one and zero. The numerical boundary cases illuminate intuition: If the other player chooses the strategy she has always chosen before, then the history and recent history vectors are exactly the same, and $S_i(t) = 0$ (player i is not at all surprised) so $\phi_i(t) = 1$ (player i does not decay the lagged attraction at all, since what other players did previously is presumed to be informative). If the other player has been choosing a single strategy in $n-1$ previous periods, then suddenly switches to a brand new strategy, the decay rate is $\phi_i(t) = \frac{(2n-1)}{n^2}$. Hence, $\phi_i(t) = 0.36$ when $n=5$ and $\phi_i(t) = 0.19$ when $n=10$. As n grows larger, $\phi_i(t)$ converges to zero (player i decays the lagged

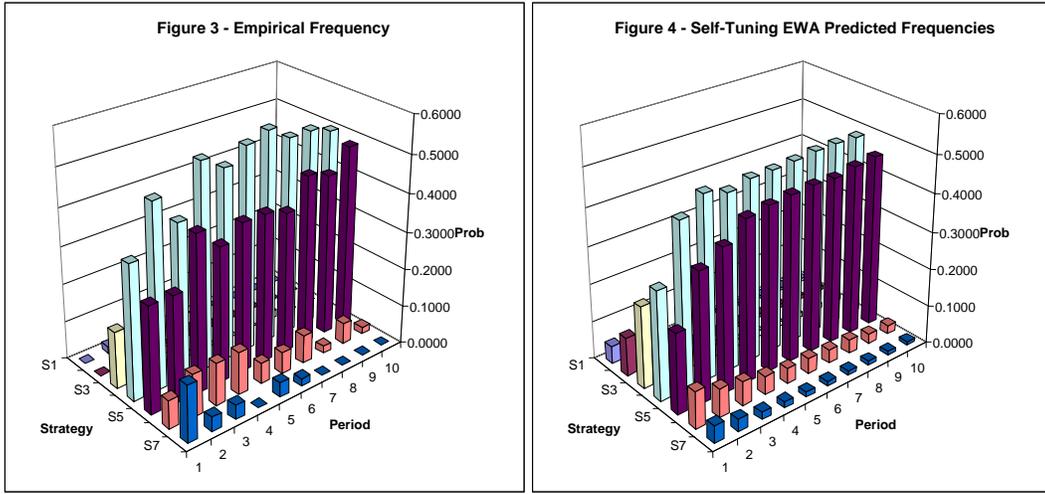
attraction completely and ‘starts over’). In general, the longer the previous stable history, the bigger is the surprise when another player switches, and the lower the decay rate $\phi_i(t)$.

The attention function $\xi_{ij}(t)$

The parameter $\xi_{ij}(t)$ is the weight on foregone payoffs. Presumably this is tied to the attention subjects pay to alternative payoffs, ex-post. Subjects who have limited attention are likely to focus on strategies that would have given higher payoffs than what was actually received, because these strategies represent missed opportunities. To capture this property, define $\xi_{ij}(t) = 1$ if $\pi_i(s_{ij}, s_{-i}(t)) \geq \pi_i(t)$ and 0 otherwise. That is, subjects reinforce chosen strategies (where the payoff inequality is always an equality) and all unchosen strategies which have better payoffs (where the inequality is strict) with a weight of one. They reinforce unchosen strategies with equal or worse payoffs by zero.

Note that this $\xi_{ij}(t)$ can transform the self-tuning rule into special cases over time. If subjects are strictly best-responding (ex post), then no other strategies have a higher ex-post payoff so $\xi_{ij}(t) = 0$ for all strategies j which were not chosen, which reduces the model to choice reinforcement. However if they always choose the worst strategy, then $\xi_{ij}(t) = 1$, which corresponds to weighted fictitious play. If subjects neither choose the best nor the worst strategy, the updating scheme will push them (probabilistically) towards those strategies that yield better payoffs, as is both characteristic of human learning and normatively sensible. These dynamics resemble the adaptive decision making by individuals that has been extensively studied by marketers (e.g. Payne et al 1988). They suggest that consumers choose among decision rules of various complexity by using simple rules and switching to complex rules when there is a benefit from doing so. The $\xi_{ij}(t)$ function is similar. When a player is best-responding it doesn’t pay to waste attention on foregone payoffs and she does not. But when a player is responding badly, she should attend to other strategies’ payoffs, and she does.

The self-tuning EWA model is able to capture the dynamics of the median-action game described in 7.1. Figure 3 shows the empirical frequency for each of the 7 strategies across time for all the six sessions. This aggregate data show that the initial choices were relative more diffuse and converges towards NE behavior (around 4 and 5) in the later rounds. Figure 4 shows how self-tuning EWA model tracks these changes over time. Ho et al (2004) also show how the same model captures learning dynamics in other games ranging from coordination, dominance-solvable to games with a unique mixed-strategy equilibrium. So as in economics and game theory, a single learning model has generality because it can be applied to games with very different structures.



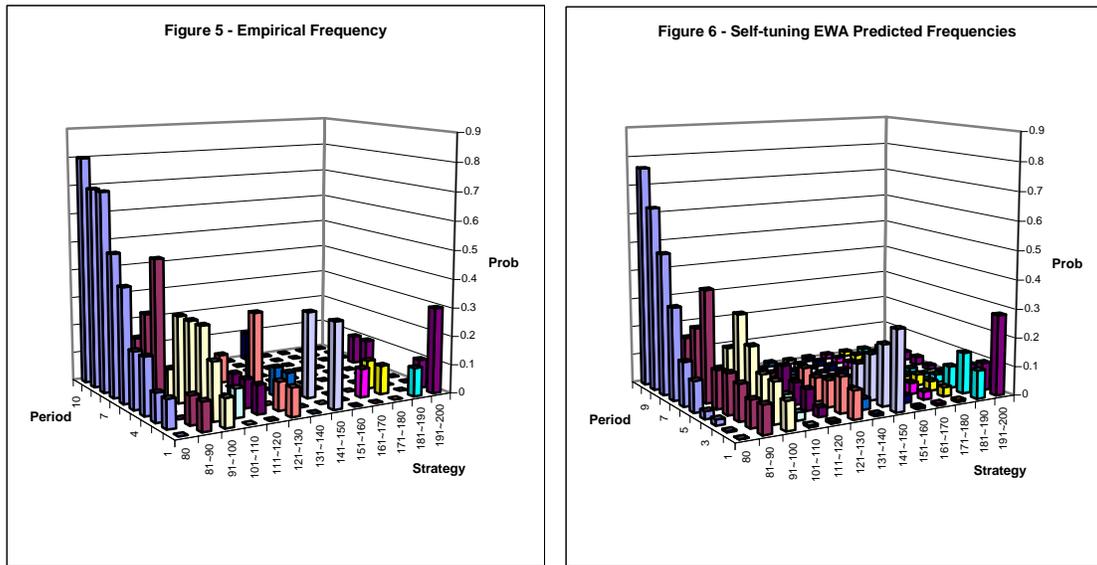
7.3 Marketing Application: Price-Matching with Loyalty

Capra et al (1999) experimentally studied a type of price-matching competition (first introduced by Basu 1994). In this game two firms must choose an integer price between marginal cost and a reservation price which is common to all customers. If the prices are equal, each firm receives the common price as its payoffs (i.e., the sales volume is normalized to one). If the prices are unequal, each of them receives the *lower* of the two prices. In addition, the firm that charges the lower price receives a bonus R and the firm that charges the higher price pays a penalty of R . Hence, this type of competition is akin to a duopoly where the two sellers promise customers that their posted prices are the lowest and offer lowest-price guarantees in the form of price-matching to underpin their claims. Since customers believe that firms that offer such guarantees have the lowest price (Jain and Srivastava 2000), the firm that charges the lowest price earns a goodwill reward while the high-price firm suffers a reputation loss when customers compare their posted prices (see Ho and Lim 2004 for a generalization of this game). As long as $R > I$, the unique NE for this game is that equilibrium prices will be at marginal costs. The proof is simple: If another firm is expected to choose a price P , then by undercutting P by one unit, the lower price firm earns $P - I + R$, which is more than it can earn by matching P (and therefore earning P) or overpricing (earning $P - R$). Because both firms always want to undercut by one unit, there is a “race to the bottom”, in theory. Strikingly, this prediction is also invariant to the size of the penalty/reward factor R .

The authors tested the NE prediction of this game where R was varied at six different levels (5, 10, 20, 25, 50, 80). Subjects grouped randomly in pairs were told to simultaneously

name a price from 80 to 200. Each subject played the game for 10 rounds for each level of R . Figure 5 shows empirical frequencies for $R=50$.

A wide range of prices are posted in the early rounds. Prices gradually fall, between 91-100 in rounds 3-5, 81-90 in rounds 5-6, towards the NE prediction of 80 in later rounds. As shown in Figure 6, the self-tuning EWA model captures the data well.⁴⁶ It is able to capture, given the initial dispersion of prices, the dynamics of equilibration and the high proportion of NE outcomes in the later rounds (around 80% choose the NE in the last round).



8. MARKETING APPLICATIONS

In Sections 2-7, we described a marketing-related application of each of the new models we presented. Our goal was to show exactly how each new model was used and its application could lead to new insights. This section extends and speculates on how these new tools can be more broadly applied to determining and influencing the four main marketing-mix decisions (i.e., decisions on price, product, promotion and place).

8.1 Pricing

The concepts of reference price and asymmetric price effects, which have been ascribed to reference-dependence and loss aversion, is one of the most well-known empirical regularities in marketing (Winer 1986, Mayhew and Winer 1992, Hardie et al 1993, Kalyanaram and Winer

⁴⁶ In Figure 6, the initial attractions are “burned in” from the initial period data rather than computed from the CH model. Using the CH fit, which is a mixture of random prices (from 0-step thinkers) and prices at 200-2R-(K-1) from K-step thinkers, gives a worse initial fit for the first couple of periods but picks up the learning path well after that.

1995). Yet with the exception of Greenleaf (1995), there has been virtually no analytical research on how past prices or the magnitude of loss aversion affects firms' pricing strategies, particularly in the contexts of competition and inter-temporal pricing. For example, does reference-dependence predict that prices should move in cycles and decline within each cycle as in the case of a monopoly? More importantly, can marketing researchers formulate pricing strategies for reference-dependent customers for new products or for those products in which customers do not have price histories? Koszegi and Rabin (2004) have developed a nice model of reference-dependent preferences that assumes that the customer's reference point is not based on past prices, but on whether she *expects* to purchase a certain product on a particular shopping trip. Their framework provides us with a natural way to view advertising and other practices as a means to influence customer expectations of purchase and willingness-to-pay. Marketing practices like free 30-day trials, and test drives of cars, can be seen as attempts to shift the reference point, so that *not* buying the product is perceived as a loss.

Heidhues and Koszegi (2005) study the influence of reference-dependence on pricing. They show that if consumers are sufficiently loss-averse, then firms should not adjust prices continuously as costs change (prices are "sticky"). The reason is that cutting a price when costs shrink creates a low reference point for future purchases, which makes it hard to get consumers to accept a future price increase even if costs change; so firms should just keep prices fixed. There are many studies of price stickiness, which often ascribe stickiness to exogenous "menu costs"; their paper derives stickiness *endogenously* from loss-averse consumer preferences and optimal firm behavior. They also show that price-cost markups vary in a counter-cyclical fashion, shrinking in booms and growing in recessions.

Another area of research is to study how firms should price products in the presence of customers who are hyperbolic discounters. In a recent paper, Oster and Morton (2004) studied the prices of about 300 magazines in the US. They found that the ratio of subscription to newsstand prices are higher for "investment" magazines like *Forbes* compared to "leisure" magazines like *Entertainment Weekly*. This can be explained by the fact that the sophisticated time-inconsistent customers require only relatively small subscription discounts to induce them to subscribe to the "investment" magazines.

Another research question concerns the pricing of durable goods by a monopoly. One question is the optimal price path for a rational firm that sells to hyperbolically-discounting customers. Another interesting case is when decision-makers at the firm are hyperbolic discounters (a strong possibility as many firm leaders face pressure from the stock market to deliver short-term profits) or when customers mistakenly believe that the firm's managers have hyperbolic time preferences.

Finally, *all* the existing models of price competition in marketing we are aware of use NE as a tool to generate descriptive and prescriptive predictions. Applying alternative solution concepts such as QRE or CH models represents a potentially fruitful area of research as they more accurately represent the presence of latent payoffs or cognitive limitations of decision-makers respectively. A recent paper in the field of industrial organization (Baye and Morgan 2004) has applied the QRE to Bertrand competition and showed that above marginal-cost prices and even persistent price dispersion can exist even when search costs are zero (such as when customers can use engines to search lowest prices on the Internet).

8.2 Product

Many customers are offered the option of buying extended product or service warranties when they purchase consumer durables. These extended warranties are marketed as ways to augment the basic warranty (usually provided free-of-charge by the manufacturer for a limited period) in order to “insure” the value of the new purchase. They are either offered by the manufacturer of the product or by third-party companies, and either increases the scope or length of the basic warranty (or both). One of the open research questions is why customers buy these warranties in the first place – while many articles in the practitioner literature have suggested that the additional coverage is not worth the price of the premiums, this has not been studied formally. Are customers driven by extreme risk aversion or is it more likely loss aversion (as we believe) that drives the demand for such warranties? Specifically, are customers loss-averse with respect to losing their newly owned products (with the value of the product inflated through the endowment effect) such that they are willing to buy or even pay extra for the extended warranties? Or do they overweight the small probabilities of breakdown (as in the prospect theory $\pi(p)$ weighting function)?

The Bass diffusion model is one of the most widely-used marketing tools to predict and explain the rate of product adoption of a product at an aggregate level. Chatterjee and Eliashberg (1990) provided a theoretical underpinning to the Bass diffusion model by using a rational framework to model the adoption process of individual potential customers. We believe that this “micro-modeling” approach can be naturally extended to incorporate bounded-rationality on the part of the consumers. For example, the (β, δ) model of Section 4.2 can be adapted to capture not only customers who seek immediate gratification and make early purchases but also customers who have decided to buy but keep putting off purchases.

Perhaps the most straightforward application of behavioral economics in the area of product management is to model customers who are hyperbolic discounters in the consumer goods industry (see Hoch and Loewenstein 1991 for early work). For example, decisions of

product assortment and inventory control in a category can be made based on the purchase behavior of segments of naïfs and sophisticates. For goods that are harmful in the long run, sophisticates will prefer small package sizes to aid self-control (Werthenbroch 1998). Certain types of products are probably also more prone to impulse purchase (i.e., they trigger low values of β or limbic activity) and hence are ideally placed near a cash register. Furthermore, how should a manager determine the optimal mix of different packaging sizes given the existence of the various customer segments?

8.3 Promotion

One of the key questions in promotions is the frequency and depth of price promotions. Narasimhan (1988) and Raju et al (1990) have shown that when two stores have their own loyal customers and compete for “switchers” by setting prices, the NE solution is characterized by a mixed-strategy profile that has been interpreted as price promotion (Rao et al 1995). Results of experimental games by Choi and Messinger (2005) and the first two authors of this paper⁴⁷ suggest however that while NE is a fairly good predictor of mean prices, it predicts the *distribution* of prices poorly, and it is the distribution determines the depth and frequency of price promotions. One of the regularities in Choi and Messinger (2005) that cannot be explained by NE is that prices start high initially, decline over time, then increase abruptly again, with this cycle repeating over the experimental rounds. Closer examination of the data reveals that in each period, subjects appear to condition their price on those that had been chosen by their rivals in the previous round. They try to undercut their rival’s most recent price, a dominant strategy if their rival’s price remains unchanged.⁴⁸ This is one of the clearest demonstration of the breakdown of the mutual consistency assumption of NE, for subjects fail to account for the fact that their rivals would have anticipated these actions *fully* and revised their prices accordingly. Both the CH and EWA models could be applied to these dynamic processes and could potentially predict the pattern of prices better than NE, by taking into account limited strategizing by decision makers, and how they learn from experience.

Another popular promotional vehicle is the use of price rebates to induce customers to accelerate purchase. Customers would usually have to pay the “full price” for the product, but can receive a cash refund if they redeem the rebate by mailing in their proof of purchase within a stipulated period of time. An interesting issue is how rebate redemption behavior will change if firms vary the length of time they allow customers to redeem the rebate and the monetary size

⁴⁷ The data is available from the authors of this paper upon request.

⁴⁸ The experiments in Choi and Messinger (2005) involved pairs of subjects that set prices in a repeated setting.

of these rebates if the customers consist of naïf and sophisticated hyperbolic discounters. For instance, should the profit-maximizing firm extend the redemption period and increase the size of the rebate to attract the naïf customers, knowing that they are more likely to delay or forget (if it is a function of delay) to redeem the rebate and end up paying the full price? And how should this be reconciled with the desires of the sophisticates for tighter redemption deadlines or even “instant” rebates as forms of commitment devices? These questions have yet to be formally analyzed.

8.4 Place

The study of marketing distribution channels has been heavily dominated by theoretical modeling. Compared to areas like consumer behavior, there is relatively little empirical research on channels. The study of channels presents marketing researchers unique opportunities to combine empirical and especially experimental work (Anderson and Coughlan 2002) with the modeling tools in behavioral economics. Existing models of channel relationships can be generalized following careful empirical studies that challenge the predictions of orthodox economic theory.

For example, firms in an independent channel cannot capture all the possible profits in the channel by using a linear price contract (this is the double marginalization (DM) problem mentioned in section 2.3). Pioneering work by Jeuland and Shugan (1983) and Moorthy (1987) showed theoretically that more complex nonlinear contracts such as quantity discounts and two-part tariffs can eliminate the DM problem. One of the fundamental questions that arise from this stream of research is: given that there are many forms of quantity discount contracts that are used in practice, should some contracts be preferred over others? Using standard experimental economics methodology, Lim and Ho (2004) compared two popular quantity discount contracts, the two-block and three-block declining tariffs, which are theoretically equally good in solving the DM problem. They found that the three-block tariff yields higher channel profits than the two-block tariff. Using elements of reference-dependence and the EWA model (which captures the fact that players in a game care about counterfactual or forgone payoffs), the authors showed that the results can be consistent with an equilibrium where the downstream firm in the channel cares about the (additional) counterfactual profits it would have earned if the lower marginal price in the adjacent block were to be applied to the current block of the tariff. This research also leads to the more general and wide-open question of whether the structures of different nonlinear contracts induce different forms of preferences, or if there exists a single generalized model can account for the relative performances of different contracts.

In a different behavioral approach to DM, Cui et al (2004) notes that the DM problem may not create inefficiency, in practice, even with a linear price contract if channel members are not purely self-interested. The authors show that if firms are sufficiently inequity averse (as in Fehr and Schmidt 1999), the efficient outcome that is reached by a vertically-integrated monopoly can be achieved with separate firms, so that complex nonlinear contracts are superfluous.

Another potentially rich area for applying the models of behavioral economics is in sales-force management. Questions that can be studied include how the structure of individual or team compensation plans changes, or how sales territories can be divided, if salespersons have social preferences and care about the payoffs of both the firm and/or other salespersons, or are intrinsically motivated (see Camerer and Malmendier forthcoming).

9. RESEARCH IMPLICATIONS

In this paper, we show how the standard utility function of decision makers like customers or managers can be generalized to capture three well-established empirical regularities about people: they care both about absolute and changes in outcomes relative to some reference point; they care both about their own and others' payoffs, and they exhibit a taste for immediate gratification. Recent advances in neuroscience suggest that these empirical regularities can be traced to the way our brain operates. These utility modifications have been increasingly popular in economics because they have one or two additional parameters over the standard models, which make them amenable to empirical testing and estimation.

We also provide three alternatives to the standard NE solution concepts in competitive situations. The QRE model generalizes the standard equilibrium concepts by allowing decision makers to better-respond instead of best-respond. The CH model relaxes the mutual consistency assumption of NE to allow decision makers to encounter surprises (since their beliefs of what others will do may not be the same as what others actually do). The self-tuning EWA model captures how decision makers might respond to experience over time when they compete in repeated interactions. These empirical alternatives have been shown to predict behavior significantly better in controlled laboratory settings and have shown some promise to perform well in field settings. Again, each alternative comes with only one additional parameter relative to NE can be effectively used for predicting market behavior.

We would like to emphasize that the proposed models are by no means the final or complete models that capture all documented behavioral regularities. In fact, research in psychology and economics is still actively delineating the boundaries of these behavioral regularities (e.g. see Novemsky and Kahneman 2005 on loss aversion). It is also important to

continue to test the predictions of standard models. More tests give more facts about when and how the standard models fail (and when they succeed), which are useful because behavioral alternatives are still being proposed and refined. Tests which include rational models as special cases (which most models in sections 2-7 do) also give estimates of the values of the additional parameter(s) introduced in the behavioral models, and can also serve as a check of whether these behavioral models are well-specified (e.g. showing whether the parameter values obtained from the data are psychologically plausible). The recent trend towards testing well-specified theories in the laboratory (e.g. Amaldoss et al 2000, Ghosh and John 2000, Srivastava et al 2000, Ho and Zhang 2004, Lim and Ho 2004) is a promising avenue for future research and often makes it easier to clearly separate rational and behavioral predictions.

Marketing is inherently an applied field. Its rich sources of field data are as useful as those in any discipline for testing behavioral models. A lesson can be learned from the phenomenal success of conjoint analysis. Conjoint analysis is the most celebrated marketing research tool in marketing because its proponents understand that it is important for it to be precise, general, and empirical (see Leeflang et al. 2000 and Bradlow et al (2004a, 2004b)). In particular, it is crucial to subject the existing models to serious empirical testing and fine-tune them through simple extensions. This is how we see the accumulation of knowledge occurring through behavioral economics models. The basic point of this paper is that the empirical power of the standard models can be enhanced by adopting these modifications, and combines the best of psychological insight with the power of mathematical formality.

9.1 Implications for behavioral researchers

We believe it is crucial for behavioral researchers to continue to document robust violations of standard models. However, showing the existence of an important behavioral regularity is only the necessary first step towards its wide applicability in marketing. To receive wide applicability, it is necessary that the regularity be precisely specified in a formal model. This formal specification process requires an active collaboration between the behavioral and quantitative researchers. Such collaboration, whether in the form of co-authorship or mutual influence from a common understanding of facts and modeling language, is a promising area for future research.

This paper also implies that it is important for the behavioral researchers to demonstrate important behavioral regularities in the field. The field experimentation approach allows one to test the idea that “bounded rationality” may not survive in the marketplaces because the latter rewards rationality more.

9.2 Implications for empirical researchers

Empirical researchers in marketing have been very successful in testing standard economic models using the field datasets. We believe our proposed revised utility functions should be used in future empirical testing because they nest the standard models. The empirical tests will provide useful information as to when and how the standard models fail. We believe this is an extremely fertile area of research.

The C-H and QRE models can be used to study off-equilibrium path behavior in the framework of New Empirical Industrial Organization, which is an active area of research for marketing. These revised models offer two advantages. First, they allow researchers to empirically check a foundational assumption of the field – that is, firms always play pure equilibrium strategies and never encounter surprises and make mistakes. We speculate that this standard assumption is likely to be true in mature industries and less so in developing industries. To the extent that this assumption is problematic, the new approaches provide a way to handle this inadequacy. Second, these new methods of game-theoretic analysis allow us to include and study off-equilibrium path behavior. This could significantly change the estimated demand model and the implied price elasticities.

There is room for applying the EWA learning models to capture how managers and customers learn over time. For instance, Ho and Chong (2003) apply the EWA model to predict consumers' product choices and show that it outperforms several existing models including Guadagni and Little (1983) in an extensive dataset involving more than 130,000 purchases across 16 product categories.

9.3 Implications for analytical researchers

We believe there is a huge opportunity for analytical modelers to incorporate the revised utility and game-theoretic functions in their modeling work. For example, it will be fruitful to investigate how a firm's marketing mix actions will change when it faces a group of customers who have reference-dependent preferences, care about fairness, and are impatient. How would the market structure and degree of competition vary as a result of these changes? We believe incorporating these changes might provide explanations to many seemingly market paradoxes that cannot be explained using standard economic models. For example, Rotemberg (2005) show that if customers care about fairness, it is optimal for firms to engage in temporary sales events and they should announce their intention of increasing prices before actually doing so.

Similarly, we believe the analytical modelers can apply both the CH and the QRE models to analyze how firms might compete in a specific market setting. These models will

allow us to capture behavior that would otherwise be suppressed by the stringent requirements of no surprises and zero mistakes. Both seem particularly promising in modeling rapidly changing product markets and in markets where firms may not have a sufficient knowledge of actual demand and supply conditions.

Besides the six topics we have narrowly focused on, there are many more rich questions about applicability of behavioral economics to market-level outcomes which marketing modelers are well-equipped to study. A rapidly-emerging question is how firms should make marketing mix choices when consumers exhibit various types of bounds on rationality. As Ellison (2005) notes, if consumers exhibit various biases relative to rational choice, from the firm's point of view these biases will have the same practical importance as product differentiation—except an identical product might be “differentiated” by idiosyncrasies in consumer cognition, rather than in tastes. This insight suggests familiar models might be adapted to study the behavioral economics of firm marketing behavior in the face of consumer rationality limits.

REFERENCES

- Ainslie, G. (1975), "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control", *Psychological Bulletin*, 82 (4), 463-496.
- Amaldoss, W., Meyer R., Raju J. and A. Rapoport (2000), "Collaborating to Compete A Game-Theoretic Model and Experimental Investigation of the Effect of Profit-Sharing Arrangement and Type of Alliance," *Marketing Science*, 19 (2), 105-126.
- Anderson, E. and A. Coughlan (2002), "Channel Management: Structure, Governance and Relationship Management", in Weitz, B. and R. Wensley (eds.), *Handbook of Marketing*, London: Sage Publications.
- Ariely, D. and K. Wertenbroch (2002), "Procrastination, Deadlines, and Performance: Self-control by Precommitment", *Psychological Science*, 13 (3), 219-224.
- Babcock, L. and G. Loewenstein (1997), "Explaining Bargaining Impasse: The Role of Self-Serving Biases", *Journal of Economic Perspectives*, 11 (1), 109-126.
- Barberis, N. and R. Thaler (2003), "A Survey of Behavioral Finance", in Constantinides, G., Harris, M. and R. Stultz (eds.), *Handbook of the Economics of Finance*, Elsevier Science, North Holland, Amsterdam.
- Basu, A., Lal, R., Srinivasan, V. and R. Staelin (1985), "Salesforce Compensation Plans: An Agency Theoretic Perspective", *Marketing Science*, 4 (4), 267-291.
- Basu, K. (1994), "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory", *American Economic Review*, 84 (2), 391-395.
- Bateman, I., Kahneman, D., Munro, A., Starmer, C. and R. Sugden (forthcoming), "Testing Competing Models of Loss Aversion: An Adversarial Collaboration", *Journal of Public Economics*.
- Baye, M. and J. Morgan (2004), "Price-Dispersion in the Lab and on the Internet: Theory and Evidence", *RAND Journal of Economics*, 35 (3), 449-466.
- Benartzi, S. and Thaler, R. (1995), "Myopic Loss Aversion and the Equity Premium Puzzle", *Quarterly Journal of Economics*, 110 (1), 73-92.
- Benzion, U., Rapoport, A. and J. Yagil (1989), "Discount Rates Inferred from Decisions: An Experimental Study", *Management Science*, 35, 270-284.
- Bhatt, M., and C. Camerer (forthcoming), "Self-Referential Thinking and Equilibrium as a State of Mind: fMRI Evidence from Games," *Games and Economic Behavior*.
- Bolton, G. and A. Ockenfels (2000), "ERC: A Theory of Equity, Reciprocity and Competition", *American Economic Review*, 90, 166-193.
- Bradlow, E., Hu, Y. and T. Ho (2004a), "Modeling Behavioral Regularities of Consumer Learning in Conjoint Analysis", *Journal of Marketing Research*, 41, 392-396.
- Bradlow, E., Hu, Y. and T. Ho (2004b), "A Learning-based Model for Imputing Missing Levels in Partial Conjoint Profiles", *Journal of Marketing Research*, 41, 369-381.

Camerer, C. (1988), "An Experimental Test of Several Generalized Utility Theories", *Journal of Risk and Uncertainty*, 2, 61-104.

Camerer, C. (1995), "Individual Decision Making", chapter 8 in Kagel, J. and A. Roth (eds.), *Handbook of Experimental Economics*, Princeton: Princeton University Press.

Camerer, C. (1999), "Behavioral Economics: Reunifying Psychology and Economics", *Proceedings of the National Academy of Sciences*, 96, 10575-10577.

Camerer, C. (2001), "Prospect Theory in the Wild: Evidence from the Field", in Kahneman, D. and A. Tversky (eds.), *Choices, Values and Frames*, Cambridge: Cambridge University Press.

Camerer, C. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton: Princeton University Press.

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

Camerer, C. and T. Ho (1998), "EWA Learning in Coordination Games: Probability Rules, Heterogeneity, and Time Variation", *Journal of Mathematical Psychology*, 42, 305-326.

Camerer, C. and T. Ho (1999), "Experience-Weighted Attraction Learning in Normal Form Games", *Econometrica*, 67, 837-874.

Camerer, C. and D. Lovo (1999), "Overconfidence and Excess Entry: An Experimental Approach", *American Economic Review*, 89, 306-318.

Camerer, C. and U. Malmendier (forthcoming), "Behavioral Economics of Organizations", in Diamond, P. and H. Vartiainen (eds.), *Economic Institutions and Behavioral Economics*, Yrjö Jahnsson Foundation, Princeton University Press.

Camerer, C., Babcock, L., Loewenstein, G. and R. Thaler (1997), "Labor Supply of New York City Cab Drivers: One Day at a Time", *Quarterly Journal of Economics*, 112, 407-442.

Camerer, C., Ho, T. and J. Chong (2002), "Sophisticated EWA Learning and Strategic Teaching in Repeated Games", *Journal of Economic Theory*, 104, 137-188.

Camerer, C., Issacharoff, S., Loewenstein, G. and T. O'Donoghue (2003a), "Regulation for Conservatives: Behavioral Economics and the Case for Asymmetric Paternalism", *University of Pennsylvania Law Review*, 151, 1211-1254.

Camerer, C., Loewenstein, G. and M. Rabin (eds.) (2003b), *Advances in Behavioral Economics*, Russell Sage Foundation, Princeton University Press.

Camerer, C., Ho, T. and J. Chong (2004), "A Cognitive Hierarchy Model of Games", *Quarterly Journal of Economics*, 119, 861-898.

Capra, M., Goeree, J., Gomez, R. and C. Holt (1999), "Anomalous Behavior in a Traveler's Dilemma?", *American Economic Review*, 89, 678-690.

Charness, G. and M. Rabin (2002) "Understanding Social Preferences with Simple Tests", *Quarterly Journal of Economics*, 117, 817-869.

- Chatterjee, R. and J. Eliashberg (1990), "The Innovation Diffusion Process in a Heterogeneous Population: A Micromodeling Approach", *Management Science*, 36 (9), 1057-1079.
- Chua, Z. and C. Camerer (2004), "Experiments on Intertemporal Consumption with Habit Formation and Social Learning", unpublished research paper.
- Chen, K., Lakshminarayanan, V. and L. Santos (2005), "The Evolution of Our Preferences: Evidence from Capuchin-Monkey Trading Behavior", Working Paper, Yale School of Management.
- Choi, S. and P. Messinger (2005), "Edgeworth Promotional Cycles in the Loyal-Switcher Game", Working Paper, Marketing Department, University of Alberta.
- Cui, H., Raju, J. and J. Zhang (2004), "Fair Channel", Working Paper, Marketing Department, Wharton School, University of Pennsylvania.
- DeGeorge, F., Patel, J. and R. Zeckhauser (1999), "Earnings Management to Exceed Thresholds", *Journal of Business*, 72 (1), 1-33.
- Della Vigna, S. and U. Malmendier (2004), "Contract Design and Self-Control: Theory and Evidence", *Quarterly Journal of Economics*, 119, 353-402.
- Dufwenberg, M. and G. Kirchsteiger (2004), "A Theory of Sequential Reciprocity", *Games and Economic Behavior*, 47, 268-298.
- Ellison, G. (2005), "Bounded Rationality in Industrial Organization", presented at 2005 World Congress of the Econometric Society, www.econ.ucl.ac.uk/eswc2005.
- Erev, I. and A. Roth (1998), "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique Mixed-Strategy Equilibria", *American Economic Review*, 88, 848-881.
- Farber, H. (2004), "Reference-Dependent Preferences and Labor Supply: The Case of New York City Cab Drivers", Working Paper #497, Industrial Relations Section, Princeton University.
- Farber, H. (2005), "Is Tomorrow Another Day? The Labor Supply of New York City Cab Drivers", *Journal of Political Economy*, 113 (1), 46-82.
- Fehr, E. and K. Schmidt (1999), "A Theory of Fairness, Competition and Cooperation", *Quarterly Journal of Economics*, 114, 817-868.
- Fehr, E., Klein, A. and K. Schmidt (2004), "Contracts, Fairness and Incentives", Discussion Paper 2004-07, Department of Economics, University of Munich.
- Friedman, J., Palfrey T. and J. Thisse (1995), "Random Choice Behavior in Continuum Models", unpublished research note.
- Frederick, S., Loewenstein, G. and T. O'Donoghue (2002), "Time Discounting and Time Preference: A Critical Review", *Journal of Economic Literature*, 40, 351-401.
- Genesove, D. and C. Mayer (2001), "Loss Aversion and Seller Behavior: Evidence from the Housing Market", *Quarterly Journal of Economics*, 116, 1233-1260.

- Ghosh, M. and G. John (2000), "Experimental Evidence for Agency Models of Salesforce Compensation", *Marketing Science*, 19 (4), 348-365.
- Greenleaf, E. (1995), "The Impact of Reference Price Effects on the Profitability of Price Promotions", *Marketing Science*, 14, 82-104.
- Gruber, J. and B. Koszegi (2001), "Is Addiction Rational? Theory and Evidence", *Quarterly Journal of Economics*, 116, 1261-1305.
- Guadagni, P. and J. Little (1983), "A Logit Model of Brand Choice Calibrated on Scanner Data", *Marketing Science*, 2, 203-238.
- Gul, F. and W. Pesendorfer (2001), "Temptation and Self-Control", *Econometrica*, 69, 1403-1435.
- Hannan, R., Kagel, J. and D. Moser (forthcoming), "Partial Gift Exchange in an Experimental Labor Market: Impact of Subject Population Differences, Productivity Differences and Effort Request on Behavior", *Journal of Labor Economics*.
- Hardie, B., Johnson, E. and P. Fader (1993), "Modeling Loss Aversion and Reference Dependence Effects on Brand Choice", *Marketing Science*, 12 (4), 378-394.
- Harper, D. (1982), "Competitive Foraging in Mallards: "Ideal Free" Ducks", *Animal Behavior*, 20, 575-584.
- Healy, P. (2004), "Fairness, or Just Gambling on It? An Experimental Analysis of the Gift Exchange Game", Social Science Working Paper 1183, Caltech.
- Heidhues, P. and B. Koszegi (2005), "The Impact of Consumer Loss Aversion on Pricing", Working Paper, Department of Economics, UC Berkeley.
- Ho, T., Camerer, C. and K. Weigelt (1998), "Iterated Dominance and Iterated Best-Response in p -Beauty Contests", *American Economic Review*, 88, 947-969.
- Ho, T. and N. Lim (2004), "A Theory of Lowest-Price Guarantees", Working Paper, Marketing Department, Haas School of Business, UC Berkeley.
- Ho, T. and J. Chong (2003), "A Parsimonious Model of SKU Choice," *Journal of Marketing Research*, 40, 351-365.
- Ho, T. and J. Zhang (2004), "Does the Format of Pricing Contracts Matter?" Working Paper, Marketing Department, Haas School of Business, University of California at Berkeley.
- Ho, T., Camerer, C. and J. Chong (2004). "The Economics of Learning Models: A Self-tuning Theory of Learning in Games", Working Paper, Haas School of Business, UC Berkeley.
- Hoch, S. and G. Loewenstein (1991), "Time-Inconsistent Preferences and Consumer Self-Control", *Journal of Consumer Research*, 17 (4), 492-507.
- Holcomb, J. and P. Nelson (1992), "Another Experimental Look at Individual Time Preference", *Rationality and Society*, 4 (2), 199-220.
- Jain, S. and J. Srivastava (2000), "An Experimental and Theoretical Investigation of Price-Matching Refund Policies", *Journal of Marketing Research*, 37, 351-62.

- Jeuland, A. and S. Shugan (1983), "Managing Channel Profits", *Marketing Science*, 2 (3), 239-272.
- Kalyanaram, G. and R. Winer (1995), "Empirical Generalizations from Reference Price and Asymmetric Price Response Research", *Marketing Science*, 14 (3), G161-G169.
- Kahneman, D. (1988), "Experimental Economics: A Psychological Perspective", in Tietz, R., Albers, W., and R. Selten, (eds.), *Bounded Rational Behavior in Experimental Games and Markets*, New York: Springer-Verlag.
- Kahneman, D. and A. Tversky (1979), "Prospect Theory: An Analysis of Decision Under Risk", *Econometrica*, 47, 263-291.
- Kahneman, D. and A. Tversky (1992), "Advances in Prospect Theory: Cumulative Representation of Uncertainty", *Journal of Risk and Uncertainty*, 5, 297-324.
- Kahneman, D., Knetsch, J. and R. Thaler (1990), "Experimental Tests of the Endowment Effect and the Coase Theorem", *Journal of Political Economy*, 98, 1325-1348.
- Kivetz, R., Netzer, O. and V. Srinivasan (2004), "Alternative Models for Capturing the Compromise Effect", *Journal of Marketing Research*, 41, 237-257.
- Knutson, B. & R. Peterson (2005), "Neurally Reconstructing Expected Utility", *Games and Economic Behavior*, 52, 305-315.
- Koszegi, B. and M. Rabin (2004), "A Model of Reference-Dependent Preferences", Working Paper, Department of Economics, University of California at Berkeley.
- Laibson, D. (1997), "Golden Eggs and Hyperbolic Discounting", *Quarterly Journal of Economics*, 112 (2), 443-477.
- Leeflang, P., Wittink, D., Wedel, M. and P. Naert (eds.), (2000), *Building Models for Marketing Decisions*, New York: Kluwer Academic Publishers.
- Lim, N. and T. Ho (2004), "Do Quantity Discounts Coordinate the Channel? Evidence from Multi-Block Tariffs", Working Paper, Marketing Department, Haas School of Business, University of California at Berkeley.
- List, J. (2003) "Does Market Experience Eliminate Market Anomalies?", *Quarterly Journal of Economics*, 118, 41-71.
- Little, J. (1970), "Models and Managers: The Concept of a Decision Calculus", *Management Science*, 16 (8), B466-B485.
- Loewenstein, G. (1987), "Anticipation and the Valuation of Delayed Consumption", *Economic Journal*, 97, 666-684.
- Loewenstein, G. and D. Prelec (1992), "Anomalies in Intertemporal Choice: Evidence and An Interpretation", *Quarterly Journal of Economics*, 107, 573-597.
- Loewenstein, G. and D. Prelec (1993), "Preferences for Sequences of Outcomes", *Psychological Review*, 100 (1), 91-108.

- Loewenstein, G. and R. Thaler (1989), "Anomalies: Intertemporal Choice", *Journal of Economic Perspectives*, 3, 181-193.
- Mayhew, G. and R. Winer (1992), "An Empirical Analysis of Internal and External Reference Price Effects Using Scanner Data", *Journal of Consumer Research*, 19, 62-70.
- McClure, S., Laibson, D., Loewenstein, G. and J. Cohen (2004), "Separate Neural Systems Value Immediate and Delayed Monetary Rewards", *Science*, 306, 503-507.
- McFadden, D. (1999), "Rationality for Economists", *Journal of Risk and Uncertainty*, 19, 73-105.
- McKelvey, R. and T. Palfrey (1995), "Quantal Response Equilibria for Normal Form Games", *Games and Economic Behavior*, 10, 6-38.
- Miravete, E. (2003). "Choosing the Wrong Calling Plan? Ignorance and Learning", *American Economic Review*, 93, 297-310.
- Moorthy, S. (1985), "Using Game Theory to Model Competition", *Journal of Marketing Research*, 22, 262-82.
- Moorthy, S. (1987), "Managing Channel Profits: Comment", *Marketing Science*, 6 (4), 375-379.
- Nagel, R. (1995), "Unraveling in Guessing Games: An Experimental Study", *American Economic Review*, 85, 1313-1326.
- Narasimhan, C. (1988), "Competitive Promotional Strategies", *Journal of Business*, 61 (4), 427-449.
- Novemsky, N. and D. Kahneman (2005), "The Boundaries of Loss Aversion", *Journal of Marketing Research*, 42, 119-128.
- Ochs, J. (1995), "Games with Unique Mixed-Strategy Equilibria: An Experimental Study", *Games and Economic Behavior*, 10, 174-189.
- Odean, T. (1998), "Are Investors Reluctant to Realize their Losses?", *Journal of Finance*, 53, 1775-1798.
- O'Donoghue, T. and M. Rabin (1999a), "Doing It Now or Later", *American Economic Review*, 89 (1), 103-124.
- O'Donoghue, T. and M. Rabin (1999b), "Incentives for Procrastinators", *Quarterly Journal of Economics*, 114, 769-816.
- O'Donoghue, T. and M. Rabin (1999c), "Addiction and Self- Control", in Elster, J. (ed.), *Addiction: Entries and Exits*, Russell Sage Foundation.
- O'Donoghue, T. and M. Rabin (2000), "The Economics of Immediate Gratification", *Journal of Behavioral Decision Making*, 13 (2), 233-250.
- O'Donoghue, T. and M. Rabin (2001), "Choice and Procrastination", *Quarterly Journal of Economics*, 116, 121-160.
- O'Donoghue, T. and M. Rabin (2002), "Addiction and Present-Biased Preferences", Working Paper, Department of Economics, Cornell University.

- O'Donoghue, T. and M. Rabin (2003), "Self Awareness and Self Control", in Baumeister, R., Loewenstein, G. and D. Read, (eds.), *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*, Russell Sage Foundation.
- Oster, S. and F. Morton (2004), "Behavioral Decision Making: An Application to the Setting of Magazine Subscription Prices", Working Paper, Yale School of Management.
- Payne, J., Bettman, J. and E. Johnson (1988), "Adaptive Strategy Selection in Decision Making", *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14, 534-552.
- Pender, J. (1996), "Discount Rates and Credit Markets: Theory and Evidence from Rural India", *Journal of Development Economics*, 50 (2), 257-296.
- Phelps, E. and R. Pollak (1968), "On Second-Best National Saving and Game-Equilibrium Growth", *Review of Economic Studies*, 35 (2), 185-199.
- Plott, C. and K. Zeiler (2005), "The Willingness to Pay/Willingness to Accept Gap, the "Endowment Effect", Subject Misconceptions and Experimental Procedures for Eliciting Valuations", *American Economic Review*, 95 (3), 530-545.
- Prelec, D. and G. Loewenstein (1991), "Decision Making over Time and under Uncertainty: A Common Approach", *Management Science*, 37, 770-786.
- Putler, D. (1992), "Incorporating Reference Price Effects into a Theory of Consumer Choice", *Marketing Science*, 11, 287-309.
- Rabin, M. (1993), "Incorporating Fairness into Game Theory and Economics", *American Economic Review*, 83, 1281-1302.
- Rabin, M. (1998), "Psychology and Economics", *Journal of Economic Literature*, 36, 11-46.
- Rabin, M. (2002), "A Perspective on Psychology and Economics", *European Economic Review*, 46, 657-685.
- Raju, J., V. Srinivasan and R. Lal (1990), "The Effects of Brand Loyalty on Competitive Price Promotional Strategies", *Management Science*, 36, 276-304.
- Rao, R., Arjunji, R. and B. Murthi, (1995) "Game Theory and Empirical Generalizations Concerning Competitive Promotions", *Marketing Science*, 14 (3), G89-G100.
- Rapoport, A. and D. Seale (forthcoming), "Coordination Success in Non-Cooperative Large Group Market Entry Games", in Plott, C. and V. Smith, (eds.), *Handbook of Experimental Economic Results*.
- Rapoport, A., Seale, D. and E. Winter (2002), "Coordination and Learning Behavior in Large Groups with Asymmetric Players", *Games and Economic Behavior*, 39, 111-136.
- Rotemberg, J. (2005), "Fair Pricing", Working Paper, Harvard Business School.
- Rubinstein, A. (2003), "'Economics and Psychology'? The Case of Hyperbolic Discounting", *International Economic Review*, 44, 1207-1216.
- Samuelson, P. (1937), "A Note on the Measurement of Utility", *Review of Economic Studies*, 4, 155-161.

- Schultz, W. and A. Dickinson, (2000), "Neuronal Coding of Prediction Errors", *Annual Review of Neuroscience*, 23, 473-500.
- Selten, R. (1998), "Features of Experimentally Observed Bounded Rationality", *European Economic Review*, 42, 413-436.
- Simonson, I. (1989), "Choice Based on Reasons: The Case of Attraction and Compromise Effects", *Journal of Consumer Research*, 16 (2), 158-174.
- Simonson, I. and A. Tversky (1992), "Choice in Context: Tradeoff Contrast and Extremeness Aversion", *Journal of Marketing Research*, 29, 281-295.
- Srivastava, J., Chakravarti, D. and A. Rapoport (2000), "Price and Margin Negotiations in Marketing Channels: An Experimental Study of Sequential Bargaining under One-Sided Uncertainty and Opportunity Cost of Delay", *Marketing Science*, 19 (2), 163-184.
- Thaler, R. (1980), "Toward a Positive Theory of Consumer Choice", *Journal of Economic Behavior and Organization*, 1, 39-60.
- Thaler, R. (1981), "Some Empirical Evidence on Dynamic Inconsistency", *Economic Letters*, 8, 201-207.
- Tovar, P. (2004), "The Effects of Loss Aversion on Trade Policy and the Anti-Trade Bias Puzzle", Working Paper, Economics Department, University of Maryland.
- Tversky, A. and D. Kahneman (1981), "The Framing of Decisions and the Psychology of Choice", *Science*, 211 (4481), 453-458.
- Tversky, A. and D. Kahneman (1991), "Loss Aversion in Riskless Choice: A Reference-Dependent Model", *Quarterly Journal of Economics*, 106, 1039-1061.
- Tversky, A. and I. Simonson (1993), "Context-Dependent Preferences: The Relative Advantage Model", *Management Science*, 39 (10), 1179-1189.
- Van Damme, E. (1999), "Game Theory: The Next Stage", in Gerard-Varet, L, Kirman, A. and M. Ruggiero (eds.), *Economics beyond the Millennium*, Oxford University Press.
- Van Huyck, J., Battalio, R. and R. Beil (1991), "Strategic Uncertainty, Equilibrium Selection and Coordination Failure in Average Opinion Games", *Quarterly Journal of Economics*, 106, 885-911.
- Von Neumann, J. and O. Morgenstern (1944), *The Theory of Games and Economic Behavior*, Princeton: Princeton University Press.
- Weber, M. and C. Camerer (1998), "The Disposition Effect in Securities Trading: An Experimental Analysis", *Journal of Economic Behavior and Organization*, 33, 167-184.
- Wertebroch, K. (1998), "Consumption Self-Control by Rationing Purchase Quantities of Virtue and Vice", *Marketing Science*, 17 (4), 317-337.
- Winer, R. (1986), "A Reference Price Model of Demand for Frequently-Purchased Products", *Journal of Consumer Research*, 13, 250-256.

Zauberman, G. and J. Lynch (2005), "Resource Slack and Discounting of Future Time versus Money", *Journal of Experimental Psychology: General*, 134 (1), 23-37.