

Behavioral Models of Managerial Decision-Making

Avi Goldfarb, Rotman School of Management, University of Toronto

Teck-Hua Ho, Haas School of Business, University of California, Berkeley

Wilfred Amaldoss, Fuqua School of Business, Duke University

Alexander L. Brown, Department of Economics, Texas A&M University

Yan Chen, School of Information, University of Michigan

Tony Haitao Cui, Carlson School of Management, University of Minnesota

Alberto Galasso, Rotman School of Management, University of Toronto

Tanjim Hossain, Rotman School of Management, University of Toronto

Ming Hsu, Haas School of Business, University of California, Berkeley

Noah Lim, Wisconsin School of Business, University of Wisconsin-Madison

Mo Xiao, Eller College of Management, University of Arizona

Botao Yang, Marshall School of Business, University of Southern California

September 2011

Abstract

This paper reviews the literature that applies behavioral economic models to managerial decisions. It organizes the literature into research that focuses on alternative utility functions and research that focuses on non-equilibrium models. Generally, behavioral models have seen less application to manager decisions than to consumer decisions and therefore there are many opportunities to develop new theoretical models, new laboratory experiments, and new field applications. The application of these models to field data is particularly underdeveloped.

Keywords: behavioral economics, managerial decision-making

1. Introduction

How do managers make choices? The dominant paradigm in empirical and theory work in economics is to assume that manager choices are made by fully rational decision-makers. These models often assume managers seek to maximize the present value of current and future earnings, solve a dynamic optimization problem, and play a Bayesian Nash Equilibrium. An increasing amount of research, however, has documented that these (and other) standard assumptions are often violated. In their place, several formal models of alternative assumptions have been developed and tested.

The general literature on bounded rationality and alternative utility functions has been reviewed by Ho, Lim, and Camerer (2006) and DellaVigna (2009). It highlights a growing body of work that has emphasized consumer analysis inside and outside the lab. For management implications, this literature has emphasized how rational firms can exploit the bounded rationality of consumers. For example, DellaVigna and Malmendier (2006) design optimal contracts in the presence of time-inconsistent consumer preferences, and Grubb (2009) discusses exploiting consumer overconfidence. When relevant, we briefly highlight papers on consumer biases that are closely related to the discussion. In discussing the literature on consumer biases, we are not comprehensive and we refer readers to the other reviews mentioned above. There is also a rich literature in behavioral finance that discusses biases in investing behavior. Again, while we touch on a handful of these papers when they are directly relevant to the discussion, we are not comprehensive and we refer readers to other reviews such as Baker, Ruback, and Wurgler (2007). What distinguishes our review is the focus on biases in managerial decision-making.

We organize the paper around two broad themes: alternative utility functions and non-equilibrium models. Alternative utility functions include reference dependence, hyperbolic discounting, and social preferences such as fairness and altruism. We emphasize that social concerns can be particularly important for understanding behavior by managers. Non-equilibrium models involve boundedly rational approaches to manager choices. These boundedly rational approaches include limited thinking steps in games and inability to make optimal decisions. For each of these behavioral regularities under these two broad themes, we discuss the existing theory and laboratory evidence and the small but growing literature on

evidence from the field. Because the application of formal behavioral models to managerial decisions is a relatively new area of inquiry, we also discuss several promising avenues for future research. Table 1 overviews the state of existing research on non-rational managerial behavior according to our classification of the various behavior regularities.

Table 1: State of Existing Research by Key Behavioral Regularity

Behavioral Regularity		Theory	Experimental	Field	Sparse
Alternative Utility Functions	Loss aversion	x	x	x	
	Fairness and other social preferences	x	x		
	Social comparison	x	x		
	Context effects				x
	Projection bias				x
	Time inconsistency				x
	Herding				x
	Inattention				x
Self control				x	
Non-equilibrium Models	Strategic thinking in games	x	x	x	
	Competition neglect	x	x	x	
	Overconfidence	x	x	x	
	Learning and equilibration	x	x		
	Biological basis	x	x		
	Dynamic decision-making				x
	Mix rational and non-rational players				x

2. Alternative Utility Functions

There is a rich literature on alternative utility functions by consumers that discusses loss aversion (e.g. Heidhues and Koszegi 2008), context effects (e.g. Orhun 2009), projection bias (e.g. Simonsohn 2010b; Conlin, O'Donoghue, and Vogelsang 2007), time inconsistency (DellaVigna and Malmendier 2006; Oster and Scott Morton 2005), fairness (e.g. Anderson and Simester 2010), herding (Simonsohn and Ariely 2008), and other ways in which consumers may not conform to a standard utility model. Much of this literature explores how rational managers can best-respond to consumers with non-standard utility functions; however, there is surprisingly little about non-standard utility functions by managers.

This section explores such models where managers may display behavior that suggests a non-standard utility function. Most of the existing work in this area has emphasized fairness and other social preferences and therefore we emphasize these areas and discuss the potential for future work related to other non-standard utility functions such as self control, inattention, and reference dependence.

2.1 Fairness

Research in behavioral economics and psychology has suggested that the standard utility framework may be systematically biased in particular settings. For example, research has shown that people use reference points in making decisions (e.g. Lim and Ho, 2007; Ho and Zhang 2008; Farber 2008) and that people tend to discount the near future disproportionately relative to the distant future (e.g. Thaler 1981; DellaVigna and Malmendier 2006).

Of particular importance for managerial decision making is that some people do not simply maximize their own payoff; they also care about others' payoffs. There is a long literature documenting the importance of fairness and equity (Güth, Schmittberger, and Schwarze 1982; Kahneman, Knetsch, and Thaler 1986; Loewenstein, Thompson, and Bazerman 1989; Camerer 2003; Anderson and Simester 2010). It has been shown that these social concerns in the form of fairness/equity have a significant impact on decision making. Specifically, Rabin (1993) formalizes the notion of procedural fairness to explain the fact that people are nice to those who are nice to them and punish those who are not. Fehr and Schmidt (1999) model distributive fairness by incorporating people's aversion to inequality. Ho and Su (2009) develop peer-induced fairness to explain the phenomenon that people may look to their peers as a reference to evaluate their own payoffs.

Fairness and equity are particularly important in understanding manager decisions because such social preferences can affect optimal firm strategies significantly. Cui et al. (2007), for instance, study how fairness may affect channel coordination. They find that a manufacturer in a dyadic channel can use a simple wholesale price above her marginal cost to maximize both channel profit and channel utility when the retailer has strong concerns of fairness. Thus, a linear pricing contract is able to coordinate the channel. Studies by Demirag et al. (2009) and Katok and Pavlov (2009) suggest that the channel can also

be coordinated with linear pricing contract when demand functions are non-linear or when there is information asymmetry between channel members. Cui and Mallucci (2010) study the consequences of fairness on channel members' decision making in situations involving choices on both investments and prices. They analytically show that firms are more willing to invest in a channel in order to affect the formation of the equitable payoff, when firms are fair-minded and make pricing decisions to assure their pecuniary payoffs reflect their respective contributions to the channel. Cui and Mallucci embed the fairness utility function of channel members in a quantal response model, in which firms may make mistaken decisions due to bounded rationality. The estimation results confirm the existence of significant fairness concerns among channel members and suggest that the principle of strict egalitarianism fits the data better than the principle of liberal egalitarianism and the principle of libertarianism. Chen and Cui (2010) incorporate consumers' concerns of peer-induced price fairness into a model of price competition to explain the extensive adoption of uniform pricing by firms. They find that uniform pricing induced by consumer concern for fairness, and emerged in equilibrium, can help mitigate price competition and hence increase firms' profits. In addition, an individual firm may not have incentive to unilaterally mitigate consumer concern for price fairness to its own branded variants, which suggests the long-run sustainability of the uniform pricing strategy. As a result, fairness concerns provide a natural mechanism for firms to commit to uniform pricing which enhances their profits.

The decisive impact of fairness concerns on managers' decision making and firms' performance is also investigated extensively in other empirical and managerial studies. For example, Fehr, Alexander, and Schmidt (2007) find that, contrary to the traditional theory, unenforceable bonus contracts outperform explicit incentive contracts in a moral-hazard context when there are fair-minded players. An empirical study by Scheer, Kumar, and Steenkamp (2003) suggests that Dutch dealers react negatively to both advantageous and disadvantageous inequality they experience. Through a series of large-scale field experiments, Anderson and Simester (2004, 2008) find that consumer fairness concerns may have a significant impact on a firm's profitability.

Although there has been some nascent research studying how social preferences including fairness and equity affect managers' decision making, there are many open questions that need further attention. For instance, will concerns of procedural fairness affect the interactions of channel members differently than the concerns of distributive fairness? Does playing fairly always benefit a firm and the society, or may it be detrimental to the society and the firm itself? In addition, procedural fairness and distributive fairness are currently the two most studied theories of fairness (Rabin 1993; Fehr and Schmidt 1999). It is worth investigating how a recently developed theory of fairness, peer-induced fairness, may affect managers' decision making in certain contexts (Ho and Su 2009). Relevant research questions include: how should a manufacturer design and propose contracts to multiple retailers who may not be competing in the same market but have concerns of peer-induced fairness; and how a principle designs a compensation plan to effectively motivate agents who may care about not only their own absolute pecuniary payoff but how much their compensation is relative to other agents' payoffs (Cui, Raju, and Shi 2010; Lim 2010).

In addition to studying how modifying manager's utility formulation has a direct effect on their decisions, there is an opportunity to explore how consumer notions of fairness and the tendency to engage in social comparison have an indirect effect on managerial decision making. For example, Amaldoss and Jain (2005a; 2005b) show consumer's desire for uniqueness and conformism affect firm's prices. Amaldoss and Jain (2008; 2010) explore how the notion of reference group can affect a firm's product line decisions. Specifically, they show that the presence of reference group effects can motivate firms to add costly features which provide limited or no functional benefit to consumers. Furthermore, reference group effects can induce product proliferation. In certain other situations, firms may offer a limited edition. While they have examined the impact of consumer reference groups on managerial decisions, future research can examine how the notion of reference groups among managers affects decisions.

2.2 Other Social Preferences

There is also experimental evidence in sales compensation research that suggests that utility functions that incorporate social preferences can explain behavior better. Lim, Chen, and Ahearne (2010)

examine whether sales managers should incorporate a group-based component into the compensation plans of salespeople. Conventional wisdom based on economic theory argues against the use of group incentives because salespeople, who are assumed to make effort decisions that maximize their own pecuniary payoffs, will not internalize the benefits that accrue to other members in the group (so that there is “free-riding”). In other words, standard economic models predict that a commission pay plan based solely on a salesperson’s own sales will induce greater effort than a plan that consists of a commission rate that is based on both the salesperson’s own sales and the sales of other salespeople. However, if salespeople have social preferences, salespeople may put in greater effort when there is a group incentive.

This paper tests these competing predictions by conducting a field experiment in an undergraduate sales program with student subjects who were responsible for fund-raising activities. The results show that sales were higher when there was some group-based compensation, thereby suggesting the existence of social preferences. Next, this paper implements an incentive-aligned laboratory experiment to examine the environmental conditions under which group-based incentives might be optimal. Specifically, they manipulate whether the subjects know who the members of their team. They find that when subjects do not know the members of their team, pure individual incentive-based compensation works best. In contrast, when the members do know their team members, effort is highest under a hybrid of individual and team incentives. This shows that, consistent with the existence of social preferences, incorporating a component of group-based incentives (without increasing the component of group pay such that it becomes a pure revenue sharing plan) can be optimal, but only in environments where social preferences are more likely to be “operational”. At a broader level, this research suggests that whether non-standard utility functions (such as those that incorporate social preferences) predict actual behavior better can depend critically on the context (Simonson and Tversky 1993; Lim 2010) – a fruitful avenue of research would be to isolate the environmental and institutional factors that can affect the existence and strength of social preferences.

2.3 Recommendations for Future Work

As discussed above, there is a growing body of laboratory-based evidence on the impact of social preferences on managerial behavior. Perhaps surprisingly, this literature has not explored how social preferences affect manager behavior in the field, with the notable exception of Zhang and Zhu's (2010) paper documenting the role of social recognition for contributions to Wikipedia. Therefore, we view this as a rich opportunity for future research.

Also surprisingly, there is little work on how other alternative utility functions (besides social preferences) might affect managerial behavior. The literature has documented that consumers are likely to be subject to behavioral biases in many settings and firms are run by human beings who may be subject to similar biases. As any other individual, managers are affected by self-control problems (DellaVigna and Malmendier, 2006), inattention (Hossain and Morgan, 2006), context effects (Orhun 2009), and reference dependence (Genesove and Mayer, 2001). While there are notable exceptions such as Ho and Zhang (2008) and Farber (2008) on reference dependence, there are several remaining opportunities for future research projects.

3. Non-equilibrium Models

This section discusses models of out-of-equilibrium behavior and applies them to the decisions of managers. First, we review the literature on strategic thinking in games and the relaxation of the assumption of consistency of beliefs. Next, we discuss learning and when we should expect markets to reach equilibrium. Third, we discuss managerial overconfidence and how it can lead to non-equilibrium outcomes. We classify overconfidence under non-equilibrium models rather than non-standard utility functions because it is about incorrect beliefs rather than utility. Finally, we discuss areas where existing theory and laboratory experiments can be fruitfully applied to understanding managerial decisions.

3.1 Strategic Thinking in Games

Perhaps the most developed area of research that applies formal behavioral models to managerial decision-making is the work on strategic thinking in games. In this subsection we begin by discussing the theory and experimental evidence and then summarize the growing evidence from the field.

3.1.1 Theory and Experimental Evidence of Strategic Thinking in Games

Research has shown that subjects often do not play equilibrium even in extremely simple economic environments (for a comprehensive review, see Camerer 2003). This is so because standard equilibrium models embody three assumptions and these assumptions are often not all empirically valid. The assumptions are: 1) players form beliefs about what others will do, 2) they maximize payoffs based on their beliefs, and 3) players' beliefs of others coincide with others' actual behavior. Several models have been proposed to relax the third assumption because players often have a hard time correctly guessing others' behavior.

To provide a concrete example, we discuss Hossain and Morgan's (2011) model of platform competition using the cognitive hierarchy model proposed by Camerer, Ho, and Chong (2004). Their model builds on prior work by Nagel (1995), Stahl and Wilson (1995) and Ho, Camerer, and Weigelt (1998). Many markets share the following features—agents of two types interact with one another on a platform which may serve to match agents and to facilitate transactions. Examples of platforms in two-sided markets abound in our daily lives and include dating services, stock exchanges, and credit cards. A central question in platform competition revolves around market concentration. In practice, we see that the online dating market includes many platforms with significant market shares while the online auction market in North America has virtually tipped to eBay. Theoretical models explain this heterogeneity in market concentration using the concepts of market size and market impact effects—while an agent benefits from an increase of agents of the opposite type of agents in the platform they choose, she is harmed by an increase of agents of her own type because they compete with her.

Existence of multiple equilibria is an issue in a general class of models of platform competition when agents are completely rational. For example, Ellison and Fudenberg (2003) show that, when agents of a given type are homogeneous in their preferences for the two platforms and can choose to subscribe to only one, almost any breakdown of market shares between the two platforms may arise in equilibrium. However, Hossain and Morgan (2011) show that, when agents follow the cognitive hierarchy model of Camerer, Ho, and Chong (2004), equilibrium analysis of such a model leads to sharp theoretical

predictions. In a market where agents can choose only one of two platforms, there exists a unique equilibrium where the market virtually tips to the risk-dominant platform. This result is consistent with the experimental findings of Hossain, Minor, and Morgan (2011). On the other hand, when players are allowed to choose both platforms, the market virtually tips to the Pareto-dominant platform in the unique equilibrium under the cognitive hierarchy model. On the other hand, when platforms are horizontally differentiated, multiple platforms obtain positive market share in the unique equilibrium.

In the model of Hossain and Morgan (2011), the platforms do not act strategically. An obvious extension of the model can be to allow platforms to choose their pricing and exclusivity strategies optimally with both rational and boundedly rational managers of platforms. The model can also be extended to allow for agents who have heterogeneous preferences for the two platforms. Experimental results of Hossain, Minor and Morgan (2011) suggest that the main driving force behind coexistence of multiple platforms is horizontal differentiation of platforms. It will be useful to explore how theoretical models with heterogeneous agents fare under the cognitive hierarchy model. Finally, the cognitive hierarchy model can be used in analyzing a variety of other coordination games.

3.1.2 Field Evidence of Strategic Thinking in Games

A broad literature documents that individual managers affect firm behavior and performance. Bertrand and Schoar (2003) track managers as they move across firms and show that by manager fixed effects explain many firm practices. Kaplan, Klebanov, and Sorenson (2008) show that success relates to manager skills. Similarly, Bloom and Van Reenen (2007) show that success relates to managerial style.

A logical next step is to see whether manager ability matters in a strategic (game-theoretic) setting. In such settings, the identification of model primitives usually relies on the imposition of equilibrium assumptions. This means observed play in data could be due to deviations from equilibrium or to unobserved heterogeneity in the payoffs to different strategy choices. Therefore, it is challenging to identify deviations from these equilibrium assumptions in a non-laboratory setting where the payoffs are typically unobserved.

One potential empirical approach is to look for a field setting where payoffs can be inferred. For example, Hortacsu and Puller (2008) examine bidding behavior in electricity auctions where marginal cost data exist. They use this data to generate a Nash equilibrium prediction and show that larger and more experienced firms bid closer to the Nash equilibrium prediction.

An alternative empirical approach is to apply a model that has been well-tested in laboratory settings and then look for evidence of internal and external validity. Goldfarb and Yang (2009) extend Camerer, Ho, and Chong's (2004) "Cognitive Hierarchy" model to an empirical setting. As such, Goldfarb and Yang develop a new type of structural model based on an alternative solution concept. They apply this model to a technology adoption game played by internet service providers. They show that more strategic firms (as measured in 1997) were more likely to survive the dot-com crash. The model is used to simulate the impact of strategic thinking on the market and the simulations show that strategic thinking by retailers slows the diffusion of the new technology.

Goldfarb and Xiao (2011) provide further insight into the incidence of heterogeneity in management ability in a high stakes game. They examine the entry decisions of competitors into local telephone markets following the 1996 Telecom Act. They show patterns in the data that suggest something other than Bayesian Nash: more experienced and better educated managers tend to enter less competitive markets. This pattern motivates estimation of a cognitive hierarchy model and the results show that manager characteristics are key determinants of measured managerial ability. Furthermore, measured ability rises sharply after the industry shakeout.

Understanding the ability to think through iterated steps may inform many key settings in industrial organization including car sales, monopoly regulation, and information disclosure. For example, there is an important result in the information disclosure literature that higher quality firms have incentives to reveal their true quality to distinguish themselves from lower quality competitors. Grossman (1981) and Milgrom (1981) show that non-disclosure is completely unraveling because consumers correctly infer that only the lowest quality firms wouldn't reveal. Brown, Camerer, and Lovo (2011a) argue that unraveling might break down if either consumers or firms do not fully engage in the iterated process of quality

inference and best response. They examine the decision of movie studios to show their movies to critics prior to opening and document that movies that are not shown to critics (“cold-opened movies”) have 15-25% higher revenue than comparable movies that are critically reviewed. In their companion paper, Brown, Camerer, and Lovallo (2011b) impose a laboratory-tested behavioral model, the cognitive hierarchy model, on the data; their estimates support the idea that consumers do very few iterations of strategic thinking. In contrast, movie studios appear to be best-responding to consumers that are playing a fully Bayesian-Nash strategy. That is, studios act as if they believe consumers complete the fully iterative process. It appears that movie studios would therefore benefit by cold opening more movies. Given the widespread application of iterating thinking in information disclosure settings, we believe there is great potential in applying Brown, Camerer, and Lovallo’s (2011a, 2011b) framework to other settings.

Another aspect of these alternative solution concept models is competition neglect. In particular, firms may not fully account for competitors’ actions. There is an increasing body of evidence that competition neglect is not unusual. For example, Simonsohn (2010a) studies seller behavior on eBay. He finds that sellers overcrowd auction closing times into peak bidding hours. Furthermore, he provides evidence that this is not driven by inattention: it is the sellers that specifically chose the closing time and the professional sellers that crowd most into the peak times. Moore and Cain (2007) also provide evidence that people do not fully consider competitors in games and that such competition neglect may explain prior results that suggest overconfidence.

More generally, Camerer and Malmendier (2007) note that bounded rationality is likely to be most important in manager decisions when (i) decisions are infrequent and do not deliver clear feedback, (ii) the manager does not specialize in that type of decision, and (iii) managers are protected from market pressure and competition. One main challenge in future work in this area is to find such settings that also offer rich enough data for empirical application. To meet this challenge, researchers need to be creative.

3.2 Learning

Although the above models extend the scope of traditional economic models of managerial decision-making, they are agnostic on the process of how markets might reach equilibrium. Likewise,

empirical applications of these models also assume that equilibrium has either already been reached, or if there are multiple equilibria, equilibrium selection has already taken place. Data from laboratory experiments, however, show a more nuanced picture. In most experiments, at least for the first few rounds, observed behavior is far from (any) equilibrium (Roth and Erev 1995; Camerer 2003). If equilibration occurs, the equilibrium is reached through some iterative adaptive process, which we will follow convention and refer to as *learning*. The growing theory and experimental literature on learning in games has examined when and how non-standard strategies in games persist. Focusing on different methods of learning, this literature has shed light on some long-standing issues in equilibrium models of non-cooperative games.

3.2.1 Learning and Equilibrium Selection

As discussed earlier, learning has been a particularly fruitful approach in predicting which equilibrium will be selected in games with multiple Nash equilibria. In the class of potential games, for example, several belief-based learning models converge to the potential maximizing equilibrium (Monderer and Shapley 1996). Combining the classical belief-based learning approach with group-contingent social preferences, Chen and Chen (2011) propose a theoretical model of social identity and derive conditions under which social identity changes equilibrium selection by changing the potential function. Specifically, when people feel more altruistic towards their ingroup members (Chen and Li 2009), they will exert higher equilibrium effort in the sense of first-order stochastic dominance in a minimum effort game. To test the ability of this model to predict behavior in the laboratory, they design an experiment using the minimum effort games of Van Huyck, Battalio, and Beil (1990) with parameter configurations where learning would result in convergence to the inefficient, low-effort equilibrium absent of group identity (Goeree and Holt 2005). In their near-minimal treatments, they show that while matching subjects with ingroup or outgroup members has some effect on the effort levels chosen, they are not statistically distinguishable from the control, where no groups are induced. On the other hand, when they enhance the groups by allowing them to communicate with group members in solving a simple task before playing the minimum effort game, they find that matching subjects with ingroup members has a

statistically significant positive effect on the subjects' provided efforts. Thus, consistent with the model, ingroup matching significantly increases coordination and efficiency. This paper contributes to the theoretical foundations of social identity by demonstrating that, by using a simple group-contingent social preference model, it is possible to endogenize the exogenous norms in the original Akerlof and Kranton social identity models (2000, 2005) and to reconcile the theory with experimental findings in a number of coordination games.

The results also have practical implications for organization design. Organizations are more frequently encountering the issue of integrating a diverse workforce, and motivating members coming different backgrounds to work towards a common goal. Chen and Chen (2011) demonstrate that creating a deep sense of common identity can motivate people to exert more effort to reach a more efficient outcome. A successful application of this idea comes from Kiva (<http://www.kiva.org/>), a person-to-person microfinance lending site, which organizes loans to entrepreneurs around the globe. In August 2008, Kiva launched its lending teams program which organizes lenders into identity-based teams. Any lender can join a team based on her school, religion, geographic location, sports, or other group affiliations. The lending teams program significantly increases the amount of funds raised.

There are several directions for future research. A possible next step in this line of research would be to extend this result to other coordination games. For example, the provision point mechanism is a simple public goods mechanism but with multiple equilibria. To select the most efficient equilibrium, one could use the same induced identity method used in this study. Chen and Chen's (2011) model predicts that successful coordination to higher levels of public goods can be achieved systematically even with a weak method of increasing other-regarding preferences. Another direction is to evaluate the effects of identity-based teams outside the lab through field studies in fundraising or online communities.

3.2.2 Biological Basis of Learning

In numerous laboratory experiments, the general finding is that, for nontrivial games, players gradually reach equilibrium over time through some process of adaptation, typically referred to as learning (Camerer 2003, Chapter 6). A number of models of learning in games have been proposed, particularly reinforcement and belief-based models, as well as hybrid models such as experienced weighted attraction (EWA) (see Camerer and Ho 1999 and Ho, Camerer, and Chong 2007).

Hsu, Mathewson, and Zhu (2010) build upon this literature by studying the neural mechanisms underlying strategic learning using functional magnetic resonance imaging (fMRI). This is a potentially fruitful endeavor as the neural mechanisms of learning have been revolutionized from the discovery of a class of neurons, namely the dopamine neurons, that appear to implement the temporal difference (TD) form of reinforcement learning. Derived from behavioral psychology and machine learning literatures, at the core of TD learning is the computation and updating of a reward prediction error (RPE), whereby organisms (in this case players) learn from the discrepancy between what is expected to happen and what actually happens (Sutton and Barto 1981). This includes a number of recent papers implicating such dopaminergic regions in decisions under risk and uncertainty (e.g. Fiorillo et al 2003). More importantly from the perspective of strategic learning, this literature offers a set of biologically plausible formal models of behaviour that has the potential to directly connect behavioral observations of learning dynamics in games on the one hand (Roth and Erev 1995; Camerer 2003), and the neural observations of the brain dynamics on the other.

Specifically, Hsu et al. (2010) used an asymmetric version of the patent race game, first studied experimentally in Rapoport and Amaldoss (2000), to search for regions of the brain involved in computation of expected payoffs and prediction errors that can be used to guide behavior. The large strategy space of this game improves the recovery of key parameters in learning models relative to smaller games typically used in such studies (Wilcox 2006).

Reinforcement and belief-based models were implemented in the manner consistent with temporal difference models (Sutton and Barto 1981) and fitted to the neural data. The results show evidence of

reinforcement and belief-based learning signals in the manner predicted by EWA learning. Somewhat surprisingly, these distinct signals are represented in both overlapping and distinct brain regions. There are a number of potential extensions to this study. First, despite the empirical success of the aforementioned models, the model fits are far from perfect. Belief-learning in particular, relies on strong functional form assumptions regarding the construction of beliefs, which has long been assumed in standard models to be unobservable. Recent studies using proper scoring rules to elicit beliefs, however, found substantial improvements in fit (Nyarko and Schotter, 2002). As shown out by Rutstrom and Wilcox (1996), however, the act of elicitation itself may well bias the learning dynamics. In contrast, direct extraction of beliefs from neural activity presents the possibility of an unbiased measurement of beliefs. More generally, neural data can discipline behavioral models of learning by providing direct data regarding the causal mechanisms behind decision-making and learning.

3.3 Overconfidence

Standard microeconomic models assume that players are on average correct about the distribution of states of the world. However, a large number of studies in applied psychology document that individuals tend to overestimate their ability thus violating this assumption (Svenson 1981; Cooper et al. 1988). We categorize this bias under “non-equilibrium models” because it is about incorrect beliefs in managerial decisions rather than a different underlying utility function.

The experimental literature on overconfidence provides several indications that CEOs are particularly likely to be affected by this bias (Camerer and Lovallo 1999; Camerer and Malmendier 2007). In a series of papers, Malmendier and Tate (e.g. 2005, 2008) use field data to show that CEO overconfidence has a significant impact on a variety of important strategic decisions. They construct a measure of overconfidence using the personal investments of CEOs and classify as overconfident those CEOs that do not exercise their stock options, even when the underlying stock price is very high. As discussed earlier, Moore and Cain (2007) argue that some of the results on overconfidence with respect to competition can be explained by competition neglect.

Galasso and Simcoe (2011) exploit Malmendier and Tate's measure in order to examine the impact of CEO overconfidence on innovation. They find that the arrival of an overconfident CEO is correlated with an increase in corporate innovation measured either as number of patents, citations received, or R&D expenditure. This correlation is robust to controlling for firm fixed effects and for a number of time-varying firm and CEO characteristics. They also find that CEO overconfidence has a greater impact on innovation under intense product market competition. Overall, their results suggest that overconfident CEOs are more likely to take their firms in a new technology direction.

There are likely many other ways in which managerial overconfidence might impact corporate performance. Two papers that make progress in this direction are Galasso (2010) on overconfidence and patent litigation and Watanabe (2010) on overconfidence and medical malpractice litigation. Still, given the current evidence on managerial overconfidence, further work on the consequences of such overconfidence is needed.

3.4 Underexplored Areas: Dynamic Decision-making and Mixing Rational and Non-rational firms

There are two areas that have been relatively unexplored that we think are particularly ripe for further analysis: Dynamic decision-making and the mixing of rational and non-rational firms.

So far the literature has been focused on static games. Even the learning literature has largely focused on learning to play static games. Looking forward, it would be useful to think about biases in dynamic games. In dynamic games, there is a further complication because people have been shown to be boundedly rational in how forward-looking they are (e.g. Loewenstein and Prelec 1992; Soman et al 2005). It is possible that many firms may also be boundedly rational in this way. For example, Che, Sudhir, and Seetharaman (2007) show that firms appear to only look ahead one period in setting prices for both breakfast cereals and ketchup. Accommodating this possibility would then combine non-standard (myopic) utility models with non-standard solution methods.

Another area of potential impact is competition between rational and non-rational firms. De Long et al (1990) showed that in financial markets, irrational traders can have a large impact on asset prices and such irrational traders can earn higher returns than rational investors. This spawned a very large literature

in behavioral finance on the role of naïve and non-rational investors in financial markets. There is also a fairly rich theory literature going back to at least the 1980s that consider heterogeneous abilities by firm-like players in games including Stahl (1993) and Haltiwanger and Waldman (1985). Importantly, these models emphasize that either rational or naïve players can be dominant, depending on the specific nature of the game. We believe there is potential for important contributions in understanding the evolution of industry structure that accommodates the possibility of both rational and non-rational firms, perhaps in a dynamic setting. Aguirregabiria and Magesan (2011) provide a tool for such analysis.

4. Conclusion

In conclusion, applying behavioral biases to managers is an important and growing area of study. We have identified several particularly promising areas for future work, which we summarize below:

- a) Theory and lab research on the impact of fairness on a broader range of managerial decisions, including welfare analysis.
- b) Theory and lab research on social preferences, particularly in coordination games.
- c) Research on the effect of social preferences on managerial behavior using data from the field in order to understand the broader applicability of the laboratory-generated results.
- d) Research on how alternative utility functions, aside from social preferences, might affect managerial behavior. Examples include self control, context effects, inattention, and reference dependence.
- e) Research that applies the computational and equilibrium selection advantages of alternative solution concepts such as cognitive hierarchy to help solve coordination games, in theory and in structural empirical work.
- f) Field work that examines the conditions under which we observe bounded rationality by managers in games, including disclosure games, entry games, technology adoption games, and others.
- g) Theory and lab work on the biological basis of economic behavior, which can in turn help discipline existing theory and inspire new models.
- h) Field work on the role of overconfidence in manager decisions and firm performance.

- i) Field work on the degree to which managers can solve dynamic games.
- j) Theory and (especially) field work on the consequences of mixing rational and non-rational firms.

While there has been substantial progress recently, there is much more work to be done to understand when and how behavioral biases apply to managerial decision-making.

References

- Aguirregabiria, V., and A. Magesan (2011). Identification and Estimation of Dynamic Games when Players' Beliefs are Not in Equilibrium. Working paper, University of Toronto.
- Akerlof, G A., R E. Kranton (2000). Economics and Identity, *Quarterly Journal of Economics* 115(3): 715–753.
- Akerlof, G A., R E. Kranton (2005). Identity and the Economics of Organizations. *Journal of Economic Perspectives*, 19(1): 9–32.
- Amaldoss, W., S. Jain (2005a). Conspicuous Consumption and Sophisticated Thinking. *Management Science* 51(10): 1449-1466.
- Amaldoss, W., S. Jain (2005b). Pricing of Conspicuous Goods: A Competitive Analysis of Social Effects. *Journal of Marketing Research*, 42: 30-42.
- Amaldoss, W., S. Jain (2010). Reference Groups and Product Line Decisions: An Experimental Investigation of Limited Editions and Product Proliferation, *Management Science* 56(4): 621-644.
- Amaldoss, W., S. Jain (2008). Trading Up: A Strategic Analysis of Reference Group Effects. *Marketing Science* 27(5): 932-942.
- Anderson, E., D. Simester (2004). Long-run effects of promotion depth on new versus established customers: Three field studies. *Marketing Science* 23(1): 4-20.
- Anderson, E., D. Simester (2008). Does demand fall when customers perceive that prices are unfair? The case of premium pricing for large sizes. *Marketing Science* 27(3): 492-500.
- Anderson, E., D. Simester (2010). Price stickiness and customer antagonism. *Quarterly Journal of Economics* 125(2): 729-765.
- Baker, M., R. Ruback, and J. Wurgler (2007), “Behavioral Corporate Finance: A Survey,” in Espen Eckbo, ed. *Handbook in Corporate Finance: Empirical Corporate Finance*, Elsevier/North Holland, ch. 4: 145-188.
- Bertrand, M. and A. Schoar (2003). Managing with Style: The Effect of Managers on Firm Policies, *Quarterly Journal of Economics* 118: 1169-1208.

- Bloom, N. and J. Van Reenen (2007), "Measuring and Explaining Management Practices Across Firms and Countries", *Quarterly Journal of Economics*, 122(4): 1351-1408.
- Brown, A. L., Camerer, C. F. and Lovallo, D. (2011a). To Review or Not to Review? Limited Strategic Thinking at the Movie Box Office. Working paper, Texas A&M University.
- Brown, A. L., Camerer, C. F. and Lovallo, D. (2011b). Estimating Structural Models of Equilibrium and Cognitive Hierarchy Thinking in the Field: The case of withheld movie reviews. Working paper, Texas A&M University.
- Camerer, C. and Ho, T-H. (1999). Experience-Weighted Attraction Learning in Normal Form Games," *Econometrica*, 67: 837-874
- Camerer, C. F. (2003). Behavioral game theory: Experiments in strategic interaction. Princeton University Press.
- Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004). A Cognitive Hierarchy Model of Games. *Quarterly Journal of Economics* 119(3): 861-898.
- Camerer, C. and D. Lovallo (1999). Overconfidence and Excess Entry: An Experimental Approach. *American Economic Review* 89(1): 306-318.
- Camerer, C. and U. Malmendier (2007). Behavioral Economics and its Applications, In *Behavioral Economics of Organizations*. Princeton University Press.
- Che, H., K. Sudhir, P.B. Seetharaman (2007). Bounded Rationality in Pricing Under State-Dependent Demand: Do Firms Look Ahead? How Far Ahead? *Journal of Marketing Research* 44(3): 434-49.
- Chen, R. and Y. Chen (2010). The Potential of Social Identity for Equilibrium Selection. Forthcoming in *American Economic Review*, October 2011.
- Chen, Y., T.H. Cui (2010). The benefit of uniform price for branded variants. Working Paper.
- Chen, Y., S. Li (2009). Group Identity and Social Preferences. *American Economic Review* 99(1): 431–57.
- Conlin, M., T. O'Donoghue, and T. Vogelsang (2007). Projection Bias in Catalog Orders. *American Economic Review* 97(4): 1217-1249.
- Cooper, A. C., C. A. Woo, W. Dunkelberg (1988). Entrepreneurs perceived chances for success. *Journal of Business Venturing* 3: 97–108.
- Cui, T.H., P. Mallucci (2010). Quantal response equilibrium in a fair channel – an experimental investigation. Working Paper.
- Cui, T.H., J.S. Raju, M. Shi (2010). S-shaped Incentive Schemes and Pay Caps. Working Paper.

- Cui, T.H., J.S. Raju, Z.J. Zhang (2007). Fairness and channel coordination. *Management Science* 53(8): 1303-1314.
- DellaVigna S. (2009). Psychology and Economics: Evidence from the Field, *Journal of Economic Literature* 47: 315-372.
- DellaVigna, S. and U. Malmendier (2006). Paying Not to Go to the Gym. *American Economic Review* 96(3): 694-719.
- DeLong, J.B. A. Shleifer, L.H. Summers, and R.J. Waldmann (1990). Noise Trader Risk in Financial Markets. *Journal of Political Economy* 98(4): 703-738.
- Demirag, O. C., Y. F. Chen, J. Li (2009). Channel coordination under fairness concerns and nonlinear demand. Working Paper.
- Ellison, G, and D Fudenberg (2003): Knife-Edge or Plateau: When Do Market Models Tip?, *Quarterly Journal of Economics*, 118(4): 1249-1278.
- Farber, H.S. (2008). Reference-dependent preferences and labor supply: The case of New York City taxi drivers. *American Economic Review* 98(3): 1069-1082.
- Fehr, E., K. Alexander, K.M. Schmidt (2007). Fairness and contract design. *Econometrica* 75(1): 121-54.
- Fehr, E., K.M. Schmidt (1999). A theory of fairness, competition and co-operation. *Quarterly Journal of Economics* 114(3): 817-868.
- Fiorillo, C, P Tober and W Schultz (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299: 1898-1902.
- Galasso A. (2011). Broad Cross-License Negotiations, working paper, University of Toronto.
- Galasso A. and T. Simcoe (2011). CEO overconfidence and innovation, *Management Science* 57: 1469-1484.
- Genesove, D, and C Mayer (2001). Loss Aversion and Seller Behavior: Evidence from the Housing Market. *Quarterly Journal of Economics*, 116(4): 1233-60.
- Goeree, J K. and C A. Holt (2005). An Experimental Study of Costly Coordination,” *Games and Economic Behavior* 51(2): 349-364.
- Goldfarb, A., M. Xiao (2011). Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets. Forthcoming, *American Economic Review*.

- Goldfarb, A., B. Yang (2009). Are all managers created equal? *Journal of Marketing Research* 46(5): 612-622.
- Grossman, S. J. (1981). The Informational Role of Warranties and Private Disclosure About Product Quality. *Journal of Law and Economics* 24: 461-483.
- Grubb, M. (2009). Selling to Overconfident Consumers. *American Economic Review* 99(5): 1770-1807.
- Güth, W., R. Schmittberger, B. Schwarze (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization* 3(4): 367-388.
- Haltiwanger, J. and M. Waldman (1985). Rational Expectations and the Limits to Rationality: An Analysis of Heterogeneity. *American Economic Review* 75(3): 326-340.
- Heidhues, P. and B. Koszegi (2008). Competition and Price Variation when Consumers are Loss Averse, *American Economic Review* 98: 1245-1268.
- Ho, T-H., Camerer, C., and Weigelt, K. (1998). Iterated Dominance and Iterated Best Response in Experimental P-Beauty Contests, *American Economic Review*, 88: 947-969.
- Ho, T-H, N Lim, C. Camerer (2006). Modeling the psychology of consumer and firm behavior with behavioral economics. *Journal of Marketing Research* 43(3): 307-331.
- Ho, T-H., Camerer, C. and Chong, J-K. (2007). Self-tuning Experience-Weighted Attraction Learning in Games, *Journal of Economic Theory* 133(1): 177-198.
- Ho, T-H, X Su (2009). Peer-induced fairness in games. *American Economic Review*, 99(5): 2022-2049.
- Ho, T-H. and Zhang, J-J. (2008). Designing Price Contracts for Boundedly Rational Customers: Does the Framing of the Fixed Fee Matter? *Management Science*, 54(4): 686-700.
- Hortacsu, A. and S. Puller (2008). Understanding Strategic Models of Bidding in Deregulated Electricity Markets: A Case Study of ERCOT, *RAND Journal of Economics* 39: 86-114.
- Hossain, T., D. Minor, and J. Morgan (2011). Competing Matchmakers: An Experimental Analysis,” Forthcoming, *Management Science*.
- Hossain T. and J. Morgan (2006). ...Plus Shipping and Handling: Revenue (Non)-Equivalence in Field Experiments on eBay. *Advances in Economic Analysis & Policy*, 6(2).
- Hossain, T., and J. Morgan (2011). When Do Markets Tip? A Cognitive Hierarchy Approach,” Working Paper, University of Toronto.
- Hsu, M., K. Mathewson, and L. Zhu (2010). Neural Computations underlying Strategic Learning, Manuscript, University of California, Berkeley.

- Kahneman, D, J L. Knetsch, R. Thaler (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review*. 76(4): 728-741.
- Kaplan, S. N., M. Klebanov, and M. Sorenson (2008), "Which CEO characteristics and abilities matter?" NBER working paper #14195.
- Katok, E., V. Pavlov (2009). Fairness and coordination failures in supply chain contracts. Working Paper.
- Lim, N. and Ho T-H. (2007). Designing Price Contracts for Boundedly Rational Customers: Does the Number of Blocks Matter? *Marketing Science*, 26(3): 312-326.
- Lim, N. (2010). Social Loss Aversion and Optimal Contest Design. *Journal of Marketing Research*, forthcoming.
- Lim, N, H Chen and M Ahearne (2010), "Do Group Incentives Work? Evidence from the Field and Lab," Marketing Department Working Paper, University of Wisconsin-Madison.
- Loewenstein, George, and Drazen Prelec (1992). Anomalies in Intertemporal Choice: Evidence and an Interpretation. *Quarterly Journal of Economics* 107: 573-98.
- Loewenstein, G. E., L. Thompson, M. H. Bazerman (1989). Social utility and decision making in interpersonal contexts. *Journal of Personality and Social Psychology* 57(3): 426-441.
- Malmendier, U. and G. Tate (2005). CEO Overconfidence and Corporate Investment. *Journal of Finance* 60(6): 2661-2700.
- Malmendier, U. and G. Tate (2008). Who makes acquisitions? CEO overconfidence and the market's reaction. *Journal of Financial Economics* 89(1): 20-43.
- Milgrom, P. R. (1981). "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics* 12(2): 380-391.
- Monderer, D and L S. Shapley (1996). Potential Games. *Games and Economic Behavior*, 14: 124–143.
- Moore, Don A., and Daylian M. Cain (2007). Overconfidence and underconfidence: When and why people underestimate (and overestimate) the competition. *Organizational Behavior and Human Decision Processes* 103: 197-213.
- Nagel, R. (1995): "Unravelling in Guessing Games: An Experimental Study," *American Economic Review*, 85(5): 1313-1326.
- Nyarko, Y. and A. Schotter (2002). An Experimental Study of Belief Learning Using Elicited Beliefs. *Econometrica* 70(3): 971–1005.
- Orhun, A. Yesim (2009). Optimal Product Line Design When Consumers Exhibit Choice Set-Dependent Preferences. *Marketing Science* 28(5): 868-886.

- Oster, S. and F. Scott Morton (2005). Behavioral Biases Meet the Market: The Case of Magazine Subscription Prices. *Advances in Economic Analysis and Policy*, 5(1): 1.
- Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review* 83(5): 1281-1302.
- Rapoport, A. and W. Amaldoss (2000). Mixed strategies and iterative elimination of strongly dominated strategies: An experimental investigation of states of knowledge. *Journal of Economic Behavior & Organization* 42: 483–521.
- Roth, A. and I. Erev (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior* 8(1): 164-212.
- Rutstrom and Wilcox (1996). Stated Beliefs Versus Empirical Beliefs: A Methodological Inquiry and Experimental Test, Working Paper.
- Scheer, L. K., N. Kumar, J.-B. E. M. Steenkamp (2003). Reactions to perceived inequity in U.S. and Dutch interorganizational relationships. *Academy of Management Journal* 46(3): 303-316.
- Simonsohn, U. (2010a). eBay's Crowded Evenings: Competition Neglect in Market Entry Decisions. *Management Science* 56(7): 1060-1073.
- Simonsohn, U. (2010b). Weather to Go to College. *Economic Journal*. 120(543): 270-280.
- Simonsohn, U. & Ariely, D. (2008). When Rational Sellers Face Non-Rational Consumers: Evidence from Herding on eBay. *Management Science* 54(9): 1624-1637
- Simonson, I. A. Tversky (1993). Context-Dependent Preferences. *Management Science* 39(10): 1179-89.
- Soman, Dilip, G. Ainslie, S. Frederick, X. Li, J. Lynch, P. Moreau, A. Mitchell, D. Read, A. Sawyer, Y. Trope, K. Wertenbroch, G. Zauberman (2005). The Psychology of Intertemporal Choice. *Marketing Letters*, 16 (3,4): 347-360.
- Stahl, D.O. (1993). Evolution of Smart_n Players. *Games and Economic Behavior* 5: 604-617.
- Stahl, D O., and P Wilson (1995). On Players' Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior* 10(1): 213-254.
- Sutton, R. S. and A. G. Barto (1981). "Toward a modern theory of adaptive networks: expectation and prediction." *Psychol Rev* 88(2): 135-70.
- Svenson, O. (1981). Are we all less risky and more skilful than our fellow drivers? *Acta Psychologica* 47(2): 143-148.

- Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters* 8(3): 201-207.
- Van Huyck, J B., R C. Battalio, and R O. Beil (1990). Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure, *American Economic Review* 80: 234–248.
- Watanabe Y. (2010). Learning and Bargaining in Dispute Resolution: Theory and Evidence from Medical Malpractice Litigation, working paper.
- Wilcox, N. (2006). Theories of Learning in Games and Heterogeneity Bias. *Econometrica*, 74(5): 1271-1292.
- Zhang, X., F. Zhu (2010). Group Size and Incentives to Contribute: A Natural Experiment at Chinese Wikipedia. Forthcoming *American Economic Review*.