Modeling Risky Economic Decision-Making with Bounded Rationality

Kay-Yut Chen and **Tad Hogg**
HP Labs
Palo Alto, CA

Abstract

Models of economic decisions often assume people maximize a preference or *utility* function. While this assumption provides insights into a wide range of economic behavior, in some cases its predictions contradict observed behavior. We illustrate this situation with human-subject experiments of decisions involving risk. We show how the addition of bounded rationality to the economic model can explain the behavior, and we estimate the parameters for our experimental subjects showing they divide into distinct groups with different risk preferences. This combined model illustrates how standard economic models can be extended to a wider range of human decision-making behavior.

Introduction

Human decisions often involve tradeoffs among competing goals and uncertainty in the outcomes. Economics provides a framework to understand a wide variety of such decisions, based on the concept of utility maximization. This model postulates a person has *utility* values associated with possible states of the world, and situations with higher utilities are preferred to those with lower ones. Any monotonic transformation of a given utility function gives the same ordering of preferences, so the utility is neither unique nor comparable between people. Nevertheless, for situations primarily involving exchanges in market contexts, the utilities are often expressed in terms of the monetary value a person would pay for various items.

The utility framework provides key insights into many economic scenarios with reasonable predictions in some contexts. However, this model makes strong assumptions of rationality, each decision-maker is able to evaluate and maximize a known (by the decision-maker) utility function. In many situations, these assumptions do not accurately describe how people make decisions (Camerer 1995). Instead, human behavior in economic contexts can show many effects contrary to predictions of the standard economic framework. Such effects include those arising from social preferences (how people deal with each other), individual bounded rationality (how people make mistakes in decisions) and uncertainty (Kahneman & Tversky 1979). For example, fairness is an important social preference in some decision-making contexts (Guth, Schmittberger, & Schwarze 1982). Another effect is loss aversion, where people tend to weigh the possibility of a loss more heavily than that of a gain in making decisions (Kahneman & Tversky 1979). When the utility is difficult to evaluate or maximize, cognitive limitations or bounded rationality can lead people to make suboptimal choices (Newell & Simon 1972). Such choices can arise for a variety of reasons, including simply not understanding or being aware of the options, mistakes in evaluating the outcomes, and various biases based on how the choices are described (Kahneman & Tversky 1979).

In this paper, as one economic context where the standard framework has difficulty, we examine modeling risk behavior and bounded rationality (Camerer 1995). Risk behavior has been studied in various applications (Cox, Smith, & Walker 1988; Chen & Plott 1998), with primarily focus on how uncertainty modifies preferences, in the context of preference maximization. From experiments with lottery choices, we show that no consistent risk preferences in the standard economic framework can explain the observed decisions. We designed the experiment to minimize social preference effects and loss aversion, thereby allowing us to focus on bounded rationality as a key behavioral effect. Thus, it is unlikely a modification of risk preferences can explain the observations. Instead, we turn to a bounded rational explanation using one model of bounded rationality, the quantal response framework. This framework also applies to strategic games, such as auctions (Goeree, Holt, & Palfrey 2002; 2003).

The next section describes the quantal response model of bounded rationality. We then describe how this model applies to the specific situation of decisions under risk, our experiment design and its results. This discussion shows how the quantal response model introduces bounded rationality to account for our observations. We generalize this result to show why the standard utility framework cannot account for our observations.

Quantal Response Model

One approach to model bounded rationality in the utility framework is the quantal response model (McKelvey & Palfrey 1995). This model of behavior considers utility-maximizing agents who occasionally make mistakes in identifying their best choice. In particular, this models people behaving *as if* they evaluate the utilities of their choices with the addition of random errors. The net result is a probabilistic model in which people have higher probabilities to select choices with higher utilities, but do not always do so. While this model can employ various assumptions on the distribution of mistakes, a common approach takes the probability a person selects choice with utility U to be proportional to $\exp(\gamma U)$. Here γ characterizes the likelihood a person selects choices with less than the best utility: $\gamma = 0$ corresponds to random choices and $\gamma \to \infty$ when the person almost always selects the highest utility choice. While γ characterizes the mistakes, its quantitative value depends on the arbitrary scaling of the utility function. Thus more significant quantities from the model involve scale-invariant combinations of γ and the utility, such as the probability a person selects the choice with highest utility.

In the quantal response model, the likelihood of mistakes among two choices increases as the difference in their utilities decreases. That is, mistakes tend to arise among choices of similar utility where extensive analysis may be required to identify the better option. This is an example of a more general challenge in mapping continuous values (e.g., utility preferences among options) to discrete outcomes (e.g., choosing one of those options) when the difference in values is small (Lamport 1986).

Example Application: Decisions with Risk

One situation leading to challenging decisions is when outcomes have a mix of properties leading to conflicting preferences. One such example is decisions involving risk, where an outcome could have a large payoff with a low probability. In this situation, a person must make a tradeoff between a desirable feature (high payoff) and an undesirable one (low chance for the payoff). Thus decisions with risk provide a prototypical modeling example involving a variety of behavioral effects (Kahneman & Tversky 1979) while also being readily studied theoretically and experimentally.

Risk involves making decisions in the face of uncertainty about outcomes and the choices others (e.g., competitors) may make. In this context, we say someone is risk averse when they prefer a certain payment to a risky choice with the same expected value. Difficulties in modeling behavior under uncertainty arise from how individuals act when confronted with risk (Kahneman & Lovallo 1993), misperceptions of randomness (Wagenaar 1972) and aversion to ambiguity (Camerer & Weber 1992). Furthermore, risky alternatives suffer from framing issues: people tend to be risk averse with formulations emphasizing positive outcomes, but risk seeking when emphasizing poor outcomes (Kahneman & Tversky 1979).

While there is substantial research into how people behave under uncertainty and mechanisms that adjust for risk attitudes (e.g., in adjusting group predictions (Chen, Fine, & Huberman 2003)), little work exists in how to design mechanisms to influence that behavior. Developing models of actual decision-making under risk is a key component for designing such mechanisms. For example, a quantitative model of how people make decisions involving risk is useful as basis for improving organizations by adjusting risk preferences (Hogg & Huberman 2008) or altering signalling with deliberately risky decisions to hide lower skill levels compared to peers (Siemsen 2008). Such models could also have broader public policy implications to help people better manage risk (Shiller 2003).

lottery	payoff 1	payoff 2	probability of payoff 1
A (risky)	1000	0	$1/2$
B (moderate)	650	250	$1/2$
C (safe)	400	N/A	1

Table 1: Lottery payoffs in "experimental dollars" and probabilities. The lotteries decrease in both expected value and variance from top to bottom. Lottery C has no risk.

To illustrate bounded rationality and the quantal response model, we consider a simple example of decisions with risk: lotteries (Holt & Laury 2002). A lottery consists of two possible payoffs and an associated probability for each outcome. For example, a lottery could pay either \$6, with probability 0.6, or \$2, with probability 0.4. The remainder of this section describes a set of lottery experiments, a utility-based model of the observed behavior and estimated parameters for the model. In this case, bounded rationality is a key ingredient to account for the observations.

Experiments

Experiments identify behavioral aspects of human economic decision-making in controlled situations (Kagel & Roth 1995). In an experiment, people are recruited to play the roles of decision makers. Subjects receive a full description of the experiment with no deception, and are paid according to their profits in the experiment.

Experiments with lotteries are particularly simple since the lottery choices do not involve interactions among people. Lotteries thus represent simple risk choices where payoffs and probabilities are precisely known, unlike, e.g., deciding whether to trust someone to fulfill a contract (Chen, Hogg, & Wozny 2004), reciprocate offers in the trustee game (Keser 2003) or combine monetary with social rewards such as status (Loch, Huberman, & Stout 2000; Wedekind & Milinski 2000).

We ran several email-based experiments using members of the Stanford community. In our experiments, we sent a selection of lottery pairs to participants, who then indicated which of each pair they preferred in an return email. Although each experiment asked about multiple pairs of lotteries, participants were informed that only *one* of the pairs, randomly selected, would actually pay off. That is, after collecting the responses, we randomly selected one of the lotteries they selected and determined the payout from just that lottery. We chose this design to reduce portfolio effects, where, for example, if all lottery pairs were performed, a risk averse person might pick a few high-payoff low-probability lotteries, comfortable in the knowledge that other, lower risk, choices would at least give some payoff.

choices	number of participants
A,B	51
A,C	18
C,B	14
C,C	24

Table 2: Number of people making each choice among two pairs of lotteries: picking A or C, and B or C. This experiment involved at total of 107 people.

Table 1 shows some of the lotteries we used, with payoffs given in "experimental dollars", converted at a preannounced exchange rate of 20 to $1. In one set of experiments we used these lotteries as two pairs, each consisting of a choice between a risky lottery and the safe option, i.e., the first choice was between lotteries A and C, the second between B and C. Table 2 shows the number of times people made each of the four possible selections for these pairs. The most popular choices were for the two risky lotteries, indicating a degree of risk-seeking behavior in this situation. Another large group select the safe choice (lottery C) in both cases, indicating risk-averse preferences. For simplicity in illustrating the use of the quantal response and utility framework, we focus just on these observations though our experiments involved a larger set of lotteries. In particular, the payoffs in these lotteries provide opportunities for outcomes which the standard utility framework cannot explain.

Decision Model

Two major approaches to modeling decisions under risk are the expected utility framework in economics and the mean-variance utility model popular in finance (Morone 2007). The expected utility model considers the utility of a risky choice as given by the expected value of the utilities of the payoffs associated with the possible outcomes. These utilities need not be linear in the payoffs, thereby allowing the model to account for risk averse or risk seeking preferences through utilities that are concave or convex, respectively. In contrast, the mean-variance utility model considers a linear combination of expected payoff and its variance. In either case, these models are only suitable for comparing choices involving a limited range of payoffs (Rabin & Thaler 2001), as is the case for our lottery experiments.

For the experiments reported in Table 2, the mean-variance framework provides a better fit to the observations, so we adopt this model here. In this model, the utility for a decision with expected payoff E and variance V is

$$U_r = E - rV \tag{1}$$

where the parameter r characterizes the person's risk preference, e.g., positive values correspond to risk averse behavior. In this utility, the value of r depends on the scale of the payoffs since expected value is linear in the payoffs while the variance is quadratic. More generally, Eq. (1) corresponds to a linear expansion of a general functional form $U(E, V)$. Provided the decisions involve options where E and V vary over a relatively small range, the linear expansion captures this utility with a single parameter r, which characterizes how the person trades expected value against variance.

parameter	value	confidence interval
f	0.8	$0.7 - 0.9$
r_1	2×10^{-4}	$-1 \times 10^{-4} - 3 \times 10^{-4}$
r_2	0.5	$0.3 - 0.7$
γ	0.02	$0.01 - 0.04$

Table 3: Maximum likelihood parameter estimates and 95% confidence intervals.

Estimation

With experimental observations and a model with parameters, we estimate values for the parameters via maximum likelihood, and then test the adequacy of the resulting fit with a randomization test based on the difference between the observed and predicted times each choice was made based on the best fit parameters for each random sample (Clauset, Shalizi, & Newman 2007).

The simplest model is a homogeneous population. In this case we require estimates for the risk parameter r and the bounded rationality γ. The risk parameter can be either positive or negative while γ is nonnegative. A maximum-likelihood parameter estimation for this homogeneous model gives r close to zero and γ about 0.01 based on the observations of Table 2. However, a randomization test indicates the homogeneous model is unlikely to account for the observations of Table 2, with p-value less than 10^{-3}.

Instead, the relatively large number of people choosing both risky (A,B) or both safe (C,C) choices suggests members of our subject pool have diverse risk preferences. One model for such variation considers the population as consisting of two groups: a fraction f with risk parameter r_1 and the remaining fraction with risk parameter $r_2 > r_1$. For this model, Table 3 shows the maximum-likelihood parameter estimates from the experimental observations of Table 2 and their confidence intervals estimated from 1000 bootstrap samples (Cohen 1995). The parameters achieving the maximum likelihood are not unique, and we show values with the largest γ which still achieves the maximum likelihood. This corresponds to minimizing the number of choices arising from "mistakes" rather than from the preferences expressed by the utilities. The same randomization test as used for the homogeneous model shows the two-group model is consistent with the observations.

To better understand the risk parameters from this fit, from Eq. (1), lottery A has higher utility than C for risk parameter $r < R_{A,C} \equiv 1/2500 = 4 \times 10^{-4}$, and lottery B has higher utility than C for $r < R_{B,C} \equiv 1/800 = 12.5 \times 10^{-4}$. Conversely, when r is larger than these values, the utility of C is larger than A or B, respectively. The values $R_{A,C}$ and $R_{B,C}$ are above the confidence interval for r_1 and below the confidence interval for r_2 in Table 3. Thus our estimates indicate the fraction f of the population has risk preferences leading to the choices A,B having highest utility, and the remaining fraction gives highest utility to C,C. Other choices made by each group are "mistakes", including when members of the first group pick C,C or those in the second group pick A,B. With the parameters of Table 3, these cases arise in 35% of the choices, almost all involving people in the first group, with risk parameter r_1.

To illustrate the consequences of the bounded rationality, consider a person in the first group. For the choice between lotteries A and C, the corresponding utilities, from Eq. (1), are 453 and 400, respectively. So the choice maximizing utility is lottery A. With bounded rationality, the quantal response model predicts the person actually selects lottery A with probability 78%. This provides a sense of the level of consistency of the choices we observed in the experiments: in this case we interpret the remaining 22% of the choices as mistakes due to bounded rationality.

The confidence intervals on the parameters are fairly broad, indicating the number of observations is not sufficient for tight estimates on model parameters. In particular, the r_2 value, for risk-averse people, is only weakly constrained by the limited set of observations in Table 2. Because these people mainly selected both safe choices, we only have a lower bound on their risk preference and would need further lottery choices with less risky alternatives to the safe choice, or both lotteries involving some risk, to better characterize their risk preference. Such additional lotteries were included in our full set of experiments.

We also examined the expected utility framework, with risk preference corresponding to the utility associated with a payoff x given by x^r, with positive risk parameter r. Here $r < 1$ corresponds to risk averse preference. This utility, combined with the quantal response model, also gives a reasonable maximum-likelihood fit to Table 2 based on two distinct risk preferences in the population. However, the maximum likelihood is smaller than achieved with the mean-variance model, making the latter a more likely explanation, as quantified, for example, with the Akaike information criterion (James & Plank 2007).

Bounded Rationality Requirement

The behavior in our lottery experiment is not adequately described through simple utility maximization with either the mean-variance utility of Eq. (1) or the expected utility framework. A question arises is whether *any* reasonable choice of utility for lotteries could explain the observations without resorting to bounded rationality, i.e., assuming people pick the choice with highest utility. As an initial approach to this question, we demonstrate no choice of risk preference for the mean-variance utility can account for the observations under the assumption of utility maximizing behavior. A similar argument gives the same conclusion for the expected utility framework.

In the experiment, we observe some people choose lottery A over C and C over B. For a utility-maximizing decision maker, these choices require $U_r(A) > U_r(C) > U_r(B)$. In the mean-variance utility framework of Eq. (1) we claim there is no r, i.e., a consistent risk attitude, that satisfies these inequalities, as demonstrated below more generally.

Inconsistency of Utility Maximization

Consider a person making *two* choices among pairs of lotteries. The first choice is between a lottery L_1 with probability p to receive payment x and otherwise nothing, or a riskless lottery Z with sure payment z. The second choice is between a lottery L_2 with probability q of receiving y_1 and $1-q$ of receiving y_2, or the same fixed payment lottery Z as in the first decision. We take these parameters to satisfy $x > y_1 > z > y_2 > 0$. Thus, both decisions involve choosing between a risky lottery and a riskless one.

Theorem: If the decision-maker strictly prefers L_1 over the riskless lottery Z and strictly prefers Z over L_2, i.e.,

$$U_r(L_1) > U_r(Z) > U_r(L_2) \tag{2}$$

then there is no value of r such that the decision-maker the choice maximizing $U_r(L)$, given by Eq. (1), when the payoffs and outcome probabilities satisfy

$$\frac{qy_1 + (1-q)y_2 - z}{q(1-q)(y_1-y_2)^2} \geq \frac{px - z}{p(1-p)x^2} \tag{3}$$

Proof: The variance of a lottery with probably p to receive w_1 and probably $1-p$ to receive w_2 is $p(1-p)(w_1-w_2)^2$. Thus Eq. (2) implies

$$px - rp(1-p)x^2 > z \tag{4}$$

and

$$z > qy_1 + (1-q)y_2 - rq(1-q)(y_2-y_1)^2 \tag{5}$$

Eq. (4) implies $r < \frac{px-z}{p(1-p)x^2}$ and Eq. (5) implies $r > \frac{qy_1+(1-q)y_2-z}{q(1-q)(y_1-y_2)^2}$. Thus, $\frac{px-z}{p(1-p)x^2} > \frac{qy_1+(1-q)y_2-z}{q(1-q)(y_1-y_2)^2}$, which contradicts Eq. (3). Thus conditions Eq. (2) and Eq. (3) are inconsistent. QED.

Eq. (3) states the increase in expected value between L_2 and Z, compared to the variance of L_2, is larger than the corresponding ratio for the choice between L_1 and Z. Thus, the mean-variance utility model suggests a person preferring L_1 to Z would find the trade-off between expected gain and variance even more attractive for L_2 compared to Z.

The theorem applies to each decision-maker individually. So even if utilities (i.e., the value of the risk parameter r) differ among people, we would not expect any choices of rational utility maximizers to violate Eq. (3).

Application to Experimental Observations

For lotteries A, B and C in the experiments, the values in Table 1 correspond to $p = q = 1/2$, $x = 1000$, $z = 400$, $y_1 = 650$ and $y_2 = 250$, which satisfy Eq. (3) and so the observed choices for A and C are inconsistent with mean-variance utility maximization. A similar argument as given above shows this set of choices is also inconsistent with the expected-utility framework.

The theorem shows the joint choices of lottery A over C, and C over B can be considered as "inconsistent choices" because no consistent mean-variance utility formulation offers an explanation based on utility maximization. For the other three choices in Table 2, such consistent formulations are possible. In particular, with the mean-variance utility, a person with $r < \min(R_{A,C}, R_{B,C}) = 1/2500$ has highest utility for the choices A,C; a person with $R_{A,C} < r < R_{B,C}$ has highest utility for choices C,B; and a person with $r > \max(R_{A,C}, R_{B,C}) = 1/800$ has highest utility for choices C,C. Thus, while observations of these choices *could* be mistakes, as they are not the maximum utility choices, the observations *need not* be attributed to mistakes. The choices A,C, on the other hand, must arise from mistakes. Thus, from Table 2 we have *at least* 18 of the 107 observations arising from mistakes. We can use this to estimate a lower bound on the probability P people do not choose their utility-maximizing choices. In particular, the maximum likelihood estimate for P is just $18/107 = 17\%$ with 95% confidence interval of $11\% - 25\%$. From this discussion, the theorem and our observations imply we are likely to have P of at least 11%. This compares with the estimate of $P = 35\%$ of choices attributed to mistakes by the maximum likelihood fit to the two-group model given above, which considers some of the other three choices to also be mistakes.

Discussion

In summary, we conducted human subject experiments and showed, in the context of lottery choices, people are inconsistent with standard economic theory. By accounting for bounded rationality, we showed the quantal response model gives a reasonable alternate explanation. More generally, extending standard economic models to include behavioral effects (e.g., bounded rationality) broadens the scope of economic approaches to modeling human behavior.

One challenge in applying economic models to human behavior is relating the somewhat arbitrary choice of utility functions to observable behavior. The utility reflects a preference ordering so any monotonic transformation of the utilities gives the same behavior predictions. More specific models, such as the quantal response and expected utility maximization frameworks, rely not only on the ordinal but also the cardinal quantity of the utility, so they are not invariant with respect to such transformations. For example, in the quantal response model the likelihood of mistakes depends on the difference in utilities whereas rational behavior just uses the ordering. Specifically, the probability of a choice is proportional to $e^{\gamma U}$, which gives the same results when all utilities are multiplied by some factor as long as γ is correspondingly divided by the same value. Thus while using numerical utility values in a model can be useful, robust predictions must focus on scale-independent behaviors, such as choice preferences or probabilities.

An important modeling question is the extent to which our claim of the limitations of utility maximization generalizes to arbitrary utility models, beyond the mean-variance and expected utility formulations we discussed. Our claim is not true for arbitrary nonlinear preference functions, e.g., allowing arbitrary utility values for the choices. The more interesting question is whether there are conditions on the utility that are both behaviorally reasonable and consistent with the experimental observations. Such reasonableness conditions include increasing preference for higher payoffs and consistent risk preferences (either risk seeking or risk averse) over the limited range of payoffs involved in the experiments.

When the quantal response model reasonably describes behavior, a practical issue is estimating its parameter γ. Since utilities are not known a priori, the utility used for estimation may not reflect people's actual preferences, in which case the estimation procedure will consider observed deviations from the best choice according to the utility model as "mistakes", thereby decreasing the estimated value of γ. Thus, estimating γ confounds actual errors with a misspecified utility model. As one approach to this problem, estimates of γ with different utility models can suggest the most descriptive utility model. If parameters indicate one utility model gives a statistically significant higher value for γ than another, the former ascribes less of the observed behavior to "random" mistakes. Controlled experiments (as opposed to observational field data with no ability to change the available choices), can partially address this confounding of model accuracy and bounded rationality. For example, asking a person to select among the same set of choices multiple times, after a sufficient time interval or recasting choices in different terms to avoid correlated mistakes, will induce multiple independent errors for estimating γ. On the other hand, systematic correlated "mistakes" likely indicate an incorrect utility model.

There are several directions for future research. In the context of managing risk behavior, it is important to measure the risk and bounded rational parameters for an individual, as opposed to the population-level estimate discussed in this paper. However, it is difficult to obtain enough choices from each person for a precise estimation of parameters. In this case, an adaptive procedure may help (Castro, Willett, & Nowak 2005; Zheng & Padmanabhan 2006), i.e., subsequent choices presented to an individual are based on their prior choices. In this context, a model is useful not only to explain observations but also to help design further experiments by suggesting ranges of parameters to focus on, e.g., appropriate risk levels to better distinguish the parameters.

Utility-based models, such as quantal response, are not the only approach to understanding bounded rationality. Other approaches include psychologically-based models, such as estimates based on a series of simulated outcomes (Busemeyer & Townsend 1993), which also focus on decision time rather than just the outcome. An interesting question is what level of cognitive detail is necessary to model various economic behaviors. From a modeling standpoint, specific experiments can be designed to distinguish the structure and types of mistakes people make. An example is to distinguish among systematic mistakes, caused by a person using an incorrect method, and random errors, caused by imprecise calculations. Another example is to identify correlations between mistakes, which are neglected in the model we discussed. One possibility to help distinguish random mistakes from an incorrect utility model is using what people say *others* will choose (Prelec 2004), under the assumption that members of the group have some common information and preferences. Such distinctions could suggest how to improve utility models. Beyond conventional economic and psychological experiments, new technologies allow a more fine-grained observation of the decision process (Bernheim 2008; Logothetis 2008) and may lead to improved predictive models, including for decisions involving risk (Kuhnen & Knutson 2005). Moreover, increasingly common mobile devices people carry can provide information on small group interactions (Pentland 2007) which could be included as additional components of utility-based models, e.g., to allow modeling social preference effects.

The utility and quantal response framework described in this paper is useful for economic decisions with well-defined choices over a limited set of alternatives. This framework of individual choices with bounded rationality can be the basis for larger-scale stochastic or agent-based models of large groups of people. These include both conventional economic situations and broader contexts such as user behavior on the web (Huberman *et al.* 1998; Lerman 2007).

Acknowledgements

We thank Murat Kaya and Ana Meyer for helpful discussions.

References

Bernheim, B. D. 2008. Neuroeconomics: A sober (but hopeful) appraisal. Working Paper 13954, Natl. Bureau of Economic Research, Cambridge, MA.

Busemeyer, J. R., and Townsend, J. T. 1993. Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review* 100:432–459.

Camerer, C., and Weber, M. 1992. Recent developments in modeling preferences: Uncertainty and ambiguity. *J. of Risk and Uncertainty* 5:325–370.

Camerer, C. 1995. Individual decision making. In Kagel, J., and Roth, A. E., eds., *The Handbook of Experimental Economics*. Princeton Univ. Press. chapter 8, 587–703.

Castro, R.; Willett, R.; and Nowak, R. 2005. Faster rates in regression via active learning. In *Proc. of Advances in Neural Information*, 179–186. Cambridge, MA: MIT Press.

Chen, K.-Y., and Plott, C. R. 1998. Nonlinear behavior in sealed bid first price auctions. *Games and Economic Behavior* 25:34–78.

Chen, K.-Y.; Fine, L. R.; and Huberman, B. 2003. Predicting the future. *Information Systems Frontiers* 5(1):47–61.

Chen, K.-Y.; Hogg, T.; and Wozny, N. 2004. Experimental study of market reputation mechanisms. In *Proc. of the 5th ACM Conference on Electronic Commerce (EC'04)*, 234–235. ACM Press.

Clauset, A.; Shalizi, C. R.; and Newman, M. E. J. 2007. Power-law distributions in empirical data. arxiv.org preprint 0706.1062.

Cohen, P. R. 1995. *Empirical Methods for Artificial Intelligence*. Cambridge, MA: MIT Press.

Cox, J. C.; Smith, V. L.; and Walker, J. M. 1988. Theory and individual behavior of first-price auctions. *Journal of Risk and Uncertainty* 1:61–99.

Goeree, J. K.; Holt, C. A.; and Palfrey, T. R. 2002. Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory* 104:247–272.

Goeree, J. K.; Holt, C. A.; and Palfrey, T. R. 2003. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior* 45:97–113.

Guth, W.; Schmittberger, R.; and Schwarze, B. 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3:367–388.

Hogg, T., and Huberman, B. A. 2008. Solving the organizational free riding problem with social networks. In Lerman, K., et al., eds., *Proc. of the AAAI Symposium on Social Information Processing*, 24–29.

Holt, C. A., and Laury, S. K. 2002. Risk aversion and incentive effects. *American Economic Review* 92:1644–1655.

Huberman, B. A.; Pirolli, P. L. T.; Pitkow, J. E.; and Lukose, R. M. 1998. Strong regularities in World Wide Web surfing. *Science* 280:95–97.

James, A., and Plank, M. J. 2007. On fitting power laws to ecological data. arxiv.org preprint 0712.0613.

Kagel, J., and Roth, A. E., eds. 1995. *The Handbook of Experimental Economics*. Princeton Univ. Press.

Kahneman, D., and Lovallo, D. 1993. Timid choices and bold forecasts: a cognitive perspective on risk taking. *Management Science* 39:17–31.

Kahneman, D., and Tversky, A. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–292.

Keser, C. 2003. Experimental games for the design of reputation management systems. *IBM Systems Journal* 42(3):498–506.

Kuhnen, C. M., and Knutson, B. 2005. The neural basis of financial risk taking. *Neuron* 47:763–770.

Lamport, L. 1986. Buridan's principle. Available at research.microsoft.com/users/lamport/pubs/buridan.pdf.

Lerman, K. 2007. User participation in social media: Digg study. In *IEEE/WIC/ACM Intl. Conf. on Web Intelligence and Intelligent Agent Technology*, 255–258.

Loch, C. H.; Huberman, B. A.; and Stout, S. 2000. Status competition and performance in work groups. *J. of Economic Behavior and Organization* 43:35–55.

Logothetis, N. K. 2008. What we can do and what we cannot do with fMRI. *Nature* 453:869–878.

McKelvey, R. D., and Palfrey, T. R. 1995. Quantal response equilibria for normal form games. *Games and Economic Behavior* 10:6–38.

Morone, A. 2007. Comparison of mean-variance theory and expected-utility theory through a laboratory experiment. Working Paper 19, Univ. of Bari.

Newell, A., and Simon, H. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

Pentland, A. S. 2007. Automatic mapping and modeling of human networks. *Physica A* 378:59–67.

Prelec, D. 2004. A Bayesian truth serum for subjective data. *Science* 306:462–466.

Rabin, M., and Thaler, R. H. 2001. Risk aversion. *J. of Economic Perspectives* 15(1):219–232.

Shiller, R. J. 2003. *The New Financial Order: Risk in the 21st Century*. Princeton University Press.

Siemsen, E. 2008. The hidden perils of career concerns in R&D organizations. *Management Science* 54:863–877.

Wagenaar, W. A. 1972. Generation of random sequences by human subjects: a critical review of the literature. *Psychological Bulletin* 77:65–72.

Wedekind, C., and Milinski, M. 2000. Cooperation through image scoring in humans. *Science* 288:850–852.

Zheng, Z., and Padmanabhan, B. 2006. Selectively acquiring customer information: A new data acquisition problem and an active learning-based solution. *Management Science* 52:697–712.