Strategic Experimentation In The Lab

By: John R. Boyce, **David M. Bruner**, & **Michael McKee**

Abstract

This paper reports the results of experimental tests of the Nash equilibrium predictions in a one-armed bandit game with information spillover. Players learn the probability that a risky prospect pays by either taking draws from the distribution themselves or ob- serving the outcome of another player's choice. Our experiment is designed to learn whether players experiment strategically, anticipating the opportunity to free-ride on others' information and doing so. While error rates exhibit a bias toward under- experimentation, we observe a significant strategic effect. Structural parameter estimates suggest the lack of experimentation observed is due to decision error and somewhat pes- simistic priors, rather than risk preferences or probability weighting.

Boyce, J. R., **Bruner, D. M.**, and **McKee, M**. (2016) Strategic Experimentation in the Lab. Manage. Decis. Econ., 37: 375–391. doi: 10.1002/mde.2723. Publisher version of record available at: https://onlinelibrary.wiley.com/doi/10.1002/mde.2723

Strategic Experimentation in the Lab

John R. Boyce[a], David M. Bruner[b] and Michael McKee[b,*]

[a]*Department of Economics, University of Calgary, Calgary, AB T2N 1N4, Canada*
[b]*Department of Economics, Appalachian State University, Boone, NC 28608, USA*

This paper reports the results of experimental tests of the Nash equilibrium predictions in a one-armed bandit game with information spillover. Players learn the probability that a risky prospect pays by either taking draws from the distribution themselves or observing the outcome of another player's choice. Our experiment is designed to learn whether players experiment strategically, anticipating the opportunity to free-ride on others' information and doing so. While error rates exhibit a bias toward under-experimentation, we observe a significant strategic effect. Structural parameter estimates suggest the lack of experimentation observed is due to decision error and somewhat pessimistic priors, rather than risk preferences or probability weighting.

1. INTRODUCTION

Most decisions of chance involve a considerable degree of uncertainty or ambiguity (Knight, 1921; Ellsberg, 1961). Consider the following: a prospective business owner deciding whether a neighborhood is a good location for a new establishment; a taxpayer deciding whether to engage in evasion; and a consumer deciding whether to purchase a new product. In each case, the likelihood of success is unknown to the individual. Often, the only means of resolving this uncertainty is through costly experimentation (Gittins, 1979). It is not always the case, however, that one need be the proverbial guinea pig. In the presence of spillovers, it is possible that agents can free-ride on others' information, creating an incentive to strategically delay experimentation themselves (Guzman and Ventura, 1998; Bolten and Harris, 1999).[1] The business owner can observe the experience of similar businesses in the vicinity. The taxpayer can observe the fate of friends, coworkers, and family members. The consumer can observe the experience of others who decided to purchase the new product. The behavioral question is whether individuals actually engage in this sort of strategic experimentation, by free-riding on others' information when the opportunity is anticipated. This is an open empirical question and has important implications regarding the public reporting of information, such as entrepreneur success rates or consumer reports, that can mitigate such strategic delay.

To address this question, this paper reports the results of a laboratory experiment designed to examine whether individuals behave in a manner consistent with the Bayesian-Nash prediction in a one-armed bandit problem with information spillover. We consider a two-player representation, where each player may simultaneously choose either a safe or risky option in each of two stages. The risky option has two states of nature, success or failure, with unknown chances common to both players. If the risky option is selected by either player, the outcome is observed by both players. Thus, prior to making their second choice, each player will have an information set consisting of the result from zero, one, or two draws from the distribution of uncertain outcomes associated with the risky choice. As Bolten and Harris (1999)

*Correspondence to: Department of Economics, Appalachian State University, Boone, NC 28608, USA. E-mail: mckeemj@appstate.edu

show, the information spillover changes the incentive to engage in costly experimentation via the *free-ride effect*. That is, a marginal player should strategically delay experimentation if they anticipate the other player will provide them with information.

The experiment induces player heterogeneity in the opportunity cost of experimentation. Players are assigned one of three safe options: low, middle, or high. Players who face a low (high) safe option have a dominant strategy to choose the risky (safe) option in the first stage. These dominant strategy players are always paired with a middle safe option player, who has a conditional best-response to *experiment strategically*. That is, these players should only experiment with the risky option in the first stage, and provides public information, when paired with a high safe option player.[2] When paired with a low safe option player, however, they should free-ride on the other player's provision of information. Hence, the same amount of information should be generated in either game; only one player should ever experiment with the bandit first stage.

The first-stage behavior of subjects in the experiment indicates a strategic response to the information spillover; there is a significant reduction in experimentation when the opportunity to free-ride is anticipated. Still, consistent with public goods experiments (Andreoni, 1995), we observe too little experimentation when players should and too much experimentation when players should not. While this is to be expected in a binary choice setting with player decision error (Andreoni, 1995), these error rates are asymmetric, exhibiting a bias towards under-experimentation.[3] This is consistent with previous evidence of ambiguity aversion (Ellsberg, 1961; Hogarth and Einhorn, 1990; Camerer and Weber, 1992; Viscusi and Magat, 1992; Fox and Tversky, 1995; Wakker *et al*., 1997; Anderson, 2012). Nonetheless, when only one player experiments (which occurs roughly two-thirds of the time in the data), it is the player predicted to do so at least 96% of the time. Thus, although players experiment with the bandit less than predicted, we still find evidence that players behave strategically, by free-riding when possible in the presence of an information spillover.

To our knowledge, this research is the first to provide evidence that *individuals* experiment strategically. Hendricks and Kovenock (1989) and Caplin and Leahy (1998) provide evidence that *firms* exhibit such behavior in the exploration for petroleum and new retail markets, respectively. Dixit and Pindyck (1999) suggest the same may be true for research and development. While the present findings indicate such strategic free-riding on information is not restricted firm behavior, individual's experiment significantly less than classic decision theory predicts. On the one hand, this may be beneficial when the risky activity is detrimental to social welfare, such as with tax evasion. On the other hand, it suggests insufficient information will be generated for welfare enhancing risky activities, such as entrepreneurship and/or new product adoption. Policies that promote public reporting of information, such as entrepreneur success rates and/or consumer reports, could alleviate information shortages.

Analysis of second-stage decisions suggests individuals would be responsive to such policies. Overall, subjects choose the Bayesian expected income maximizing choice 82.9% of the time. Still, the results indicate two sources of potential bias to Bayesian updating: (i) the strength of the information signal and (ii) previous behavior. Similar to recently reported behavior (Holt and Smith, 2009; Poinas *et al*., 2012), we find the probability of a player choosing the risky option is positively correlated to the proportion of successes observed, despite controlling for Bayesian updating. This is consistent, however, with behavioral models that allow for decision error (Fechner, 1860/1966; Luce, 1959; Smith and Walker, 1993). Still, the probability of choosing the risky option in the second stage is also positively correlated with doing so in the first stage, which is also consistent with previously reported behavior (Charness and Levin, 2005; Hey and Panaccione, 2011).[4] Nonetheless, controlling for these two factors, we find significant evidence of Bayesian updating. In each treatment, the largest increase in the proportion of players that choose the risky option in the second stage occurs after observing the critical proportion of successes in the first stage. Hence, it is unlikely the lack of experimentation is due to a failure to process information.

Finally, we explore the extent to which risk preferences, probability weighting (a popular method of modeling ambiguity aversion), and/or pessimistic priors (another approach to modeling ambiguity aversion) can explain why players experimented less than predicted. As such, the analysis adds to the growing literature regarding estimation of prior beliefs (McKelvey and Page, 1990; Nyarko and Schotter, 2002; Holt and Smith, 2009; Karni, 2009; Offerman *et al*., 2009; Andersen *et al*., 2010; Hao and Houser, 2012). This literature has primarily focused on incentive-compatible mechanisms to elicit prior beliefs, such as scoring rules, which require risk neutral preferences for truthful revelation (Offerman *et al*., 2009).[5] The present analysis adopts an alternative strategy similar to that of Offerman *et al*. (2009). The experimental design employs a standard multiple price list (MPL) to elicit risk attitudes (Holt and Laury, 2002), which are subsequently used to calibrate estimates of prior beliefs.[6] The estimation strategy combines first-stage and second-stage decisions from the bandit game (with and without information) with choices from the risk preference elicitation MPL (with known and unknown probabilities) to obtain structural parameter estimates of risk preference, probability weighting, prior beliefs, and stochastic decision error.

Overall, neither risk preference nor probability weighting appear to be the reason for the lack of experimentation observed in first-stage choices of the bandit game. When significantly different from neutrality, the estimated constant relative risk aversion (CRRA) parameter implies risk-seeking preferences. While there is evidence of weighting *known* probabilities, consistent with previous research (Tversky and Kahneman, 1992; Wu and Gonzalez, 1996; Prelec, 1998; Stott, 2006), we find no evidence of significant weighting of *unknown* probabilities, which is also consistent with previous findings (Budescu *et al*., 2011). Rather, the lack of experimentation in the first stage of the bandit game appears to be mainly the result of slightly pessimistic priors and decision error, as these parameter estimates are significantly different from the values assumed in the theoretical framework. This is consistent with ambiguity aversion in prospective reference theory (Viscusi, 1989).

The remainder of the paper is organized as follows. In the next section, we present a model of a one-armed bandit problem with information spillovers and derive a solution to the game used in the experiment. In Section 3, we present the experimental design and formal hypotheses of expected behavior. In Section 4, we present the results of the data analysis. Finally, we summarize and discuss our findings in Section 5.

2. INFORMATION SPILLOVER IN A BANDIT PROBLEM

Suppose two players, i and j, play a game against nature in two stages, $t=1,2$. The payoffs to each player are the sum of first-stage and second-stage returns. In each stage, the players simultaneously choose whether or not to take a safe option, paying $S_i>0$, or to take a risky option, a sample of n Bernoulli trials that result in outcomes, $x_1,\ldots,x_n$, where each $x_i\in\{0,1\}$ pays R for a success, such that $nR>S_i$ and zero for a failure. The likelihood a player observes $X=\sum_{i=1}^{n} x_i$ successes in n trials is given by $f_n(x_1,\ldots,x_n|\theta)\propto\theta^X(1-\theta)^{n-X}$, where the probability of a success for each draw from the risky option is θ, such that $0<\theta<1$, which is unknown to the players. Hence, players must make their decisions in each stage based upon their subjective beliefs (Savage, 1954) and any available information.[7]

Let $\xi(\theta)$ denote the distribution of each player's prior beliefs about θ. Assume that $\xi(\theta)$ is a beta distribution with parameters α and β, so that $\xi(\theta)\propto\theta^{\alpha-1}(1-\theta)^{\beta-1}$ (DeGroot, 1970, p. 40).[8] Since players are completely uninformed about the value of θ, it is natural to assume $\alpha=\beta=1$ for each player, which implies that $\xi(\theta)$ is a uniform distribution over the interval zero to one.[9]

As each player may draw n samples, $N\in\{0,n,2n\}$ denotes the number of draws observed in the first stage of the game. Hence, the possible information sets in the second stage are denoted $I_0=\{\varnothing\}$, $I_1(X)=\{n,X\}$, $I_2(X)=\{2n,X\}$, where the subscript on each information set represents the number of players who have chosen the risky option in the first stage. The posterior distribution of beliefs is thus given by $\xi(\theta|x_1,\ldots,x_N)\propto\theta^{\alpha+X-1}(1-\theta)^{\beta+N-X-1}$, which is also beta distribution with parameters $\alpha+X$ and $\beta+N-X$ (DeGroot, 1970). Let $p_t=E[\theta|I]$ denote each players' common expectation of θ in stage t, given the information, I, available at that point. Thus $p_1=\frac{\alpha}{\alpha+\beta}$ and $p_2(I)=\frac{\alpha+X}{\alpha+\beta+N}$.[10] Hence, the expectation of the uninformed prior and posterior distributions are $p_1=\frac{1}{2}$ and $p_2=\frac{1+X}{2+N}$, respectively.

Player i chooses the risky option in the first stage if and only if

$$Rnp_1-S_i+EVMS_i(N)>0, \tag{1}$$

where $EVMS_i(N)$ is the expected value of a message service that provides N draws from the risky option to a player with safe return S_i. The difference in the first two terms in (1) is the expected opportunity cost of obtaining information in the first stage. This equals the expected return from choosing the risky option less the forgone return to choosing the safe option. Thus, the decision rule is to choose the risky option if the expected value of the message service is greater than the expected cost of obtaining the information. However, the expected value of the message service depends on how much information will be provided.

Suppose player j has chosen the safe option in the first stage. Then, the expected value of the message service can be written as

$$EVMS_i(n) = \sum_{X=0}^{n} f_n(X|I_0)\{\max[S_i, Rnp_2(I_1(X))] - \max[S_i, Rnp_2(I_0)]\} \geq 0 \quad (2)$$

It is well known that for any given return to the safe option relative to the risky option, the expected value of the message service is nonnegative (Hirshleifer and Riley, 1992, p. 180). Information has value because it may change the perceived optimal decision in the second stage. Thus, the expected value of the message service is the difference between the weighted average of expected returns from the optimal decision, $\max[S_i, Rnp_2(I_1(X))]$, given the information signal $I_1(X)$, weighted by the expected probability of the particular information signal $f_n(X|I_0)$, and the expected return from the optimal decision in the absence of information, $\max[S_i, Rnp_2(I_0)]$. This difference is greater than zero whenever the player potentially chooses an action with information that differs from what they would have chosen absent information and zero otherwise.[11]

In contrast, if player j has chosen the risky choice in the first stage, player i's decision is no longer a comparison between zero information and a positive amount of information. Rather, player i's choice is now between two message services, one in which n draws are provided (by the other player) and one in which $2n$ draws are provided (by both players). When player j has chosen the risky option in the first stage, the expected value of the message service,

$$EVMS_i(2n) = \sum_{X=0}^{2n} f_{2n}(X|I_0)\max[S_i, Rnp_2(I_2(X))] - \sum_{X=0}^{n} f_n(X|I_0)\max[S_i, Rnp_2(I_1(X))], \quad (3)$$

is the net expected value of the message service that provides $2n$ draws, relative to the message service that provides n draws, to a player whose safe return is S_i. The first and second summations in $EVMS_i(2n)$ are the weighted averages of the updated expected returns from the second-stage optimal decisions, given that either *both* players or *only* the other player chooses the risky option, respectively. The expected opportunity cost of obtaining the message service is measured in the safe return foregone to player i. This cost is independent of whether the other player does or does not make the risky choice. The expected value of information, however, does depend upon the choice of the other player. To see this, we turn our attention to the parameters used in the experiment.

2.1. Solution to the Game Used in the Experiment

In the experiment, the payoff to a successful draw is $R = \$5$ and choosing the risky option gives the player $n = 3$ draws. There are three types of players in the experiment, indexed by their safe returns: $S_L = \$4$, $S_M = \$8$, and $S_H = \$12$. The model predicts that the S_M player will systematically vary their behavior depending upon which type of player they are paired with, S_L or S_H. Thus, an S_M player is always paired with an either S_L or S_H player.[12] This experimental design yields four treatments consisting of the combinations of a player's own type and that of the other player $\{S_i, S_j\} \in \{\{\$4, \$8\}, \{\$8, \$4\}, \{\$8, \$12\}, \{\$12, \$8\}\}$.[13] Figure 1 displays the normal form of the first-stage games used in the experiment. The values in Figure 1 represent the sum of first-stage and second-stage expected payoffs from playing a particular strategy, given optimal second-stage behavior and an uninformed prior.[14] The S_M player's payoffs are the first number in each cell. Best-responses are in bold, italicized font. There is a unique pure-strategy Nash equilibrium in each of the games in Figure 1: the lower cost player experiments and the other free-rides.[15]

3. EXPERIMENTAL DESIGN AND HYPOTHESES

The experimental design pares down the number of confounding factors to a minimum. While one-armed bandit problems are typically studied in an infinite horizon, the two-stage game allows us to eliminate two types of confounding effects. By restricting our attention to a two-stage setting, we eliminate incentive to provide information in the current stage to *encourage* information provision by others in the future (Bolten and Harris, 1999). This allows us to cleanly test for the free-riding effect. In addition, the cognitive costs of deriving optimal stopping rules in infinite horizon bandit games may cause players to form heuristics to simplify the repeated sampling problem (Gans *et al.*, 2007). The two-stage game lessens the cognitive costs of finding the best-response.

$S_M = \$8$		$S_L = \$4$: *Safe*		*Risky*	
	Safe	16.00	11.50	***16.88***	***15.13***
	Risky	***16.38***	11.63	16.21	***15.06***

(*a*) $\{\$4, \$8\}$ Game

$S_M = \$8$		$S_H = \$12$: *Safe*		*Risky*	
	Safe	16.00	***24.00***	***16.88***	19.50
	Risky	***16.38***	***24.00***	16.21	19.52

(*b*) $\{\$12, \$8\}$ Game

Figure 1. Normal forms of the first-stage games used in the experiment, $n=3$ and $R=5$.

Furthermore, while there is considerable evidence that players are not 'perfect Bayesians'—players tend to place too little weight on prior beliefs and too much weight on new information—our experiment is designed to afford players the best shot at being 'good Bayesians'.[16] For example, it is known that when there are a small number of possible data generating processes, players form their posterior beliefs on heuristics such as whether the sample is 'representative' of one of the possibilities. Hence, we allow for a large number of possible data generating processes to minimize the influence of heuristics, such as representativeness. Moreover, when both players choose the risky option, both players receive all information simultaneously, which Hoffman *et al.* (2011) have shown to produce average beliefs closest to the Bayesian posterior. Furthermore, we employ a simple case of Bayesian updating, a binomial sample, in which the sample proportion asymptotically approaches the uninformed posterior Bayesian estimate.

Accordingly, in the experiment, paired players simultaneously make dichotomous choices across two stages: to play a lottery, with an unknown chance of winning, or not. Players simultaneously choose between a guaranteed amount, predetermined to be one of $S_i = \{\$4,\ \$8,\ \$12\}$, and a lottery with an unknown probability distribution. The lottery is framed as three draws with replacement from an urn containing 100 balls, composed of an unknown proportion θ of red balls and proportion $1-\theta$ of blue balls, where $0 \le \theta \le 1$. Each red ball pays \$5, and each blue ball pays \$0. Players are informed that θ is held constant in both stages. If either of the players in a pair chooses the lottery, both players observe the results. Hence, while players have no information regarding θ in the first stage, lottery results can be used to inform second-stage decisions, if at least on paired player chooses the lottery.

Each player participates in 20 rounds of two-stage play, employing a within-subjects design (i.e., all subjects are exposed to all treatments). In each round, players are randomly assigned to a treatment with an anonymous partner. Each treatment pair is assigned a probability of success from a uniform distribution with a mean of 0.5, in accordance with an uninformed prior. Treatments, player pairing, and the underlying probability of success are randomly assigned in each round and remain fixed for that round.[17] In each round, treatments are drawn with replacement, while player pairing and the underlying probability of success are drawn without replacement. Players know which treatment they have been assigned in the first stage of every round.

In all sessions, the instructions are read aloud, as well as presented on computer screens, to ensure common knowledge.[18] During the instruction phase, players must pass a short quiz on basic arithmetic and statistics to reduce errors due to misunderstanding.[19] In addition, risk preferences are elicited using a standard MPL, made popular by Holt and Laury (2002), which consists of a menu of 10 binary choices between their \$5 show-up fee or a lottery with an increasing (decreasing) chance of winning \$10 (\$0). Players were paid on the basis of the outcome of one randomly selected choice in the MPL task and one randomly selected round of the bandit game. Experimental sessions lasted approximately 90 min, and participant earnings averaged \$18. A total of 52 players participated in the experiment over three sessions (18 in session 1, 20 in session 2, and 14 in session 3). Because each player makes choices in each of two stages of 20 rounds of play, there are 2080 total observed risky/safe choices. The participant pool is composed of volunteer students at the North American University.[20]

3.1. Hypotheses Regarding First-Stage Choices

Let the probability that a player in treatment $\{S_i, S_j\}$ experiments with the risky option in the first stage be given by

$$Pr\left(\Delta EV_{ij} > 0\right) = \beta_{ij}\mathbf{D}_{ij} + \epsilon_{ij}, \tag{4}$$

where $\Delta \mathrm{EV}_{ij}$ denotes the sum of expected net gains from choosing the risky option, $\mathbf{D}_{ij}$ is an indicator variable equal to one if the treatment is $\{S_i, S_j\}$, and zero otherwise; and $\in_{ij}$ is an unobserved error term.

The β_{ij} parameters to be estimated correspond to the conditional mean probabilities that each $\{S_i, S_j\}$ player chooses the risky option in the first stage.[21] The Nash equilibrium to the game in Figure 1 yields the following first-stage predictions: $\beta_{4,8} = \beta_{8,12} = 1$ and $\beta_{8,4} = \beta_{12,8} = 0$. Hence, if players err, their errors are necessarily one sided: those who should play the risky (safe) strategy can only err by choosing the safe (risky) option (Smith and Walker, 1993). Thus, the error rates are $\in_{4,8} = 1 - \beta_{4,8}$ and $\in_{8,12} = 1 - \beta_{8,12}$ for the {\$4, \$8} and {\$8, \$12} treatments, respectively. Likewise, the error rates are $\in_{12,8} = \beta_{12,8}$ and $\in_{8,4} = \beta_{8,4}$ for the {\$12, \$8} and {\$8, \$4} treatments, respectively.

Notice in Figure 1, players in the {\$4, \$8} and the {\$12, \$8} treatments each have much more to lose than players in the {\$8, \$4} and the {\$8, \$12} treatments from making errors. That is, $\Delta EV_{4,8} > \Delta EV_{8,12}$ and $\Delta EV_{12,8} < \Delta EV_{8,4}$. Hence, those players with greater salience in their choices should have smaller errors relative to the Nash predictions.[22] Thus, while {\$4, \$8} and {\$8, \$12} players are each expected to choose the risky option, salience implies that $\in_{4,8} < \in_{8,12}$. Similarly, while {\$12, \$8} and {\$8, \$4} players are each expected to choose the safe option, salience implies $\in_{12,8} < \in_{8,4}$. In other words, $\beta_{8,12} < \beta_{4,8}$ (salience in risky choices) and $\beta_{12,8} < \beta_{8,4}$ (salience in safe choices). Combining the free-riding prediction, $\beta_{8,4} > \beta_{8,12}$, with the two salience predictions, then implies the following ordering of predicted proportions of risky choices:

$$\beta_{4,8} \overbrace{>}^{\text{salience in risky choices}} \beta_{8,12} \overbrace{>}^{\text{free-riding rationality}} \beta_{8,4} \overbrace{>}^{\text{salience is safe choices}} \beta_{12,8} \qquad (5)$$

As this gives a complete ordering of the expected proportions of risky choices by treatment type, we summarize it in the following hypothesis.

Hypothesis 1:

$\beta_{4,8} > \beta_{8,12} > \beta_{8,4} > \beta_{12,8}$: In the first stage, {\$4, \$8} players should choose the risky option more often than {\$8, \$12} players, who should choose the risky option more than {\$8, \$4} players, who should choose the risky option more often than {\$12, \$8} players.

Note that rejection of the hypothesis is achieved by evidence that contradicts any one of the inequalities. This can be achieved by either under-experimentation or over-experimentation. Under-experimentation occurs when both players choose the safe choice in the first stage, resulting in $\in_{8,12} > \in_{8,4}$. On the other hand, over-experimentation occurs when both players choose the risky choice, resulting in $\in_{8,12} < \in_{8,4}$. We summarize these errors in the following hypothesis.

Hypothesis 2:

(a) $1 - \beta_{8,12} > \beta_{8,4}$: In the first stage, under-experimentation results in a higher error rate for {\$8, \$12} players relative to {\$8, \$4} players; and

(b) $1 - \beta_{8,12} < \beta_{8,4}$: In the first stage, over-experimentation results in a higher error rate for {\$8, \$4} players relative to {\$8, \$12} players.

It is clear from inspection of Hypothesis 2 that (a) and (b) cannot both hold; (a) implies $\beta_{8,12} + \beta_{8,4} > 1$, while (b) implies that $\beta_{8,12} + \beta_{8,4} < 1$. Furthermore, it is possible to reject both by a finding that $\beta_{8,12} + \beta_{8,4} = 1$.

3.2. Hypotheses Regarding Second-Stage Choices

The Bayesian predictions are players with a safe option of S_i should choose the risky option in the second stage whenever $X \geq X_N^{S_i} = \frac{S_i(2+N)}{nR} - 1$ for $N \in \{n, 2n\}$.[23] Allowing for player error, however, implies the difference in expected utility between the safe and risky options is increasing in the difference $X - X_N^{S_i}$. Hence, players' second-stage behavior should depend not only on the type of information they observe, $XX_N^{S_i}$, but also on the strength of the signal, $X - X_N^{S_i}$. Let ΔEV_{it}^2 denote the expected net gain from choosing the risky option in the second stage to the S_i player in round t. Hence, the probability that a player chooses the risky option in the second stage may be written as

$$Pr(\Delta EV_{it} > 0) = \beta_{Aij}\mathbf{D}_{Xi} + \beta_{Bij}\left(\mathbf{X} - \mathbf{X}_N^{S_i}\right) + \in_i, \quad (6)$$

where $\mathbf{D}_{Xi}$ is a vector of indicator variables which equal to one if $X \geq X_N^{S_i}$ for each treatment, and zero otherwise;[24] $\mathbf{X} - \mathbf{X}_N^{S_i}$ is a vector of the differences between the observed and the critical number of successes, which is equal to zero when no draws are observed; and $\in_i$ is the unobserved error. While strict Bayesian updating implies $\beta_{Aij} = 1$ and $\beta_{Bij} = 0 \ \forall \ \{S_i, S_j\}$, allowing for decision error implies the following hypothesis.

Hypothesis 3:
$\beta_{Aij} > 0 \ \forall \ \{S_i, S_j\}$: In the second stage, all players should choose the risky option more often whenever the Bayesian condition is satisfied, $X \geq X_N^{S_i} = \frac{S_i(2+N)}{nR} - 1 \quad for \ N \in \{n, 2n\}$.
and

Hypothesis 4:
$\beta_{Bij} > 0 \ \forall \ \{S_i, S_j\}$: In the second stage, the likelihood of a player choosing the risky option is increasing in the strength of the information signal, $X - X_N^{S_i}$.

3.3. Alternative Explanations

These theoretical predictions rely on the assumptions that players are as follows: (i)risk neutral, (ii) ambiguity neutral, and (iii) have uninformed prior beliefs. In order to explore the validity of these assumptions, first-stage and second-stage choices from the bandit game (with and without information) are pooled with those in the risk preference MPL (with and without uncertainty). This within-subject variation permits joint estimation of structural parameters of risk preference and prior beliefs, using maximum likelihood methods.[25] Joint estimation of the parameters of the prior belief distribution and the utility function avoids potential misspecification bias.[26]

There is substantial evidence of heterogeneity in the risk preference of individuals (Holt and Laury, 2002; Stott, 2006; Andersen *et al*., 2008). Hence, assume player i's preferences are given by the popular CRRA utility function $U_i(\pi_t) = \frac{\pi_t^{1-r_i}}{1-r_i}$, where π_t is the return from the S_i player's choice in stage t.[27] The parameter r_i measures the risk preference of the player, where $r_i = 0$ implies risk neutrality; $r_i > 0$ implies risk aversion; and $r_i < 0$ implies risk-loving preferences.

There is also evidence that preferences are not linear in probabilities (Camerer and Ho, 1994; Prelec, 1998; Stott, 2006; Tversky and Kahneman, 1992; Wu and Gonzalez, 1996). Probability weighting is also a popular way of modeling ambiguity aversion (Einhorn and Hogarth, 1985; Schmeidler, 1989; Viscusi, 1989).[28] Therefore, the following probability weighting function proposed by Tversky and Kahneman (1992) is incorporated to allow for more flexibility in preferences:[29]

$$\omega(p_t) = \frac{\left(p_t^X(1-p_t)^{N-X}\right)^{\gamma}}{\left\{\sum_{X=0}^{N}\left(p_t^X(1-p_t)^{N-X}\right)^{\gamma}\right\}^{\frac{1}{\gamma}}}, \quad (7)$$

where p_t is the expected probability of success, given any available information, and γ represents the curvature parameter. For $0 < \gamma < 1$ $(\gamma > 1)$ respondents overweight (underweight) small probabilities and underweight (overweight) large probabilities. Weighted and objective probabilities are identical for $\gamma = 1$.

Finally, stochastic decision error is introduced to account for player's ability to make mistakes. The Fechner (1860/1966) model of stochastic choice assumes each player maximizes their stochastic subjective expected utility, such that the probability the risky option is chosen in stage t can be written as

$$Pr\left[\sum_{X=0}^{n}\{\omega(p_t)U_i(y_t)\} + EUMS_{it} - U_i(S_i) > \in_i\right], \quad (8)$$

where $y_t = XR$ is the return to player S_i from choosing the risky option in stage t, $\in_i$ is a stochastic noise parameter to account for decision error, and $EUMS_{it}$ is the subjective expected utility of the message service in stage t, which is

$$EUMS_{i1} = \sum_{X=0}^{n}\left\{\omega(p_1)\max\left[\sum_{X=0}^{n}\omega(p_2)U_i(y_2), U_i(S_i)\right]\right\} - \max\left[\sum_{X=0}^{n}\omega(p_1)U_i(y_2), U_i(S_i)\right] \quad (9)$$

for treatments $\{\$4, \ \$8\}$ and $\{\$8, \ \$12\}$, and

$$EUMS_{i1} = \sum_{X=0}^{2n}\left\{\omega(p_1)\max\left[\sum_{X=0}^{n}\omega(p_2)U_i(y_2), U_i(S_i)\right]\right\} - \sum_{X=0}^{n}\left\{\omega(p_1)\max\left[\sum_{X=0}^{n}\omega(p_2)U_i(y_2), U_i(S_i)\right]\right\} \quad (10)$$

for treatments $\{\$8, \ \$4\}$ and $\{\$12, \ \$8\}$, and zero in the second stage.[30] Alternatively, the Luce (1959) model of stochastic choice assumes the probability player i chooses the risky option in stage t is

$$\frac{\left(\sum_{X=0}^{n}\omega(p_t)U_i(y_t)+EUMS_{it}\right)^{\frac{1}{\varepsilon_i}}}{\left(\sum_{X=0}^{n}\omega(p_t)U_i(y_t)+EUMS_{it}\right)^{\frac{1}{\varepsilon_i}}+U_i(S_i)^{\frac{1}{\varepsilon_i}}}. \quad (11)$$

Estimation of both models avoids making inferences based on what Wilcox (2007) refers to as a ‘stochastic identifying restriction’.[31] The theoretical predictions assume players are risk neutral, $r_i=0$, do not distort probabilities, $\gamma=1$, and have uninformed prior beliefs such that $\alpha=\beta=1$, implying $p_0=\frac{1}{2}$. We explore the validity of these assumptions in the following analysis of the experimental results.

4. RESULTS

4.1. Analysis of First-Stage Decisions

Because players repeat the game over 20 rounds, it is worthwhile to see how the proportion of risky choices in each treatment varies over the course of the experiment. Figure 2 shows the mean proportion of risky choices in each round by treatment. The {\$4, \$8} and {\$12, \$8} trends are quite stable across rounds and are close to the predictions of one and zero, respectively. Treatments {\$8, \$4} and {\$8, \$12}, however, reflect much more volatility and are quite far from the predictions of zero and one, respectively.[32] Nevertheless, the vertical alignment is consistent with Hypothesis 1. The higher volatility in the {\$8, \$12} and {\$8, \$4} treatments is consistent with the lower salience.

To conduct a rigorous test of the predictions, we estimate the parameters of the model in Equation (4), controlling for player fixed-effects. Table 1 reports the regression results for linear probability models estimated via ordinary least squares.[33] Model (1) simply includes a set of binary predictors representing treatment effects. Thus, these estimates represent the raw sample proportions for each treatment. Model (2) estimates these treatment effects while controlling for player-specific fixed-effects. The results from the panel model are consistent with the results from the pooled regression. Table 2 reports tests of the hypotheses stated in the previous section using the regression models reported in Table 1. We focus on the versions of the Nash predictions which account for player error, as the strict

Table 1. Regression Results for First Stage Choices

Treatment	Ordinary least squares	Fixed-effects
{\$4, \$8}	0.946***	0.942***
	(0.019)	(0.022)
{\$8, \$12}	0.467***	0.469***
	(0.055)	(0.036)
{\$8, \$4}	0.221***	0.225***
	(0.040)	(0.026)
{\$12, \$8}	0.021	0.019
	(0.017)	(0.025)
R^2	0.501	0.615

The data consists of a panel of 52 players over 20 decision rounds (1040 observations). Robust standard errors are reported in parentheses. Statistical significance of t-tests that the estimated coefficient is zero is indicated by asterisks.
***Significant at 1% level.

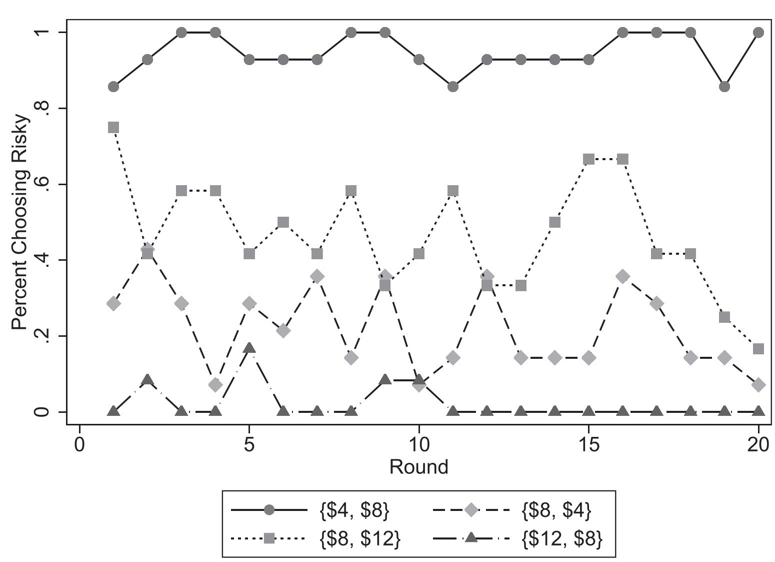

Figure 2. Proportion of risky choices in first stage across rounds by treatment.

Table 2. Hypothesis Test Results for First-Stage Choices

Hypothesis		Ordinary least squares	Fixed-effects
1a:	$H_O\colon \beta_{4,8}-\beta_{8,12}=0$	0.480***	0.473***
		(0.054)	(0.051)
1b:	$H_O\colon \beta_{8,12}-\beta_{8,4}=0$	0.245***	0.244***
		(0.052)	(0.049)
1c:	$H_O\colon \beta_{8,4}-\beta_{12,8}=0$	0.180***	0.186***
		(0.046)	(0.064)
2:	$H_O\colon 1-\beta_{8,12}-\beta_{8,4}=0$	0.354***	0.345***
		(0.077)	(0.084)

Columns (1) and (2) correspond to the models estimated in Table 2. The table reports the observed difference in the estimated coefficients from Table 2. Standard errors are reported in parentheses and statistical significance is indicated by asterisks.
***Significant at 1% level.

predictions are rejected in all cases. We summarize the results of the hypothesis tests regarding first-stage decisions subsequently.

Result 1. Both models support the full set of Nash equilibrium predictions implied by Hypothesis 1: {\$4, \$8} players choose the risky option more often than {\$8, \$12} players; {\$8, \$12} players choose the risky option more often than {\$8, \$4} players; and {\$8, \$4} players choose the risky option more often than {\$12, \$8} players.

Result 2. Both models support the under-experimentation Hypothesis 2 that the {\$8, \$4} error rate is smaller than the {\$8, \$12} error rate.

The results of the regression analysis are consistent with the theoretical predictions, accounting for decision error and the cost of such error. Random decision error in a binary decision setting, such as the bandit problem, will only lower (raise) the sample proportion away from one (zero). Smith and Walker (1993) demonstrate the importance of salience in the presence of random decision error. For a given cost of making a decision, here captured by the variance of the random error term, the probability of making the optimal decision is increasing in the difference in payoffs. Hence, because treatments {\$8, \$4} and {\$8, \$12} are the least salient (i.e., they have the lowest opportunity cost of suboptimal behavior), these exhibit the largest error rates and the most volatile behavior. Overall, error rates in first-stage choices exhibit a bias toward under-experimentation. Still, the results support the central hypothesis that players experiment strategically; there is a significant reduction in experimentation when players anticipate they can free-ride on another players' information.

4.2. Analysis of Second-Stage Decisions

Analyzing how players respond to information in their second-stage choices can provide some insight into the validity of the assumption players employ Bayesian updating in their valuation of information and may help explain the under-experimentation observed in the first stage of the game. Figure 3 shows the proportion of players from each treatment that chooses the Bayesian expected income maximizing choice in each round. Although it does not appear that subjects improve their performance relative to the Bayesian prediction over the course of the experiment, there is little room for improvement. Overall, subjects choose the Bayesian expected income maximizing choice 82.9% of the time. This proportion varies across treatments with subjects choosing the Bayesian prediction 89.1% of the time in the {\$12, \$8} treatment, 83.9% of the time in the {\$4, \$8} treatment, 81.3% of the time in the {\$8, \$12} treatment, and 77.9% of the time in the {\$8, \$4} treatment. These differences are consistent with differences in the expected net benefits of choosing the Bayesian prediction.

Figure 4 depicts the proportion of players choosing the risky option according to the number of successes observed, for the cases where one or both players chose to experiment with the risky option in the first stage. In the left panel, exactly one player in a pair experiments and in the right panel, both players in a pair chose to experiment.[34] Figure 4 reveals in each treatment the proportion of players choosing the risky option is increasing in the number of observed successes. Bayesian updating, however, generates sharp predictions regarding responses to information.

Theoretically, a player should choose the risky option if a critical number of successes are observed, $X \geq X_N^{\$S_i}$. In either panel of Figure 4, the largest increase in the proportion of players choosing the risky option occurs after observing the critical number of successes, for each treatment. The largest increase in the {\$4, \$8} treatment occurs after observing $X_3^{\$}4 = 1$, an increase of 40% (left panel), and after observing $X_6^{\$}4 = 2$, an increase of 66% (right panel). The largest increase in the {\$8, \$4} and {\$8, \$12} treatments occurs after observing $X_3^{\$}8 = 2$ when a single player chose the risky option, increases of 42% and 32% (left panel), respectively. The same is true for the {\$8, \$4} treatment when both players chose the risky

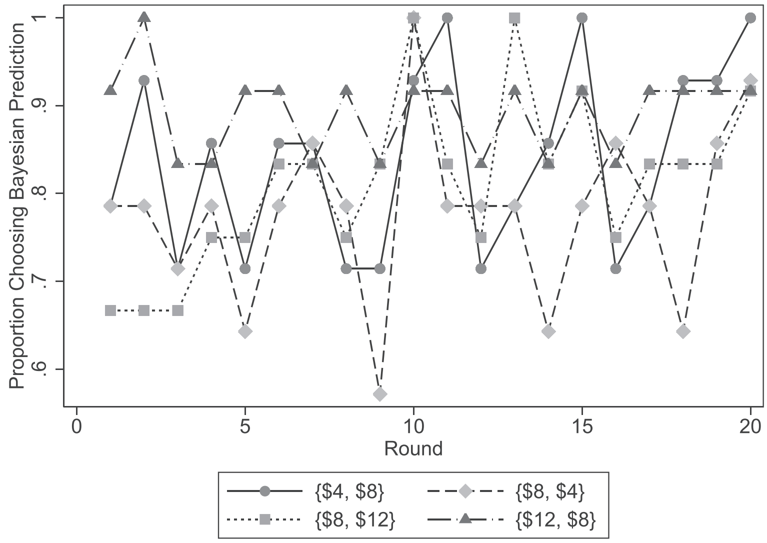

Figure 3. Proportion of Bayesian expected income maximizing choices across rounds by treatment.

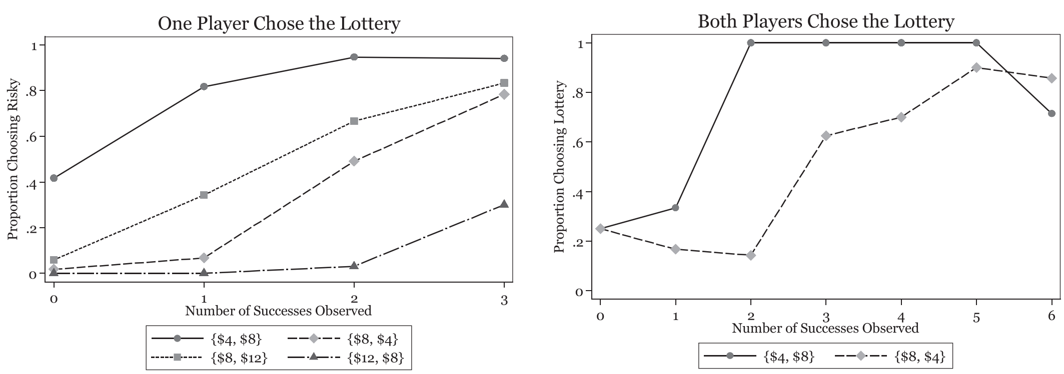

Figure 4. Proportion of risky choices by number of successes observed across treatments.

option; there is an increase of 48% (right panel) after observing $X_6^{\$8} = 4$. Finally, the only significant increase in the {$12, $8} treatment occurs after observing $X_3^{\$12} = 3$, an increase of 27% (left panel).[35]

From the aforementioned analysis, the proportion of players choosing the risky option in the first stage was highest in the {$4, $8} treatment, followed by {$8, $12} and then {$8, $4}, and lowest in the {$12, $8} treatment. Figure 4 reveals a similar pattern in the second-stage choices, regardless of the information signals. This represents persistence in player behavior.

Table 3 reports the regression results for linear probability model estimates of Equation (6) using ordinary least squares. We estimate two models that include three types of variables: (i) a set of indicator variables which is equal to one when the critical number of successes is observed, $X \geq X_N^{\$S_i}$, representing pure

Table 3. Regression Results for Second-Stage Choices

	Ordinary least squares	Fixed-effects
Bayes		
{$4, $8}	0.451*** (0.039)	0.491*** (0.041)
{$8, $4}	0.358*** (0.055)	0.322*** (0.058)
{$8, $12}	0.253*** (0.075)	0.240*** (0.073)
{$12, $8}	0.081 (0.091)	0.030 (0.098)
Salience		
{$8, $4}	0.028 (0.024)	0.017 (0.025)
{$8, $4}	0.090*** (0.020)	0.095*** (0.021)
{$8, $12}	0.156*** (0.029)	0.153*** (0.030)
{$12, $8}	0.105*** (0.017)	0.119*** (0.019)
Chose Risky	0.217*** (0.048)	0.160*** (0.055)
Constant	0.212*** (0.034)	0.244*** (0.038)
R^2	0.476	0.510

Columns (1) and (2) correspond to the models estimated in Table 2. The table reports the observed difference in the estimated coefficients from Table 2. Standard errors are reported in parentheses and statistical significance is indicated by asterisks.
***Significant at 1% level.

Bayesian predictions for each treatment; (ii) the value $X - X_N^{S_i}$ to capture the salience of the information signal by treatment;, and (iii) an indicator variable that is one if a player chose the risky option in the first stage, to account for the observed tendency to repeat that decision. The second column reports panel regression results that add player-specific fixed-effects to the pooled regression reported in the first column.

The regression coefficients are interpreted as marginal effects given the linear specification.[36] The first set of coefficients is the average increases in the proportion of players choosing the risky option when the Bayesian criterion is satisfied. The regressions indicate that the proportion of players choosing the risky option when the Bayesian criterion is satisfied increases significantly in all but the $\{\$12,\ \$8\}$ treatment. The results also suggest an increase in the strength of the information signal further encourages players to choose the predicted option, as the salience effects are positive and significant for all cases but the $\{\$4,\ \$8\}$ treatment. Players also show a significant propensity to stick with their first-stage choices. Nonetheless, the results support the Bayesian prediction that the probability of choosing the risky option in the second stage increases after observing the critical number of successes. We summarize the results of the hypotheses tests regarding second stage decisions below.

Result 3. Both models support the Bayesian hypothesis that the probability of choosing the risky option in the second stage increases after observing the critical number of successes. Three out of three β_{Aij} are statistically greater than zero.

Result 4. Both models support the hypothesis that salience has an effect upon deviations from the Bayesian-Nash predictions in the second stage. Three out of four β_{Bij} are statistically greater than zero.

This suggests the under-experimentation observed in first-stage decisions is not due to an inability to employ Bayesian updating. Still, controlling for both Bayesian updating and the salience of the information signal, we observe a significant tendency to repeat first-stage choices in the second stage. Hence, we explore whether player preferences and/or prior beliefs can explain these observed deviations from the theoretical predictions.

4.3. Estimation of Preference and Prior Belief Parameters

The estimation results associated with Equations (8) and (11) are shown in the columns 1–4 and columns 5–8 of Table 4, respectively. Columns 1 and 5 report the estimation results from the risk preference elicitation MPL. While neither models strongly rejects risk neutrality, the estimated CRRA parameter changes sign across models. Moreover, only the Luce error specification suggests significant probability distortion. The results from the MPL data suggest the estimated parameter values are sensitive to the error specification.[37] The remaining columns report the results from the joint maximum likelihood estimation of preference and prior belief parameters.

Columns 2 and 6 report the uncalibrated estimates obtained from first-stage and second-stage choices (with and without information) from the bandit game. Columns 3 and 7 report the estimates calibrated by the MPL data. These were obtained by combining the MPL data with the bandit data (with and without known probabilities) to obtain calibrated parameter estimates of the prior belief distribution. Finally, columns 4 and 8 also report calibrated parameter estimates, allowing for differential weighting of known and unknown probabilities.[38] While the CRRA parameter is significantly different from risk neutrality in the Luce error models, it is in the direction of risk seeking preferences. Hence, it does not appear that risk aversion is the culprit of the observed lack of experimentation in the first stage of the bandit game. Likewise, there is not much evidence that probability weighting played a large role either. In particular, the results in columns 4 and 8 suggest only the known probabilities, in the MPL, were substantially weighted.[39] As with the previous analysis, decision error seems to be a significant component to the observed deviations from theoretical predictions, as the estimated error parameter is always significantly different from unity.

Finally, the table reports the estimated parameters of a beta distribution of prior beliefs. All of the Fechner error models suggest beliefs were stronger than an uninformed prior (i.e., $\alpha > 1$ and $\beta > 1$). This is consistent with the observed tendency to repeat first-stage choices in the second stage of the bandit game. While the Luce models also indicate beliefs were stronger than an uninformed prior, only one parameter is significantly different from unity. Combining these parameters provides an estimate of the prior belief of success, $p_0 = \frac{\alpha}{\alpha+\beta}$. The Fechner models imply priors were significantly pessimistic, although the magnitude of this pessimism is not terribly large. One the other hand, the degree of pessimism, when significant, is larger in the Luce models. However, the estimated prior is much more sensitive to calibration and the restriction of equal probability weighting in the Luce model.

Table 4. Joint Maximum Likelihood Estimation of Structural Parameters

	Luce error				Fechner error			
Parameter	(1) MPL data	(2) Bandit data	(3) Combined restricted	(4) Combined unrestricted	(5) MPL data	(6) Bandit data	(7) Combined restricted	(8) Combined unrestricted
CRRA	−0.451*	−0.611***	−0.162	−0.296**	0.127	0.011	0.080	0.065
$H_0: \rho = 0$	(0.262)	(0.178)	(0.107)	(0.117)	(0.111)	(0.110)	(0.074)	(0.072)
Weighting	0.481***			0.538***	0.892			0.776**
Known probability	(0.075)			(0.044)	(0.293)			(0.087)
		0.905		0.953		1.051		1.050*
Unknown probability		(0.124)		(0.030)		(0.035)		(0.028)
			0.672***				0.942	
$H_0: \gamma = 1$			(0.039)				(0.059)	
Error	0.207***	0.322***	0.237***	0.248***				
$H_0: \mu = 1$	(0.052)	(0.040)	(0.028)	(0.027)				
					2.203***	2.365***	2.156***	2.129**
					(0.078)	(0.460)	(0.341)	(0.320)
Prior		1.006	1.194	1.351		1.919**	1.763**	1.990***
$H_0: \alpha = 1$		(0.207)	(0.272)	(0.221)		(0.360)	(0.384)	(0.378)
$H_0: \beta = 1$		1.738***	1.138	1.964***		2.246***	2.028***	2.272***
		(0.263)	(0.277)	(0.280)		(0.392)	(0.401)	(0.399)
$p_0 = \frac{\alpha}{\alpha+\beta}$		0.366***	0.512	0.408***		0.461**	0.465**	0.467**
$H_0: p_0 = 0.5$		(0.025)	(0.026)	(0.020)		(0.018)	(0.015)	(0.016)
Observations	520	2080	2600	2600	520	2080	2600	2600
Log likelihood	−240.060	−827.661	−1089.077	−1070.283	−237.390	−838.733	−1079.796	−1076.592

Robust standard errors are reported in parenthesis. All estimates were compared with the values reported in the first column, with significance levels of the z-tests indicated by asterisks.
*Significant at 10%; **5%; ***and 1%.

Overall, neither risk preference nor probability weighting appears to be the reason for the lack of experimentation observed in first-stage choices of the bandit game. When significantly different from neutrality, the estimated CRRA parameter implies risk-seeking preferences. There is evidence of weighting known probabilities, as the estimated parameter value is consistent with those previously reported (Tversky and Kahneman, 1992; Camerer and Ho, 1994; Wu and Gonzalez, 1996 Prelec, 1998). Also consistent with previous findings (Budescu *et al*., 2011), there is no evidence of significant weighting of unknown probabilities. Rather, the lack of experimentation in the bandit game appears to be mainly the result of decision error and slightly pessimistic priors, as these parameter estimates are significantly different from the values assumed in the theoretical framework. Pessimistic priors are not only consistent with the under–experimentation observed in first stage decisions, but also provide an explanation for the observed tendency to repeat first-stage choices in the second stage of the game.

5. DISCUSSION AND CONCLUSIONS

We began with the question: how does an information spillover in an armed bandit problem affect players' decisions to experiment? Theoretically, the presence of a spillover creates an incentive to free-ride on the information generated by others. Still, it remains an empirical question as to whether people actually engage in such strategic experimentation. As such, this paper reports the results of an experiment designed to determine whether players free-ride in a bandit problem with information spillover.

This is a fairly common decision setting, and an understanding of the information provided by players' decisions can inform policy debates concerning the public reporting of success rates for entrepreneurship, outsourcing, performance of new technology, and so on. In such settings, some individuals have private incentives to take risks, but in doing so, their behavior can inform others of the likelihood of a good (or bad) payoff. While our results indicate that individuals rationally free-ride on the information provided by others in such settings, they experiment less than predicted resulting in information shortages. Policies directed toward encouraging risk taking—such as policies encouraging the adoption of new technology—could mitigate the effects of individual aversion to experimentation. Moreover, policies that encourage public reporting of information in such settings could improve the efficiency of productive risk-taking activities.

Analysis of second-stage decisions suggests individuals would be responsive to such policies. Overall, subjects choose the Bayesian expected income maximizing choice 82.9% of the time. Hence, the lack of experimentation in the first stage of the bandit game does not appear to be due to how players process information. Still, the results indicate two sources of potential bias to Bayesian updating: (i) the strength of the information signal and (ii) previous choices.

To determine the likely reason for the observed tendency to experiment less than predicted, we explore the validity of the assumptions made to generate theoretical predictions in the analysis of the results, as Bayesian predictions are largely supported by second stage choices. We also find evidence that the strength of the information signal matters, however, which is consistent with behavioral models of decision error. In addition, we observe a significant tendency to repeat previous choices regardless of the information signal. Still, controlling for these factors, we observe a significant increase in the propensity to choose the risky option in the second stage of the bandit game when the number of successes observed is sufficient to make the expected net benefits of doing so profitable, as predicted by Bayesian updating.

Finally, we explore the extent to which risk preferences, probability weighting (a popular method of modeling ambiguity aversion), and/or pessimistic priors (another approach to modeling ambiguity aversion) can explain why players experimented less than predicted. We jointly estimate parameters of risk preference, probability weighting, and prior belief parameters using choices from a standard risk preference elicitation MPL in combination with the decisions in the bandit game. When significantly different from neutrality, the estimated CRRA parameter implies risk-seeking preferences. There is evidence of weighting *known* probabilities; however, we find no evidence of significant weighting of *unknown* probabilities. Estimated prior beliefs were somewhat pessimistic and stronger than an uninformed prior. Hence, the lack of experimentation in the bandit game appears to be mainly the result of decision error and slightly pessimistic priors.

Nonetheless, only one player experiments with the bandit, as predicted, about two-thirds of the time. Moreover, it is the predicted player at least 96% of the time. Thus, although players experiment with the bandit less than predicted, we still find evidence that players behave strategically, by free-riding when possible in the presence of an information spillover.

APPENDIX

EXPECTED PAYOFF CALCULATIONS

The following appendix provides detailed calculations of players' first-stage expected payoffs. We begin with the calculations of the expected values of the prior and posterior distributions of beliefs.

Beliefs

Given an uniformed prior, such that $p_1 = \frac{1}{2}$, the probabilities of observing X successes out of $N=3$ draws are

$$f_3(0) = \frac{1}{8}, f_3(1) = \frac{3}{8}, f_3(2) = \frac{3}{8}, \text{ and } f_3(3) = \frac{1}{8} \ .$$

Likewise, the probabilities of observing X successes out of $N=6$ draws are

$$f_6(0) = \frac{1}{64}, f_6(1) = \frac{3}{32}, f_6(2) = \frac{15}{64}, f_6(3) = \frac{5}{16}, f_6(4) = \frac{15}{64}, f_6(5) = \frac{3}{32}, \text{ and } f_6(6) = \frac{1}{64} \ .$$

When a player observes X successes out of $N=3$ draws, their posterior belief will be

$$p_2 = \frac{1}{5} \text{ for } X = 0, p_2 = \frac{2}{5} \text{ for } X = 1, p_2 = \frac{3}{5} \text{ for } X = 2, \text{ and } p_2 = \frac{4}{5} \text{ for } X = 3 \ .$$

Likewise, when a player observes X successes out of $N=6$ draws, their posterior belief will be

$$p_2 = \frac{1}{8} \text{ for } X = 0, p_2 = \frac{2}{8} \text{ for } X = 1, p_2 = \frac{3}{8} \text{ for } X = 2, p_2 = \frac{4}{8} \text{ for } X = 3, p_2 = \frac{5}{8} \text{ for } X = 4, p_2 = \frac{6}{8} \text{ for } X = 5, \text{ and } p_2 = \frac{7}{8} \text{ for } X = 6 \ .$$

We now turn our attention to the expected payoff calculations for each treatment. All calculations assume players choose the Bayesian expected value maximizing choice in the second stage.

{\$4, \$8} Treatment

Suppose that player $S_M = \$8$ has chosen the safe option. If player $S_L = \$4$ chooses the safe option in the first stage, they can expect to earn

$$\$4 + \frac{1}{8}\$0 + \frac{3}{8}\$5 + \frac{3}{8}\$10 + \frac{1}{8}\$15$$
$$= \$4 + \frac{1}{2}\$15 = \$11.50.$$

Otherwise, if they choose the risky option in the first stage, they can expect to earn

$$\frac{1}{8}\left(\$0 + \max\left[\$4, \frac{1}{5}15\right]\right) + \frac{3}{8}\left(\$5 + \max\left[\$4, \frac{2}{5}\$15\right]\right)$$
$$+\frac{3}{8}\left(\$10 + \max\left[\$4, \frac{3}{5}\$15\right]\right)$$
$$+\frac{1}{8}\left(\$15 + \max\left[\$4, \frac{4}{5}15\right]\right) = \$15.13.$$

Suppose that player $S_M = \$8$ has chosen the risky option. If player $S_L = \$4$ chooses the safe option in the first stage, they can expect to earn

$$\$4 + \frac{1}{8}\max\left[\$4, \frac{1}{5}\$15\right] + \frac{3}{8}\max\left[\$4, \frac{2}{5}\$15\right]$$
$$+\frac{3}{8}\max\left[\$4, \frac{3}{5}\$15\right] + \frac{1}{8}\max\left[\$4, \frac{4}{5}\$15\right]$$
$$= \$11.63.$$

Otherwise, if they choose the risky option in the first stage, they can expect to earn

$$\frac{1}{8}\$0 + \frac{3}{8}\$5 + \frac{3}{8}\$10 + \frac{1}{8}\$15 + \frac{1}{64}\max\left[\$4, \frac{1}{8}\$15\right]$$
$$+\frac{3}{32}\max\left[\$4, \frac{2}{8}\$15\right] + \frac{15}{64}\max\left[\$4, \frac{3}{8}\$15\right]$$
$$+\frac{5}{16}\max\left[\$4, \frac{4}{8}\$15\right] + \frac{15}{64}\max\left[\$4, \frac{5}{8}\$15\right]$$
$$+\frac{3}{32}\max\left[\$4, \frac{6}{8}\$15\right] + \frac{1}{64}\max\left[\$4, \frac{7}{8}\$15\right]$$
$$= \$15.13.$$

{\$8, \$4} and {\$8, \$12} Treatments

Suppose that player $S_L = \$4$ or $S_H = \$12$ has chosen the safe option. If player $S_M = \$8$ chooses the safe option in the first stage, they can expect to earn,

$$\$8 + \$8 = \$16.00.$$

Otherwise, if they choose the risky option in the first stage, they can expect to earn

$$\frac{1}{8}\left(\$0 + \max\left[\$8, \frac{1}{5}\$15\right]\right)$$
$$+\frac{3}{8}\left(\$5 + \max\left[\$8, \frac{2}{5}\$15\right]\right)$$
$$+\frac{3}{8}\left(\$10 + \max\left[\$8, \frac{3}{5}\$15\right]\right)$$
$$+\frac{1}{8}\left(\$15 + \max\left[\$8, \frac{4}{5}\$15\right]\right) = \$16.38$$

Suppose that player $S_L = \$4$ or $S_H = \$12$ has chosen the risky option. If player $S_M = \$8$ chooses the safe option in the first stage, they can expect to earn

$$\$8 + \frac{1}{8}\max\left[\$8, \frac{1}{5}\$15\right] + \frac{3}{8}\max\left[\$8, \frac{2}{5}\$15\right]$$
$$+\frac{3}{8}\max\left[\$8, \frac{3}{5}\$15\right] + \frac{1}{8}\max\left[\$8, \frac{4}{5}\$15\right]$$
$$= \$16.88.$$

Otherwise, if they choose the risky option in the first stage, they can expect to earn,

$$\frac{1}{8}\$0 + \frac{3}{8}\$5 + \frac{3}{8}\$10 + \frac{1}{8}\$15 + \frac{1}{64}\max\left[\$8, \frac{1}{8}\$15\right]$$
$$+\frac{3}{32}\max\left[\$8, \frac{2}{8}\$15\right] + \frac{15}{64}\max\left[\$8, \frac{3}{8}\$15\right]$$
$$+\frac{5}{16}\max\left[\$8, \frac{4}{8}\$15\right] + \frac{15}{64}\max\left[\$8, \frac{5}{8}\$15\right]$$
$$+\frac{3}{32}\max\left[\$8, \frac{6}{8}\$15\right] + \frac{1}{64}\max\left[\$8, \frac{7}{8}\$15\right]$$
$$= \$16.21.$$

{\$12, \$8} Treatment

Finally, suppose that player $S_M = \$8$ has chosen the safe option. If player $S_H = \$12$ chooses the safe option in the first stage, they can expect to earn

$$\$12 + \$12 = \$24.00.$$

Otherwise, if they choose the risky option in the first stage, they can expect to earn,

$$\frac{1}{8}\left(\$0+\max\left[\$12,\frac{1}{5}\$15\right]\right)+\frac{3}{8}\left(\$5+\max\left[\$12,\frac{2}{5}\$15\right]\right)+\frac{3}{8}\left(\$10+\max\left[\$12,\frac{3}{5}\$15\right]\right)+\frac{1}{8}\left(\$15+\max\left[\$12,\frac{4}{5}\$15\right]\right)=\$19.50.$$

Suppose that player $S_M=\$8$ has chosen the risky option. If player $S_H=\$12$ chooses the safe option in the first stage, they can expect to earn

$$\$12+\frac{1}{8}\max\left[\$12,\frac{1}{5}\$15\right]+\frac{3}{8}\max\left[\$12,\frac{2}{5}\$15\right]+\frac{3}{8}\max\left[\$12,\frac{3}{5}\$15\right]+\frac{1}{8}\max\left[\$12,\frac{4}{5}\$15\right]=\$24.00.$$

Otherwise, if they choose the risky option in the first stage, they can expect to earn

$$\frac{1}{8}\$0+\frac{3}{8}\$5+\frac{3}{8}\$10+\frac{1}{8}\$15+\frac{1}{64}\max\left[\$12,\frac{1}{8}\$15\right]+\frac{3}{32}\max\left[\$12,\frac{2}{8}\$15\right]+\frac{15}{64}\max\left[\$12,\frac{3}{8}\$15\right]+\frac{5}{16}\max\left[\$12,\frac{4}{8}\$15\right]+\frac{15}{64}\max\left[\$`1,\frac{5}{8}\$15\right]+\frac{3}{32}\max\left[\$12,\frac{6}{8}\$15\right]+\frac{1}{64}\max\left[\$12,\frac{7}{8}\$15\right]=\$19.52.$$

NOTES

1. Pure information spillovers are fundamentally different from the literature on information cascades and herding. The latter assumes only players' *actions* are observable, while the former assumes *outcomes* are observable.
2. Because it is a dominant strategy for players to choose the high safe option; this situation is strategically equivalent to the absence of spillover for players with a middle safe option, because they should not expect their partner to provide any information.
3. There is considerable evidence that suggests the individuals undervalue information (McKelvey and Page, 1990; Meyer and Shi, 1995; Banks *et al*., 1997; Gans *et al*., 2007; Anderson, 2012). Although Kraemer *et al*. (2006) report experimental results that subjects purchase too much information (appearing to overvalue it), Anderson (2012) demonstrates this is consistent with ambiguity aversion.
4. Hey and Panaccione (2011) refer to this as *resolute* behavior.
5. Hao and Houser (2012), Holt and Smith (2009), and Karni (2009) use Becker *et al*., (1964) mechanisms that do not require risk neutrality.
6. Andersen *et al*., (2010) suggest this approach, similar to the technique to jointly estimate risk and time preferences used by Andersen *et al*., (2008).
7. For simplicity, the theoretical framework used to establish testable hypotheses assumes risk and ambiguous neutral preferences. In the next section, we develop a structural equation that allows for more flexibility in preferences to account for risk and ambiguity aversion (seeking). The following section reports the estimation results regarding the validity of these assumptions.
8. Viscusi and O'Connor (1984) advocate the use of a beta distribution for Bernoulli processes.
9. We explore the validity of this assumption in the data analysis. Viscusi (1989) proposes weighting parameters in the updating function to account for ambiguity aversion.
10. Viscusi (1989) proposes weighting parameters in this updating function to account for ambiguity aversion.
11. Snow (2010) shows that the value of information that resolves uncertainty is positive and increasing in ambiguity aversion. Paradoxically, Anderson (2012) demonstrates that ambiguity aversion also lowers the value of information that reduces uncertainty. Anderson (2012) conducts an experimental test of these hypotheses and finds evidence supporting both predictions.
12. Pairing S_M players complicates the game by introducing a coordination problem.
13. Hence, 25% of players are S_L types, 50% are S_M types, and 25% are S_H types.
14. Detailed calculations of these expected payoffs are provided in the appendix.
15. Notice the cost of decision error for S_H and S_L players is much higher (approximately 23% and 40%, respectively) relative to S_M players (2.5–3.7%). The small difference in expected payoffs increases the likelihood of rejecting the Nash equilibrium predictions.
16. See Tversky and Kahneman (1971), Kahneman and Tversky (1972), Kahneman and Tversky (1973), Tversky and Kahneman (1973), Grether (1980), Grether (1992), El-Gamal and Grether (1995), Holt and Smith (2009), and Nyarko and Sopher (2006) for a discussion of these results.
17. Random treatment assignment under a within-subjects design allows the data analysis to control for potential subject-specific fixed-effects, such as risk preference.
18. Screen images and instructions are available from the authors.
19. Players could not proceed until they answered all questions correctly.
20. Participants are recruited by email via the lab's Online Recruitment System for Experimental Economics (ORSEE) (Greiner, 2004). The experiment is programmed and conducted with the software Z-Tree (Fischbacher, 2007).
21. The analysis controls for any subject-specific fixed-effects. The within-subjects design implies player fixed-effects are uncorrelated with the treatments effects.

22. The low cost of error in the $\{\$8, \$4\}$ and $\{\$8, \$12\}$ treatments is by virtue of the bandit problem itself. This only *reduces* the chances that behavior supports the Nash predictions.
23. Note $S_i = \$4$ players should also choose the risky option in the second stage when $N = 0$.
24. The dummy variable for the $\{\$4, \$8\}$ treatment also takes a value of one when $N = 0$ because the expected net benefits of the risky option are positive.
25. A detailed discussion of the estimation technique can be found in Harrison (2007).
26. For example, Andersen *et al*., (2008) demonstrate joint estimation of risk and time preference parameters removes the upward bias in estimated discount rates that assume risk neutrality.
27. The choice of a CRRA utility function is based on its popularity and its ability to explain behavior, 'under one specific payoff scale, constant relative risk aversion can provide an excellent fit for the data patterns' (Holt and Laury, 2002, p1652).
28. Camerer & Weber (1992) provide an excellent overview of the various approaches to modeling ambiguity aversion.
29. Budescu *et al*. (2011) investigate the effect of probability weighting in the estimation of prior beliefs. They do not find substantial evidence of bias in elicited beliefs due to probability distortion.
30. This assumes the other player chooses optimally.
31. Wilcox (2007), Harrison (2007), and Andersen *et al*., (2008) demonstrate that the main finding of Holt and Laury (2002), increasing relative risk aversion, is contingent on their choice of the Luce (1959) model with CRRA because the choice probability is invariant to the scale of payoffs.
32. There appears to be an 'end-game' effect in these two series.
33. We also verified that probit and logit specifications yield qualitatively similar results.
34. For 274 observations, $N = 0$; for 644 observations, $N = 3$; and for 122 observations, $N = 6$.
35. Risk neutrality predicts indifference at $X = X_3^{\$}12$. We never observe $X \geq X_6^{\$}12 = 6$.
36. We verified that probit specifications yield qualitatively similar results.
37. The difference in the estimated decision error parameters is due to the alternative specifications of the error structure in Equations (8) and (11). In either case, however, no error implies a parameter value of unity.
38. Tversky and Kahneman (1992) suggest the weighting function be allowed to differ across these domains.
39. The results in column 8 are consistent with the case made by Wakker (2001) for an (inverse) *S*-shaped weighting function for unknown (known) probabilities.

REFERENCES

Andersen SG, Harrison GW, Lau MI, Rutström EE. 2008. Eliciting risk and time preferences. *Econometrica* **76**: 583–618.

Andersen S, Fountain J, Harrison GW, Rutström EE. 2010. Estimating subjective probabilities, Available at: http://ideas.repec.org/p/exc/wpaper/2010-08.html

Anderson CM. 2012. Ambiguity aversion in multi-armed bandit problems. *Theory and Decision* **72**(1): 15–33.

Andreoni J. 1995. Cooperation in public goods experiments? Kindness or confusion. *American Economic Review* **85**: 891–904.

Banks J, Olson M, Porter D. 1997. An experimental analysis of the bandit problem. *Economic Theory* **10**: 55–77.

Becker GM, DeGroot MH, Marschak J. 1964. Measuring utility by a single-response sequential method. *Behavioral Science* **9**: 226–232.

Bolten P, Harris C. 1999. Strategic experimentation. *Econometrica* **67**: 349–374.

Budescu D, Abbas A, Lijuan W. 2011. Does probability weighting matter in probability elicitation? *Journal of Mathematical Psychology* **55**: 320–327.

Camerer CF, Ho T-H. 1994. Violations of the betweeness axiom and nonlinearity in probabilities. *Journal of Risk and Uncertainty* **8**: 167–196.

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

Caplin A, Leahy J. 1998. Miracle on sixth avenue: information externalities and search. *The Economic Journal* **108**: 60–74.

Charness G, Levin D. 2005. When optimal choices feel wrong: a laboratory study of Bayesian updating, complexity, and affect. *American Economic Review* **95**: 1300–1310.

DeGroot MH. 1970. Optimal Statistical Decisions, McGraw-Hill: New York.

Dixit A, Pindyck R. 1999. Investment Under Uncertainty, Princeton University Press: Princeton, NJ.

Einhorn HJ, Hogarth RM. 1985. Ambiguity and uncertainty in probabilistic inference. *Psychological Review* **92**: 433–461.

El-Gamal MA, Grether DM. 1995. Are people Bayesian? Uncovering behavioral strategies. *Journal of the American Statistical Association* **90**: 1137–1145.

Ellsberg D. 1961. Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics* **75**(4): 643–669.

Fechner G. 1860/1966. Elements of Psychophysics, Holt, Rinehart and Winston: New York.

Fischbacher U. 2007. z-Tree: Zurich toolbox for readymade economics experiments. *Experimental Economics* **10**: 171–178.

Fox CR, Tversky A. 1995. Ambiguity aversion and comparative ignorance. *The Quarterly Journal of Economics* **110**: 585–603.

Gans N, Knox G, Croson R. 2007. Simple models of discrete choice and their performance in bandit experiments. *Manufacturing & Service Operations Management* **9**: 383–408.

Gittins JC. 1979. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society, Series B* **41**: 148–177.

Greiner B. 2004. The Online Recruitment System ORSEE 2.0 - A Guide for the Organization of Experiments in Economics. *Working Paper Series in Economics 10, University of Cologne*.

Grether DM. 1980. Bayes rule as a descriptive model: the representativeness heuristic. *Quarterly Journal of Economics* **95**: 537–557.

Grether DM. 1992. Testing Bayes rule and the representativeness heuristic: some experimental evidence. *Journal of Economic Behavior and Organization* **17**: 31–57.

Guzman RM, Ventura G. 1998. A model of experimentation with information externalities. *Journal of Economic Dynamics and Control* **23**: 9–34.

Hao L, Houser D. 2012. Belief elicitation in the presence of naïve respondents: an experimental study. *Journal of Risk and Uncertainty* **44**: 161–180.

Harrison GW. 2007. Maximum Likelihood Estimation of Utility Functions using *Stata*, *Working Paper 06-12*, *Department of Economics*, *College of Business Administration*, University of Central Florida.

Hendricks K, Kovenock D. 1989. Asymmetric information, information externalities, and efficiency: the case of oil exploration. *RAND Journal of Economics* **20**(Summer): 164–182.

Hey JD, Panaccione L. 2011. Dynamic decision making: what do people do? *Journal of Risk and Uncertainty* **42**: 85–123.

Hirshleifer J, Riley JG. 1992. The Analytics of Uncertainty and Information, Cambridge University Press: New York

Hoffman RM, Kagel JH, Levin D. 2011. Simultaneous versus sequential information processing. *Economics Letters* **112**: 16–18.

Hogarth R, Einhorn H. 1990. Venture Theory: A Model of Decision Weights. *Management Science* **36**: 780–803.

Holt CA, Laury SK. 2002. Risk aversion and incentive effects. *American Economic Review* **92**(5): 1644–1657.

Holt CA, Smith AM. 2009. An Update on Bayesian Updating. *Journal of Economic Behavior and Organization* **69**: 125–134.

Kahneman D, Tversky A. 1972. Subjective probability: a judgment of representativeness. *Cognitive Psychology* **3**: 430–454.

Kahneman D, Tversky A. 1973. On the psychology of prediction. *Psychological Review* **80**: 237–251.

Karni E. 2009. A mechanism for eliciting probabilities. *Econometrica* **77**: 603–606.

Knight FH. 1921. *Risk*, *Uncertainty and Profit*, Houghton Mifflin: Boston.

Kraemer C, Nöth M, Weber M. 2006. Information aggregation with costly information and random ordering: experimental evidence. *Journal of Economic Behavior & Organization* **59**: 423–432.

Luce D. 1959. Individual Choice Behavior, John Wiley & Sons: New York.

McKelvey RD, Page T. 1990. Public and private information: an experimental study of information pooling. *Econometrica* **58**: 1321–1339.

Meyer RJ, Shi Y. 1995. Sequential choice under ambiguity: intuitive solutions to the armed-bandit problem. *Management Science* **41**: 817–834.

Nyarko Y, Schotter A. 2002. An experimental study of belief learning using elicited beliefs. *Econometrica* **70**: 971–1005.

Nyarko Y, Sopher B. 2006. On the informational content of advice: a theoretical and experimental study. *Economic Theory* **29**: 433–452.

Offerman T, Sonnemans J, Van de Kuilen G, Wakker PP. 2009. A truth serum for non-Bayesians: correcting proper scoring rules for risk attitudes. *The Review of Economic Studies* **76**: 1461–1489.

Poinas F, Rosaz J, Roussillon B. 2012. Updating beliefs with imperfect signals: experimental evidence. *Journal of Risk and Uncertainty* **44**: 219–241.

Prelec D. 1998. The probability weighting function. *Econometrica* **66**: 497–528.

Savage L. 1954. The foundations of statistics, revised and enlarged edition, New York: dover publications ed., John Wiley and Sons: New York. 1972

Schmeidler D. 1989. Subjective probability and expected utility without additivity. *Econometrica* **57**: 571–587.

Smith VL, Walker JM. 1993. Monetary rewards and decision cost in experimental economics. *Economic Inquiry* **31**: 245–261.

Snow A. 2010. Ambiguity and the value of information. *Journal of Risk and Uncertainty* **40**: 133–145.

Stott HP. 2006. Cumulative prospect theory's functional menagerie. *Journal of Risk and Uncertainty* **32**: 101–130.

Tversky A, Kahneman D. 1971. Beliefs in the Law of small numbers. *Psychological Bulletin* **76**: 105–110.

Tversky A, Kahneman D. 1973. Availability: A Heuristic for Judging Frequency and Probability. *Cognitive Psychology* **5**: 207–232.

Tversky A, Kahneman D. 1992. Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty* **5**: 297–323.

Viscusi WK. 1989. Prospective reference theory; toward an explanation of the paradoxes. *Journal of Risk and Uncertainty* **2**: 235–264.

Viscusi W, Magat WA. 1992. Bayesian decisions with ambiguous belief aversion. *Journal of Risk and Uncertainty* **5**: 371–387.

Viscusi WK, O'Connor CJ. 1984. Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers? *American Economic Review* **2**: 942–956.

Wakker PP. 2001. Testing and characterizing properties of nonadditive measures through violations of The surething principle *Econometrica* **69**: 1039–1059.

Wakker PP, Thaler R, Tversky A.1997. Probabilistic insurance. *Journal of Risk and Uncertainty* **15**: 7–28.

Wilcox NT. 2007. *Risk Aversion in Experiments*, Vol. 12 of *Research in Experimental Economics*, JAI Press: Greenwich, CT.

Wu G, Gonzalez R. 1996. Curvature of the probability weighting function. *Management Science* **42**: 1676–1690.