The Journal of Risk and Uncertainty, 27:2; 139–170, 2003

Preference Reversals and Induced Risk Preferences: Evidence for Noisy Maximization

JOYCE E. BERG
Department of Accounting, Henry B. Tippie College of Business, University of Iowa, Iowa City, IA 52242, USA

JOHN W. DICKHAUT
Department of Accounting, Carlson School of Management, University of Minnesota, 321 19th Avenue South, Minneapolis, MN 55455, USA

THOMAS A. RIETZ*
Department of Finance, Henry B. Tippie College of Business, University of Iowa, Iowa City, IA 52242, USA

Abstract

We combine two research lines: preference reversal research (Lichtenstein and Slovic, 1971) and research on lottery-based risk preference induction (Roth and Malouf, 1979). Our results are informative for both research lines. We show that inducing risk preferences in preference reversal experiments has dramatic effects. First, while our subjects still display reversals, they do not display the usual pattern of "predicted" reversals suggested by the compatibility hypothesis. By inducing risk averse and risk loving preferences, we can dramatically reduce reversal rates and even produce the opposite pattern of reversals. Our results are consistent with the assumption that subjects maximize expected utility with error. This provides evidence that Camerer and Hogarth's (1999) framework for incentive effects can be extended to include the risk preference induction reward scheme.

Keywords: preference reversal, risk preference induction, incentives, expected utility theory

JEL Classification: C91, D81

We combine the study of risk preference induction and preference reversal and generate results that prove interesting for both areas of study. First, our results suggest that Camerer and Hogarth's (1999) framework for understanding incentive effects can be extended from simple monetary payoff experiments to those that use binary lottery systems to induce risk preferences.[1] Second, preference reversals appear to be the result of preferences expressed with random error rather than systematic error due to task-specific characteristics. These results stand in contrast to work reported in Selten, Sadrieh, and Abbink (1999) where induction did not appear to affect reversals. Our study suggests that this is due to subjects being approximately indifferent between alternatives. In such cases, the opportunity cost of error is small and reversals virtually costless.

In preference reversal experiments (beginning with Lichtenstein and Slovic, 1971), subjects both compare gambles directly (in "paired choice tasks") and reveal their preference over gambles by pricing them separately (in "pricing tasks"). A reversal occurs when the

*To whom correspondence should be addressed.

implied preference ordering differs across these two tasks. If subjects have consistent preferences over gambles, are rational, understand the tasks and make no mistakes, we should see no such reversals. However, reversals occur frequently, are surprisingly robust and appear insensitive to incentives (see, for example, Grether and Plott, 1979). Moreover, reversals seem to be the result of systematic error, with substantially more "p-bet" reversals than "\$-bet" reversals.[2] Researchers argue that this pattern cannot be simple random noise.

We study reversals under a new incentive scheme where subjects are rewarded with lotteries rather than dollars. In theory, this incentive scheme allows us to induce preferences that eliminate any uncertainty about which gambles subjects should prefer. This would allow us to disentangle strictly utility-based explanations of preference reversal from other explanations that invoke task-specific preferences. But, "in theory" may not reflect actual subject behavior. Camerer and Hogarth (1999) argue that complex relationships exist between incentives and behavior in economics experiments and that polar views on incentives (that incentives never have an effect or that incentives always promote more rational behavior) are too simplistic. They show that examining incentive effects by measuring mean behavior frequently results in a different conclusion than when examining incentive effects by measuring the variance in behavior. Why? Incentives can affect both the mean and variance of subject responses, possibly in different manners and to different degrees. Though Camerer and Hogarth (1999) do not address the efficacy of the binary lottery incentives used for risk preference induction, one might suspect that similar complex relationships exist.

Our results, and their contrast with previous work by Selten, Sadrieh, and Abbink (1999), provide a first step in integrating the study of preference induction into the framework developed by Camerer and Hogarth (1999). In the preference reversal context, our results shed light on causes of preference reversal. Our results are most consistent with the idea that choices under incentives that induce risk preferences are the result of maximization with errors, what we will call "noisy maximization." In theory, inducing risk neutrality should make subjects indifferent between the gambles we use, while inducing risk aversion or risk loving preferences should result in a clear preference for one gamble or the other. On average, observed choices and prices are consistent with this. However, there is apparent noise, or response variance, that causes reversals. The rate and pattern of these reversals change with the induced risk preferences in a manner consistent with noisy maximization.

In the next section, we describe risk preference induction, preference reversals and their joint study. Then, we describe our experimental methods and procedures. Next, we present our results and end with conclusions and discussion.

1. Risk preference induction and preference reversals

1.1. Inducing risk preferences

In theory, inducing risk preferences (Berg et al., 1986; Roth and Malouf, 1979) lets experimenters test virtually any prediction that is derived from expected utility theory.[3] The experimenter does not need to know each subject's native utility function. Instead, the experimenter pays subjects using an artificial commodity that is later converted into the probability of winning a prize. If subjects make choices consistent with expected utility maximization and prefer more money to less, then this payoff scheme induces subjects to behave as if they have the (predetermined) risk preferences defined by the conversion function. If the conversion is concave (convex, linear), then risk aversion (loving, neutrality) is induced in the artificial commodity.

Early studies on induction (e.g., Berg et al., 1986; Roth and Malouf, 1979) demonstrate some success using the technique in bargaining games, individual choices and value elicitation. However, consistent with Camerer and Hogarth's (1999) arguments, the relationship between the technique and behavior appears to be complex. In two person bargaining games, Roth and Malouf (1979) discuss how knowledge of the other player's lottery prize levels and expected values affect the equilibrium achieved. In an individual choice setting, Berg et al. (1986) report a mean shift in preference that is consistent with predictions, but also note that there is variation in subject behavior. In first price sealed bid auctions, the technique predicts well only if the predicted bids under risk induction are adjusted to capture subjects' desires to win the auction (Rietz, 1993) and the technique has more predictive power at the auction price level than at the individual level (Cox and Oaxaca, 1995). In a recent review article, Berg, Dickhaut, and Rietz (2003) suggest that the performance of the technique is sensitive to the prior experience of subjects in the task as well as the magnitude of payoffs: as the opportunity cost of error increases, the performance of the technique improves. This observation is more fully examined in Prasnikar (2001), who demonstrates a strong relationship between level of monetary payoffs and performance of induction.

1.2. The preference reversal task

Figure 1 shows the timeline for the typical subject in a preference reversal experiment (for example, Lichtenstein and Slovic, 1971). There are six pairs of gambles. First, three pairs are presented to the subject who must state which gamble in each pair is preferred (in some studies, indifference is admissible). Then, prices for each gamble in all six pairs are elicited. Finally, the last three pairs are presented to the subject. The gambles in each pair have approximately the same expected value as one another. One gamble, the "p-bet," has a high probability of winning a low amount while the other, the "$-bet," has a low probability of winning a large amount. Grether and Plott (1979) and many subsequent studies add actual monetary payoffs to the task to accommodate typical economic concerns about the nature of subject incentives, wealth effects, order effects and other potential confounds.

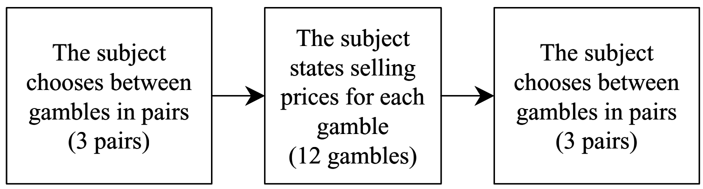

Figure 1. Timeline for the typical subject in a preference reversal experiment.

Reversals, or inconsistencies in the ranking of gambles across these two types of measurement, are typically classified as "predicted" or "unpredicted" reversals. If, within a pair, a subject chooses the p-bet over the $-bet, but prices the $-bet higher, it is called a "predicted" reversal. If the $-bet is chosen over the p-bet while the p-bet is priced higher, it is called an "unpredicted" reversal. Lichtenstein and Slovic (1971) first reported that predicted reversals significantly outnumber unpredicted reversals. Others confirmed this finding (e.g., Hamm, 1979; Lichtenstein and Slovic, 1971; Schkade and Johnson, 1989). Simply adding real dollar payoffs to the gambles has little effect (e.g., Lichtenstein and Slovic, 1973; Grether and Plott, 1979; Reilly, 1982; Berg, Dickhaut, and O'Brien, 1985). Numerous studies have sought an explanation of the phenomenon (e.g., Goldstein and Einhorn, 1987; Kahneman and Tversky, 1979; Karni and Safra, 1987; Loomes, Starmer, and Sugden, 1989; Loomes and Sugden, 1983; Segal, 1988). In addition to replicating and finding the preference reversal result, several authors have tried to eliminate it by instituting a market-like mechanism, such as arbitraging subjects (Berg, Dickhaut, and O'Brien, 1985) or employing an auction structure (Cox and Grether, 1996). Such attempts have met with some success. Selten, Sadrieh, and Abbink (1999) and our study, unlike market-based studies, are designed to examine whether preference reversal phenomena can be mitigated by attempting to reduce within and between subject variability by adopting the preference induction technique.

1.3. The preference reversal task with induced preferences

Because we induce risk preferences, our subjects choose between and price gambles with payoffs in probabilities of winning yet another binary lottery. The second lottery (the induction lottery) pays off in dollars. Induction lotteries can be designed so that all subjects should choose the same gamble in each pair during the choice task. Further, for a given gamble in the pricing task, all subjects should state the same value.

How will risk preference induction affect behavior in preference reversal experiments? We start with the conclusions of Selten, Sadrieh, and Abbink (1999) who study risk preference induction in settings where behavioral anomalies such as preference reversals are commonly observed. They conclude:

> "Our studies seem to indicate that the subjects' attitudes towards binary lottery tickets are not fundamentally different from those towards money. Both kinds of stimuli seem to be processed in a similar way. This results in similar patterns of behavior. In as far as there is a difference, it goes in the opposite direction from what would be expected on the basis of von Neumann-Morgenstern utility theory."

They draw this conclusion because they observe little or no change in traditional behavioral biases including "the reference point effect," "the preference reversal effect," and "violations of stochastic dominance" when they induce risk neutrality.

Our study takes a more systematic look at the preference reversal effect when risk preferences are induced. We frame our investigation using three potential explanations for preference reversals. The first is that subjects make errors, which creates noise that sometimes causes preference reversals. The idea is simple: subjects can make errors in either the choice task, the pricing task or both. When subjects are nearly indifferent between gambles, even small errors could produce reversals. And, for a given propensity for error, the closer subjects are to indifference, the greater the reversal rate. We call this explanation the "noisy maximization" hypothesis.

A primary objection to the noisy maximization hypothesis is the commonly observed pattern of more p-bet ("predicted") than $-bet ("unpredicted") reversals. To explain this pattern, researchers (e.g., Bostic, Herrnstein, and Luce, 1990; Tversky, Slovic, and Kahneman, 1990) have argued for the "compatibility hypothesis." According to the argument, reversals result from a tendency for the unit of assessment to influence the outcomes of judgment tasks. In pricing tasks, subjects evaluate the gambles in terms of money, which is compatible with the payoffs of the lottery and this compatibility focuses subjects on the possible payoffs of the gambles. In choice tasks, no such compatibility exists.[4] Thus, the compatibility hypothesis predicts a stronger preference for $-bets in the pricing task than in the choice task, resulting in systematic p-bet reversals.

Risk preference induction can help distinguish between these potential explanations. If induction does create the intended preferences, then we can create either known indifference across gambles or clear preferences across gambles. Inducing risk neutrality would then allow us to investigate the pattern of reversals when there is no preference across gambles. In contrast, inducing risk aversion would make all subjects prefer the lower variance p-bet. Similarly, inducing risk loving would make all subjects prefer the higher variance $-bet. If the noisy maximization hypothesis holds, moving subjects away from indifference should decrease the probability that small errors will result in reversals.

However, risk preference induction could introduce more response variance (noise), resulting in an increase in the reversal rate. Induction depends on the reduction of compound lotteries, an assumption of expected utility theory. Luce (2000, pp. 46–47), among others, argues that subject behavior does not conform to the compound lottery axiom. If so, or if subjects find the increased complexity of induction burdensome, the frequency of errors may rise. When subjects are nearly indifferent between gambles, this increase in error rate could result in increased reversals. Further, if subjects cannot reduce compound lotteries, inducing risk aversion and risk loving will not reliably create a preference for one gamble or the other. Thus, induction may actually increase reversals arising from errors.

Finally, if the compatibility hypothesis captures the underlying cause of reversals, then induction may have no effect on reversals at all. Tversky, Slovic, and Kahneman (1990, p. 211) state the compatibility hypothesis as follows:

> "The weight of any aspect (for example, probability, payoff) of an object of evaluation is enhanced by compatibility with the response (for example, choice, pricing)."

Compatibility exists in the usual pricing task because gamble payoffs are expressed in dollars and prices are also expressed in dollars. Thus, Tversky, Slovic, and Kahneman (1990, p. 211) argue compatibility implies that "payoffs will be weighted more heavily in pricing than in choice." Under preference induction, gamble payoffs are expressed in terms of "points" and prices are also expressed in terms of points. Thus, the payoff/pricing compatibility effect should be the same under risk preference induction as under monetary payoffs.

To study the competing hypotheses, we induce risk averse, risk neutral and risk loving preferences. Noisy maximization and the compatibility hypothesis make different predictions under these three conditions. If the noisy maximization hypothesis accounts for preference reversals, inducing risk neutrality should result in a high reversal rate (because the cost of error in the choice task is low when subjects are nearly indifferent between gambles), but no systematic pattern of reversals. Inducing risk aversion results in a strict preference for the p-bet, so reversal rates should decrease (the opportunity cost of error increases). Further, the reversal rate conditional on choosing the p-bet should be low (because the p-bet should be preferred) while the reversal rate conditional on choosing the $-bet should be high (reflecting corrections of errors in choice). Inducing risk loving results in a strict preference for the $-bet, so reversal rates should decrease (the opportunity cost of error increases). Further, the reversal rate conditional on choosing the $-bet should be low (because the $-bet should be preferred) while the reversal rate conditional on choosing the p-bet should be high (reflecting corrections of errors in choice).

If the compatibility hypothesis explains preference reversals, we should observe little if any effect on reversal rates or the pattern of reversals under all three conditions. However, because the risk induction technique may affect preferences in the choice task, examining the pattern of reversals will require that we examine conditional reversal rates rather than absolute reversal rates. The compatibility hypothesis predicts that the reversal rate conditional on choosing the p-bet will be higher than the reversal rate conditional on choosing the $-bet across all of our conditions.

Finally, if subjects cannot reduce compound lotteries, we should find no systematic results in our data. Error rates and, hence, reversals, could rise in all three of our settings, consistent with the apparently less rational behavior documented in Selten, Sadrieh, and Abbink (1999).

2. Methods and procedures

2.1. Task description

In our study, as in traditional preference reversal experiments, subjects make 18 decisions: 3 paired choice decisions, followed by 12 pricing decisions, followed by the 3 remaining paired choice decisions. Traditionally, when the experiment incorporates monetary rewards, one of the 18 decisions is picked randomly and payoffs are determined by the outcome of that decision. In our experiments, each decision results in a payoff. The subject receives points based on the outcome of the decision and those points determine a lottery that is played immediately for a monetary prize.[5]

The payoffs for gambles in our study are stated in units of an artificial commodity, points, which have no value outside the experiment. We use six pairs of gambles, where gambles within a pair have the same expected point payoff. These gambles are shown in Table 1.

Value is induced on points via a transformation function (call this G) that converts points into the probability of winning the monetary prize. Subjects do not see the algebraic form of this function. Instead, as in Berg et al. (1986), the transformation function is presented to subjects in the form of a "prize wheel" that is used to determine whether they win the

Table 1. Pairs of gambles used in the experiment.

Pair	Type	Probability of winning	Points if win	Points if lose	Expected points (risk neutral certainty equivalent)	Risk averse certainty equivalent	Risk loving certainty equivalent
1	P	35/36	9	2	8.81	8.71	8.86
	$	11/36	27	1	8.94	4.09	17.33
2	P	33/36	14	2	13.00	12.13	13.43
	$	9/36	40	4	13.00	6.56	27.90
3	P	32/36	15	14	14.89	14.88	14.89
	$	12/36	36	4	14.67	7.55	26.54
4	P	30/36	23	5	20.00	16.52	21.59
	$	18/36	40	0	20.00	6.19	33.81
5	P	27/36	26	22	25.00	24.82	25.15
	$	18/36	39	11	25.00	16.89	33.11
6	P	29/36	13	3	11.06	10.01	11.74
	$	7/36	37	5	11.22	6.90	23.16

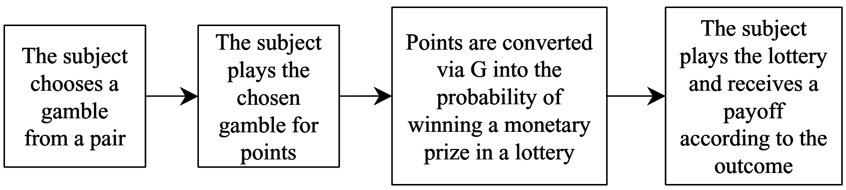

Figure 2. Timeline for the paired choice task with induced risk preferences.

monetary prize. In this way, all subjects see the transformation function before they make any decisions.

Figure 2 shows the timeline for implementing the induction procedure in a paired choice decision. Berg et al. (1986) have shown that, when this reward mechanism is used, a subject who is an expected utility maximizer will act as if he has a utility function specified by G for experimental points.[6] Wealth effects do not affect the preferences induced. As long as outcomes from choices are independent, the inducing technique will induce the same risk preferences on each choice in a sequence.

In the paired choice task, the subject is predicted to choose the gamble that maximizes the probability of winning the monetary prize given the transformation function specified by the experimenter. If the transformation function converting points to the probability of winning the preferred prize is concave, then the subject will behave as if risk averse in experimental points. If the transformation function is convex, then the subject will behave as if risk loving in experimental points. If the transformation function is linear, then the

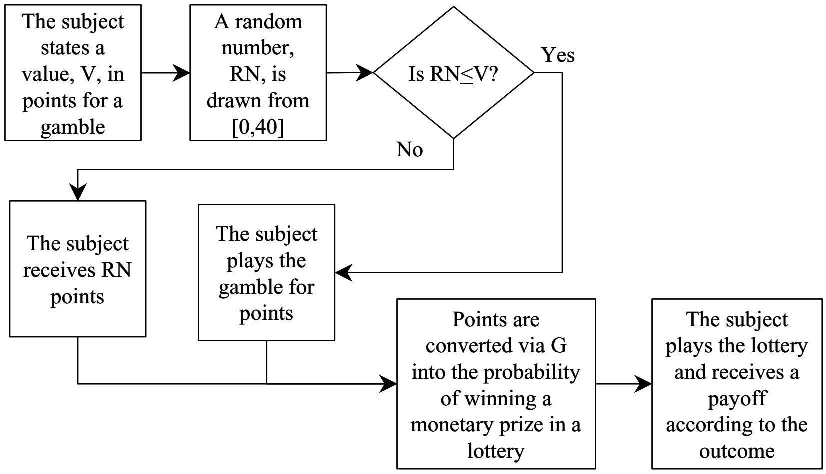

Figure 3. Timeline for the pricing task with induced risk preferences.

subject will behave as if risk neutral. Thus, our design creates three types of subjects: those that are risk averse and should prefer the p-bet, those that are risk loving and should prefer the $-bet and those that are risk neutral and should be indifferent between the two bets.

A similar modification of the traditional preference reversal task is made in the pricing decision. Figure 3 shows the timeline for a typical pricing decision with induced risk preferences.

As in the paired choice task, subjects are predicted to make a decision (in this case, state a point value) that maximizes the probability of winning the monetary prize. In theory, this stated value should be the certainty equivalent of the gamble under the utility function induced by G. Because we restrict subjects to state an integer value, subjects in our experiment are predicted to round prices to the highest integer less than the certainty equivalent of the gamble.

2.2. *Risk preferences induced*

Our design varies both the risk preferences induced and the level of monetary incentives. Subjects were induced to be either risk averse with $\mathrm{G}(w) = -\mathrm{e}^{-0.11w}$, risk loving with $\mathrm{G}(w) = \mathrm{e}^{0.11w}$ or risk neutral with $\mathrm{G}(w) = w$, where w represents the number of points earned in the task.[7] The induced certainty equivalents for each gamble are shown in Table 1. Monetary prizes were varied so that each risk group contained two subgroups: one with low monetary incentives and one with high monetary incentives. In the low incentives conditions, the risk averse group's prizes were $1, the risk loving group's prizes were $3 and the risk neutral group's prizes were $2. In the high incentives conditions, monetary prizes were doubled: the risk averse group's prizes were $2, the risk loving group's prizes were $6 and the risk neutral group's prizes were $4. Within an incentive level treatment, these prizes result in approximately equal expected payoffs across induction treatments.

Subjects in these experiments consisted primarily of undergraduate students enrolled in business school classes at the University of Iowa and the University of Minnesota. The experiments were run in several sessions with students randomly assigned to one of the six treatment groups (high and low incentives with induced risk aversion, risk loving and risk neutral preferences).

3. Predictions and results

3.1. Predictions

Because there are no wealth effects in experimental points associated with a reward structure that transforms points to lotteries, whenever the same transformation function is used in the paired choice task and the pricing task, choices in the paired choice task and stated values in the pricing task should be consistent: subjects should choose the gamble with the highest certainty equivalent in the pair and price the gambles according to their certainty equivalents. Therefore, when risk averse or risk loving preferences are induced, subjects should exhibit no preference reversals if they maximize expected utility (or other criteria linear in probabilities, as discussed in footnote 6).

Camerer and Hogarth (1999) suggest that things will not be so simple. In general, incentives have complex effects: sometimes affecting the mean response, sometimes response variance, sometimes both, sometimes in complex ways. We hypothesize that induction will display similar complexity.

First consider the noisy maximization hypothesis (that preference reversals are caused by subjects who are nearly indifferent between gambles and make errors). For our gambles, inducing risk neutrality should make the subjects nearly indifferent between gambles in a pair because those gambles have nearly identical expected values. As a result, *any* choice pattern is consistent with the induced preferences. While prices should be equal for gambles in a pair, small random errors would lead to differences in prices and, as a result, possibly high apparent reversal rates. Further, if these reversals do result from random error, then there is no reason to expect that there will be more predicted than unpredicted reversals. In contrast, inducing risk averse or risk loving preferences creates a clear difference between the gambles, so we should see fewer reversals when such preferences are induced. Moreover, those reversals should display clear patterns that differ across the two risk groups. When risk aversion is induced, a stated preference for the $-bet is truly a response error. This means that the conditional reversal rate for subjects who choose the $-bet should be high, while the conditional reversal rate for subjects who choose the p-bet should be low. The opposite is true for the risk loving group. In this case, a stated preference for the p-bet is a response error, so the conditional reversal rate should be high for p-bet choices and low for $-bet choices.

Now consider the compatibility hypothesis. Under this hypothesis, there will be a tendency to overweight point payoffs in the pricing task, resulting in $-bets being priced relatively higher. This will result in reversals, particularly if subjects are nearly indifferent in the paired choice task. So, we predict reversal rates in our risk neutral condition that are consistent with other preference reversal experiments. Because risk aversion and risk loving induce clear preferences across the gambles, we make no prediction about the overall level of reversals in those treatments. However, the compatibility hypothesis does make predictions about the pattern of reversals that we should observe. Because $-bets tend to be overpriced, we should observe higher conditional reversal rates for p-bet choices than for $-bet choices in all three of our risk preference conditions.

Finally, consider the effect of a general failure of the compound lottery axiom. Such a failure would increase response errors and, likely, increase reversal rates. Evidence of clear shifts in preference across induction treatments would be evidence against the failure of the compound lottery axiom.

3.2. Overview of results

Our evidence is most consistent with the noisy maximization hypothesis. Although inducing risk neutral preferences results in reversal rates that are similar to the previous literature, the systematic pattern of more "predicted" than "unpredicted" reversals is eliminated. Instead, conditional error rates in this setting are nearly identical. Inducing risk averse preferences not only significantly decreases the reversal rate, but it also produces a pattern of reversals opposite of that predicted by the compatibility hypothesis. The conditional reversal rate is higher for $-bet choices than for p-bet choices. Inducing risk loving preferences clearly shifts preferences toward the $-bet. While reversal rates do not fall dramatically, risk loving results in more unpredicted than predicted reversals overall. Nevertheless, the conditional reversal rate for p-bet choices is higher than for $-bet choices, as predicted by noisy maximization.

3.3. Reversal rates

For comparison, we first present a subject level analysis summarizing reversal rates across subjects (as in Selten, Sadrieh, and Abbink, 1999) and then present a decision level analysis summarizing reversal rates across decisions (as in Lichtenstein and Slovic, 1971; Grether and Plott, 1979).

3.3.1. Subject level data. Table 2 shows the comparison between the Selten, Sadrieh, and Abbink (1999) data and our data when the subject is the unit of analysis.[8] Selten, Sadrieh, and Abbink classify subjects into one of four categories: (1) no reversals, (2) more predicted than unpredicted reversals, (3) more unpredicted than predicted reversals, (4) equal number of predicted and unpredicted reversals. We present our data summarized in the same way. This results in four data sets under induced risk neutrality. Our data sets differ by the level of payoffs. Selten, Sadrieh, and Abbink's differ by the amount of information available to subjects.[9] We also have two data sets each under induced risk aversion and induced risk loving, differing by the level of payoffs.

Table 2, Panel A, shows several striking differences between the sets of data. First, inducing risk aversion significantly increases the number of subjects who display no reversals

Table 2. Reversal occurrences by subject across preference induction treatments.

Panel A: Percent of subjects with no reversals, more predicted than unpredicted reversals, more unpredicted than predicted reversals and equal numbers of types of reversals across treatments

	Selten, Sadrieh, and Abbink's (1999) risk neutral treatment		Our risk averse treatment		Our risk neutral treatment		Our risk loving treatment	
Number of subjects exhibiting:	Without summary statistics	With summary statistics	High incentives	Low incentives	High incentives	Low incentives	High incentives	Low incentives
No reversals	4 (8%)	5 (21%)	10 (42%)	10 (43%)	2 (8%)	5 (21%)	3 (14%)	3 (13%)
More predicted reversals	36 (75%)	16 (67%)	8 (33%)	9 (39%)	8 (31%)	7 (29%)	3 (14%)	4 (17%)
More unpredicted reversals	3 (6%)	2 (8%)	4 (17%)	4 (17%)	12 (46%)	9 (38%)	12 (57%)	13 (57%)
Equal numbers of reversals	5 (10%)	1 (4%)	2 (8%)	0 (0%)	4 (15%)	3 (13%)	3 (14%)	3 (13%)
Total	48	24	24**	23**	26**	24**	21**	23**

**Observations with indifference in choices but different dollar prices for pairs excluded.

Panel B: χ^2 Tests for significant differences across treatments

		Selten, Sadrieh, and Abbink's (1999) risk neutral treatment	Our risk averse treatment		Our risk neutral treatment		Our risk loving treatment	
		With summary statistics	High incentives	Low incentives	High incentives	Low incentives	High incentives	Low incentives
Selten, Sadrieh, and Abbink's (1999) risk neutral treatment	Without summary statistics	3.00	15.55*	17.25*	19.15*	17.07*	27.63*	27.04*
	With summary statistics	–	5.33	5.27	12.84*	8.98*	17.42*	16.75*
Our risk averse treatment	High incentives	–	–	2.04	9.94*	3.86	10.09*	10.05*
	Low incentives	–	–	–	13.26*	6.82	13.71*	13.46*
Our risk neutral treatment	High incentives	–	–	–	–	1.85	2.11	1.54
	Low incentives	–	–	–	–	–	2.34	2.03
Our risk loving treatment	High incentives	–	–	–	–	–	–	0.09

*Significant at the 95% level of confidence.

at all (43% of subjects overall under risk aversion versus 13 and 14% under risk neutrality and risk loving preferences, respectively). Second, our data show significant reductions in predicted reversals across the board. In our data, the percentage of subjects who display more predicted than unpredicted reversals is 28% overall relative to Selten, Sadrieh, and Abbink's (1999) 72% overall. Third, inducing risk loving preferences significantly increases the number of subjects who display more unpredicted than predicted reversals (57% overall under risk loving versus 21% under risk neutrality[10] and 17% under risk aversion). The χ^2 tests in Table 2, Panel B show frequent significant differences across induction treatments (risk neutral versus risk aversion versus risk loving) and show no significant differences across the type of information presented to subjects (Selten, Sadrieh, and Abbink, 1999) or the level of incentives alone (in our experiments). Thus, inducing different preferences has a decided effect on the patterns of choice. Most striking is the reversal of the normal pattern of reversals (from mostly predicted to mostly unpredicted) as one moves from inducing risk aversion to inducing risk loving preferences.

3.3.2. Decision level data. Traditionally, psychologists and economists (e.g., Lichtenstein and Slovic, 1971; Grether and Plott, 1979) have reported overall response rates. Specifically, using individual decisions for gamble pairs as the unit of measure, they present the percentages of non-reversals, predicted reversals and unpredicted reversals.

Table 3, Panel A contrasts results from the Lichtenstein and Slovic (1971) experiment with incentives, the Grether and Plott (1979) experiments with incentives, Selten, Sadrieh, and Abbink's (1999) experiments and ours. Again, there are striking differences. First, the overall reversal rate (the sum of the "predicted" and "unpredicted" reversal rates) across the Lichtenstein and Slovic and Grether and Plott treatments remains remarkably constant (ranging from 33 to 37%) across differences in value elicitation methods. The Selten, Sadrieh, and Abbink treatments result in slightly lower reversal rates (from 21 to 33%). The reversal rates in our risk neutral and risk loving treatments fall within the range of the other research (ranging from 24 to 36%). However, the systematic pattern of "predicted" reversals disappears when risk neutrality is induced. The pattern is reversed (with more "unpredicted" than "predicted" reversals) when risk loving preferences are induced. Our risk averse treatments result in significantly lower reversal rates (16% and 16%) compared to all other treatments.

Table 3, Panel B presents χ^2 tests for differences in reversal rates across the studies shown in Panel A (the italicized cells indicate comparisons between treatments in the same study). These tests show no significant differences as a result of implementing incentive compatible monetary incentives (Grether and Plott, Experiment 1), changing the means of eliciting values (Grether and Plott, Experiment 2), changing from monetary incentives to induced risk neutrality (Selten, Sadrieh, and Abbink monetary versus binary lottery treatments) or changing information structure (Selten, Sadrieh, and Abbink summary statistics versus none). In our data, there is only one significant difference between high and low incentive levels (under induced risk neutrality). The primary within study effect in our data is the clear significance of changing from induced risk aversion, to risk neutrality, to induced risk loving preferences.

Table 3. Reversal rates across risk treatments.

Panel A: Rates of no reversals, predicted reversals, unpredicted reversals and indifference across gamble pairs

Experiment		Percentage of choices				
		No reversal (%)	Predicted reversals (%)	Unpredicted reversals (%)	Indifference (%)	Total choices
Lichtenstein and Slovic (1971)	Experiment 3 (Incentives)	63	32	5	NA	84
Grether and Plott (1979)	Experiment 1 (Incentives)	62	25	8	5	276
	Experiment 2 (Selling prices)	58	23	11	8	228
	Experiment 2 (Equivalents)	58	27	9	6	228
Selten, Sadrieh, and Abbink (1999) monetary incentives	Without summary statistics	76	15	6	3	192
	With summary statistics	75	19	6	0	96
Selten, Sadrieh, and Abbink (1999) risk neutral treatment	Without summary statistics	66	21	6	7	192
	With summary statistics	66	27	6	1	96
Our risk averse treatment	High incentives	83	10	6	1	144
	Low incentives	80	11	5	4	138
Our risk neutral treatment	High incentives	53	15	20	12	156
	Low incentives	50	11	13	26	144
Our risk loving treatment	High incentives	61	8	22	8	132
	Low incentives	51	9	27	13	144

(*Continued on next page.*)

Table 3. (*Continued*).

Panel B: χ^2-Tests for differences between data sets

Study	Data set description	Data set number	2	3	4	5	6	7	8	9	10	11	12	13	14
Lichtenstein and Slovic (1971)	Experiment 3 Incentives	1**	1.94	4.31	1.85	9.64*	4.31	2.69	0.61	17.65*	14.43*	15.40*	12.23*	25.69*	29.65*
Grether and Plott (1979)	Experiment 1 Incentives	2	–	*3.82*	*0.77*	10.42*	8.19*	1.91	3.42	20.79*	15.19*	22.71*	48.13*	28.20*	44.83*
	Experiment 2 Selling prices	3	–	–	*2.03*	16.48*	13.22*	4.08	8.57*	26.21*	19.59*	8.97*	26.57*	16.93*	25.13*
	Experiment 2 Equivalents	4	–	–	–	14.81*	11.17*	3.08	4.93	25.75*	19.39*	17.67*	38.63*	25.59*	37.52*
Selten, Sadrieh, and Abbink (1999)	Monetary	5	–	–	–	–	*3.56*	*6.15*	*6.94*	3.40	1.40	29.92*	47.96*	26.07*	45.41*
	Monetary with summary statistics	6	–	–	–	–	–	*7.53*	*3.05*	5.28	6.33	23.58*	35.96*	23.23*	34.83*
	Risk neutral	7	–	–	–	–	–	–	*5.22*	15.41*	9.30*	18.97*	32.96*	24.02*	37.04*
	Risk neutral with summary statistics	8	–	–	–	–	–	–	–	12.67*	11.78*	21.88*	35.92*	26.76*	36.49*
Berg, Dickhaut, and Rietz	Risk averse high incentives	9	–	–	–	–	–	–	–	–	*1.72*	*33.52**	*47.10**	*23.86**	*42.08**
	Risk averse low incentives	10	–	–	–	–	–	–	–	–	–	*27.63**	*38.40**	*20.87**	*37.04**
	Risk neutral high incentives	11	–	–	–	–	–	–	–	–	–	–	*12.51**	*4.64*	*4.23*
	Risk neutral low incentives	12	–	–	–	–	–	–	–	–	–	–	–	*18.42**	*15.22**
	Risk loving high incentives	13	–	–	–	–	–	–	–	–	–	–	–	–	*3.13*
	Risk loving low incentives	14	–	–	–	–	–	–	–	–	–	–	–	–	–

*Significant at the 95% level of confidence.
**2 degrees of freedom for comparisons to Lichtenstein and Slovic and 3 degrees of freedom for all other comparisons.

Across studies, we see that Lichtenstein and Slovic's results do not differ from any of the Grether and Plott treatments or from three of the four Selten, Sadrieh, and Abbink treatments. However, there are also differences across studies. Some of these differences are expected. For example, we expect other treatments to differ significantly from our induced risk aversion and induced risk loving preferences. We cannot explain why Selten, Sadrieh, and Abbink's monetary treatments differ significantly from Grether and Plott's monetary treatments while their induced risk preference treatments do not.

3.3.3. Conditional reversal rates. Even more striking evidence on the effects of induction comes from examining conditional reversal rates. Table 4 gives the entire pattern of responses for our data and allows the calculation of conditional reversal rates.

In the risk neutral treatment, p-bet choices displayed reversals 27% of the time while \$-bet choices displayed reversals 37% of the time. These conditional rates do not differ significantly from each other at the 95% level of confidence (Pearson's $\chi^2(1) = 3.31$). The data in this treatment are consistent with subjects being on average indifferent across gambles. Consistent choice and pricing preference for the p-bet happens 36% of the time, while consistent choice and pricing preferences for the \$-bet happens 27% of the time.[11] In addition, indifference between the gambles is indicated by one or both metrics 19% of the time. This evidence is consistent with weak if any preferences across the bets. With weak preferences, high reversal rates can be caused by relatively small response variance.

In the risk averse treatment, p-bet choices displayed reversals 11% of the time while \$-bet choices displayed reversals 76% of the time. These conditional rates differ significantly from each other at the 95% level of confidence (Pearson's $\chi^2(1) = 61.38$). The p-bet was chosen and priced (weakly) higher than the \$-bet 82% of the time, while the \$-bet was chosen and priced (weakly) higher only 2% of the time. This evidence is consistent with preferences for the p-bet and small error rates.

In the risk loving treatment, p-bet choices displayed reversals 59% of the time while \$-bet choices displayed reversals 30% of the time. These conditional rates differ significantly from each other at the 95% level of confidence (Pearson's $\chi^2(1) = 13.03$). Consistent choice and pricing responses occurred just under 60% of the time for \$-bets, compared to a rate of 6% for p-bets. This evidence is consistent with preferences for the \$-bet and somewhat larger error rates than in the risk averse treatment.

Figure 4 shows the conditional error rates under each treatment along with confidence intervals derived from the normal approximation to the binomial distribution for each. Within treatment differences in conditional error rates are consistent with the χ^2 statistics discussed above.[12] The reversal patterns across treatments also differ significantly from each other. Consider the difference in conditional reversal rates defined as the reversal rate conditional on choosing the p-bet (i.e., predicted reversals) minus the rate conditional on choosing the \$-bet (i.e., unpredicted reversals). The difference in conditional reversal rates under risk loving (29.0%) is significantly higher than under risk neutrality (−10.2%, with a z-statistic of 3.92) and significantly higher than under risk aversion (−65.1%, with

Table 4. Choice task versus pricing task outcomes.

	P-Bet priced higher	$-Bet priced higher	Indifference in the pricing task
Panel A: Risk averse preferences			
P-Bet chosen	226	29	6
	*80.1%**	*10.3%*	*2.1%*
	(82.2%)*	(10.5%)	
$-Bet chosen	16	4	1
	5.7%	*1.4%*	*0.4%*
	(5.8%)	(1.5%)	
Indifference in the choice task	0	0	0
	0.0%	*0.0%*	*0.0%*
Panel B: Risk neutral preferences			
P-Bet chosen	95	40	12
	31.7%	*13.3%*	*4.0%*
	(38.9%)	(16.4%)	
$-Bet chosen	49	60	22
	16.3%	*20.0%*	*7.3%*
	(20.1%)	(24.6%)	
Indifference in the choice task	6	8	8
	2.0%	*2.7%*	*2.7%*
Panel C: Risk loving preferences			
P-Bet chosen	12	24	5
	4.3%	*8.7%*	*1.8%*
	(4.9%)	(9.7%)	
$-Bet chosen	68	143	19
	24.6%	*51.8%*	*6.9%*
	(27.5%)	(57.9%)	
Indifference in the choice task	1	3	1
	0.4%	*1.1%*	*0.4%*

*Percentages of total responses appear in *italics*.
**Percentages of non-indifference responses appear in parentheses ().

a z-statistic of 7.47). The difference in conditional reversal rates under risk neutrality is significantly higher than under risk aversion (−10.2% versus −65.1%, with a z-statistic of 4.98).

Thus, the conditional reversal rate evidence is consistent with the predictions of the noisy maximization hypothesis and inconsistent with the compatibility hypothesis.

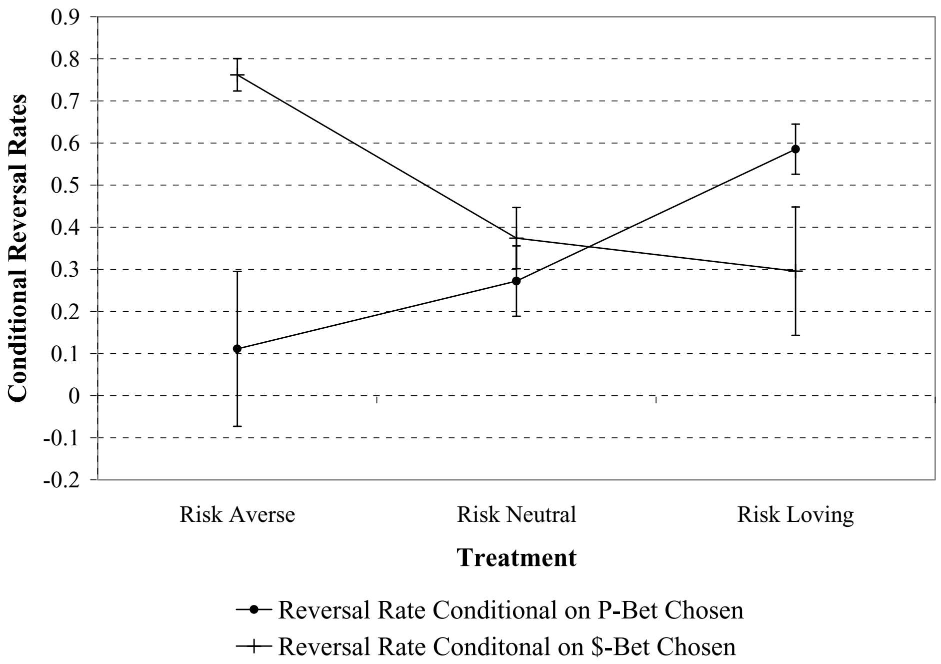

Figure 4. Conditional reversal rates across treatments and 95% confidence intervals.

3.4. Evidence from certainty equivalents

Comparing the certainty equivalents of each gamble with the prices given by subjects for the gambles provides additional evidence for the noisy maximization hypothesis. If induction worked perfectly, then the price for each gamble should equal its certainty equivalent. Figure 5 shows the 25th percentile, median and 75th percentile prices submitted for each gamble along with that gamble's certainty equivalent for each risk preference treatment.[13] With few exceptions, the graphs show distributions of prices that, on average, fall close to the induced certainty equivalents with some noise. Consistent with a higher reversal rate, the noise appears greatest under risk loving. On average, subjects over-priced gambles by 2.77 points under risk neutrality, over-priced by 1.44 points under risk aversion and under-priced by 2.00 points under risk loving preferences. The average (across gambles) standard deviation in stated prices was 5.95 points under risk aversion, 6.50 under risk neutrality and 8.57 under risk loving preferences. The average (across gambles) absolute distance between the certainty equivalent and the average price was 0.44 standard deviations under risk aversion, 0.31 standard deviations under risk neutrality 0.40 standard deviations under risk loving preferences. Thus, overall, the prices submitted by subjects are consistent with inducing the theoretical preferences with some noise.

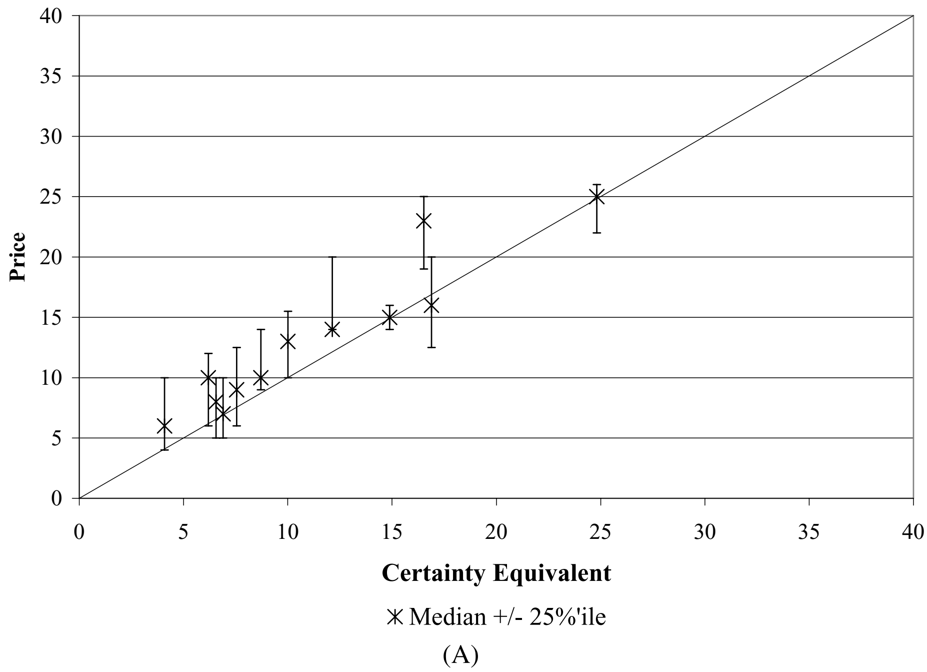

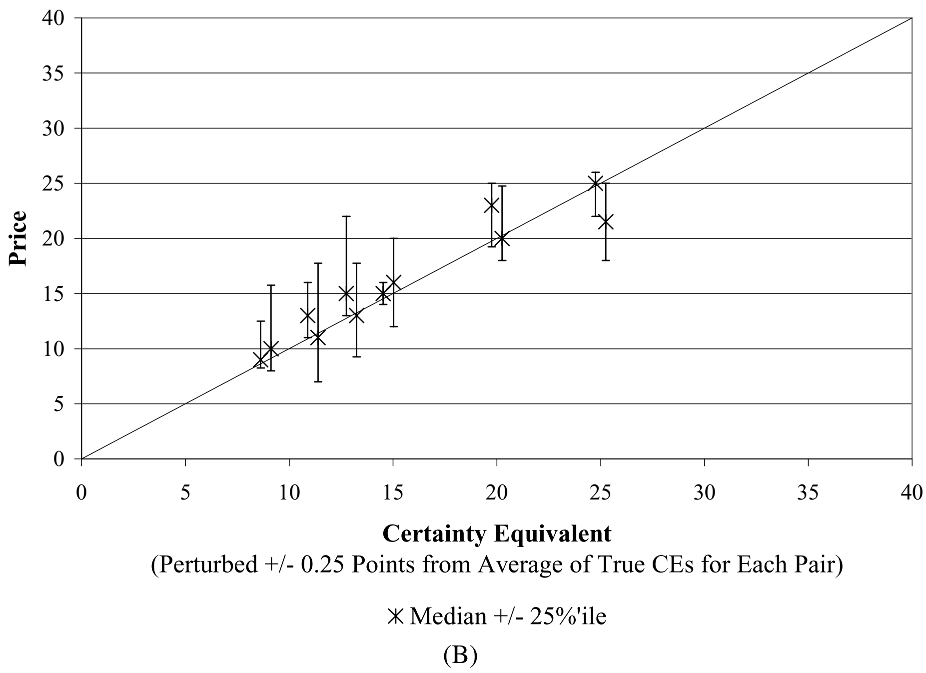

Figure 5. Certainty equivalents and prices for gambles; Panel A: Risk averse treatment, Panel B: Risk neutral treatment, and Panel C: Risk loving treatment. (*Continued on next page*.)

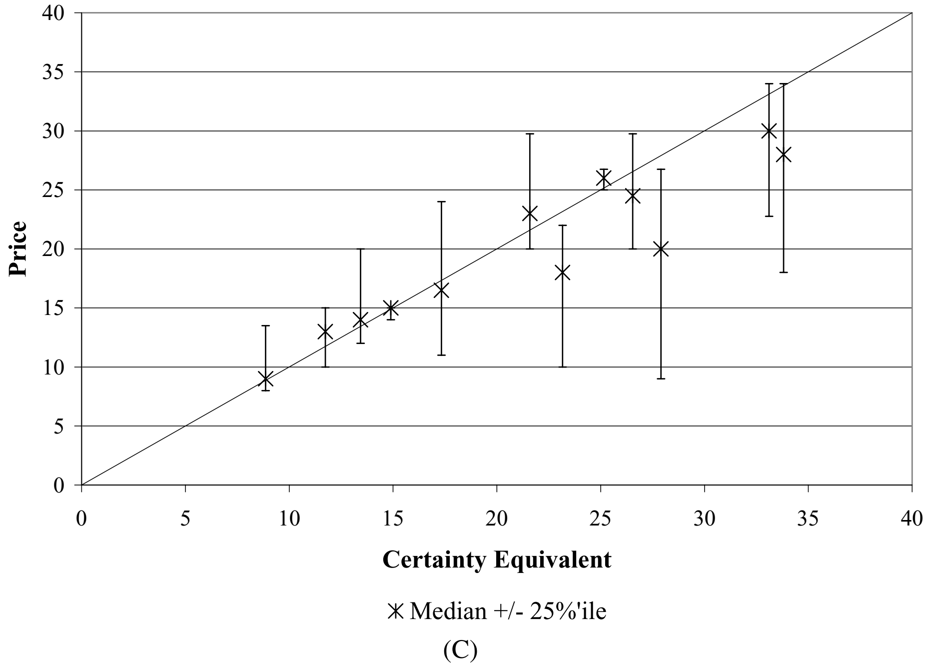

Figure 5. (*Continued*).

3.5. The two-error rate model

Perfect performance of the induction technique in the preference reversal task would imply that there would be no reversals in the risk averse and risk loving data. Table 2 shows this is not the case. The evidence presented so far suggests that individuals generally follow an expected utility maximization model but are subject to response variance or noise. Lichtenstein and Slovic (1971) first introduced this idea using a "two-error rate" formulation. However, the model failed to fit their data, leading them to conclude that the reversals they saw were systematic rather than random deviations from expected utility theory. If induction eliminates the systematic nature of errors, we should be able to use their model to explain our data.

To develop the model, let "q" represent the percentage of subjects who prefer the p-bet according to their underlying preference ordering for gambles, "r" represent the error rate in the choice task and "s" represent the error rate in the pricing task.[14] If we assume that errors in the choice task and the pricing task are random (that is, error rates do not differ across bets or subjects) and independent (that is, making an error in the choice task does not affect the probability of making an error in the pricing task), then the pattern of observations generated in a preference reversal experiment should conform to Figure 6, where a, b, c and d represent the percentage of observations that fall in each cell. The four cells represent all combinations of preferences indicated by the choice and the pricing tasks for a particular

	P-Bet Priced Higher	\$-Bet Priced Higher
P-Bet Chosen	$(q)(1\text{-}r)(1\text{-}s) + (1\text{-}q)(r)(s)$ *a*	$(q)(1\text{-}r)(s) + (1\text{-}q)(r)(1\text{-}s)$ *b*
\$-Bet Chosen	*c* $(q)(r)(1\text{-}s) + (1\text{-}q)(1\text{-}r)(s)$	*d* $(q)(r)(s) + (1\text{-}q)(1\text{-}r)(1\text{-}s)$

where:
q = percentage of subjects whose underlying preference ordering ranks the P-Bet higher
r = error rate in the paired-choice task
s = error rate in the pricing task

Figure 6. Two-error rate model.

pair of gambles: (a) the proportion of comparisons where the p-bet was both chosen and priced higher than the \$-bet, (b) the proportion where the p-bet was chosen but the \$-bet was priced higher (c) the proportion where the \$-bet was chosen but the p-bet was priced higher and (d) the proportion where the \$-bet was chosen and priced higher.

If behavior conforms to the two-error rate model, then these proportions are also functions of q, r and s as defined in Figure 6. Solving for q, r and s gives the following equations.[15]

$$q(1-q) = \frac{ad - bc}{(a+d) - (b+c)}. \tag{1}$$

$$r = (a + b - q)/(1 - 2q) \tag{2}$$

and

$$s = (a + c - q)/(1 - 2q). \tag{3}$$

There are at most two sets of parameters that satisfy these equations. In one set, the estimated percentage of subjects that prefer the p-bet is consistent with risk aversion (that is, $\hat{q} > .5$). The other set is consistent with risk loving ($\hat{q} < .5$).[16] In our analysis, we assume when we induce risk aversion (loving) q will be greater (less) than 0.5.

Table 5 shows the overall reversal rates for each data set along with estimates of q, r and s. The Selten, Sadrieh, and Abbink (1999) data differs a bit from the rest of the data because half of their p-bets have a significantly higher expected value than the \$-bets and half have a significantly lower expected value. So, we would hypothesize q's close to 0.5 for this data. The two-error-rate model estimates are consistent with this hypothesis. The estimate of q is 0.5 when $q(1-q) = (ad - bc)/(a - b - c + d) = 0.25$. In Selten, Sadrieh, and Abbink's binary lottery treatment with feedback on statistics, q cannot be estimated because $q(1-q) = 0.28$. In their other treatments q is estimated to be 0.68, 0.69 and 0.50 exactly.

Table 5. Reversal rates, estimated preference rates and estimated error rates from the two-error rate model*.

Study	Data set description	Reversal rate* (%)	Estimated q (Preference for p-bet)	Estimated r (Error rate in choice task)	Estimated s (Error rate in pricing task)
Lichtenstein and Slovic (1971)	Experiment 3 Incentives	37	No real root	NA	NA
Grether and Plott (1979)	Experiment 1 Incentives	35	0.88	0.68	0.92
	Experiment 2 Selling Prices	37	0.84	0.68	0.87
	Experiment 2 Equivalents	38	0.90	0.65	0.90
Selten, Sadrieh, and Abbink (1999)	Monetary	22	0.68	0.77	1.03
	Monetary with summary statistics	25	0.50	NA**	NA**
	Risk neutral	30	0.69	0.67	1.10
	Risk neutral with summary statistics	34	No real root	NA**	NA**
Berg, Dickhaut, and Rietz	Risk averse high incentives	16	0.99	0.07	0.11
	Risk averse low incentives	17	0.99	0.06	0.12
	Risk neutral high incentives	40	0.66	0.34	0.18
	Risk neutral low incentives	32	0.61	0.24	0.15
	Risk loving high incentives	33	0.07	0.11	0.28
	Risk loving low incentives	41	−0.10***	0.19	0.36

*For experiments in which subjects were permitted to respond "indifferent," percentage reversals is calculated using only non-indifference responses.
**Estimates for r and s are undefined when the estimated q is 0.50.
***Sampling error can produce negative estimated probability values outside of the valid zero to one range when the true q is close to the limits.

In our risk neutral treatment, we would expect q to be close to 0.5 if subjects choose randomly because the gambles have approximately the same expected value (a maximum difference of \$0.011 after taking the induction lottery into account). As in the Selten, Sadrieh, and Abbink data, subjects show a slight preference for the p-bet, with estimated q's of 0.66 and 0.61.

In our risk averse and risk loving treatments, the model's estimates suggest that subjects have a strong preference for one type of bet or another (all the estimates of q are near 0 or 1). If the induction technique worked perfectly, q should be 1 for risk averse preferences (in our data it is estimated at 0.99 and 0.99) and q should be 0 for risk loving preferences (in our data it is estimated at −0.10 and 0.07).[17]

The implied task error rates are striking.[18] With induced risk neutrality, the subjects are essentially indifferent between gambles. We would expect any randomness due to errors or biases to generate high reversal rates because small effects can easily swing preferences of nearly indifferent subjects between gambles. In the data without induction or when Selten, Sadrieh, and Abbink induce risk neutrality, the choice task error rate (r), the pricing task error rate (s) or both need to exceed 67% to fit the data. In stark contrast, when we induce risk neutrality on equal expected value gambles, the estimated task error rates fall dramatically (to 34% or below).[19] When we induce risk loving, estimated error rates are 36% or below. Most dramatically, when we induce risk averse preferences, both error rates fall to 12% or below.

4. Conclusions

Preference reversal tasks traditionally have been used to demonstrate the failure of expected utility theory. We show that, by altering the incentive mechanism, it is possible to alter the preference reversal phenomenon, providing data that suggests subjects maximize expected utility with errors. This is accomplished by using the risk preference induction technique, which itself depends on expected utility maximization (or at least the ability to reduce compound lotteries).

What is the best interpretation of these results along with those from the prior literature? We believe that extending Camerer and Hogarth's (1999) ideas to inducing preferences provides an explanation. Camerer and Hogarth differentiate the effects of incentives on mean behavior from the effect on variance of behavior. A similar distinction helps differentiate the effects of induction. The intended effect of induction is to create uniformly risk neutral, risk averse or risk loving preferences. That is, induction is intended to change mean behavior. But, the induction process increases the number of steps subjects must undertake to perform the task and therefore increases task difficulty. With induction, subjects must reduce compound lotteries as well as make choices in the economic context under study. This can add noise to responses. Selten, Sadrieh, and Abbink (1999) argue that "background risk" (the increased variance of ultimate payoffs resulting from using a lottery) increases the number of expected utility violations. Again, the argument is that the lottery mechanism can add noise. In other words, induction can also change the response variance in the data.

If subjects are already approximately risk neutral, then the observable effects of inducing risk neutrality could all lie in the increase in response variance. Inducing risk neutrality *should* result in near indifference across gambles and, in fact, subjects do not display strong preferences across gambles. Subjects *should* price each gamble in a pair the same. While they do so on average, there is considerable variance. This response "noise" results in reversals at about the same rate as in prior literature. However, the usual pattern of more predicted than unpredicted reversals disappears under induced risk neutrality. Choices and reversal patterns seem consistent with purely random errors.

If one induces risk aversion in subjects who are largely risk neutral, then one should observe a change in the typical responses (i.e., a change in mean behavior) that can swamp any increase in response variance that results. That is, inducing risk aversion *should* result in a strong preference for the less risky p-bet and higher prices for the p-bet. Indeed, subjects generally choose the p-bet. While prices equal the induced certainty equivalents on average, there is noise. The strength of preference overcomes the noise on average leading to few reversals. When reversals occur, conditional rates are consistent with noise causing them. When subjects choose the less risky p-bet (which they should prefer), they seldom reverse. When they choose the more risky $-bet (an apparent mistake) they frequently reverse.

Similarly, if one induces risk loving in subjects who are largely risk neutral, then one should observe a change in the typical responses (i.e., a change in mean behavior) that can swamp any increase in response variance that results. That is, inducing risk loving *should* result in a strong preference for the more risky \$-bet and higher prices for the \$-bet. Indeed, subjects generally choose the \$-bet. While prices equal the induced certainty equivalents on average, there is noise. The strength of preference does not overcome the noise as easily as in the risk averse case. The result is a higher reversal rate. Nevertheless, when reversals occur, conditional rates are consistent with noise causing them. When subjects choose the more risky \$-bet (which they should prefer), they seldom reverse. When they choose the less risky p-bet (an apparent mistake) they frequently reverse.

Far from dismissing preference induction, as Selten, Sadrieh, and Abbink (1999) would have us do, the evidence calls for much more research to fully understand the tradeoffs and expected effects of using risk preference induction in experiments. Further study on incentive levels, types of induced risk preferences and other variations of the risk preference technique will be useful in understanding the tradeoffs involved in using the technique.

Appendix: Instructions[20]

This is an experiment in individual decision making. As a participant in this experiment, you will have opportunities to play for eighteen \$4.00 prizes. Whether or not you receive a particular \$4.00 prize will be determined by spinning the spinner on your prize wheel. If the spinner stops in the area designated as the WIN area on your prize wheel, then you will receive the \$4.00 prize. If the spinner stops in the area outside the WIN area, then you will receive nothing.

For example, suppose the WIN area of your prize wheel is designated as 0 through 5. Then, if the spinner stops on a number *less than or equal to* 5, you will receive the \$4.00 prize. If the spinner stops on a number greater than 5, you will receive nothing. Although the WIN area on your prize wheel will vary, it will always be determined by starting at zero and moving clockwise.

Now suppose that the WIN area on your prize wheel is designated as 0 through 30. Please spin the spinner to determine whether you would have received the \$4.00 prize or not.

So far, you have discovered that a spin on your prize wheel will determine whether or not you receive a \$4.00 prize. However, you need to know how the WIN area on your prize wheel is determined before you can complete the experiment. The markings on the circumference of your prize wheel denote points, and you will receive points for making decisions. There are 18 decision items in this experiment. When a decision is made, the WIN area on your prize wheel will be designated as the area between 0 and the number of points you receive as a result of the decision. Then the spinner on your prize wheel will be spun to determine whether you receive the \$4.00 prize. Points do *not* accumulate from decision to decision.

Each decision you make will involve one or more bets. These bets will be indicated by pie charts as shown below. When a bet is played, one ball will be drawn from a bingo cage that contains 36 red balls numbered 1, 2, . . . , 36. The ball drawn determines the point outcome of the bet. This point outcome will designate the upper boundary of the WIN area on your prize wheel. For example, suppose you are playing the bet below. If the red ball drawn was less than *or equal to* 10, you would receive 30 points. If the red ball drawn was greater than 10, you would receive 5 points.

Now let's use the bet shown below as a practice item.

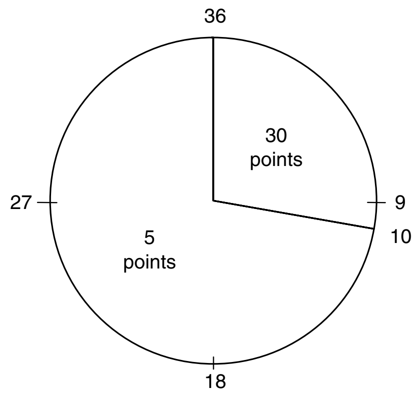

The number that the experimenter drew from the cage of red balls is ____.
This means that I would receive ____ points as a result of this bet.
Therefore the WIN area on my prize wheel is designated as 0 through ____.
Now, spin the spinner. As a result of my spin I would have received \$4.00/nothing (circle the correct word).

{*Page break.*}

Part 1:

In this part you will be asked to consider several pairs of bets. For each pair you should indicate which bet you prefer to play or indicate that you are indifferent between them. After each decision, you will have an opportunity to play for a \$4.00 prize using the following procedure:

1. The bet you indicate as preferred will be played and you will receive the points indicated by its outcome. If you check "Indifferent" the bet you play will be determined by a coin toss.
2. The WIN area of your prize wheel will be designated as the area from 0 through the number of points which you have received. You will spin the spinner to determine whether you win the \$4.00 prize.

{*Three paired choice tasks follow.*}

Part 2:

In this part you are given several opportunities to play bets to obtain points. For each bet you must indicate the smallest number of points for which you would give up the opportunity to play the bet.

After each decision, you will have an opportunity to play for a \$4.00 prize using the following procedure:

1. A ball will be drawn from a bingo cage containing 41 green balls numbered 0, 1, 2, . . . , 40. If the number on this green ball is less than *or equal to* the number you have specified, you will keep the bet and play it. You will receive the points indicated by the outcome of the bet. If the number on the green ball is greater than the number you have specified, you will give up the bet and in exchange receive the points equal to the number on the *ball*.
2. The WIN area of your prize wheel will be designated as the area from 0 through the number of points which you have received. You will spin the spinner to determine whether you win the \$4.00 prize.

It is in your best interest to be accurate; that is, the best thing you can do is be honest. If the number of points you state is too high or too low, then you are passing up opportunities that you prefer. For example, suppose you would be willing to give up the bet for 20 points but instead you say that the lowest amount for which you would give it up is 30 points. If the green ball drawn at random is between the two (for example 25) you would be forced to play the bet even though you would rather have given it up for 25 points.

On the other hand, suppose that you would give it up for 20 points but not for less, but instead you state your amount as 10 points. If the green ball drawn at random is between the two (for example 15) you would be forced to give up the bet for 15 points even though at that amount you would prefer to play it.

{*Page break.*}

Practice Item 1: Suppose you have the opportunity to play the bet shown below. What is the *smallest* number of points for which you would give up this opportunity? Remember that the WIN area on your prize wheel will be designated as the area from 0 through the number of points you receive as a result of your decision.

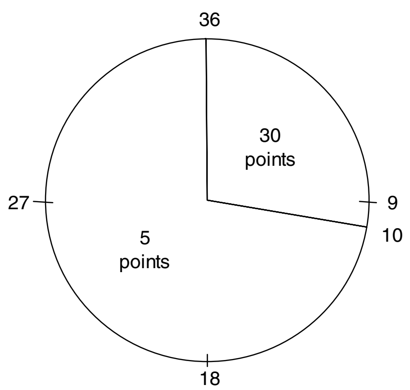

Decision ____.

In order to determine the WIN area on your prize wheel which results from this decision, you need to know two things:

(1) The ball drawn from the cage of green balls.
(2) The ball drawn from the cage of red balls.

The two examples in this practice item fix these draws so that you can concentrate on how your decision and the results of the draws will determine your WIN area.

{*Page break.*}

Example 1: Use your decision in the Practice Item 1 and suppose the green ball drawn at random is 2.

The number on the green ball is	(a) greater than my indicated amount.
	(b) less than or equal to my indicated amount.
Therefore, I would	(a) receive points equal to the number on the ball and not play the bet.
	(b) play the bet and receive points according to its outcome.

Will you be playing the bet? ____ (yes/no) If your answer is YES, you will need to know the outcome of the bet before you can determine the WIN area of your prize wheel. If your answer is NO, you do not need to know the outcome of the bet to determine the WIN area. Suppose the red ball drawn to determine the outcome of the bet was 18.

Based on my point decision above, and the results of the draws from the bingo cages, the number of points I would have is ____.

This means that the WIN area of my prize wheel would cover the numbers 0 through ____.

If the spinner stopped on the number 5, I would (circle the correct words) win/not win the $4.00 prize.

If the spinner stopped on the number 40, I would (circle the correct words) win/not win the $4.00 prize.

{*Page break.*}

Complete this page only if you would have been playing the bet to receive points.

Suppose the red ball drawn to determine the outcome of the bet was 10 instead of 18.

Based on my point decision above, and the results of the draws from the bingo cages, the number of points I would have is ____.

This means that the WIN area of my prize wheel would cover the numbers 0 through ____.

If the spinner stopped on the number 5, I would (circle the correct words) win/not win the $4.00 prize.

If the spinner stopped on the number 40, I would (circle the correct words) win/not win the $4.00 prize.

Stop here and wait for the experimenter to tell you to go on to Example 2.

{*Page break.*}

Example 2: Now use your decision in the Practice Item 1 and suppose instead that the green ball drawn at random is 38.

The number on the green ball is	(a) greater than my indicated amount. (b) less than or equal to my indicated amount.
Therefore, I would	(a) receive points equal to the number on the ball and not play the bet. (b) play the bet and receive points according to its outcome.

Will you be playing the bet? ____ (yes/no) If your answer is YES, you will need to know the outcome of the bet before you can determine the WIN area of your prize wheel. If your answer is NO, you do not need to know the outcome of the bet to determine the WIN area. Suppose the red ball drawn to determine the outcome of the bet was 18.

Based on my point decision above, and the results of the draws from the bingo cages, the number of points I would have is ____.

This means that the WIN area of my prize wheel would cover the numbers 0 through ____.

If the spinner stopped on the number 5, I would (circle the correct words) win/not win the \$4.00 prize.

If the spinner stopped on the number 40, I would (circle the correct words) win/not win the \$4.00 prize.

{*Page break.*}

Complete this page only if you would have been playing the bet to receive points.

Suppose the red ball drawn to determine the outcome of the bet was 10 instead of 18.

Based on my point decision above, and the results of the draws from the bingo cages, the number of points I would have is ____.

This means that the WIN area of my prize wheel would cover the numbers 0 through ____.

If the spinner stopped on the number 5, I would (circle the correct words) win/not win the \$4.00 prize.

If the spinner stopped on the number 40, I would (circle the correct words) win/not win the \$4.00 prize.

Stop here and wait for the experimenter to tell you to go on to the next practice item.

{*Page break.*}

Practice Item 2: Suppose you have the opportunity to play the bet shown below. What is the *smallest* number of points for which you would give up this opportunity? Remember that the WIN area on your prize wheel will be designated as the area from 0 through the number of points you receive as a result of your decision.

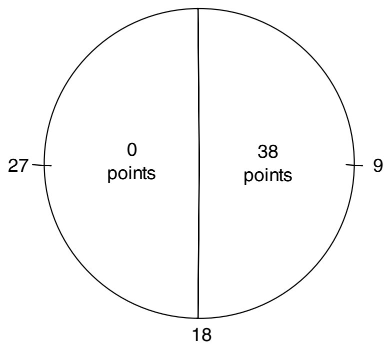

Decision ____.
The green ball drawn at random is ____.
The number on this green ball is (a) greater than my indicated amount.
(b) less than or equal to my indicated amount.
Therefore, I would (a) receive points equal to the number on the ball and not play the bet.
(b) play the bet and receive points according to the outcome.
The red ball drawn to determine the outcome of the bet was ____.
Based on my decision above, and the results of the draws from the bingo cages, the number of points I would have is ____.
This means that the WIN area of my prize wheel would cover the numbers 0 through ____.
My spinner stopped on the number ____.
Therefore I would have (circle the correct words) won/not won the $4.00 prize.

{*Page break.*}

Practice Item 3: Suppose you have the opportunity to play the bet shown below. What is the *smallest* number of points for which you would give up this opportunity? Remember that the WIN area on your prize wheel will be designated as the area from 0 through the number of points you receive as a result of your decision.

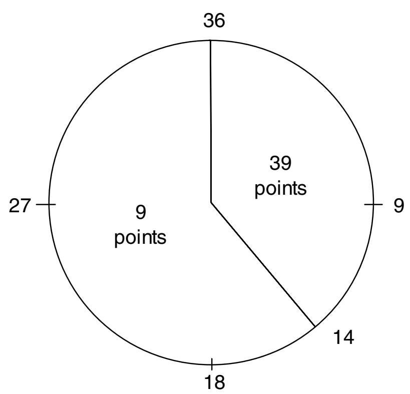

Decision ____.
The green ball drawn at random is ____.
The number on this green ball is (a) greater than my indicated amount.
(b) less than or equal to my indicated amount.
Therefore, I would (a) receive points equal to the number on the ball and not play the bet.
(b) play the bet and receive points according to the outcome.

The red ball drawn to determine the outcome of the bet was ____.
Based on my decision above, and the results of the draws from the bingo cages, the number of points I would have is ____.
This means that the WIN area of my prize wheel would cover the numbers 0 through ____.
My spinner stopped on the number ____.
Therefore I would have (circle the correct words) won/not won the $4.00 prize.

{*Twelve pricing tasks follow.*}

Part 3:

This part is exactly like Part 1. You will be asked to consider several pairs of bets, and for each pair you should indicate which bet you prefer to play or indicate that you are indifferent between them. After each decision, you will then have an opportunity to play for a $4.00 prize using the following procedure:

1. The bet you indicate as preferred will be played and you will receive the points indicated by its outcome. If you check "Indifferent" the bet you play will be determined by a coin toss.
2. The WIN area of your prize wheel will be designated as the area from 0 through the number of points which you have received. You will spin the spinner to determine whether you win the $4.00 prize.

{*Three paired choice tasks follow.*}

Acknowledgments

We thank Daniel Kahneman, Chandra Kanodia, Duncan Luce, Kevin McCabe, an anonymous referee, the editor and workshop participants at the University of Minnesota and the Experimental Science Association Meetings for their helpful comments and discussions.

Notes

1. The binary lottery risk preference induction mechanism is an alternative means of giving subjects incentives that allows control of risk preferences. It was introduced by Roth and Malouf (1979) and extended by Berg et al. (1986).
2. In preference reversal studies, subjects are typically presented with two gambles that have the same expected value: a low variance gamble that has a high probability of getting a high payoff but a small high payoff dollar amount (the p-bet) and a high variance gamble that has a low probability of getting a high payoff, but a big high payoff dollar amount (the $-bet). Preferences are elicited through two tasks: a paired-choice task and an individual-bet pricing task. A p-bet reversal occurs when a subject chooses the p-bet in the choice task, but prices the same bet lower in the pricing task. A $-bet reversal occurs when the opposite happens. These reversals are often termed "predicted" and "unpredicted," respectively.
3. We note that Schotter and Braunstein (1981) introduce a different means of manipulating subjects' risk preferences. They pay subjects in points and transform the point payoffs into dollars using a concave function to increase subjects' risk aversion in points. We do not study this method here because, unlike the lottery method, wealth effects and native preferences can confound results.

4. Bostic, Herrnstein, and Luce (1990) demonstrate that by making the pricing task more like a choice task (using the PEST procedure), they can reduce the number of observed reversals.
5. We do this to maintain consistency with other experiments that investigate risk preference induction, for instance, Berg et al. (1986). This differs from the payoff frequency in Selten, Sadrieh, and Abbink (1999) where subjects receive points after every decision, but those points accumulate for two decisions with the induction lottery played only after every second decision.
6. The set of assumptions needed are actually a subset of those required for expected utility maximization. Non-expected utility specifications in which subjects display preferences that are linear in probabilities will also lead to this result.
7. The functions used to map points into the 360 degrees of the prize wheel were:

$$\text{Risk Averse: Degrees} = 364.4784 - 364.4784 \times \mathrm{e}^{(-0.11 \times W)}$$
$$\text{Risk Loving: Degrees} = -4.47478 + 4.47478 \times \mathrm{e}^{(0.11 \times W)}$$
$$\text{Risk Neutral: Degrees} = 9w$$

8. We have excluded observations with reported indifference across pairs in choices but reported different valuations in the pricing task because it is unclear whether these are predicted or unpredicted reversals according to Selten, Sadrieh, and Abbink's (1999) classification scheme.
9. In one treatment, Selten, Sadrieh, and Abbink's (1999) subjects were allowed to see mean and variance statistics on each gamble if they requested them. In the other, these statistics were not available.
10. The risk neutral percentage comes from aggregating across all risk neutral treatments in our data and Selten, Sadrieh, and Abbink (1999).
11. Indifference in the pricing task is counted as a consistent choice because the pricing grid may not be fine enough to capture small differences in valuations.
12. The difference in reversal rates under risk neutrality is -0.10, which is insignificant at the 95% level of confidence (z-statistic $= -1.82$). The difference in reversal rates under risk aversion is -0.65 with "predicted" reversals significantly exceeding "unpredicted" reversals (z-statistic $= -6.85$). The difference in reversal rates under risk loving is 0.29 with "unpredicted" reversals significantly exceeding "predicted" reversals (z-statistic $= 3.51$).
13. In the risk neutral treatment, certainty equivalents for the two gambles in a pair are approximately equal. To avoid overlapping lines and an uninterpretable graph, we perturbed the certainty equivalents for each pair by $+/-.25$ points from the average of the certainty equivalents for the pair.
14. Technically, s is the chance that deviations from true values revealed in the pricing tasks are sufficiently large to cause a reversal of the ordering of the gambles from true preferences.
15. We will be unable to calculate q whenever the denominator in Eq. (1) is zero. This occurs only when $r = 1/2$ or $s = 1/2$. Similarly, when $q = 1/2$, the denominators in Eq. (2) and Eq. (3) are zero, so r and s cannot be calculated. As long as $q \neq 1/2$, $r \neq 1/2$, and $s \neq 1/2$, then all three parameters can be calculated from a, b, c, and d.
16. These sets of parameters are related as follows. Let $\{q_1, r_1, s_1\}$ represent the parameter set obtained when we choose the solution to Eq. (1) that is consistent with risk aversion (that is, $q_1 > .5$). Then the second set of parameters, $\{q_2, r_2, s_2\}$, satisfies:

$$q_2 = 1 - q_1$$
$$r_2 = 1 - r_1 \quad \text{and}$$
$$s_2 = 1 - s_1.$$

If $\hat{q} = .5$, then there is no second solution to Eq. (1). In this case, we cannot estimate error rates since the denominators in Eqs. (2) and (3) are zero.

17. Sampling error can produce negative estimated q's (outside of the valid zero to one range) when the true q is close to the limits.
18. As noted earlier, no real root exists for the Lichtenstein and Slovic experiment. In related work (Berg, Dickhaut, and Rietz, 2002), we attribute this to the nature of incentives used in this experiment.
19. Note that even though task error rates fell, preference reversal rates in this condition did not fall.

20. These are the risk neutral, high incentives instructions. Only the amounts of the prizes and the prize wheels changed for other treatments.

References

Berg, Joyce E., Lane A. Daley, John W. Dickhaut, and John R. O'Brien. (1986). "Controlling Preferences for Lotteries on Units of Experimental Exchange," *Quarterly Journal of Economics* 101, 281–306.

Berg, Joyce E., John W. Dickhaut, and John R. O'Brien. (1985). "Preference Reversal and Arbitrage," In Vernon Smith (ed.), *Research in Experimental Economics*. Greenwich: JA Press, Vol. 3, pp. 31–72.

Berg, Joyce E., John W. Dickhaut, and Thomas A. Rietz. (2002). "Preference Reversals: The Impact of Truth-Revealing Incentives," Working Paper, the University of Iowa.

Berg, Joyce E., John W. Dickhaut, and Thomas A. Rietz. (2003). "On the Performance of Lottery Procedures for Controlling Risk Preferences," Forthcoming in Charles R. Plott and Vernon Smith (eds.), *Handbook of Experimental Economics Results*. North Holland, Amsterdam.

Bostic, Raphael, Richard J. Herrnstein, and R. Duncan Luce. (1990). "The Effect on the Preference-Reversal Phenomenon of Using Choice Indifferences," *Journal of Economic Behavior and Organization* 13, 193–212.

Camerer, Colin F. and Robin M. Hogarth. (1999). "The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework," *Journal of Risk and Uncertainty* 19, 7–42.

Cox, James C. and David M. Grether. (1996). "The Preference Reversal Phenomenon: Response Mode, Markets and Incentives," *Economic Theory* 7, 381–405.

Cox James C. and Ronald L. Oaxaca. (1995). "Inducing Risk-Neutral Preferences: Further Analysis of the Data," *Journal of Risk and Uncertainty* 11, 65–79.

Goldstein, William M. and Hillel J. Einhorn. (1987). "Expression Theory and the Preference Reversal Phenomena," *Psychological Review* 94, 236–254.

Grether, David M. and Charles R. Plott. (1979). "Economic Theory of Choice and the Preference Reversal Phenomenon," *American Economic Review* 69, 623–638.

Hamm, Robert M. (1979). "The Conditions of Occurrence of the Preference Reversal Phenomenon." Unpublished doctoral thesis, Harvard University.

Kahneman, Daniel and Amos Tversky. (1979). "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica* 47, 263–291.

Karni, Edi and Zvi Safra. (1987). "'Preference Reversal' and the Observability of Preferences by Experimental Methods," *Econometrica* 55, 675–685.

Lichtenstein, Sarah and Paul Slovic. (1971). "Reversals of Preference Between Bids and Choices in Gambling Decisions," *Journal of Experimental Psychology* 89, 46–55.

Lichtenstein, Sarah and Paul Slovic. (1973). "Response-Induced Reversals of Preference in Gambling: An Extended Replication in Las Vegas," *Journal of Experimental Psychology* 101, 16–20.

Loomes, Graham, Chris Starmer, and Robert Sugden. (1989). "Preference Reversal: Information-Processing Effect Or Rational Non-Transitive Choice?," *Economic Journal* 99, 140–151.

Loomes, Graham and Robert Sugden. (1983). "A Rationale for Preference Reversal," *American Economic Review* 73, 428–432.

Luce, R. Duncan. (2000). *Utility of Gains and Losses*. Mahway, NJ: Lawrence Erlbaum Associates.

Prasnikar, Vesna. (2001). "How Well Does Utility Maximization Approximate Subjects' Behavior? An Experimental Study." Working paper, Carnegie Mellon University.

Rietz, Thomas A. (1993). "Implementing and Testing Risk Preference Induction Mechanisms in Experimental Sealed Bid Auctions," *Journal of Risk and Uncertainty* 7, 199–213.

Reilly, Robert J. (1982). "Preference Reversal: Further Evidence and Some Suggested Modifications in Experimental Design," *American Economic Review* 72, 576–584.

Roth, Alvin E. and Michael W.K. Malouf. (1979). "Game-Theoretic Models and the Role of Bargaining," *Psychological Review* 86, 574–594.

Schkade, David A. and Eric J. Johnson. (1989). "Cognitive Processes in Preference Reversal," *Organizational Behavior and Human Decision Processes* 44, 203–231.

Schotter, Andrew and Yale M. Braunstein. (1981). “Economic Search: An Experimental Study,” *Economic Inquiry* 19, 1–25.

Segal, Uzi. (1988). “Does the Preference Reversal Phenomenon Necessarily Contradict the Independence Axiom?,” *American Economic Review* 78, 233–236.

Selten, Reinhard, Abdolkarim Sadrieh, and Klaus Abbink. (1999). “Money Does Not Induce Risk Neutral Behavior, But Binary Lotteries Do Even Worse,” *Theory and Decision* 46, 211–249.

Tversky, Amos, Paul Slovic, and Daniel Kahneman. (1990). “The Causes of Preference Reversal,” *American Economic Review* 80, 204–217.