

The Impact of Insurance Prices on Decision-Making Biases: An Experimental Analysis

Susan K. Laury
Georgia State University

and

Melayne Morgan McInnes
University of South Carolina

December 2001

Abstract: This paper tests whether the use of endogenous risk categorization by insurers enables consumers to make better-informed decisions even if they do not choose to purchase insurance. We do so by adding a simple insurance market to an experimental test of optimal (Bayesian) updating. In some sessions, no insurance is offered. In others, actuarially fair insurance prices are posted, and a subset of subjects is allowed to purchase this insurance. We find significant differences in the decision rules used depending on whether or not one observes insurance prices. Although the majority of choices correspond to Bayesian updating, the incidence of optimal decisions is higher in sessions with an insurance option. Most subjects given the option to purchase actuarially fair insurance choose to do so, however fewer subjects purchase insurance when the probability of a loss is higher.

Please send correspondence to:

Melayne Morgan McInnes
Department of Economics
The Darla Moore School Business
The University of South Carolina
Columbia, SC 29208
mcinnes@darla.badm.sc.edu

We wish to thank Kin Blackburn, Helen Doeringhaus, Scott Harrington, Glenn Harrison, and Charlie Holt. Funding was provided by The Darla Moore School of Business Small Research Grant, University of South Carolina Research and Productive Scholarship grant, and the National Science Foundation (SBR-9753125 and SBR-0094800).

The Impact of Insurance Prices on Decision-Making Biases: An Experimental Analysis

1. INTRODUCTION

Individual decisions about whether to wear a seatbelt or purchase a smoke detector can affect the probability of experiencing a loss and making a claim on insurance. Bond and Crocker (1991) show that when consumption is correlated with risks, efficiency can be increased by endogenous categorization, classifying risks on the basis of the insured's consumption decisions. In Bond and Crocker's model, individuals and insurers both understand the relationship between consumption decisions and the consumer's expected losses. In many instances, however, individuals are unlikely to understand the risks associated with their decisions, resulting in inefficient choices. We test whether endogenous risk categorization increases efficiency in a different way than proposed by Bond and Crocker. With endogenous categorization, insurance prices contain information that may allow consumers to make better-informed decisions. This paper describes an experiment that tests whether showing subjects actuarially fair insurance prices reduces deviations from optimal (Bayesian) decision-making.

Many studies of individual decision-making have shown that people often use simple rules-of-thumb, or heuristics, when solving complicated problems¹. Because of bounded rationality, time constraints, or other limitations, individuals cannot or do not utilize all relevant information. The reliance on heuristics may produce biases in individual decisions. In one frequently cited example, Grether (1980) finds evidence that people give less weight to base rates than Bayes' Rule requires when solving statistical problems that require them to weigh prior information and sample information.

In many naturally occurring contexts, however, individuals are able to insure themselves against risk. In these situations, they do not have to rely merely upon prior probabilities of loss or sample information. They also observe insurance prices before determining what action to take and whether or not to purchase insurance against potential loss. If the relative prices of insurance contracts correctly reflect the relative probability of loss, the price of each contract

¹ See Camerer, 1995, Davis and Holt, 1993, and Kleindorfer, Kunreuther, and Schoemaker, 1993, for surveys of these studies. For studies of decision-making in insurance markets, see Camerer and Kunreuther, 1989; McClelland, Schulze and Coursey, 1993; and Shapira and Venizia, 1999.

may be an important signal of the risk of each choice or activity. Consider, for example, the decision of whether to quit smoking. There is a wealth of information available about the health consequences of smoking, but Viscusi (1990) has shown that individual assessments of the risks from smoking are often biased. In the United Kingdom, the Stalwart Insurance Company offers higher annuity payments for smokers than for non-smokers (Investors Chronicle, 1995). For someone who understands what this higher monthly payment means, this is a dramatic signal of the mortality risk associated with smoking. A second example is the decision of whether to strap down a water heater in an earthquake prone area. These measures are not frequently adopted (Palms et al, 1990) suggesting that people do not believe such measures are effective (Kunreuther, 1996). Reduced premiums on insurance policies that specify strapping down the water heater may help to convey this information about effectiveness. We address the question of whether companies who sell endogenously-priced insurance also create a byproduct - information - that allows people to better understand risk and therefore to make better decisions.

This paper describes an experiment that introduces a simple insurance option into El-Gamal and Grether's (1995) test of Bayesian updating. We focus on the information that insurance markets create and disseminate in the form of prices. In particular, we address the question of whether allowing subjects to purchase actuarially fair insurance reduces deviations from Bayesian updating. We are interested in whether decisions are more consistent with Bayes' Rule when subjects are given the option to purchase insurance. We examine whether those who are not allowed to purchase, but who can observe posted prices for the insurance, also make better decisions than those who never observe insurance prices.

Using an environment in which biases are likely to arise allows us to study whether the introduction of insurance reduces the *prevalence* of these biases or the *kinds* of biases that are observed. We use an experimental design, described in the next section, in which others have found systematic deviations from Bayesian updating. Sections 3 and 4 present an overview and statistical analysis, respectively, of the decision rules used by individuals in different insurance treatments. Section 5 looks at the decision to purchase insurance, and the final section concludes.

2. THE DECISION-MAKING ENVIRONMENT

Procedures and Subjects

We began with the procedures used by Grether (1980) and El-Gamal and Grether (1995) in which subjects guessed from which cup a sample was drawn, and could maximize expected earnings by making decisions in accordance with Bayes' Rule. Next, we introduced a simple insurance market in which subjects could purchase a policy that fully reimbursed them for the penalty assessed when an incorrect decision is made. The posted prices of insurance for each cup accurately reflected the Bayesian posterior odds that each cup was used; therefore, the cost of insurance was lower for the cup with the highest posterior odds.

The experiment was implemented in the following manner.² Two cups were used, marked A and B. The contents of the cups were as follows (where O denotes an orange ball and W denotes a white ball):

Cup A:	O	O	O	O	W	W
Cup B:	O	O	O	W	W	W

Subjects were given a \$23 endowment and told that a sample would be drawn with replacement, from one of the two cups (determined by the roll of a six-sided die). For example, at the start of Round 1 subjects were told that Cup A would be used if the die throw was 1, 2, 3, 4, or 5; Cup B would be used if the die throw was a 6. After announcing this prior probability, the die was thrown behind a screen and the appropriate cup was selected based on this die roll.³ The cup's contents were placed into an opaque bucket and the experimenter then drew six balls with replacement from the bucket, showing and publicly recording each draw as it occurred. Subjects then guessed which cup was used. If the subject was correct, the subject kept the \$23 endowment; otherwise, \$14 was deducted, and final earnings were \$9. These payoffs correspond to a situation in which either choice may result in a loss, but the probability of incurring a loss is endogenously determined by the particular choice made.⁴ This process was repeated 16 times without feedback. Earnings were determined by one decision chosen at random after all decisions were made. A student monitor verified that all of the procedures were properly

² Instructions for both of our treatments are available on the web at <http://dmsweb.badm.sc.edu/mmcinnes/>

³ There were 16 rounds in the experiment and the prior probability for Cup A followed a predetermined, but unannounced, sequence: 5/6, 2/3, 1/3, 5/6, 1/3, 2/3, 5/6, 2/3, 1/3, 5/6, 1/3, 2/3, 5/6, 2/3, 1/3, 5/6.

⁴ We framed earnings in terms of losses to be consistent with the natural interpretation of insuring against a loss.

followed as described in the instructions.

In sessions that included an insurance option, participants were told that in some rounds they would be able to purchase insurance that would fully reimburse them for the \$14 loss resulting from an incorrect guess. The instructions read: "the insurance prices were calculated in such a way that, if we repeated the experiment enough times, on average the insurance would reimburse to you exactly the amount that you paid for the insurance."⁵ After the sample was drawn, insurance prices were posted for each cup. Those subjects who could purchase insurance made two choices: which cup was used, and whether to purchase insurance for this choice. About half of the subjects were allowed to purchase insurance in each round (but they were not told in advance in which rounds they would be able to purchase the insurance). One group of subjects was able to purchase in the first 5 rounds but not in the remaining 11, while the other group was able to purchase in the last 11 rounds but not in the first five. Therefore, those who could not purchase insurance in the majority of rounds were still given prior experience with looking at insurance prices and deciding whether or not to purchase insurance. Moreover, all subjects were required to write down the insurance prices to be sure that all observed them.

Overall, 88 subjects participated in 11 sessions (an additional 11 people served as monitors, one in each session); 28 participants were in five control (no insurance) sessions, while the remaining 60 subjects participated in six insurance sessions (32 of these subjects were given the option to purchase insurance in the last 11 rounds). All sessions were conducted at the University of South Carolina. Participants were recruited from undergraduate business classes and through flyers posted across the campus. About half of the subjects were business or economics majors. They ranged in age from 17 to 47, with an average age of 21.

Decision Rules and Heuristics

A specific example demonstrates how Bayes' rule may be used to compute the posterior probability that a given cup was used, as well as how heuristics may lead to a different (or

⁵ Although actuarially fair insurance is not prevalent outside of the lab, we wanted to reduce confusion or skepticism about how insurance prices were calculated.

identical) decision.⁶ Suppose that Cup A is used with probability 2/3, and Cup B is used with probability 1/3. If the sample draw is 3 orange balls and 3 white balls, this is "representative" of Cup B because the draw exactly matches the contents of this cup. If the subject under-weights the base probability that Cup A was twice as likely to be used as Cup B, the representative heuristic (also known as base right bias) may lead to the conclusion that Cup B was used. However, the correct Bayesian posterior probability that Cup A was used is 0.584.

Table 1. Summary of Decision Rules

Draw	Prob(A draw)	Bayes' Rule	Representative Heuristic	Conservative Heuristic
Prob (A) = 1/3				
6 orange, 0 white	0.737	A	A	B
5 orange, 1 white	0.584	A	A	B
4 orange, 2 white	0.413	B	A	B
3 orange, 3 white	0.260	B	B	B
2 orange, 4 white	0.149	B	B	B
1 orange, 5 white	0.081	B	B	B
0 orange, 6 white	0.042	B	B	B
Prob (A) = 2/3				
6 orange, 0 white	0.918	A	A	A
5 orange, 1 white	0.849	A	A	A
4 orange, 2 white	0.737	A	A	A
3 orange, 3 white	0.584	A	B	A
2 orange, 4 white	0.413	B	B	A
1 orange, 5 white	0.260	B	B	A
0 orange, 6 white	0.149	B	B	A
Prob (A) = 5/6				
6 orange, 0 white	0.966	A	A	A
5 orange, 1 white	0.934	A	A	A
4 orange, 2 white	0.875	A	A	A
3 orange, 3 white	0.778	A	B	A
2 orange, 4 white	0.637	A	B	A
1 orange, 5 white	0.467	B	B	A
0 orange, 6 white	0.305	B	B	A

Alternatively, the conservative heuristic predicts that subjects anchor too strongly on the base rate and fail to update their probabilities to reflect sample information. Suppose instead that the sample draw was 2 orange and 4 white balls. In this case, the Bayesian posterior probability that Cup A was used is 0.413, so Cup B should be chosen. However, the conservative heuristic predicts that Cup A may be chosen because of the higher base probability, even when fewer than 2 orange balls are drawn.⁷

⁶ See Miller, 1956, and Newell and Simon, 1972, for early work on heuristics; Kleindorfer, Kunreuther, and Schoemaker, 1993, provide a more recent review.

⁷ See Edwards, 1968; Eger and Dickhaut, 1982; Armstrong, 1985; McKelvey and Page, 1990; and Griffin and Tversky, 1992.

Table 1 summarizes these predictions. An individual should select Cup A whenever the Bayesian posterior probability is greater than 0.5 and should select Cup B otherwise. The representative heuristic tells us that Cup A will be chosen when the observed draw looks most like Cup A (2 or fewer white), and Cup B will be chosen if it looks like more Cup B (3 or more white). The conservative heuristic predicts that one does not update the prior probability to a great enough extent based on the sample information, and so the Cup with the highest prior is chosen. In addition, individuals may employ heuristics other than these; however, following El-Gamal and Grether (1995), we limit our attention to this prominent subset of potential rules.

Finally, consider how the insurance prices reflect the Bayesian posterior probabilities. Table 1 shows that when Cup A is used with a prior probability of $2/3$ and 3 orange and 3 white balls are drawn, the posterior probability that Cup A was used is 0.584. The price of insurance for those selecting cup A is equal to the expected loss of $\$14(0.416) = \5.82 . A subject who chooses Cup A must pay \$5.82 to fully insure against a \$14 loss, and will earn $\$23 - \$5.82 = \$17.18$, whether or not Cup A was used. If the subject chooses Cup B, the insurance will cost \$8.18. The fact that insurance costs more when Cup B is chosen reflects the fact that there is a higher expected loss associated with choosing Cup B than with Cup A. The Appendix contains the insurance prices for each prior and draw combination.

3. OVERVIEW OF INDIVIDUAL CHOICES

In this section we summarize the behavioral patterns that emerged, focusing on whether decisions in sessions in which insurance was offered differed from those in sessions without insurance. We consider rounds in which use of the representative or conservative heuristic leads to a different choice than Bayesian updating would predict. For example, when 3 orange and 3 white balls are drawn and the prior is $2/3$, both Bayes' Rule and the conservative heuristic predict Cup A will be chosen (see Table 1). If a subject instead chooses Cup B, this is identified as a representative decision. In rounds where all three rules agree but the subject chooses the opposite cup, we label this decision simply as "other."

Figure 1. Observed Frequency of Non-Bayesian Decisions

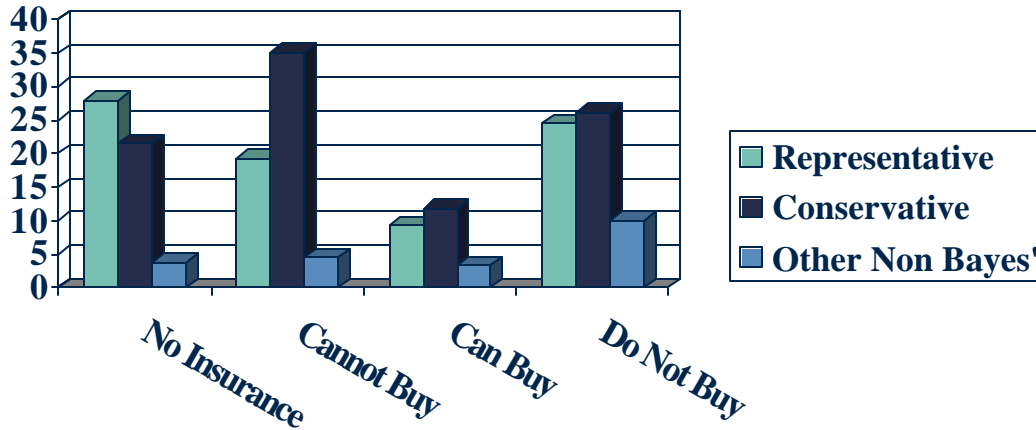


Figure 1 shows the aggregate incidence of these choices. In sessions in which no insurance was offered (columns labeled “No Insurance”) there were 136 draws in which a representative decision could be identified; of these, 38 (27.9%) were consistent with the representative heuristic. The incidence of conservative decisions was lower: just 21.6% (out of 51 possible times). El Gamal and Grether (1995) also find that the most common form of non-Bayesian decision-making is representative followed by conservative. This pattern is reversed for subjects who observed insurance prices. Those in the second group of columns observed insurance prices but were not allowed to purchase the insurance (combining observations from all rounds). Those in the “Can Buy Insurance” columns observed the insurance prices and were allowed to purchase insurance. About 75 percent of subjects who were given the opportunity to purchase insurance chose to do so. The 25% who chose not to purchase are in the “Do Not Buy,” columns. In every insurance treatment conservative decisions are made more frequently than representative decisions. This suggests that observing insurance prices may alter decisions even if individuals choose not to buy or are not allowed to purchase insurance.

Subjects who are allowed to purchase insurance make fewer representative and conservative decisions. In this sense, we find evidence that endogenously priced insurance does improve efficiency in our experiments. However, anyone who purchases insurance should

always choose the cup predicted by Bayes' Rule.⁸ Given this, a more stringent test of the effect of insurance prices would be to compare those who chose not to insure ("Do Not Buy" column) with those who do not observe insurance prices. Representative and conservative heuristics are used with about the same frequency among subjects who never observed insurance prices and those who chose not to buy, while the rate at which subjects made decisions inconsistent with any of these rules was about double for those who chose not to buy insurance. Therefore, the insurance market did not improve decision-making among those who chose not to purchase insurance.

One explanation for this is the self-selection involved in the decision not to purchase insurance. If a person believes that the "wrong cup" (from the standpoint of Bayes' rule) is more likely to have been used, then one may perceive that the actuarially fair insurance is overpriced, given one's own subjective beliefs. For example, if several draws that look like Cup A have been observed, a draw of 3 orange and 3 white may confirm one's belief that Cup B is "due" even if the subject understands that the objective posterior probability of B is below one-half. In this event, the subject may choose not to purchase insurance and instead select the cup believed to be more likely.

In order to avoid these selection issues, we also compare those who observe insurance prices but are not allowed to purchase (Cannot Buy) with those in the no insurance sessions. Here, the results are mixed. Representative decision-making by those who observe insurance prices is reduced relative to that in the no insurance sessions (19.3 percent of those who observed prices made representative errors, compared with 27.9 percent in the no insurance sessions). However, there is a much higher incidence of conservative decision-making among this group (35.1 percent for those who observed prices, versus 21.6 percent for those who did not). It appears that posting prices for actuarially fair insurance increased the focus on base rates (and therefore the use of the conservative heuristic). To explore this more carefully, the following section presents estimation of the decision-rules used in each treatment.

⁸ When one purchases insurance, earnings are equal to \$23 (the initial endowment) minus the price of insurance for the cup that is chosen. To maximize earnings, a subject who purchases should always choose the cup with the lower price, which is the cup that has the highest Bayesian posterior probability. Overall, subjects who purchased insurance chose consistently with Bayes' rule about 98 percent of the time.

4. ESTIMATION OF DECISION RULES

We estimate the decision rules employed by subjects in these experiments using the cut-point analysis described in El-Gamal and Grether (1995). The idea is that one uses a cutoff rule, which is a function of the prior and the draw, when deciding whether to choose Cup A or Cup B. This is a natural class of decisions rules to consider that includes both Bayes' Rule and the conservative and representative heuristics. For example, when the prior probability of using Cup A is $1/3$ Bayes' Rule predicts that a subject will choose Cup A if more than four orange balls are drawn (i.e., when the posterior probability on A rises above 50 percent, as shown in Table 1). So, under this prior, a sample of 4 orange balls is the "cut-point" after which the subject will switch from choosing B to A. Similarly, when the prior for A is $2/3$, the subject should choose Cup A when more than 2 orange balls are drawn, and when the prior is $5/6$, Cup A should be chosen when more than 1 orange ball is drawn. This is shown in the first row of Table 2.

Table 2. Predicted Cut Points

	Cut-Point for Pr(A) = 1/3	Cut-Point for Pr(A) = 2/3	Cut-Point for Pr(A) = 5/6
Bayes' Rule	4	2	1
Representative	3	3	2, 3
Conservative	5, 6	-1, 0, 1	-1, 0

A cut-point specifies that Cup A is chosen if the number of orange balls drawn is greater than the specified number for a given prior.

Taking all combinations of draws for the three priors yields a total of 512 possible cutoff rules. Table 2 shows the cut-points for the three most prominent rules: Bayesian updating, the representative heuristic, and the conservative heuristic. The strong form of representative decision-making predicts that one will choose Cup A when more than 3 orange balls are drawn, regardless of the prior. A weaker form of the representative heuristic considers the prior but still gives too much weight to the sample. For example, if the prior for A is $5/6$, Bayes' Rule predicts that Cup A will be chosen if more than 1 orange ball is drawn. If we observe subjects using a cut-point of 2 rather than 1, this behavior is consistent with over-weighting the sample information and is identified as a representative decision rule. The last row shows the

conservative cut-points. At an extreme, a conservative subject would ignore the sample and always choose Cup A when the prior for A is greater than 1/2. A weaker form of conservative decision-making predicts that the prior will receive too much weight but the sample draw will also influence the choice. For example, if the prior on A is 1/3, a cut-point of 5 or 6 is consistent with the conservative heuristic that can be distinguished from Bayes' Rule.

We use the observed choices made by subjects to infer the actual cut-point employed for each of the three priors used in our experiments. Of course, none of these decision rules will be applied with absolute precision. We therefore assume that with probability g , the subject "trembles" when applying the rule and makes a random choice between A and B. Thus, a subject's decision will agree with the rule with probability $1-g/2$ and disagree with the rule with probability $g/2$. The error rate g is assumed to be the same for all subjects in all rounds. Given a subject's decisions in each round, (x_1, x_2, \dots, x_t) , the likelihood of the rule c and error rate g is given by

$$f(x_1, x_2, \dots, x_t) = \left(1 - \frac{g}{2}\right)^{X_c} \times \left(\frac{g}{2}\right)^{t-X_c}$$

where t is the number of decisions and X_c is a count of the number of times the subject's decisions agree with the rule c .

Using maximum likelihood procedures, we first estimated the single rule that best fits our subjects' decisions.⁹ Results are shown in the first two rows of Table 3. Recall, almost everyone who purchased insurance in these experiments chose consistently with Bayes' Rule; hence, there are few observations on subjects who were able to purchase insurance but chose not to. Therefore, we limit our attention to those who never observed insurance prices ("No Insurance") and those who observed prices but were not allowed to purchase insurance.

In both treatments, the estimated cutoff rule is consistent with Bayes' Rule when the prior is 1/3 and 2/3. We can distinguish between these two treatments based on the estimates for the cut-point used when the prior on A is 5/6. For those in the "No Insurance" treatment, this cut-point is consistent with a representative heuristic. For those who observe insurance prices but

⁹ The estimation algorithm checks all possible rules. The Gauss code that was used for all of the following estimation is available from the authors upon request. While the standard regularity conditions for maximum likelihood are not met, El-Gamal and Grether (1995) state that the procedure yields consistent estimates and that the Likelihood Ratio test statistic is asymptotically distributed as Chi-squared.

are not allowed to buy, the estimated cut-point does not distinguish between a conservative and a Bayesian cutoff rule. Notice, however, that Bayes' Rule and the conservative heuristic dictate that the same cup will be chosen, except when five or six white balls are drawn. Since we never observed such a draw when the prior on A was 5/6, the data are unable to distinguish between these rules.¹⁰ The estimate of epsilon (see Table 4) shows that the rules are applied with about the same precision in both treatments (about 87 percent of the time). The estimates are roughly comparable El-Gamal and Grether's, which ranged from .22 to .31 in sessions in which subjects were financially motivated for their decisions.

Table 3. Estimated Decision Rules^a

	Rule 1			Rule 2			Number of Subjects Classified to Each	
	1/3	2/3	5/6	1/3	2/3	5/6	Rule 1	Rule 2
One Rule								
No Insurance	B	B	R	-	-	-	28	
Observe Prices, Cannot Buy	B	B	B/C	-	-	-	60	
Two Rules								
No Insurance ^b	B	B	R	R	C	B/C	23	5
				C	C	B/C	17	11
Observe Prices, Cannot Buy	B	B	B/C	C	C	R	32	16 ^c

^a Cutoff rules consistent with Bayes' Rule, the Representative heuristic, and the Conservative heuristic are denoted, respectively, B, R, C.

^b For the second rule chosen, two different sets of rules produced the same likelihood, but a different number of people were classified to each. Therefore, both sets of rules are shown.

^c The remaining 12 subjects were tied between these two rules.

It may be too restrictive to assume that all subjects in a given treatment use the same decision rule. Therefore, we introduce heterogeneity between subjects by estimating the two most prominent decision rules. For each pair of rules, we classify the subject to the rule that best predicts this person's actual decisions. Results from this estimation are given in the bottom rows of Table 3.

For each pair of estimated rules, the rule that provides the best fit for the majority of subjects is the same one that was obtained when we estimated a single rule (the entries for "Rule

¹⁰ Neither of these draws occurred in any of the 66 samples using a prior on A of 5/6. This was true both in the insurance and no insurance treatments.

1" are identical when either one or two decision rules are estimated). The existence of ties complicates the identification of the second most prominent decision rule. In the "No Insurance" treatment, many decision rules result in an equal likelihood, but most contain common elements. For all of the tied rules, a conservative heuristic is estimated for the prior of 2/3 and the data do not distinguish between a conservative heuristic and Bayes' Rule for the prior of 5/6. As a whole, this shows the strongest support for the conservative heuristic when the prior favors Cup A. When the prior on A is 1/3, the Bayesian cut-point is eliminated, but the data do not distinguish between the representative and conservative heuristics. When this cut-point is estimated to be representative, 5 of 28 subjects are classified to Rule 2. When this cut-point is conservative; however, more than double this number (11 subjects) are classified to Rule 2. Across the three priors, this supports a conservative heuristic as the second-most prominent decision rule used by those who do not observe insurance prices.

Turning to those who observe insurance prices but cannot buy insurance, the second most prominent decision rule is consistent with the conservative heuristic except when the prior probability on A is 5/6, when a representative decision rule is favored. Not surprisingly, the estimated error rates are somewhat smaller when we estimate two decision rules than when we impose a single rule on all of our subjects.

Table 4. Likelihood Statistics for Decision Rules

	Number of Rules	Log Likelihood	Epsilon	Information Criterion	Sample Size
No Insurance	1	-166.8	.25	-173.77	448
Observe Prices, Cannot Buy	1	-184.9	.27	-191.82	440
Pooled	1	-355.7	.26	-362.61	888
No Insurance	2	-146.1	.20	-179.35	448
Observe Prices Cannot Buy	2	-139.0	.18	-194.45	440
Pooled	2	-296.5	.20	-371.36	888

The log-likelihood for each of these estimates is shown in Table 4. We would expect a better fit for the data when more than one decision rule is allowed. Therefore we also consider the Information Criterion proposed by El-Gamal and Grether (1995) to choose the optimal number of decision rules.¹¹ This criterion adjusts the likelihood by including a penalty for the additional decision rules. The Information Criterion for both treatments is slightly lower when two decision rules are considered, suggesting that estimation of a single decision rule is sufficient.¹²

Next, we test for significant differences in decision rules used between treatments. In Table 5 we report Chi-squared test statistics for a likelihood ratio test of homogeneity between these two treatments. The null hypothesis is that the epsilon and decision rules used are identical between treatments. When a single decision rule is estimated, we reject the null hypothesis at the 10 percent level. When two decision rules are estimated, we are able to reject the null hypothesis at any standard level of confidence. Thus, the introduction of information on insurance prices has a significant effect on subjects' decision rules, even when these subjects are not allowed to purchase the insurance. Because the estimates of epsilon (the precision with which the decision rule is applied) are very similar between the two treatments, this difference appears to be accounted for by the different decision rules employed by these subjects.

Table 5. Test of Homogeneity between Treatments

Number of Rules	Restricted Log Likelihood	Chi-squared	Degrees of Freedom	p value
1	-355.7	7.97	4	0.093
2	-296.5	22.85	7	0.002

Overall, this estimation finds that subjects may use different heuristics depending on the prior probability that a given cup is used. For priors closer to one-half, Bayes' Rule provides a good fit for the observed choices. However, when the prior on Cup A is 5/6, we observed differences between treatments. When we restrict attention to a single decision rule (as supported by a Chi-squared test), those who do not observe insurance prices use a representative

¹¹ For k rules, the Information Criterion = $\log(\text{maximal likelihood}) - 3k\log(8) - k\log(2) - n\log(k)$.

¹² We also estimated the three best decision rules; the Information Criterion was lower than that for one rule for both treatments, and resulted in many ties.

heuristic. However, for those who observe insurance prices but cannot buy insurance, we cannot distinguish between the use of Bayes' rule and conservative decision-making.

5. THE INSURANCE PURCHASE DECISION

We next turn our attention to the factors that influence the decision of whether to purchase insurance. Recall that the insurance offered to subjects was actuarially fair (that is, the cost to the subject of purchasing insurance was equal to the expected loss), and this fact was emphasized in the instructions. A risk-preferring individual should never purchase this insurance, and a risk neutral individual should be indifferent between purchasing and not. However, one who is risk averse is predicted to always purchase this insurance. Pooling across subjects and decision rounds, insurance was purchased 74 percent of the time. Only three of 60 subjects never purchased insurance. By contrast, 17 of 60 subjects always purchased insurance.

Individual risk attitudes, especially when potential losses are at stake, are debated in the literature. In a classic paper, Kahneman and Tversky (1979) present evidence in favor of risk seeking behavior when people are presented with lotteries over (hypothetical) losses. Kunreuther *et al.* (1978) give further support to the prevalence of risk-seeking in the domain of losses, showing that in actual insurance markets, few people buy highly subsidized earthquake and flood insurance. However, several studies have also found that the purchase of full insurance is quite popular, which would indicate risk *aversion* when potential losses are at stake. McClelland, Schulze and Coursey (1993) show that individuals bid several times the expected loss to obtain full insurance. Moreover, people frequently shun lower priced insurance with a deductible in favor of full insurance (Johnson *et al.*, 1993, and Shapira and Venizia, 1999).

The prevalence of insurance purchases is consistent with some degree of risk aversion.¹³ About 60 percent of subjects purchased insurance in more than 2/3 of the rounds in which they were given the opportunity. This is roughly comparable to the proportion of subjects identified as being risk averse (66 percent) by Holt and Laury (2000) using a lottery over gains.

¹³ Increasing the scale of the potential loss (rather than the likelihood of loss) increased the frequency with which insurance was purchased. In pilot experiments, choosing the incorrect cup resulted in losing \$6 of a \$9 initial endowment (compared with losing \$14 of the \$23 endowment in the experiments reported here). In these pilot sessions, insurance was purchased only about 50 percent of the time. This finding is similar to Bosch-Domenech and Silvestre (1999), who report that willingness to purchase actuarially fair insurance against losses increases in the scale of the loss.

We next consider the circumstances in which individuals choose to purchase insurance in our experiments. There is an inverse correspondence between the posterior probability and the price of insurance: it costs a subject very little to insure when the posterior probability that a cup was used is very high.¹⁴ Figure 2 shows the percentage of subjects who chose to purchase insurance as a function of the posterior probability of the “more likely” cup being used (that is, the cup with a posterior probability that is greater than 50 percent). Fewer subjects bought insurance when the posterior probability was close to 50 percent. When the posterior probability strongly favored one cup over the other, the price of insurance was lower, and more subjects purchased insurance. In these situations, there is less uncertainty and subjects are more likely to be able to “solve the problem” and correctly choose the Cup that is favored by the posterior probability. Because the posterior odds strongly favor this choice, it also carries lower risk.¹⁵ In fact, most non-Bayesian choices occur when the posterior probability is close to one-half. A random effects regression of the percentage of subjects who buy insurance in a given round on the posterior probability of the most likely cup yields:


$$\text{percent buying} = \begin{array}{ccc} .27 & + & .63 * \text{posterior} \\ (.10) & & (.13) \end{array}$$

where the numbers in parentheses are the standard errors of the coefficient estimates. The coefficient for the posterior probability is significantly different than zero at any standard level of confidence.

¹⁴ In insurance settings, the “price” of insurance is typically interpreted to be the difference between the expected loss and the premium. Because the premium charged is just equal to the expected loss in our design, the “price” of insurance is zero. Throughout this paper, we instead use price to mean the amount a subject paid to purchase insurance.

¹⁵ When subjects choose not to insure, the risk associated with their cup choice can be measured as the standard deviation of the payoff conditional on their choice. If the posterior probability of the chosen Cup is p , then the standard deviation is proportional to $\sqrt{p(1-p)}$. The standard deviation is maximized at $p=0.5$. The posterior probability for the more likely cup varies between 0.5 and 1.

Figure 2. Percent Purchasing Insurance



Although many subjects purchase insurance (supporting the existence of risk aversion over losses), this inverse relationship between the probability of a loss and buying insurance is consistent with one of prospect theory's behavioral implications (Kahneman and Tversky, 1979, and Tversky and Kahneman, 1992). Specifically, risk-seeking behavior is predicted for high probabilities of loss, and risk-averse behavior may be observed when the probability of loss is low. This is consistent with the fact that insurance is bought more frequently when the posterior probability is high (and therefore the probability of loss is low).

At the conclusion of each experimental session, subjects were asked to complete a brief demographic questionnaire. While not the focus of our experiments, it is useful to consider the effect of socioeconomic characteristics on the insurance purchase decision. Despite evidence from other studies (Holt and Laury, 2000, and Eckel et al., 1998) that gender and other demographic characteristics may influence risk attitudes, the frequency with which an individual purchased insurance was not significantly related to any of these characteristics. OLS regressions including characteristics such as age, gender, major, year in school, and the purchase of insurance outside of the lab (such as auto, renters, or health insurance) were not significant when estimated as a group or in univariate regressions. We should note, however, that our design does not allow us to infer the *degree* of risk aversion. If one is even slightly risk averse, that person should always purchase insurance. Therefore, if both men and women are risk averse, all should purchase insurance, regardless of whether one group is significantly more risk averse than the other.

6. CONCLUSIONS

Studies of individual decision-making have shown that people often use simple rules-of-thumb, or heuristics, when solving complicated problems. The end-result can be systematic biases in decisions. An open question is whether the information provided by market forces can improve decisions in these contexts. The preponderance of evidence is that markets reduce, but do not eliminate, these biases.¹⁶ This paper instead focuses on a different question: whether the market for insurance yields, as a byproduct, information that may reduce biases in making related decisions. If insurers are better informed than consumers about the relationship between consumption choices and risk, then endogenous risk categorization may help to inform consumers about risk.

We began by replicating El-Gamal and Grether's experiment, in which we found evidence of decision-making biases. We followed this by a treatment in which some subjects were given the option to purchase insurance at actuarially fair prices. Other subjects observed the prices but were not allowed to purchase the insurance. Those who were allowed to purchase made decisions that were largely consistent with Bayesian updating. However, the interpretation of this result is complicated by the fact that nearly 75 percent of these subjects chose to purchase insurance. It is not clear whether the insurance prices themselves improved decisions or if subjects simply wanted to insure and therefore chose the less expensive (and therefore "correct") option. In either case, subjects who were able to purchase insurance made better decisions (from the standpoint of Bayes' Rule) than did those in our other treatments.

When all priors are combined, those who observe the price of insurance but are not allowed to buy are less likely to make decisions consistent with a representative heuristic than those who do not observe prices. However, they make conservative decisions more often. Maximum likelihood estimation of the decision rules used by subjects not allowed to purchase insurance (separated by treatment depending on whether they observed prices), suggested that estimating a single decision rule for all subjects in a given treatment was sufficient. Both sets of these subjects made decisions consistent with Bayesian updating for the 1/3 and 2/3 priors.

¹⁶ For implementations using asset markets, see Camerer, 1987 and 1990; Duh and Sunder, 1986; and Anderson and Sunder, 1995. An exception is Ganguly, Kagel, and Moser (2000) who looked at the role of market context in an asset market experiment. They compared a treatment in which uncertainty was presented in a context-rich narrative with one in which it was presented in a "ball and urn" design and found that the market context increased the prevalence of biases.

When the prior probability of using Cup A was $5/6$, we find a representative rule best describes choices for those who do not observe insurance prices. However, for those who observe prices (but cannot buy), decisions are consistent with both Bayes' Rule and the conservative heuristic.

Finally, we find that most subjects who are able to purchase insurance choose to do so, although only $1/3$ of subjects buy insurance in all rounds. A larger percentage of subjects buy insurance when the posterior probability of that a cup was used is higher. In other words, more subjects buy insurance when the probability of a loss is lower.

The notion that price can serve as an informative signal is not new: many studies have shown that price may signal the quality of a good (in a market context) or inside information (in asset markets).¹⁷ In this paper, we consider whether people use the price of insurance associated with an underlying outcome to infer information about the likelihood of that outcome. Our results indicate that subjects in the insurance session made choices that result in lower expected losses than did those in the sessions with no insurance market.

¹⁷ For price as a signal of quality, see Bagwell and Riordan (1991) and Milgrom and Roberts (1986). For price as a signal of inside information in asset markets, see Forsythe and Lundholm (1990), Banks (1985), and Plott and Sunder (1982).

References

- Anderson, M.J. and S. Sunder. "Professional Traders as Intuitive Bayesians." *Organizational Behavior and Human Decision Processes*, November 1995, 64(2), 185-202.
- Armstrong, J.S. *Long-Range Forecasting: From Crystal Ball to Computer*. New York: John Wiley, 1985.
- Bagwell, K. and M.H. Riordan. "High and Declining Prices Signal Product Quality." *American Economic Review*, March 1991, 81(1), 224-239.
- Banks, Jeffrey S. "Price-Conveyed Information Versus Observed Insider Behavior: A Note on Rational Expectations Convergence." *Journal of Political Economy*, 1985, 93(4), 807-815.
- Bond, E.W. and K.J. Crocker. "Smoking, Skydiving and Knitting: The Endogenous Categorization of Risks in Insurance Markets with Asymmetric Information." *The Journal of Political Economy*, 1991, 99, 177-199.
- Bosch-Domenech, A. and J. Silvestre. "Does Risk Aversion or Attraction Depend on Income? An Experiment." *Economics Letters*, 1999, 65, 265-73.
- Camerer, C.F. "Do Biases in Probability Judgment Matter in Markets? Experimental Evidence." *American Economic Review*, 1987, 77, 981-97.
- Camerer, C.F. "Do Markets Correct Biases in Probability Judgment? Evidence from Market Experiments." In L. Green and J. Kagel, eds., *Advances in Behavioral Economics, Vol. 2*, Greenwich, Conn.: JAI Press, 1990, 126-72.
- Camerer, C.F. "Individual Decision Making." In J. Kagel and A. Roth, eds., *Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press, 1995, 587-616.
- Camerer, C.F. and H. Kunreuther. "Experimental Markets for Insurance." *Journal of Risk and Insurance*, 1989, 2, 265-300.
- Davis, D.D. and C.A. Holt. *Experimental Economics*. Princeton, NJ: Princeton University Press, 1993.
- Duh, R.R. and S. Sunder. "Incentives, Learning, and Processing of Information in a Market Environment: An Examination of the Base Rate Fallacy." In S. Moriarty, ed., *Laboratory Market Research*, Norman, Okla.: University of Oklahoma Press, 1986.
- Eckel, C., P. Grossman, N. Lutz, and V. Padmanabhan. "Playing it Safe: Gender Differences in Risk Aversion." 1998 Discussion Paper, Virginia Tech.

- Edwards, W. "Conservatism in Human Information Processing." In B. Kleinmuntz, ed., *Formal Representation of Human Judgment*. New York: John Wiley, 1968, 17-52.
- Eger, C. and J. Dickhaut. "An Examination of Conservative Information Processing Bias in an Accounting Framework." *Journal of Accounting Research*, 1982, 20, 711-23.
- El-Gamal, M.A. and D.M. Grether. "Are People Bayesian? Uncovering Behavioral Strategies." *Journal of the American Statistical Association*, 1995, 90(432): 1137-45.
- Forsythe, R. and Lundholm R. "Information Aggregation in an Experimental Market." *Econometrica*, March 1990, 58(2), 309-347.
- Ganguly, A.R., J.H. Kagel, and D.V. Moser. "Do Asset Market Prices Reflect Traders' Judgement Biases?" *Journal of Risk and Uncertainty*, May 2000, 20(3), 219-45.
- Grether, D.M. "Bayes' Rule as a Descriptive Model: The Representativeness Heuristic." *Quarterly Journal of Economics*, November 1980, 95(3), 537-57.
- Griffin, D. and A. Tversky. "The Weighing of Evidence and the Determinants of Confidence." *Cognitive Psychology*, 1992, 24, 411-35.
- Holt, C.A. and S.K. Laury. "Risk Aversion and Incentive Effects." 2000 Discussion Paper, University of Virginia.
- Investor's Chronicle*, "How Your Health Eats Into Your Wealth." Aug. 18, 1995, v. 113, Number 1438, p. 33.
- Johnson, E.J., J. Hershey, J. Meszaros, and H. Kunreuther, "Framing, Probability Distortions, and Insurance Decisions." *Journal of Risk and Uncertainty*, 1993, 7, 35-52.
- Kahneman, D. and A. Tversky. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 1979, 47(2), 263-91.
- Kleindorfer, P.R., H.C. Kunreuther, and P.J.H. Schoemaker. *Decision Sciences: An Integrative Perspective*. Cambridge: Cambridge University Press, 1993.
- Kunreuther, H., "Mitigating Disaster Losses Through Insurance." *Journal of Risk and Uncertainty*, 12, 1996, 171-187.
- Kunreuther, H., R. Ginsberg, L. Miller, P. Sagi, P. Slovic, B. Borkin and N. Katz. *Disaster Insurance Protection: Public Policy Lessons*. New York: Wiley, 1978.
- McClelland, G.H. W.D. Schulze, and D.L. Coursey. "Insurance for Low-Probability Hazards: A Bimodal Response to Unlikely Events." *Journal of Risk and Uncertainty*, 1993, 7, 95-116.

- McKelvey, R.D., and T. Page. "Public and Private Information: An Experimental Study of Information Pooling." *Econometrica*, 1990, 58, 1321-39.
- Milgrom, P. and J. Roberts. "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, 1986, 94, 796-821.
- Miller, G.A. "The Magic Number Seven, Plus or Minus Two." *Psychological Review*, 1956, 63(2).
- Newell, A. and H.A. Simon. *Human Problem Solving*. Englewood Cliffs, N.J.: Prentice-Hall, 1972.
- Palms, R. M. Hodgson, R.D. Blanchard, and D. Lyons. *Earthquake Insurance in California: Environmental Policy and Individual Decision-Making*. Boulder: Westview Press, 1990.
- Plott, C.R. and S. Sunder. "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models." *Journal of Political Economy*, 1982, 90(4), 663-698.
- Shapira Z. and I. Venezia. "Experimental Tests of Self-Selection and Screening in Insurance Decisions." *The Geneva Papers on Risk and Insurance Theory*, 1999, 24, 139-158.
- Tversky, A. and D. Kahneman. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*. 1992, 5(4), 297-323.
- Viscusi, W. Kip. "Do Smokers Underestimate Risks?" *Journal of Political Economy*, 1990, 98(6), 1253-1269.

Appendix
Predictions and Insurance Prices

Draw (# orange)	Posterior on A	Bayes' Rule	Representative	Conservative	Insurance Price, A	Insurance Price, B
Prior Probability of Using Cup A = 1/3						
6	.737	A	A	B	3.68	10.32
5	.584	A	A	B	5.82	8.18
4	.413	B	A	B	8.22	5.78
3	.260	B	B	B	10.36	3.64
2	.149	B	B	B	11.91	2.09
1	.081	B	B	B	12.87	1.13
0	.042	B	B	B	13.41	0.59
Prior Probability of Using Cup A = 2/3						
6	.918	A	A	A	1.14	12.86
5	.849	A	A	A	2.12	11.88
4	.737	A	A	A	3.68	10.32
3	.584	A	B	A	5.82	8.18
2	.413	B	B	A	8.22	5.78
1	.260	B	B	A	10.36	3.64
0	.149	B	B	A	11.91	2.09
Prior Probability of Using Cup A = 5/6						
6	.966	A	A	A	0.48	13.52
5	.934	A	A	A	0.93	13.07
4	.875	A	A	A	1.74	12.26
3	.778	A	B	A	3.10	10.90
2	.637	A	B	A	5.08	8.92
1	.467	B	B	A	7.46	6.54
0	.305	B	B	A	9.73	4.27