THE EFFECT OF INFORMATION ON DEMAND FOR HIGH-LOAD INSURANCE

Marc A. Ragin

July 2015

Abstract

High-load insurance products, such as extended warranties and rental car insurance, often are criticized for the large profits they generate at the expense of potentially uninformed consumers. A popular solution is to require sellers to disclose more information, allowing consumers to make an informed and rational economic decision. Whether or not such disclosure actually helps consumers to make economically rational choices, however, remains an open question. We conduct a laboratory experiment to determine whether disclosure affects demand for high-load insurance and, if so, what information has the greatest effect. Following a paid real-effort task, we present experiment subjects with a potential loss of their earnings and offer high-load insurance to cover the loss. Subjects receive one of three additional information disclosures: (1) the true probability of loss, (2) the expected loss or "cost of goods sold" for the insurer, or (3) the insurer's profit on the transaction. Ultimately, none of the disclosures has any significant effect on demand, supporting existing evidence in other fields that disclosure is not effective in changing behavior. These results demonstrate that mandated disclosure is not an appropriate solution to the problem of high-load insurance and that alternative policies should be considered.

1. Introduction

Certain consumer insurance products are known to be expensive relative to the value they provide. Purchasing high-load products such as extended warranties, credit insurance, cell phone insurance, and rental car insurance is often considered to be a financial mistake and inconsistent with rational economic behavior. Consumer advocates and regulators often denounce the profits companies make selling these products (Birnbaum, 1999, 2001), while economists have documented potential behavioral biases leading consumers to make such a purchase. Despite these criticisms, however, the the market for high-load insurance remains robust—American consumers paid $30.5 billion for extended warranties and $7.8 billion for cell phone insurance in 2013 (Arnum, 2014).

Critics of these insurance products have suggested a number of regulatory responses intended to encourage consumers to make “better” decisions. One of the most common recommendations is to require various forms of disclosure, a recommendation that rarely specifies *what* information should be disclosed or *why* this information would change consumer behavior. This is the primary motivation for our study—we conduct an experiment to determine (1) whether disclosure actually changes demand for high-load insurance, and if so, (2) what information has the greatest effect on demand. In our experiment, we disclose additional information to subjects who are offered a high-load insurance product, evaluating how purchase behavior differs between treatments.

The “load” for an insurance product is the ratio of price to expected loss, and the size of the load gives a sense of the relative value of the coverage. Larger loads indicate that coverage is of relatively low value. In responding to a listener question on extended warranties, personal finance radio host Dave Ramsey noted that the cost of goods sold (i.e. the insurer’s expected loss) for an average extended warranty is 12% of the price charged (Ramsey, 2011). This implies a load of 8.3 (1.00 ÷ 0.12). Baker and Siegelman (2013) conducted several back-of-the-envelope load calculations, estimating that rental car insurance has a load of 11.3 and extended warranties have a load of 10.0. The average loads for credit life and disability insurance (which pay a specified debt if the insured dies or is disabled) were 3.4 and 5.9, respectively, in 2014, according to the National Association of Insurance Commissioners. There is no threshold for what constitutes a high load, but insurance products in competitive markets often have loads less than two—in 2014, average loads were 1.87 for homeowners insurance and 1.54 for personal auto insurance. While relatively high loads do not necessarily equate to high profits earned by firms, such high loads imply that the products are of low value to consumers.

Existing academic literature offers conflicting perspectives in recommending disclosure concerning relatively expensive financial products. A number of researchers have recommended disclosure as an "asymmetrically paternalistic" solution—disclosure may encourage consumers to make more rational choices while inflicting minimal costs to sellers and firms. Bar-Gill and Ferrari (2010), Camerer et al. (2003), Nalebuff and Ayres (2003), and Schwarcz (2010) each recommended that sellers of high-load insurance products disclose additional information to consumers. Camerer et al. (2003) provided a succinct explanation of disclosure as an asymmetrically paternalistic policy in the context of extended warranties:

> The fact that [extended warranties] are enormously profitable to retailers implies that they are costly to buyers...Perhaps people who buy warranties do not realize how slight the chance is the product will break within the warranty period, or the fact that the small loss they have to pay for repairs out-of-pocket can be easily

> absorbed into the hedonic ups and downs of everyday life...If disclosure reduces warranty purchases by reminding consumers of the low chance of product breakage, then purchasing the warranty would have been a mistake rather than a preference. If informed consumers continue to purchase the warranties, then it is quite possible that they have a good reason to do so, however unfathomable that decision may seem to an economist. (1253-1254)

While a common policy recommendation, disclosure has also been denounced as ineffective and unnecessary. Baker and Siegelman (2013) specifically criticized the recommendations of Camerer et al. (2003), citing the large amount of empirical evidence that disclosure is ineffective in encouraging rational economic choices. Baker and Siegelman even went so far as to state, "It strikes us as ironic that Camerer, Issacharoff, Loewenstein, O'Donohue and Thaler—all distinguished behavioral economists who have made careers out of demonstrating that most of us are less-than-fully rational most of the time—suggest disclosure as the preferred regulatory solution." Ben-Shahar and Schneider (2011) succinctly stated that mandating disclosure effectively assumes that people are rational when they are not, so required disclosure is almost certain to fail at improving decisions.

We conduct a lab experiment to examine the question of whether (and what) additional information might affect willingness to pay for insurance. Subjects begin by earning money in a real-effort task, and then face a potential loss to their earnings. We present subjects with a menu of insurance options to cover this potential loss. Within the insurance menu, we provide one of three additional information disclosures, each of which is motivated by theoretical models of insurance demand. For subjects assigned to Treatment 1 (T1), we disclose the true probability of loss, which is intended to address excess demand due to ambiguity aversion or a miscalculation of the loss probability. To subjects in Treatment 2 (T2), we disclose the insurer's expected loss in an attempt to establish a reference price. Subjects in Treatment 3 (T3) are provided the insurer's profit, which may evoke a sense of relative (un)fairness that affects willingness to pay. Each treatment is compared to the other two treatments, as well as to a control group with no additional information.

An example might help to illustrate the policy implications of this research. Imagine mapping the information disclosed in our experiment into the following consumer purchase scenario: You have just purchased a $1,500 Samsung 55" LED Smart TV, and the clerk asks you if you would like to buy a two-year extended warranty. She provides you with the required pamphlet of information on their warranty product. In addition to the warranty price of $179, it provides one of three different types of information:

T1) Historically, about 2.5% of these TVs fail in the first three years of ownership.

T2) Out of the premium you pay, the insurance company sets aside $30.15 to pay for losses.

T3) Out of the premium you pay, $156.33 is dedicated to overhead, marketing, and profits.

We investigate which, if any, of these disclosures is likely to change purchase behavior.

While similar disclosures may affect demand for other relatively expensive financial products (such as payday loans, title loans, or rent-to-own stores), insurance is an ideal setting to test the effectiveness of disclosure for a number of reasons. First, insurance is a relatively familiar product to most individuals in developed countries. Other financial products, such as mutual funds, may not be familiar to subjects in our experiment. Second, insurance can be studied in relative isolation in the lab, as there is no scarcity, resale value, or substitutes. Third, it is simple to manipulate the cost of goods sold for insurance by changing the size and probability of loss. This allows for easy adjustment of parameters relevant to the subject's decision. Lastly, the insurance products subjects evaluate in our experiment are unique to the experiment; therefore, subjects should not have any obvious reference price priors. Even though our experiment focuses on insurance, the results can be applied to regulation of many financial products.

Ultimately, we find that *none* of these information disclosures significantly affects willingness to pay for high-load insurance. Critics of disclosure often argue that disclosure is ineffective at changing behavior, and the results of our experiment support this criticism. Even in a controlled laboratory setting, subjects had difficulty understanding or believing the information provided—less than one-third of subjects who were *given* the probability of loss provided the correct probability of loss less than five minutes later. Learning the insurer's profit did not affect feelings of price fairness, and most subjects were willing to pay profit margins in excess of 50%. Overall, subjects did not appear to use the information disclosed, either because they failed to understand the information or deemed it irrelevant to their purchase decision. These findings provide further criticism of disclosure as an effective solution to the problem of high-load insurance demand. In light of this evidence, other regulatory solutions, such as alternative markets or rate regulation, should be prioritized.

2. Related Research

2.1. Demand for high-load insurance

The purchase of high-load insurance implies extreme levels of risk aversion under the standard expected utility models of demand for insurance. Evidence of this inconsistency has been found in studies of telephone wire insurance (Cicchetti & Dubin, 1994), extended warranties (Jindal, 2014), and low deductibles in home and auto insurance (Barseghyan et al., 2013; Cohen & Einav, 2007; Sydnor, 2010). Most of these studies have proposed alternative models and/or behavioral biases to explain this demand.

The probability of loss is an essential consideration when purchasing insurance, but the true probability of loss is almost never known to a consumer. According to the predictions of ambiguity aversion models, those consumers may seek certainty in insurance when faced with this ambiguity. Hogarth and Kunreuther (1995) conducted an experiment examining demand for extended warranties under ambiguity, finding that a higher proportion of subjects purchased an extended warranty under ambiguity (63%) than under known probabilities and loss sizes (46%). Alternatively, a consumer may misestimate the true probability or amount of loss, making their purchase decision based on their subjective expected utility. Murray (1972) determined that incorrect loss probability estimates were a primary driver of demand for low deductibles in auto insurance. Chen et al. (2009), on the other hand, found that consumers who purchased extended warranties actually had fairly accurate estimates of product failure rates—survey respondents estimated that video equipment, game consoles, and phones had failure rates of 12%, 14%, and 22%, respectively, while the true failure rates were 9%, 18%, and 26%. Huysentruyt and Read (2010) determined that the estimated probability of product failure was not a significant predictor of willingness to pay for extended warranties. In our experiment, we disclose the true probability of loss to subjects in Treatment 1 to counteract the potential effects of ambiguity.

A consumer also may consider the “cost of goods sold” for insurance (i.e. the expected loss the insurer covers) when making his purchase decision. Bearden et al. (2003) investigated how invoice prices in car sales affected value for consumers, finding that higher invoice prices increased acquisition and transaction value for consumers. Bolton and Alba (2006) found that vendor costs significantly influence perceptions of price fairness (and implicitly, expected satisfaction and purchase intention). The authors found that consumers are receptive to changes in alignable costs (i.e. cost of goods sold), but are less willing to absorb vendor costs for services rather than for tangible goods. This aversion to paying for service costs likely applies to a financial institution such as an insurance company. We disclose the insurer’s expected loss to subjects in Treatment 2, which may help subjects realize how large the premium is relative to the risk being transferred.

According to dual entitlement theory (Kahneman et al., 1986), the perceived fairness of a product’s price is influenced by the seller’s profits. Consumers who consider the price unfair under dual entitlement may reduce their demand. Campbell (1999) determined that both relative profit and the firm’s motive for changing prices influenced consumer perceptions of price fairness. In addition, she found that perceived unfairness was associated with lower shopping intentions. Urbany et al. (1989) examined consumer response to changes in ATM fees, finding that cost-justified fee increases were perceived as more fair. The authors determined that even “unfair” fee increases were unlikely to change consumer behavior, primarily due to switching costs. Maxwell et al. (1999) provided subjects in an experiment information on a seller’s pricing process and found that “socially acceptable” pricing strategies led consumers to feel positive about prices. Socially unacceptable pricing elicited feelings that prices were less fair. In a series of experiments, Bolton et al. (2003) found that price fairness was positively associated with seller costs and negatively associated with seller profits, consistent with dual entitlement theory. The authors also found that profits reported in dollar (rather than percentage) terms evoked a stronger response in subjects. To evoke potential feelings of unfairness that may affect demand, we disclose the insurer's profits to subjects in Treatment 3.

2.2. Disclosure

One of the greatest challenges of disclosure is providing information that consumers are able to use. Day (1976) pointed out that disclosures are only effective when the information is (1) readily available, (2) easy to understand, and (3) relevant to the consumer's decision. These three requirements are difficult to ensure, as most transactions that would require disclosure are complex even without the disclosed additional information. We prioritize availability and understanding in the design of our lab experiment and investigate whether the information is deemed by subjects to be relevant.

The information disclosed in our experiment has been recommended or utilized in a number of papers. Both Camerer et al. (2003) and Nalebuff and Ayres (2003) recommended disclosing loss probabilities in the form of historical claim information. To give consumers a sense of the value of coverage relative to the premium, Bar-Gill and Ferrari (2010) proposed disclosing the insurer's expected loss (a proposal supported by Schwarcz, 2010). Birnbaum (1999, 2001) has attempted to raise awareness about the relatively low value of credit life insurance by pointing out the credit life insurance industry's profits.

In a proposal to address the problem of low demand for catastrophe insurance, Kunreuther and Pauly (2004) developed a theoretical basis for all three of our disclosure treatments. Their study focused on the search costs consumers must incur to evaluate the probability of loss. The consumer makes an insurance purchase decision based on the ratio of the premium per dollar of coverage (p) to the probability of loss (π), p/π. This load tells him how large the premium is relative to the actuarially fair rate, helping to determine whether insurance would increase his expected utility. The true probability can be learned by incurring search cost Z, which includes both the cost of paying attention and the cost of seeking out additional information, and reduces wealth whether insurance is purchased or not. In the case of high-load insurance, Z is relatively large and the perceived p/π is relatively low, leading consumers to puchase coverage without bothering to search. Implied in this outcome is that consumers either overestimate search costs or loss probability, and that a more accurate estimate of loss probability would have led them to reject the coverage. Ultimately, the authors concluded that "better information about probabilities as well as about the level of insurer profits and their pricing decisions could help to motivate better insurance purchasing behavior." The authors recommended disclosing the true probability of loss, the expected loss, and/or the insurer's profit to help buyers assess the value of an insurance product.

The finance literature has a multitude of studies focusing on disclosure. Most have found disclosure to be ineffective, including those focusing on mutual fund fees (Barber et al., 2005; Choi et al., 2009, among many others) and mortgage terms (Lacko & Pappalardo, 2010). On the other hand, Durkin (2002) illustrated that awareness of credit card interest rates improved substantially after implementation of the Truth-in-Lending Act (1968). In a promising field experiment, Bertrand and Morse (2011) provided payday lending customers with a clearer explanation of the cost of the loan (a $300 loan accumulated $270 in fees over three months) and found that customers exposed to such disclosure were 11% less likely than others to use payday lending services in the following four months. Ben-Shahar and Schneider (2011) provided an extensive list of state and federal laws mandating disclosure in various contexts, including insurance. The authors also included a survey of empirical studies of disclosure, the majority of which indicate that disclosees either ignore or misunderstand the information disclosed.

Interaction between a discloser and a disclosee creates a social bond that could be awkward to break if the information changes the buyer's decision. Sah et al. (2013) provided an excellent series of experiments showing how disclosure of conflicts of interest might generate conflicting emotions for consumers, decreasing trust in the seller but increasing social pressure to comply (the "panhandler effect"). This effect was apparent in the loss ratio disclosure experiment of De Meza et al. (2010), where sellers were instructed to persuade buyers to overlook disclosed information and purchase insurance despite high loads. This is an important consideration in our experiment. Because we want to examine the value of the *information alone*, we disclose the additional information via text on the insurance menu and not via a human discloser. This avoids the potential social complexities that interactive verbal disclosure might create.

3. Experiment design

3.1. Overview

We conduct an experiment with subjects in a lab setting. Our experiment includes five stages. In Stage 1, subjects complete a real effort task (categorizing grocery store items) for twenty minutes to earn $20. Once all subjects have completed the task, they move to Stage 2 (the "insurance task"), where the money they earned in Stage 1 may be lost based on the draw of a red chip from a clear jar of red and white chips. Subjects in this stage are presented with a menu of insurance options that they may buy to cover the risk. The menu specifies the size of the loss, the amount of insurance, and the price of insurance. Each subject is randomly placed into one of four additional information conditions that may influence their purchase decision, though they are not aware of this selection. The control condition is no additional information. Subjects in Treatment 1 are provided the true probability of loss (the percentage of chips that are red). Subjects in Treatment 2 are provided the insurer's expected loss for each insurance option. Subjects in Treatment 3 are provided the insurer's profit. With this information, subjects make several insurance decisions, after which they move to Stage 3 (the "estimation task"). During this stage, subjects provide estimates of the percent of red chips in the jar. Stage 4 is a task to elicit each subject's appetite for risk (the "risk aversion task") . In Stage 5, subjects answer a series of questions to test their understanding of insurance pricing and then complete a demographics questionnaire. A chart in Appendix 1 illustrates this process.

3.2. Detailed description

As subjects arrive, there are two large clear jars of red and white chips displayed prominently at the front of the room. Subjects walk past these jars to receive their $5 show-up fee, which is not subject to loss in the experiment. Each subject signs the participant log for the show-up fee, chooses a computer in the lab, and signs a form consenting to the remainder of the experiment. Subjects must stay through the end of the experiment to keep any further earnings (in excess of the $5 show-up fee).

Printed instructions for Stage 1 are at each station. This stage is a "real effort" task to categorize grocery store items based on receipt codes and each subject will be paid $20 for his/her work. We (the experimenters) read the Stage 1 instructions aloud, answer any questions, and then give subjects twenty minutes to complete the task. At the end of twenty minutes, each subject is paid with a $10 bill, a $5 bill, and five $1 bills and asked to count the money and keep it on their desk. We then distribute instructions for Stage 2.

At the start of Stage 2, we provide a short overview by telling subjects that some or all of the $20 they just earned may be lost based on the draw of a red chip from one of the jars at the front of the room. Before continuing with the instructions, we pass the jars around the room so that each subject may examine them. Each jar contains a total of 3,200 chips. A draw from Jar X has a 15% probability of loss, while a draw from Jar Y has a 5% probability of loss. Though we do not state so explicitly, it is clear to all that the chance of loss from Jar X is higher than the chance of loss from Jar Y. We state that subjects may come to the front of the room and examine the jars at any time during the experiment, and do not need to ask or raise their hand to do so.

Once the jars have been passed around to all subjects, we continue reading aloud the instructions with subjects following along on their printed versions. Subjects are told they will face a number of scenarios which vary the amount of the loss and the jar from which the chip is drawn. The subjects will evaluate four scenarios varying the amount of loss ($10 or $20) and the jar (X or Y). We explain that at the end of the experiment, we will randomly select one of the scenarios to be binding, and we will play out the chosen loss scenario with red chips designating a loss of the specified amount and white chips designating no loss.

We explain that subjects will have an opportunity to purchase partial or full insurance, but they are not required to buy it. The insurance is priced based on the size and chance of loss, so the price of the insurance will change between scenarios. We ask subjects to imagine that insurance is offered by “Acme Insurance Company,” who charges them a premium in exchange for paying losses that may occur. Acme sets aside a portion of the premium to pay for losses and keeps the rest as profit (implicitly ignoring expenses such as overhead, which might unnecessarily complicate the experiment). The company sets premiums by multiplying the chance of loss times the amount of insurance provided, and adds profit to that amount. Specifically:

$$Premium = [(\%\ Chance\ of\ Loss) \times (Amount\ Insured)] + Profit$$

There is no risk that Acme will go bankrupt or otherwise be unable to cover losses. All of these details (including the pricing formula) are provided to subjects in the printed instruction sheet on their desk.

In each scenario, subjects choose from a menu of insurance levels. Figure 1 provides a sample menu. Subjects may purchase no insurance or "coinsurance" on 20%, 40%, 60%, 80%, or 100% of the loss. These levels of insurance are specified as dollar amounts (e.g. for a $10 loss, they may purchase $2 of insurance, $4 of insurance, etc.). Loads for these insurance products increase as the coverage increases—20% coverage has a load of 1, 40% coverage has a load of 2, etc. up to a load of 5 for full insurance. Table 1 provides the premiums for each coverage option in each of the four scenarios.

After the instructions have been read and any questions answered, subjects may begin the insurance task (Stage 2). Subjects start this stage with a practice scenario and then continue to

Table 1: Insurance premiums for each loss scenario

Scenario	Total Loss	Loss Prob	Coverage				
			20%	**40%**	**60%**	**80%**	**100%**
20X	$20	15%	$0.60	$2.40	$5.40	$9.60	$15.00
10X	$10	15%	$0.30	$1.20	$2.70	$4.80	$7.50
20Y	$20	5%	$0.20	$0.80	$1.80	$3.20	$5.00
10Y	$10	5%	$0.10	$0.40	$0.90	$1.60	$2.50
Load			1	2	3	4	5
Implied profit margin			0%	50%	67%	75%	80%

the four insurance scenarios. Once each subject has completed the scenarios, she is asked to rate how important each piece of information was in making her decision and how fair she found Acme's pricing. Each is rated on a seven-point Likert scale.

After all subjects complete Stage 2, we provide the password to continue to Stage 3 where they complete a series of estimation tasks. These tasks are intended to elicit subjects' estimates of the probability of loss from each jar. Each subject is asked to estimate the percentage of red chips in Jar X and Jar Y using a slider between 0% and 50%. At the end of the experiment, we will randomly select one of the two responses (Jar X or Jar Y) and pay an additional $1 if the estimate for that jar is correct.

We follow the estimation task in Stage 3 with a risk aversion task in Stage 4. This stage replicates part of the risk aversion experiment in Holt and Laury (2002, 2005), where each subject makes a series of ten choices between two gambles, a "safe" Gamble A (win $2.00 or $1.60) and a "risky" Gamble B (win $3.85 or $0.10). The first choice includes only a 10% chance of winning the high outcome for both gambles, so all but the most risk-seeking subjects will choose Gamble A. Subsequent questions increase the chance of winning the high outcome

Figure 1: Sample menu of insurance choices

Potential Loss	$20
Jar	X

You will lose $20 if you draw a red chip from Jar X and lose $0 if you draw a white chip from Jar X.

Insurance product	A No Insurance	B	C	D	E	F Full Insurance
In the event of loss... Acme covers You lose	 $0 $20	 $4 $16	 $8 $12	 $12 $8	 $16 $4	 $20 $0
Premium (price of insurance)	$ -	$0.60	$2.40	$5.40	$9.60	$15.00
Treatment info (if any) goes here	*Info*	*Info*	*Info*	*Info*	*Info*	*Info*

Which insurance product would you like to buy?

- ○ A: Pay $0 for no insurance
- ○ B: Pay $0.60 for $4 of insurance
- ○ C: Pay $2.40 for $8 of insurance
- ○ D: Pay $5.40 for $12 of insurance
- ○ E: Pay $9.60 for $16 of insurance
- ○ F: Pay $15.00 for $20 of insurance

[Treatment Info]						
[T1] Chance of loss	15%	15%	15%	15%	15%	15%
[T2] Portion of premium set aside for losses	$0.00	$0.60	$1.20	$1.80	$2.40	$3.00
[T3] Portion of premium kept as profit	$0.00	$0.00	$1.20	$3.60	$7.20	$12.00

by 10%, so every subject will eventually switch to choosing Gamble B as the probability of winning $3.85 approaches 100% in the last choice. Risk aversion is approximated based on the number of times the subject chooses the safe Gamble A. As in the previous stages, one of their ten choices will be randomly selected to play out at the end of the experiment.

Stage 5 includes a number of follow-up questions. First, subjects answer three “understanding” questions intended to gauge how well each subject was able to use the given pricing formula to calculate the information contained in other information treatments. For example, the subject may be given the probability and size of loss and asked to calculate the expected loss. In another, the subject might be asked to calculate the profit given the premium and expected loss. One of these questions is randomly selected at the end of the experiment, and subjects who correctly answered the selected question are paid $1.

Subjects then complete a questionnaire asking about demographic and financial information, including their real-life experience with insurance. The final question asks the subject what she thought we were trying to find out in the experiment, allowing for free response in a text entry box. This is to determine whether there is any systematic bias or confounding factor overlooked by the experimenter.

Following Stage 5, each subject individually goes to the experimenter to play out their randomly-selected choices. The subject first pays the premium she chose in the selected insurance scenario from Stage 2. The relevant jar (X or Y) is then covered and she blindly draws a chip. If the chip is red, the subject pays the amount of loss less any insurance coverage. If the chip is white, there is no loss payment required. The subject then plays out the selected risk aversion gamble from Stage 4. Finally, the subject is paid $1 for any correct probability estimates she made in Stage 3 and is paid an additional $1 if she answered the selected Stage 5 "understanding" question correctly. Once all payments have been made and totaled, the subject signs the participant log and leaves.

3.3. Treatments

The insurance task in Stage 2 is the primary focus of this experiment. Any differences in the distribution of responses between the control condition and the treatment conditions will provide insight into *whether* additional information influences choices. Differences between treatments address *what* information helps to encourage more rational choices. The information provided was the only difference between treatments. For each treatment, the additional information is in the same location on the insurance menu, in a row below the premium as shown in Figure 1. Subjects in the control condition were not provided any additional information.

Subjects in Treatment 1 are provided the true probability of loss (i.e. the percent of red chips in the jar). This is framed simply as the “Chance of loss” in the insurance menu. This treatment is intended to mitigate demand due to mistaken calculation, subjective expected utility, and/or ambiguity aversion. Disclosure of the loss probability was a primary focus in Kunreuther and Pauly (2004), though their focus was on consumers who underpurchase catastrophe insurance. An equivalent policy implementation might be providing historical claim information, as recommended in Camerer et al. (2003) and Nalebuff and Ayres (2003).

Subjects in Treatment 2 are provided the expected loss or “cost of goods sold” for Acme Insurance Company. This disclosure is framed as the “Portion of the premium set aside for losses” and displayed as a dollar amount. This information is intended to adjust reference points, as it may help subjects realize how large the premium is relative to the risk being transferred. The low seller costs relative to price also may evoke feelings of unfairness, which may decrease demand. This disclosure is in line with the policy recommendations of Bar-Gill and Ferrari (2010) and Schwarcz (2010).

Subjects in Treatment 3 are provided Acme’s profit on each level of insurance. Similar to Treatment 2, this is displayed as a dollar amount and labeled “Portion of premium kept as profit.” This treatment is intended to make subjects feel that the pricing is increasingly unfair as the levels of insurance increase. Dual entitlement theory predicts that this treatment will make the higher coinsurance policies less desirable, as the large relative profits (up to 80% profit margins) earned by the insurer may be seen as unfair by insurance purchasers. Disclosing the insurer's profits mirrors the activism of Birnbaum (1999, 2001). Academic research has shown that consumers find high profits unfair (Bolton et al., 2003; Kalapurakal et al., 1991, etc.), though this sense of unfairness does not necessarily translate to willingness to pay.

3.4. Recruitment and selection

Subjects were 134 students from a large public university, recruited from a voluntary online participant pool. The recruitment posting was titled "Insurance Study" and stated that most participants were expected to earn $10-25 depending on the choices they made (ultimately payments ranged from $13.70 to $32.25, with an average payment of $25.56). There were fourteen rounds of the experiment conducted over four days in April 2014, with between seven and fifteen subjects per round. Rounds were conducted at 9:00am, 11:00am, 1:00pm, and 3:00pm, and each round took 75-90 minutes depending on the number of subjects participating (randomly choosing and playing out the gambles at the end took several minutes per subject).

Most of the subjects were undergraduates, with approximately 27% freshmen, 20% sophomores, 18% juniors, 20% seniors, and 15% graduate or professional students. Males comprised 48% of subjects. The average age of subjects was 22, though some were as young as 18 and as old as 57. We did not collect any information on race, religion, or nationality.

4. Data

4.1. Probability estimates

The probability of loss is an essential consideration in insurance purchases and the disclosure to Treatment 1 subjects was intended to provide perfect information on the probability of loss. Subjects in the other information treatments also could have calculated the probability with relatively simple arithmetic. After making their insurance choices, subjects were asked to provide their estimated probability of loss from each jar. We call this estimate the subjective probability estimate (*s*). Figure 2 illustrates errors in these estimates, with reference lines for accurate estimates (error = 0). Subjects typically overestimated the probability—Jar X was overestimated by 5.2% on average, while Jar Y was overestimated by 5.44%. Subjects in Treatment 1 provided the most accurate estimates, with a mean estimate error of 2.5% and 4.1% for Jar X and Jar Y, respectively. Estimate errors from subjects in Treatment 3 were higher than any of the other groups.

Surprisingly, a majority of Treatment 1 subjects provided incorrect estimates (80% for Jar X and 69% for Jar Y), even though they were *explicitly told* the correct answer in the preceding stage. We suggest three possible reasons for this. First, subjects in this treatment may not have believed the information provided in the insurance task, perhaps thinking they were being tricked by the experimenter. Second, they may have overlooked or ignored the probabilities when making their insurance choices. Third, subjects may have forgotten which probability applied to which jar, or forgotten the probabilities altogether. These (mis)estimates provide an initial indication that information disclosed (even in its simplest form) may not be understood or used by disclosees.

4.2. Insurance choices

Subjects made insurance purchase decisions in four independent scenarios: 20X and 10X (each with a 15% probability of loss) and 20Y and 10Y (each with a 5% probability of loss). The primary task in each scenario was for subjects to choose one of six coverage levels (0% to

Figure 2: Distribution of errors in estimating probability of loss, by treatment

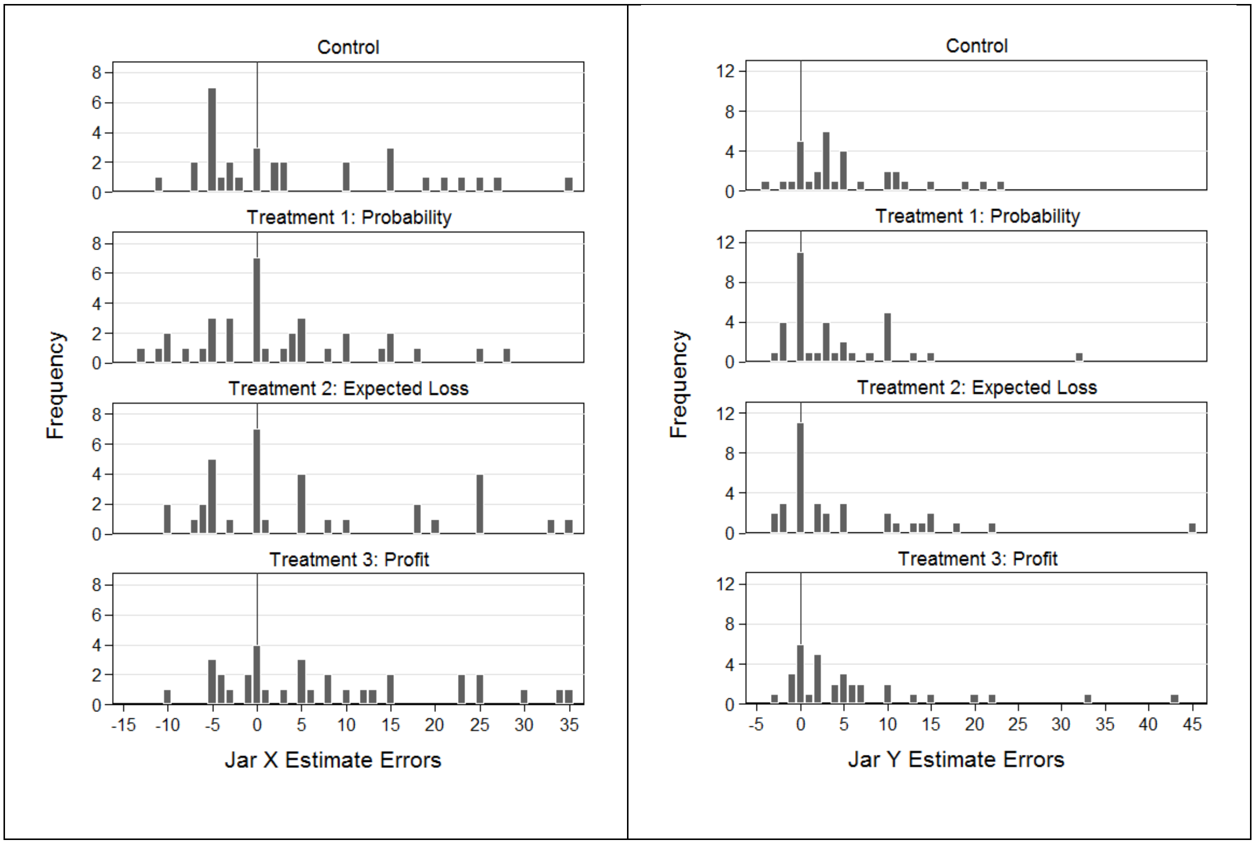

100% of the potential loss, in 20% increments). Histograms of the choices made are provided in Figure 3. Each set of histograms represents a scenario (20X in the top left, 20Y in the top right, etc.), and there are four plots within each scenario for each of the four information conditions. The x-axis ranges from no insurance on the left to full insurance on the right, and the y-axis represents the number of subjects who chose the given level of coverage.

The treatments within each scenario appear to have a similar distribution of choices, particularly for the 20X scenario. There does not appear to be a treatment that consistently results in higher or lower choices than the other treatments, though Treatment 3 (Profit) tends to look slightly different from the others. Very few subjects purchased full insurance on draws from Jar

Figure 3: Histograms of coverage choices, by scenario and treatment

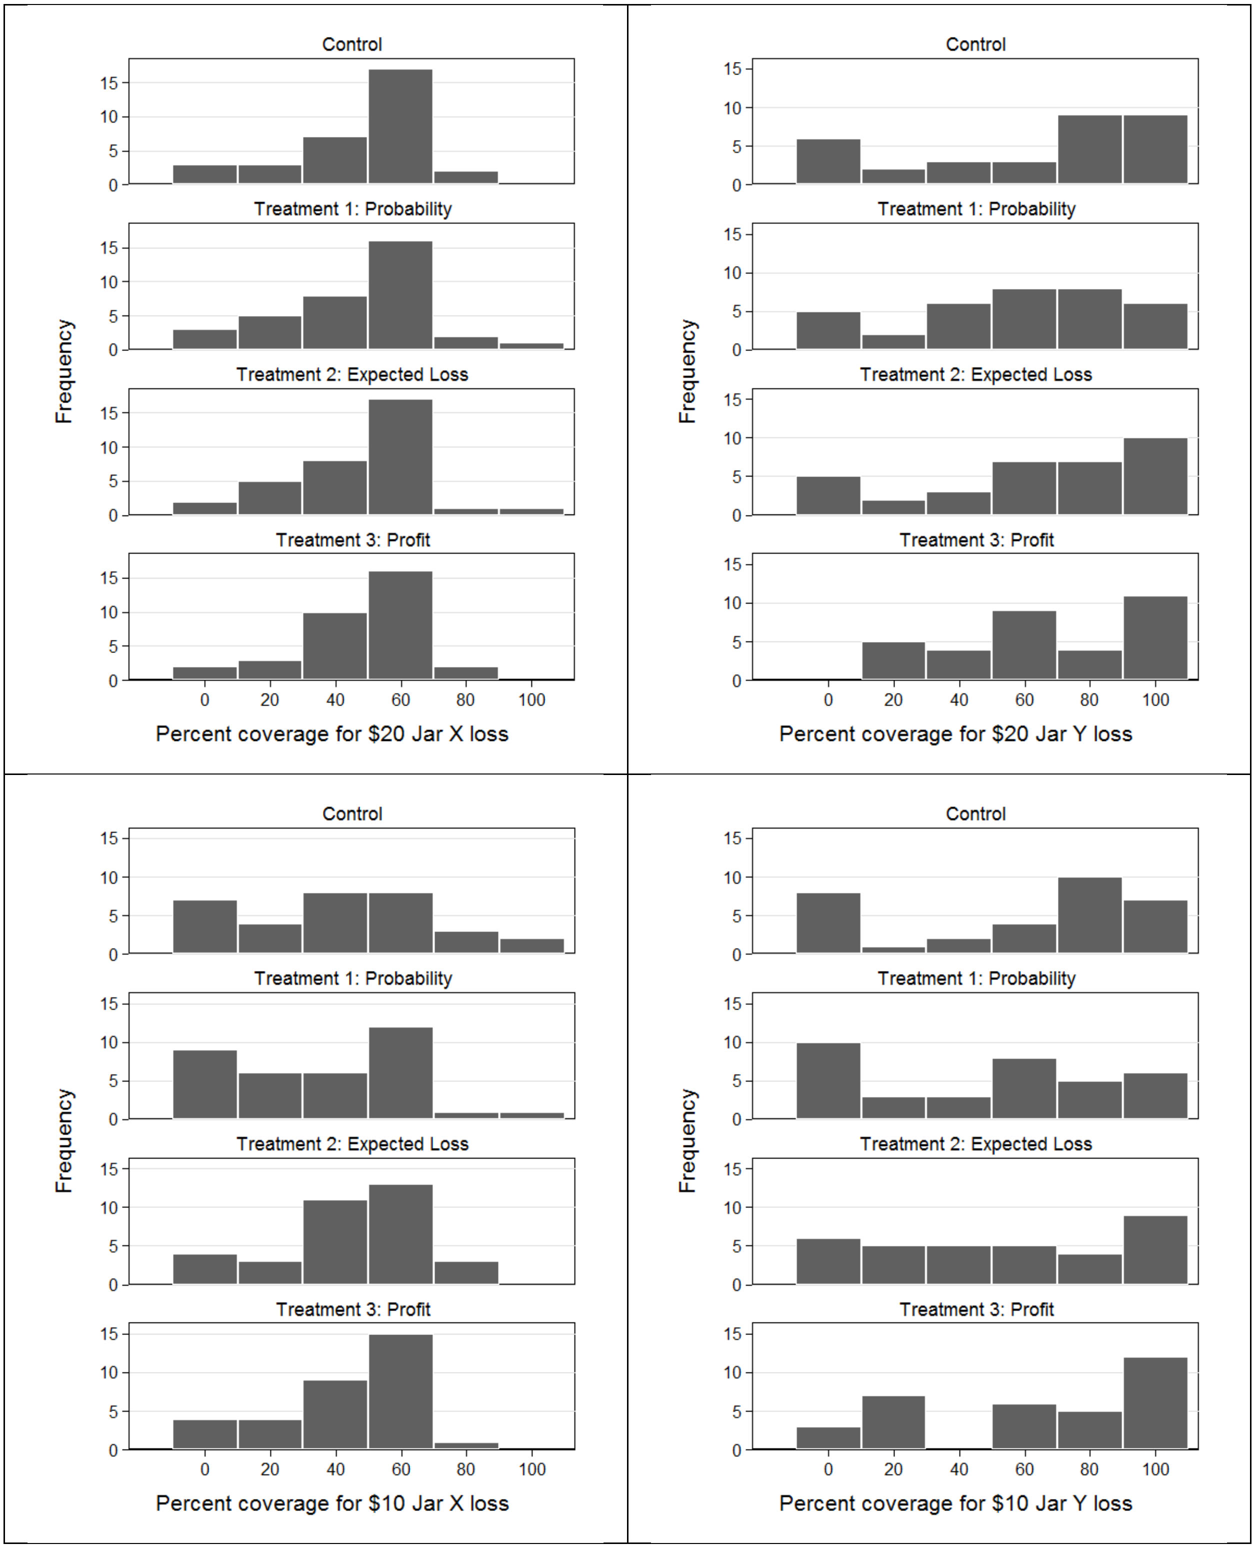

X, likely because full insurance cost 75% of the insured amount ($15 to insure a $20 loss and $7.50 to insure a $10 loss). Many more purchased full insurance for Jar Y scenarios. These choices are positively related to estimates of loss probability—the correlation between each scenario's coinsurance choice and its respective probability estimate was 0.25 for 20X, 0.27 for 10X, 0.23 for 20Y, and 0.22 for 10Y.

These probability estimates obviously are important in making choices, so we create a "subjective load" variable to incorporate the subject's beliefs regarding the chance of loss. By multiplying a subject's selected coinsurance level by his estimated chance of loss *s*, we obtain the subject's expected value of coverage. Dividing the premium by this expected value gives the subjective load (*SUBJLOAD*) paid by the subject, as illustrated in Equation 1. In other words, the subjective load is the load the subject *believed* he was paying.

$$SUBJLOAD = \frac{Prem\ Paid}{(s \times Coverage)} \tag{1}$$

We use SUBJLOAD as our primary choice variable of interest. This load measure is a more appropriate representation of each subject's willingness to pay for coverage, as it does not assume the subject knows the true probability of loss. The point estimates given by Treatment 1 subjects show that even if the true probability of loss is "known," it may not be the probability used to make decisions. Figure 4 illustrates the distribution of these loads. These histograms are the same as Figure 3, but the x-axis represents the subjective loads ranging from 1-10. There were three subjects who paid subjective loads greater than 10, but including them in the histograms made it difficult to see the rest of the distribution. These subjects paid subjective loads of 37.5 (10X and 20X), 15 (20X), and 15 (10Y). These high subjective loads stem from a combination of low probability loss estimates and high coverage choices.

Figure 4: Distribution of subjective loads paid in each scenario, by treatment

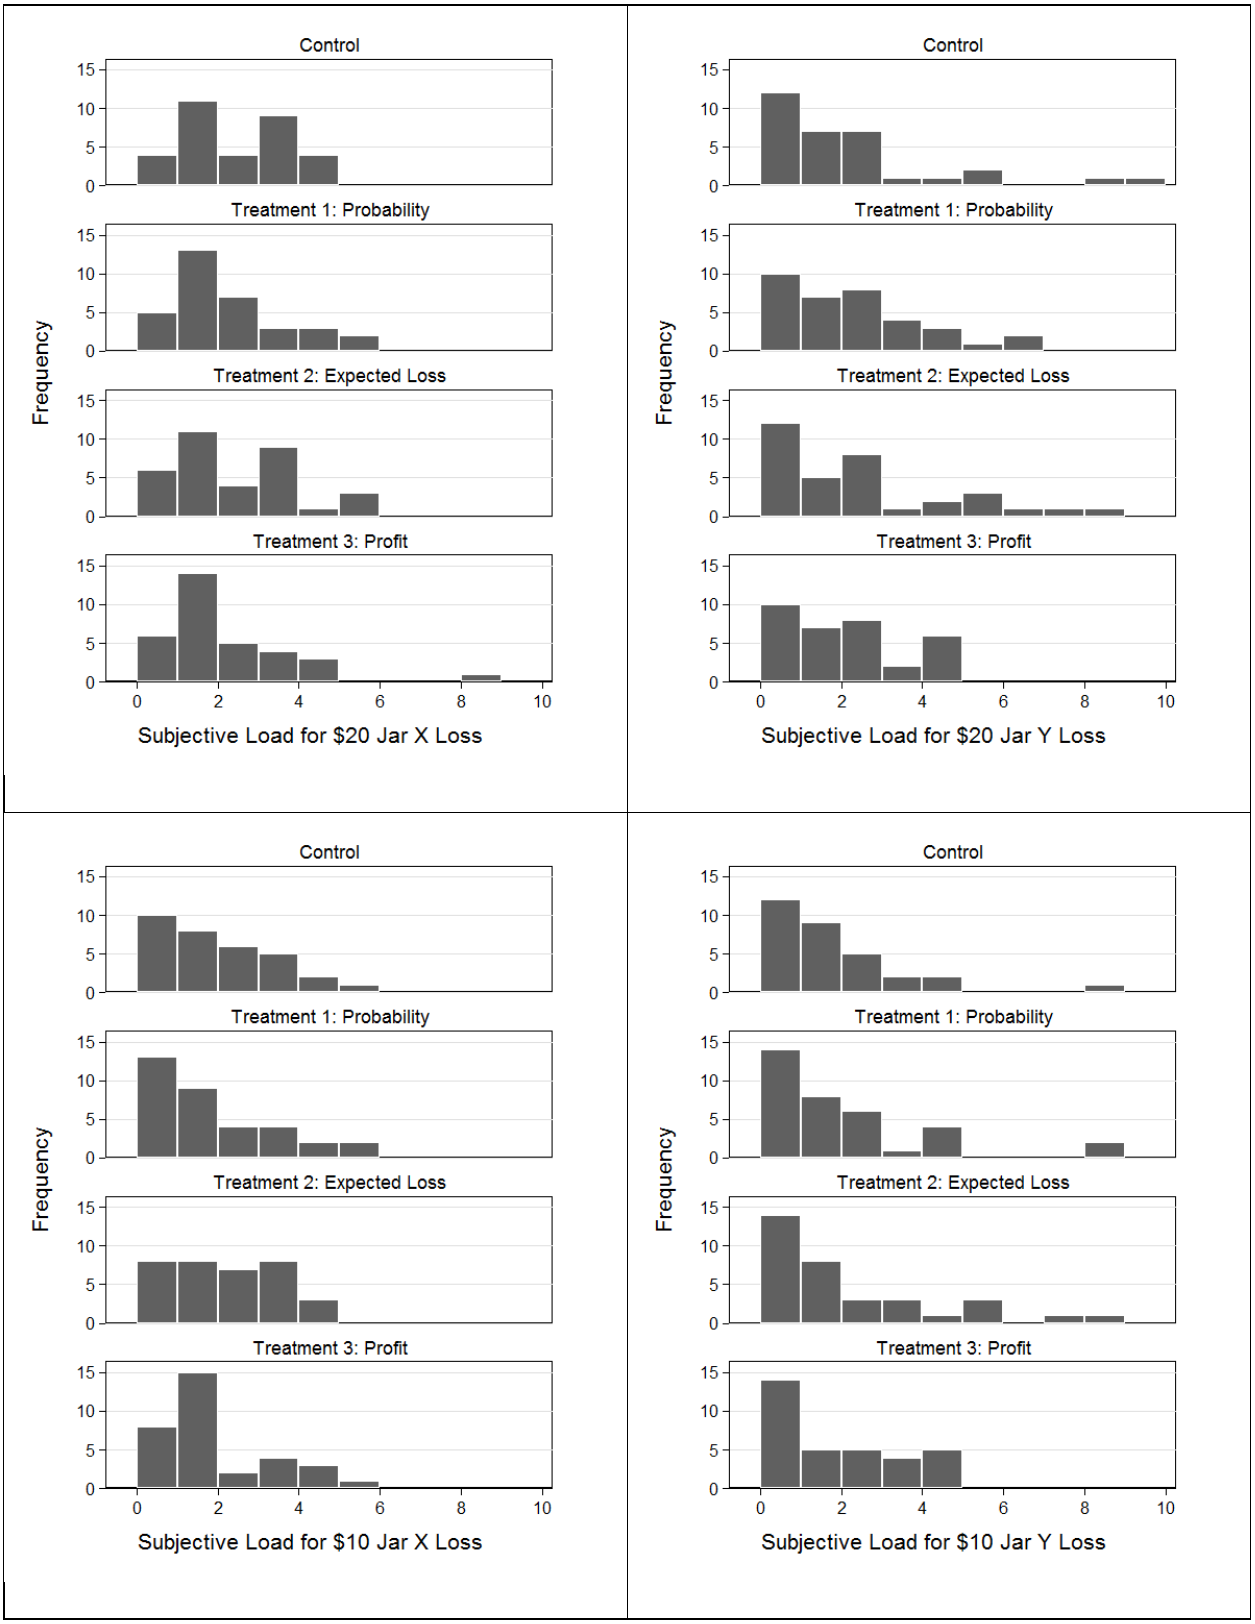

On average, subjects were willing to pay between two and three times what they believed was the expected loss. Many subjects believed they were paying loads under 2, in line with auto and homeowners insurance. Subjects paid higher subjective loads for $20 potential losses than for $10 potential losses, and slightly higher subjective loads for Jar X draws than Jar Y draws (with the exception of Treatment 3 subjects, who consistently paid more for Jar Y draws). Subjects seemed to think they were paying much lower loads than they actually were—the subjective load was smaller than the actual load for approximately 72% (56%) of subjects evaluating Jar X (Jar Y). One contributing factor to this is that subjects often overestimated the probability of loss—66% of subjects overestimated Jar X and 87% overestimated Jar Y.

4.3. Other response data

Is it possible that subjects simply did not understand the information provided? To test their understanding and use of information during the experiment, subjects answered three "understanding" questions at the end of the experiment. These questions were intended to gauge subjects' ability to calculate information available in other treatments. During this task, the Stage 2 instructions (with the pricing formula) were still in front of each subject. Subjects answered each question by moving a slider between 0 and 100. On average, subjects answered 1.45 of the three questions correctly. Subjects in Treatment 2 did relatively well in these questions, on average answering 1.97 questions correctly (all other subjects averaged 1.27 correct). These results suggest that subjects had at least some understanding of how premiums and profits were determined.

Subjects were asked to rate the "fairness" of Acme's pricing strategy on a seven-point Likert scale from "Very unfair" (with a value of -3) to "Very fair" (with a value of $+3$).

Coverage above 20% was intended to seem very unfair with regards to profits. Most subjects felt relatively neutral on the subject, giving an average fairness rating near zero. Subjects in Treatment 3 found the pricing to be slightly *more* fair (mean rating of 0.5, not significant), which conflicts with the predictions of dual entitlement theory. Subjects may have only considered the fairness of their selected insurance product, though profit margins for the average purchase were more than 66%. It also is possible that subjects looked at dollar profits rather than profit margins, or did not think about profits at all in evaluating fairness.

Responses to the demographic and financial questions generally were as expected for the given student population. All of the subjects had a checking account and most had savings accounts (92%) and credit cards (65%). Only 4% had auto loans. The most common monthly housing payment was between $500 and $750 (47% of subjects). Many subjects had some degree of experience with auto insurance (47%), cell phone insurance (41%), and/or extended warranties (45%). Fewer had purchased renters (25%) or homeowners (13%) insurance, likely because they had coverage under their parents' policies.

5. Methodology

To answer the questions of *whether* new information affects insurance choices, we compare the distribution of choices from the control condition to each of the treatment conditions. We also determine *what* information may influence choices by comparing the additional information treatments to each other. We use the subjective load (*SUBJLOAD*) described in Section 5 as the primary variable of interest and use a number of methods to test for differences in these values between treatments.

5.1. Univariate tests

We use the Wilcoxon rank-sum test (WRST) to compare distributions of *SUBJLOAD*, as the test requires only that the observations are random samples and are mutually independent. Both of these assumptions are appropriate for our data, as subjects were randomly selected into each of the treatment groups. The WRST does not assume normally distributed data. The WRST first pools all observations from two samples (*j* = 1, 2). It then assigns a rank $R_{i,j}$ to each observation *i* based on size. The test statistic *T* is the sum of the ranks for the observations in the first sample (*j* = 1):

$$T = \sum_{i=1}^{n_1} R_{i,1} \tag{2}$$

The first and second moments of this test statistic are then used to calculate a z-score $\left(z = \frac{T - E(T)}{\sqrt{Var(T)}}\right)$, which can be compared to a normal distribution. Since the WRST is based on ranked observations rather than means, it is not sensitive to outliers (an important consideration for subjective load). The WRST may have issues with Type I error if the variance between treatments is heterogeneous.

5.2. Regression analysis

It is possible that other factors may influence choice, confounding the effect of treatments. To control for these factors, we regress *SUBJLOAD* on treatment variables and control variables. *SUBJLOAD* is a continuous measure ranging from 0 (no insurance) to infinity (as the expected loss goes to zero). It cannot be negative, so we specify a Tobit model which is left-censored at zero. The *TRT* variables are indicators of the treatment assigned to subject *i*, where *TRT* is equal to 1 if the subject was in that treatment and 0 otherwise. Each of these treatments is evaluated relative to the control condition, which is the omitted category. A significant coefficient on any *TRT* variable will indicate that the information treatment affected choices, and a negative coefficient will indicate that the additional information decreased willingness to pay for the high-load insurance in our experiment.

$$SUBJLOAD = \alpha_j + \beta_1 TRT1_i + \beta_2 TRT2_i + \beta_3 TRT3_i + \boldsymbol{\beta}\mathbb{X}_{\boldsymbol{i}} + u_i \qquad (3)$$

The matrix $\mathbb{X}_{\boldsymbol{i}}$ includes control variables for estimates of the probability of loss, price fairness ratings, risk aversion (number of safe bets in the Holt & Laury task), the number of understanding questions correct, sex, age, monthly housing payment (a proxy for income), GPA, and indicator variables for credit cards and savings accounts. It is most important to control for loss probability estimates, since we ultimately want to determine whether the implied load itself changed due to the disclosures. The coefficient on probability estimates will be negative—if someone grossly underestimated the probability of loss, their implied load from any insurance purchase would be large. We also expect that risk aversion (as measured by the safe bets made) and price fairness ratings will be positively related to *SUBJLOAD*. We specify robust standard errors in this model.

6. Results

We compare distributions of the choice variables using the WRST and find that, of the 24 tests conducted and reported in Table 2, the null hypothesis of equal distributions cannot be rejected in any case. For ease of interpretation, we report the statistic ρ, which is an estimate of the probability that a draw from the first population will be larger than a draw from the second population (for the null hypothesis of no difference between distributions, ρ = 0.5). The significance of the tests, however, is determined based on the Z-score described in Section 5.1. According to the WRST, the distributions of choice variables are not significantly different between treatments, regardless of the additional information provided. Type I error is not a problem in the WRST results, as we did not reject any null hypotheses. We did, however, conduct Welch t-tests and Kolmogorov-Smirnov tests for robustness, without any difference in results.

We report the results for the Tobit model in Table 3. None of the treatments is significantly related to *SUBJLOAD*, providing further evidence that the disclosures in our experiment did not affect demand for insurance, even after controlling for other factors. While Treatment 3 has a consistently negative sign, none is even close to statistically significant (the lowest p-value is 0.29).

As expected, the Probability Estimates are negatively related to the subjective load due to how the load is calculated. Consistent with the influence of risk aversion on demand for insurance, the level of risk aversion (Safe Bets) is positively associated with the subjective load (though not significant in all cases). Feelings of price fairness do not appear to be consistently related to demand. Subjects who reported having a credit card tended to pay lower loads, while subjects who had a better understanding of the pricing models paid higher loads in the Jar Y scenarios. The Jar Y scenarios were more “reasonably” priced in pure dollar terms, particularly with respect to marginal premiums for increasing levels of coverage. We conduct a number of robustness checks without any difference in results. These include using log(*SUBJLOAD*) as the

Table 2: Wilcoxon rank-sum test results of the difference in distribution of SUBJLOAD between treatments

	20X	**20Y**	**10X**	**10Y**
Control - T1	0.512	0.444	0.533	0.508
Control - T2	0.523	0.465	0.444	0.486
Control - T3	0.568	0.437	0.513	0.445
T1 - T2	0.514	0.514	0.414	0.476
T1 - T3	0.563	0.493	0.466	0.440
T2 - T3	0.544	0.475	0.568	0.475

Note: The statistic ρ is reported in each cell. This statistic is an estimate of the probability that a random draw from the first population will be larger than a random draw from the second population. Significance is determined based on the Z-score as outlined in Section 5.1.

dependent variable in an OLS model, using coinsurance choices as the dependent variable in an ordered logit model, and specifying a GLS random effects model with subject as the panel variable and choice as the time variable.

Overall, these results provide evidence that disclosure is not effective in changing willingness to pay for high-load insurance. Even after controlling for other factors that may affect demand, the additional information disclosed in our study did not significantly affect the decisions made by the subjects. While we made every attempt to provide potentially useful information in a clear but unobtrusive manner, subjects did not appear to understand, remember, or use the information provided.

Table 3: Tobit model regression of subjective load on treatment and control variables

Dep var: *SUBJLOAD*	**20X**	**10X**	**20Y**	**10Y**
Treatment 1	0.913 (0.736)	0.358 (0.818)	-0.063 (0.525)	-0.507 (0.702)
Treatment 2	0.119 (0.542)	0.226 (0.600)	-0.073 (0.583)	-0.344 (0.750)
Treatment 3	0.061 (0.475)	0.070 (0.595)	0.480 (0.451)	0.380 (0.569)
Probability Estimate	-0.128*** (0.041)	-0.091** (0.042)	-0.103*** (0.020)	-0.100*** (0.026)
Price Fairness	0.120 (0.183)	0.254 (0.182)	-0.263* (0.150)	-0.233 (0.200)
Safe Bets	0.263* (0.149)	0.199 (0.160)	0.045 (0.079)	0.208* (0.109)
Credit Card	-0.283 (0.496)	-0.400 (0.529)	-0.334 (0.371)	-0.823* (0.426)
Understanding	-0.385 (0.360)	-0.225 (0.360)	0.234* (0.140)	0.548** (0.185)
N	134	134	134	134
Log-Likelihood	-349.446	-326.622	-276.786	-281.444
AIC	672.130	641.308	547.902	558.606

Robust standard errors in parentheses. Stars *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. A number of control variables were included in the analysis, but coefficients were not significant and did not have consistent signs across scenarios. These variables included age, sex, GPA, housing payment (a proxy for income), and indicator variables for savings accounts and auto loans held by the subject.

7. Limitations

The results of this analysis indicate that the additional information provided has no effect on insurance demand, but there may be alternative explanations for the null result. It is possible that the dollar values were too small to evoke careful consideration by subjects. In three of the four scenarios, full coverage was available for $7.50 or less. Of the twenty total insurance options available across scenarios, only five cost more than $5.00. These premiums were large relative to expected loss, but were small in total dollar amounts (less than 25% of subjects' earnings). Insurer profit margins were high, but the dollar amounts going to profits were relatively small. For example, the insurer made an 80% profit margin for selling full coverage in any scenario, but that equated to only $2 in the 10Y scenario—and 25% of subjects purchased full insurance in that scenario. In scenario 20X, the insurer made a $12 profit on the $15 premium for full insurance, which only two subjects purchased. This provides a sense that at least some of the dollar profit amounts were not large enough to elicit feelings of unfairness. The tendency to pay for full coverage as stakes decreased, however, is consistent with high willingness to pay for low levels of coverage (e.g. extended warranties).

Sample selection also may be a source of the null result. The students in our sample are not demographically representative of the general population. Undergraduates 18-22 years old may not have had sufficient experience purchasing insurance or other financial products to evaluate whether prices were acceptable. College students may have different risk preferences than the general population—age, marital status, and other life cycle factors have been shown to influence risk aversion.

A number of follow-up studies with alternative specifications might shed light on our results. Disclosing profit and expected loss in percentage terms (rather than dollar terms) might help to address the potential issue of scale. Telling subjects “market rates” for the policies (perhaps based on loads for life or auto insurance) might give subjects a more salient anchor point to evaluate the experiment premiums. A follow-up study with a more nationally or globally representative subject pool also may arrive at different results, particularly as income, age, education, and occupation vary. While the null results of this experiment are interesting in themselves, follow-up studies would determine whether similar results occur with consistency.

8. Conclusion

We conducted an experiment to evaluate whether information disclosures affect demand for high-load insurance. Subjects in our experiment were faced with a potential loss of earned money based on the draw of a red chip from a clear jar of red and white chips. Subjects were offered a menu of insurance options to protect their money from potential losses. All subjects were told the amount of loss, the dollar amounts covered, and the premium for each insurance product. The insurance products were designed to be relatively expensive, with the insurance costing between one and five times the expected loss (generating profit margins of 50-80% for the hypothetical insurer). Subjects also received one of four additional information treatments: no additional information (Control), the true probability of loss (Treatment 1), the expected loss (Treatment 2), or the insurer's profit (Treatment 3). The theoretical literature on insurance demand motivated the information provided in each of these treatments, and the disclosed information was provided in a clear but unobtrusive manner in an effort to avoid the social confounds inherent in person-to-person disclosure.

Ultimately, the disclosed information had no effect on demand for this insurance—the distribution of choices did not differ between information treatments. A major driver of demand was subjects' estimates of the loss probability, in line with subjective expected utility models of insurance demand. Most subjects overestimated the probability of loss—even those who were explicitly told the true values. We take this as evidence that subjects did not understand, did not trust, or did not remember the disclosures provided. While this was an unexpected result in our highly controlled experiment, it is likely that mandated disclosure in practice would create the same mental errors.

These results have policy implications for regulation of insurance and other financial products. Critics of high-load insurance have called for enhanced disclosure, publicly condemned high profit margins, and recommended against purchasing based on the small benefit relative to the cost. Our experiment provides evidence that insurance buyers are not sensitive to any of these framings. Buyers appear not to care about the expected value of coverage or the insurer's profit, instead focusing on the size of the potential loss and how much it will cost them to cover it. While disclosure is a popular policy recommendation, the results of our experiment add to the evidence that disclosure does not significantly change behavior. Alternative policies, such as enhancing market competition or regulating rates, likely would prove more effective in encouraging more economically rational consumer choices.

Appendix 1: Experiment flowchart

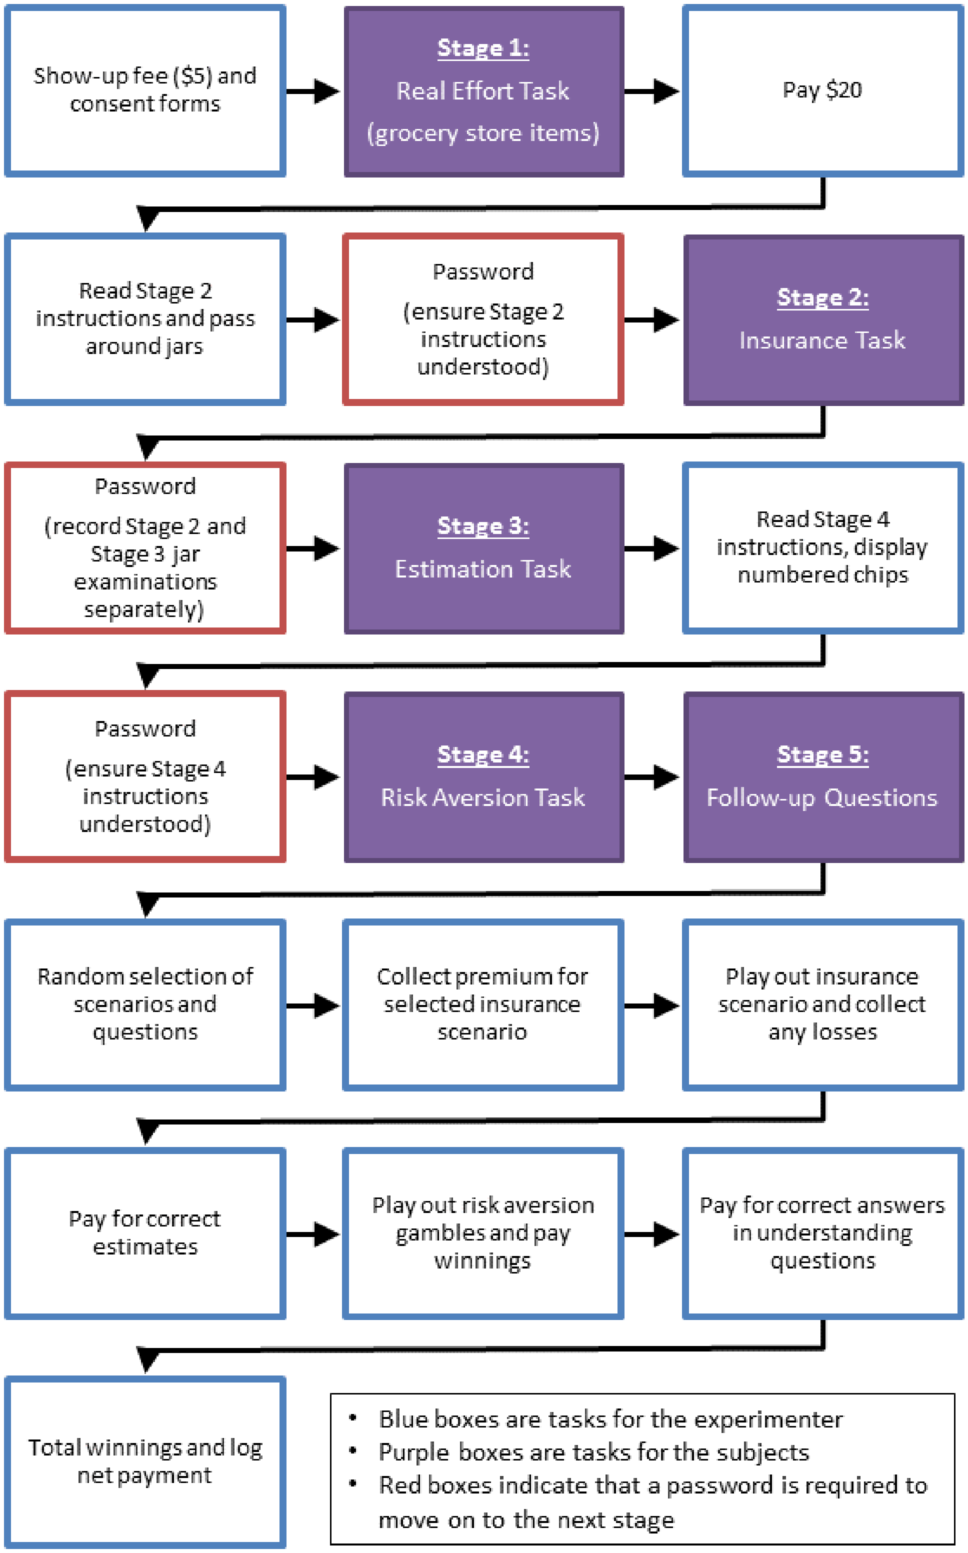

References

Arnum, E. (2014). *State of the Industry and Market Sizing*. Conference presented at the Global Warranty and Service Contract Association First Annual Conference, Chicago.

Baker, T., & Siegelman, P. (2013). “You Want Insurance With That?” Using Behavioral Economics to Protect Consumers from Add-On Insurance Products. *Connecticut Insurance Law Journal*, *20*(1), 1–60.

Barber, B., Odean, T., & Zheng, L. (2005). Out of Sight, Out of Mind: The Effects of Expenses on Mutual Fund Flows. *Journal of Business*, *78*(6), 2095–2120.

Bar-Gill, O., & Ferrari, F. (2010). Informing Consumers About Themselves. *Erasmus Law Review*, *3*(2), 93–119.

Barseghyan, L., Molinari, F., O’Donoghue, T., & Teitelbaum, J. C. (2013). The Nature of Risk Preferences: Evidence from Insurance Choices. *American Economic Review*, *103*(6), 2499–2529.

Bearden, W. O., Carlson, J. P., & Hardesty, D. M. (2003). Using Invoice Price Information to Frame Advertised Offers. *Journal of Business Research*, *56*(5), 355–366.

Ben-Shahar, O., & Schneider, C. (2011). The Failure of Mandated Disclosure. *University of Pennsylvania Law Review*, *159*(3), 647–749.

Bertrand, M., & Morse, A. (2011). Information Disclosure, Cognitive Biases, and Payday Borrowing. *The Journal of Finance*, *66*(6), 1865–1893.

Birnbaum, B. (1999). *Credit Insurance: The $2 Billion A Year Rip-Off* (pp. 1–49). Center for Economic Justice. Retrieved from http://www.cej-online.org/report.pdf

Birnbaum, B. (2001). *Credit Insurance Overcharges Hit $2.5 Billion Annually* (pp. 1–49). Center for Economic Justice. Retrieved from http://www.cej-online.org/2001credit.pdf

Bolton, L. E., & Alba, J. W. (2006). Price Fairness: Good and Service Differences and the Role of Vendor Costs. *Journal of Consumer Research*, *33*(2), 258–265.

Bolton, L. E., Warlop, L., & Alba, J. W. (2003). Consumer Perceptions of Price (Un)Fairness. *Journal of Consumer Research*, *29*(4), 474–491.

Camerer, C., Issacharoff, S., Loewenstein, G., O’Donoghue, T., & Rabin, M. (2003). Regulation for Conservatives: Behavioral Economics and the Case for “Asymmetric Paternalism.” *University of Pennsylvania Law Review*, *151*, 1211.

Campbell, M. C. (1999). Perceptions of Price Unfairness: Antecedents and Consequences. *Journal of Marketing Research*, *36*(2), 187.

Chen, T., Kalra, A., & Sun, B. (2009). Why Do Consumers Buy Extended Service Contracts? *Journal of Consumer Research*, *36*(4), 611–623.

Choi, J. J., Laibson, D., & Madrian, B. C. (2009). Why Does the Law of One Price Fail? An Experiment on Index Mutual Funds. *Review of Financial Studies*, *23*(4), 1405–1432.

Cicchetti, C. J., & Dubin, J. A. (1994). A Microeconometric Analysis of Risk Aversion and the Decision to Self-Insure. *Journal of Political Economy*, 169–186.

Cohen, A., & Einav, L. (2007). Estimating Risk Preferences from Deductible Choice. *American Economic Review*, *97*(3), 745–788.

Day, G. S. (1976). Assessing the Effects of Information Disclosure Requirements. *Journal of Marketing*, *40*(2), 42–52.

De Meza, D., Irlenbusch, B., & Reyniers, D. (2010). Disclosure, Trust and Persuasion in Insurance Markets. *IZA DP No. 5060*, 1–81.

Di Mauro, C., & Maffioletti, A. (2001). The Valuation of Insurance Under Uncertainty: Does Information About Probability Matter? *The Geneva Papers on Risk and Insurance Theory*, *26*(3), 195–224.

Durkin, T. A. (2002). Consumers and Credit Disclosures: Credit Cards and Credit Insurance. *Federal Reserve Bulletin*, *88*, 201.

Hogarth, R. M., & Kunreuther, H. (1995). Decision Making Under Ignorance: Arguing with Yourself. *Journal of Risk and Uncertainty*, *10*(1), 15–36.

Holt, C. A., & Laury, S. K. (2002). Risk Aversion and Incentive Effects. *The American Economic Review*, *92*(5), 1644–1655.

Holt, C. A., & Laury, S. K. (2005). Risk Aversion and Incentive Effects: New Data Without Order Effects. *The American Economic Review*, *95*(3), 902–904.

Huysentruyt, M., & Read, D. (2010). How Do People Value Extended Warranties? Evidence from Two Field Surveys. *Journal of Risk and Uncertainty*, *40*(3), 197–218.

Jindal, P. (2014). Risk Preferences and Demand Drivers of Extended Warranties. *Marketing Science*, *34*(1), 39–58.

Kahneman, D., Knetsch, J. L., & Thaler, R. (1986). Fairness as a Constraint on Profit Seeking: Entitlements in the Market. *The American Economic Review*, 728–741.

Kalapurakal, R., Dickson, P. R., & Urbany, J. E. (1991). Perceived Price Fairness and Dual Entitlement. *Advances in Consumer Research*, *18*(1).

Kunreuther, H., & Pauly, M. (2004). Neglecting Disaster: Why Don't People Insure Against Large Losses? *Journal of Risk and Uncertainty*, *28*(1), 5–21.

Lacko, J., & Pappalardo, J. (2010). The Failure and Promise of Mandated Consumer Mortgage Disclosures: Evidence from Qualitative Interviews and a Controlled Experiment with Mortgage Borrowers. *American Economic Review Papers and Proceedings*, *100*, 516–521.

Maxwell, S., Nye, P., & Maxwell, N. (1999). Less Pain, Same Gain: The Effects of Priming Fairness in Price Negotiations. *Psychology and Marketing*, *16*(7), 545–562.

Murray, M. L. (1972). Empirical Utility Functions and Insurance Consumption Decisions. *The Journal of Risk and Insurance*, *39*(1), 31.

Nalebuff, B., & Ayres, I. (2003). *Why Not? How to Use Everyday Ingenuity to Solve Problems Big and Small*. Harvard Business Press.

Ramsey, D. (2011, January). Never Buy Extended Warranties. Retrieved from https://www.daveramsey.com/index.cfm?event=askdave/&intContentItemId=117416

Sah, S., Loewenstein, G., & Cain, D. M. (2013). The Burden of Disclosure: Increased Compliance with Distrusted Advice. *Journal of Personality and Social Psychology*, *104*(2), 289–304.

Schwarcz, D. (2010). Insurance Demand Anomalies and Regulation. *Journal of Consumer Affairs*, *44*(3), 557–577.

Sydnor, J. (2010). (Over)insuring Modest Risks. *American Economic Journal: Applied Economics*, *2*(4), 177–199.

Urbany, J. E., Madden, T. J., & Dickson, P. R. (1989). All's Not Fair in Pricing: An Initial Look at the Dual Entitlement Principle. *Marketing Letters*, *1*(1), 17–25.