Price-Sensitive Preferences

August 17, 2009

Nina Mazar

University of Toronto, 105 St. George Street, Toronto, ON M5S3E6,

phone: 416-946-5650, fax: 416-978-5433, nina.mazar@utoronto.ca

Botond Koszegi

University of California, Berkeley, 549 Evan Hall #3880, Berkeley, CA 94720-3880, phone: 510-643-0714, fax: 510-642-6615, botond@econ.berkeley.edu

Dan Ariely

Duke University, One Towerview Road, Durham, NC 27708

phone: 919-660-7703, fax 919-681-6246, dandan@duke.edu

Price-Sensitive Preferences

ABSTRACT

A crucial assumption of neoclassical economics and the base for marketing and pricing models is that the reservation price (maximum willingness to pay) for a familiar product is independent of the non-informative prices consumers face for this product. We test this assumption using experimental procedures that are both incentive-compatible in eliciting subjects' reservation prices and allow us to vary price distributions across participants. We find that valuations are extremely sensitive to the possible prices. More importantly, our findings suggest that this sensitivity is neither due to rational inferences subjects draw from the price distribution, nor due to true price-dependent preferences; most likely, they are due to mistakes subjects make in expressing price-independent preferences. We consider the implications of our findings for simple supply-demand theory, and find that neoclassical methods may underestimate the elasticity of demand, overestimate the consumer surplus, and overestimate the dead-weight loss from taxation.

Recent research in psychology and economics questions one of the key tenets of neoclassical utility theory: That preferences are stable and independent of the external environment that decision makers face. The general assumption instead is that preferences are influenced by contextual variables. One of the most important domains for the study of such influences is pricing. In a typical purchase situation, whenever consumers intend to buy a particular good or service they encounter numerous retailers with differing market prices. Prices can be found everywhere in the immediate environment – from flyers and advertisements in newspapers and televisions to price tags in stores and on price comparison web sites – such that consumers have some expectation about a product's distribution of prices in the marketplace. Can the knowledge of the price distribution influence people's reservation prices?

The reference price literature has shown that previous or present market prices of the good in question, close substitutes, and even unrelated goods affect customers' willingness to pay for a product by serving as an anchor (e.g., Tversky and Kahneman 1974; Chapman and Johnson 2002; Monroe 1990; Simonson and Drolet 2004; see also Ariely, Loewenstein, and Prelec 2003 for arbitrary prices). However, these studies only considered the influence of a single price rather than a distribution of prices or expectations about the price distribution. In addition, with the exception of Nunes and Boatwright (2004), these studies included an explicit (hypothetical) comparative judgment, which often does not reflect real world market conditions. Finally, it is still unclear what are the possible causes for such a context-dependency. If reservation prices depend on the prices buyers expect to face in the market, demand will in general depend on supply, with important implications for marketers' pricing strategies and for pricing models. Therefore, understanding the causes for this dependency would be crucial to determine marketing, welfare, and policy implications.

The remainder of this paper is organized as follows: First, we present one experiment in which we use a multi-period market game. We find that participants' reactions for a given price depend on the distribution of prices they experienced in the market beforehand: Valuations are higher when the price distribution is skewed to the right than when it is skewed to the left. That is, whereas most economic analysis, which also forms the basis for most pricing models in marketing, assumes that exogenously given preferences determine (along with supply) the market price, the influence also runs the other direction: The price serves as a powerful contextual variable that determines expressed preferences.

In experiments 2 and 3 we replicate our findings with a different, incentive compatible procedure: The Becker-DeGroot-Marschak (1964, henceforth BDM) procedure. Similar to before, we find a dramatic violation of the independence of valuations from the price distribution. In addition, in order to identify the specific implications of these findings for economics and marketing, in experiment 2 and 3 we investigate three underlying reasons for the price dependence of preferences: rational reaction to information, true change in preferences (i.e. real context-dependent preferences), and biased (i.e. erroneous) expression of true context-independent preferences. Our results in experiment 2 rule out a rational reaction to information and favor the possibility of an error in the expression of one's true context-independent preferences. Experiment 3 provides further support for this bias-hypothesis by demonstrating how to de-bias responses. Finally, we discuss some implications of our findings for basic supply–demand analyses, illustrating how neoclassical methods may underestimate the elasticity of demand, overestimate the consumer surplus, and overestimate the dead-weight loss from taxation, and add a few concluding comments.

EXPERIMENT 1: THE MARKET GAME

In a typical purchase situation in a store, consumers face a price and have to decide whether they want to make the purchase. We set out to determine in a setting that closely resembles such a typical market decision whether reservations prices can depend upon the price distribution. For this purpose we designed a market game in which after experiencing a product's offerings in the market, participants had to answer simple yes/no questions as to whether they wanted to buy a product for various given prices. Participants used their own money, and all answers were deterministic.

Procedure

Eighty-two students from two northeastern universities in the US participated in this experiment, which was conducted in multiple sessions of approximately 15 students at a time. Each participant sat in a private booth in front of a computer screen with a program simulating a market for $0.50 gift certificates for Amazon.com. At the start of the experiment, each student was presented with a description and a picture of a $0.50 electronic gift certificate for Amazon.com. The instructions stated that the experiment consisted of two sections. In the first section they would observe the various prices for a $0.50 gift certificate available in the market and their frequency. This was meant to resemble people's experience with a product's offerings in the market. In the second section they would face the same set of prices but this time each offering would require a consequential (real) purchase decision. Both sections consisted of 50 offers, each for a gift certificate for Amazon.com worth $0.50. Consequently, the maximum number of gift certificates that participants could buy was 50. Selling prices varied from $0.10 to $0.50 in steps of $0.05 and were shown in a random order. At the end of the experiment, for each participant we calculated the sum of the prices and values of her purchased Amazon.com gift

certificates and handed out one Amazon.com gift certificate of the total value (the maximum possible value was 50 x $0.50 = $25) in exchange for the total price.

Students were randomly assigned to one of two experimental conditions that differed in the frequency at which the two extreme prices $0.10 and $0.50 were offered (left-skewed and right-skewed). In the left-skewed distribution we assigned 26 of the 50 trials to the lowest value ($0.10) and distributed the remaining 24 trials equally among the remaining eight prices ranging from $0.15 to $0.50 (three trials per price). In the right-skewed distribution we assigned 26 of the 50 trials to the highest value ($0.50) and equal frequencies to the other eight prices ranging from $0.10 to $0.45 (see Figure 1).

*** Figure 1 ***

Results & Discussion

For each participant we defined the reservations price for a $0.50 gift certificate for Amazon.com as the highest price at which participants were still buying the gift certificate. A t-test revealed that the reservation price differed significantly between participants that had experienced the left-skewed price distribution and those that experienced the right-skewed price distribution ($t(80) = 2.768$, $p = 0.007$). Participants in the market with the left-skewed price distribution (M = $0.29, SEM = 0.02) had on average an 8 ¢ lower reservation price than the participants in the market with the right-skewed distribution (M = $0.37, SEM = 0.02).

More importantly, this result is not simply due the fact that the likelihood of the occurrence of $0.1 ($0.5) was higher in the left (right)-skewed distribution. Figure 2 shows the cumulative frequency distribution of reservation prices across the two market conditions. A

follow-up analysis that excluded all $0.1 and $0.5 trials such that all participants faced the same frequency of the same prices (3 trials of each: $0.15, $0.2, $0.25, $0.3, $0.35, $0.4, $0.45) replicated the result[1]. Participants in the left-skewed market exhibited a significantly lower reservation price (M = $0.31, SEM = 0.02) than participants in the right-skewed market (M = $0.36, SEM = 0.02), ($t(76) = 2.081$, $p = 0.041$).

*** Figure 2 ***

Finally, our finding was based on a sequence of decisions across fifty trials. Therefore, we could also examine if the magnitude of the difference in reservation prices between the left and right-skewed distributions changed over time. To do so we divided the trials into two halves and compared the reservation prices for the first half of trials to the second half of trials (within-subjects). A repeated measures ANOVA showed that the effect of the price distribution (between-subjects) was significant ($F(1,80) = 8.040$, $p < 0.006$), but there was no significant effect of half ($F(1,80) = 2.448$, $p = 0.122$), and more importantly, there was no interaction ($F(1,80) = 0.010$, $p = 0.992$). In particular, the difference in mean reservation price in the left and right-skewed distribution was $0.08 in the first half of the trials, and $0.08 in the second half of the trials, suggesting that the effect of the price distributions did not change over time.

Our results might shed new light on Richard Thaler's (1985) famous "beer on the beach" example, which traditionally has been interpreted in terms of fair pricing. In his study, respondents provided their hypothetical willingness to pay for a bottle of cold beer delivered to them on the beach by a friend. Participants were told that the friend would get the beer, in one condition, from a local store and, in the other condition, from a fancy hotel. The results showed

that participants' willingness to pay was significantly higher in the latter condition. We view these results as consistent with our main finding: Because respondents expected the prices to be higher at the hotel, they responded by raising their reservation price – which, as our results show, can occur even when there aren't any fairness considerations.

We elicited students' willingness to pay for a common and popular product, a gift certificate for Amazon.com. Because participants likely understand the pleasure of consumption associated with this simple good we posited that under these conditions the independence of valuations from prices is most likely to hold. Nevertheless, the results from Experiment 1 imply that individuals' expressed preferences are not independent of the distribution of prices, as the stability of preferences across environments would imply. Our interest in this paper, however, goes beyond merely identifying one particular way in which preferences depend on contextual factors. We thought out to replicate our findings with a different incentive-compatible mechanism – the BDM procedure –, explore the source of the price-dependency, and discuss possible implications for some key economic settings.

EXPERIMENT 2: THE SOURCE OF OUR EFFECT

Becker-DeGroot-Marschak Procedure (BDM)

One alternative tool that can be used for examining the relationship between reservation prices and the price distribution that consumers expect is the BDM procedure. The BDM procedure is a widely used, incentive-compatible mechanism to measure individuals' valuations of consumption goods and other experiences (see e.g., Budescu, Shalva, and Weinberg 1988; Fox, Rogers, and Tversky 1996; Kahneman, Knetsch, and Thaler 1990; Prelec and Simester 2001). When employing this procedure, participants are informed of a distribution of prices at

which they may acquire a good (supply), and are asked to indicate their reservation price (demand). After they have done so, a random price (the realized price) is drawn from the price distribution, and the appropriate outcome is implemented: If the reservation price exceeds the realized price, the decision maker receives the good and pays the realized price; if the reservation price is below the realized price, the decision maker does not receive the good and does not pay anything. Because this mechanism precludes participants from influencing the price, it is a (weakly) dominant strategy to reveal one's true preferences. Wertenbroch and Skiera (2002) have compared the BDM procedure to other procedures for measuring consumers' willingness to pay and found that the BDM procedure reveals very reliable and valid estimates.

Our design builds upon earlier work by Bohm, Linden, and Sonnegard (1997), who demonstrated that an unrealistically high upper bound of the BDM distribution could influence sellers' elicited reservation prices for a gas card. In particular, Bohm and colleagues showed that sellers' mean reservation price elicited with the BDM procedure surpassed seller's mean price demanded in an experimental sealed-bid double auction market. While their evidence points to the importance of testing whether and how prices influence preferences, their main interest was to test for the robustness of the BDM procedure as a tool for eliciting reservation prices. In contrast, we are interested in the more fundamental relationship between preferences and market prices, or alternatively, the relationship between demand and supply. As a consequence, our work differs on several important dimensions. First and most important, we extend the scope of the investigation by examining the underlying reasons for our findings: rational inferences, context dependent preferences, or a bias. Second, we illustrate some implications of our findings for basic supply–demand analysis. Third, assuming that participants can be considered to have more real-world experiences as buyers, we elicit buyer reservation prices rather than selling prices.

Finally, participants in our studies had to spend their own money instead of house money.[2] The general procedure is described in Figure 3.

*** Figure 3 ***

Possible Sources for a Price-Dependence of Preferences

In order to explore some implications of the price dependence of preferences for economic analysis we need to know the source and nature of that dependence. We consider three possibilities: (1) uncertainty about the product's value leads to rational inferences about it from the price distribution; (2) true context-dependent preferences cause higher prices to generate higher utility from consuming the product; and (3) a bias (mistake) occurs in the expression of true price-independent preferences. Experiment 2 attempts to discriminate among these explanations and suggests that the last one is the most likely account.

First, in order to reconcile the results from Experiment 1 with neoclassical utility theory, one could argue that the stark differences in valuations across the left and right-skewed market conditions arouse because participants were somehow uncertain about the value of the gift certificate for Amazon.com. Consequently, they drew rational inferences about it from the price distributions they experienced (see Wernerfelt 1995; Prelec, Wernerfelt, and Zettelmeyer 1997; and Kamenica 2008 for similar explanations of other context effects) and adjusted their reservation price. While in a real market such inferences are surely going on, it is important to know whether the price sensitivity of valuations can be attributed entirely to those inferences. We address this question in Experiment 2 by providing participants the same information in all conditions. Specifically, by showing participants all price distributions that could be used to

determine their individual, realized price, and by selecting the relevant price distribution using a coin flip performed by the participant herself. Through this procedure we make it explicit that the particular distribution is randomly determined and hence uninformative. If rational inferences about the value of a product drive our results, this change in the experimental design should eliminate the price dependence of the elicited preferences.

Even once we demonstrate that price dependence cannot plausibly be due to informational considerations, the implications of our findings depend crucially on which of the two other (more psychological) explanations is valid. Our results may be driven by true context-dependent preferences, whereby (even random) prices affect consumers' utility from the product and thus their willingness to pay for it. That is, when consumers face higher prices, even if those prices come about randomly, they truly like and enjoy the product more and therefore are willing to pay more for it. For example, a consumer may enjoy wearing a Rolex rather than a fake just because it is more expensive, even though she and others cannot distinguish it from the fake. A second account is that the price distribution might not affect consumers' true experienced utility but rather acts as an environmental cue that leads to a mistaken expression of preferences. After all, prior research demonstrates that individuals often rely on environmental cues to guide their choices (e.g., Tversky, Sattath, and Slovic 1988; Huber, Payne, and Puto 1982; Bettman, Luce, and Payne 1998), and the price is an extremely salient such cue.

To distinguish between true context-dependent preferences and a biased expression of true context-independent preferences, we asked one third of the participants in Experiment 2 to indicate their reservation price for both price distributions (joint valuations) and the remaining two thirds of the participants to indicate their reservation price for only one of the two price distributions (separate valuations). The logic behind this manipulation is that joint valuations

should confront participants more explicitly with the influence of price distributions on their reservation prices, and to the extent that it creates an undesirable influence, they should be able to correct it (Hsee et al. 1999). Thus, when respondents make valuations jointly, the effect of the price distributions should emerge only if the cause of the difference is a true change in preferences. However, if the effect of the price distributions emerges only in the separate valuations condition, it likely is due to an error in expressing (true context-independent) preferences, not an expected influence on preferences.

To keep both the joint and separate valuations incentive compatible, we varied the timing of the coin flip relative to the timing of the valuations. In the joint valuation condition, after observing both distributions, participants indicated their reservation prices for both distributions and then flipped the coin to determine the price distribution that would be used to calculate their outcome. In the separate valuations condition, after observing both distributions, participants first flipped the coin to determine the price distribution that would be used to calculate their outcome and then they indicated their reservation price for that price distribution only.

Procedure

Sixty students at a northeastern university in the US participated in this experiment, which was conducted with one participant at a time. Each participant had a chance to buy a travel mug through the BDM mechanism but otherwise received no compensation for participation.

At the start of the experiment, each student inspected a sample mug. The experimenter then gave the participant a sheet with a description of the BDM procedure and a prominent graphical depiction of two price distributions, a left-skewed distribution and a right-skewed distribution, both of which had support from $1 to $10. The left-skewed distribution assigned a

probability of 0.5 to the lowest value ($1) and distributed the remaining probability equally in the continuous range from $1.01 to $10. The right-skewed distribution assigned a probability of 0.5 to the highest value ($10) and equal probabilities to the continuous range from $1 to $9.99 (see Figure 4). The instructions indicated to the participant that she would have a chance to buy a mug at a price to be determined randomly by one of these two price distributions and that a coin flip performed by the participant would determine the distribution to be used. In addition, the instructions for the BDM procedure made clear that if the price drawn from the distribution came out higher than the student's reservation price, she would not receive the mug and not pay anything; and that if the price drawn from the distribution was the same or lower than the student's reservation price, she would receive the mug and pay the price drawn. The description also included a sentence that emphasized the incentive compatibility of this procedure (see Appendix A).

*** Figure 4 ***

Students were randomly assigned to one of two experimental conditions: the joint or separate valuations conditions. In the joint valuations condition, after a participant finished reading the instructions, which included both distributions, she was asked to indicate for each of the two price distributions her willingness to pay for the mug should that distribution be selected by their coin flip. The participants then performed the coin flip, the experimenter used a random-number generator to determine the price from the selected distribution, and the participant's relevant choice was implemented. In the separate valuations condition, after a participant finished reading the same instructions, she flipped a coin to determine which of the two price

distributions would be used, and then indicated her willingness to pay for the mug for the selected price distribution only. After providing the evaluation, the experimenter used a random-number generator to determine the price from the selected distribution, and the participant's choice was implemented.

Results & Discussion

Using the reservation prices submitted to the BDM procedure, we analyzed the data as a 2 (right- vs. left-skewed distribution) by 2 (joint vs. separate evaluations) between-subjects ANOVA. As can be seen in Figure 5, we found a significant interaction between the type of distribution and the method of evaluation ($F(1,76) = 5.95$, $p = 0.017$). Specifically, though the difference in reservation prices in the separate evaluation conditions is significantly lower for the left-skewed distribution ($2.42) compared with the right-skewed distribution ($5.08; $F(1, 76) = 14.78$, $p < 0.001$), the difference in reservation prices in the joint evaluation condition is very small ($3.09 vs. $3.38) and statistically insignificant ($F(1, 76) = 0.19$, $p = 0.662$). That is, the results indicate a large influence of distributions when respondents provide valuations to only one of the distributions and no influence when respondents offer reservation prices for both distributions.

*** Figure 5 ***

Because our separate valuations condition replicates the qualitative results of Experiment 1 even though participants knew both price distributions, we can reject the hypothesis that our findings are solely due to rational inferences about the value of the mug drawn from the price distribution. More important, because participants in the joint valuations condition indicate very

similar reservation prices for the two distributions—apparently realizing that their willingness to pay should be independent of the price distribution— these results support the thesis that the price dependence of valuations is driven mainly by a mistake in the expression of true context-independent preferences and not by true context-dependent preferences.

An alternative account for the differing results in the joint and separate valuations conditions is that participants had indeed inherent preferences that varied across the left and right-skewed distributions (i.e. they had true price-dependent preferences) and that the mistake in the expression of their preferences occurred in the joint valuations condition (rather than in the separate valuations condition). For example, even if the procedure was incentive compatible individuals might have given consistent prices in the joint valuations condition because of demand effects or because of not being aware of their true context-dependent preferences as long as they weren't certain about the price distribution they would face. Experiment 3 was designed to further test whether the observed price-sensitivity is due to true context-dependent preferences or due to a biased expression of true context-independent preferences.

EXPERIMENT 3: DE-BIASING

Experiment 3 introduces two variants of the basic BDM design, both involving procedures that confront respondents with a sequence of prices, and ask them to make single buy / not buy decisions for each price. In the first new variant, choice based BDM (see Kahneman, Knetsch, and Thaler 1990; Fox and Tversky 1991; Bostic, Herrnstein, and Luce 1990), the full range of prices is presented in a preset order. In the second new variant, titration based BDM, the sequence is determined in an iterative manner by the respondent herself. We expected these two elicitation methods to increase respondents' attention to their inherent preferences

(Simonson 2008). Thus, if our hypothesis is correct and our participants in experiment 1 and 2 succumbed to a mistake in the expression of true context-independent preferences because of overly relying on the very salient environmental cue of prices in the market, the two new elicitations methods should reduce or eliminate any differences in reservation prices even when there is no joint evaluation (i.e. no demand effect and no uncertainty about the underlying price-distribution). However, true price-dependent preferences should not be so sensitive to how one asks for the reservation price. That is, if participants had true context-dependent preferences, the differences in reservations prices should remain with the new elicitation methods.

Procedure

One-hundred and twenty students were recruited at a northeastern university to participate in the experiment. Each participant was randomly assigned to one of six conditions in a 3 (elicitation method) by 2 (price distribution) between-subjects design. Besides the two additional elicitation methods, Experiment 3 differed from the previous experiment in two other ways. First, the product used in this experiment was a box of Godiva chocolates (Gold Ballotin 1/4 lb. with 9 pieces of confection). Second, the two price distributions had a support on the integers from $1 to $20, and assigned a probability of 1/29 to values between 2 and 19. The left skewed distribution had a probability of 10/29 for the lowest value ($1), and a probability of 1/29 for the highest value ($20). The right skewed distribution had a probability of 10/29 for the highest value ($20), and a probability of 1/29 for the lowest value ($1) (see Figure 6).

*** Figure 6 ***

The major innovation in Experiment 3 was the inclusion of two variants of the BDM method to elicit a participant's reservation price. As in previous experiments, participants were informed that a price would be drawn randomly according to a distribution illustrated graphically on the instructions. They were then asked to indicate their reservation price for the chocolates using one of the following three different but strategically essentially equivalent elicitation methods: The standard BDM, the choice based BDM, and the titration based BDM procedures.

In the standard BDM procedure participants stated their willingness to pay directly by writing down an amount, and the outcome was determined by that number and the price drawn from the price distribution (similar to Experiment 2; see Appendix B). In the choice based BDM procedure, participants answered a series of questions, each of which asked, for an integer X between 1 and 20, whether they would be willing to pay $X for the box of chocolates. After respondents answered all the questions, one of the questions was selected according to the probabilities specified by the distribution, which was given to the respondents before they answered any of the questions and was also marked next to the questions (see Appendix C). In the titration based BDM procedure, participants were asked to (eventually) write down a single reservation price, and the outcome was determined by this reservation price and the realized price exactly as in the standard BDM. However, in this condition, participants were advised to think about their reservation price in the following specific way. They were first asked to write down a number they thought was a reasonable estimate of the maximum amount they would be willing to pay for the box of chocolates. Next, they were asked to consider if the amount written was too high, too low, or just the right reservation price for the box of chocolates. If respondents answered too low or too high, they were asked to think about a new price and continue the process until they came up with the "just right" response. This final "just right" amount was

subjected to the standard BDM procedure (see Appendix D). Note that this procedure explicitly articulates the steps individuals might actually take when answering an open ended pricing question such as the one they face in a standard BDM procedure.

Results & Discussion

As can be seen in Figure 7, the standard BDM procedure replicated the findings of the previous experiments. The mean reservation price for the chocolates was dramatically lower in the left skewed distribution ($3.75) compared with the right skewed distribution ($6.85), and this difference of $3.10 was significant ($t(35) = 2.89$, $p = 0.007$).

Most important to the goals of the current experiment, the effect of the price distribution on reservation prices weakened for the two variant BDM procedures (Figure 7). In particular, the choice based BDM task revealed a significant difference ($t(31) = 2.08$, $p = 0.046$) in the mean reservation price between the two distribution conditions: the mean reservation price was lower in the left skewed distribution ($3.65) than in the right skewed distribution ($5.43). However, the difference between the two distributions was lower ($1.78) in the choice based BDM procedure than in the standard BDM procedure ($3.10). Finally, the titration based BDM procedure showed an even smaller difference between the two distributions. The mean reservation price for the chocolate was lower in the left skewed distribution ($5.18) compared with the right skewed distribution ($6.11), but this difference of $0.93 was not significant ($t(37) = 0.758$, $p = 0.453$). That is, the choice based elicitation method reduced the sensitivity of expressed valuations to the price distribution by a small amount, while the titration based elicitation method reduced it almost completely.

*** Figure 7 ***

In addition, the comparison between these variant approaches and the standard BDM can shed some light on the psychological process by which individuals decide about their reservation prices. In particular, the difference between the titration based method and the standard BDM method indicates that when respondents try to pick a price in the latter method, they do not use a procedure that resembles an iterative approach, although intuitively it seems to be a natural process for arriving at a reservation price. What respondents are doing instead is unclear, and an interesting and promising path for future research. To make the psychological processes underlying our results even more elusive, we found that on average participants went through only 1.5 iterations in the titration procedure. Thus, it seems that it is the nature of the titration task itself, and not the number of iterations, that caused the decreased sensitivity to the price distributions.

On an aggregate level, the findings of our three experiments imply that demand is not, as standard theory assumes, independent of supply. In the following section, we derive some implications of our experimental findings for an important tool of neoclassical economics, that forms the basis for most pricing models: partial equilibrium supply–demand analysis.

ECONOMIC IMPLICATIONS

Starting from undergraduate microeconomics, economists conceptualize many market phenomena using demand and supply curves that determine the equilibrium price and quantity, and they nearly always suppose that supply and demand curves are stable and independent of each other in the market. Existing research on the psychology of judgment and decision making

has already questioned the stability of preferences and hence the existence of a fixed, stable demand curve.[3] Our results indicate that beyond being unstable, the demand curve depends systematically on a key market variable of interest, the price (Bohm, Linden, and Sonnegard 1997; Simonson and Drolet 2004). We demonstrate that because consumer preferences are often inferred exactly from how consumers react to different prices, this dependence can systematically bias classical measurements of welfare.

For our theoretical argument, we modify neoclassical supply-demand analysis to incorporate the interpretation of our results that (given we have ruled out others) we take to be the most reasonable: Expressed preferences depend on market prices, but there are "true", inherent preferences relevant for welfare analysis that do not depend on prices (Simonson 2008). We start from a standard supply-demand diagram, consisting of the demand curve D1 and supply curve S1 in Figure 8. We assume that the market is initially in a kind of "long-run equilibrium", where consumers have figured out their preferences for this particular product and hence where the demand curve D1 reflects true preferences.[4] Now consider an exogenously driven shift in the supply curve from S1 to S2. This shift can occur due to an increase in producer-side taxes, transportation costs, or other input costs. Because – as we have found – price changes due to supply or tax shocks shift consumers' expressed preferences in the same direction, the resulting rise in price leads to an upward shift in the demand curve.[5] The combination of shifts in the supply and demand curves leads to a new equilibrium at point A.[6] However, a neoclassical economist – not aware of our findings – would assume that this change traces out a single demand curve in the market, inferring that the demand curve is AC. This immediately leads to our first deduction: The economist infers a lower elasticity of demand than that warranted by the true demand curve.

Our neoclassical economist's misestimation of the demand curve also leads to a misestimation of consumer welfare in various situations. If our economist extrapolates the inferred demand curve beyond A, as is commonly done, the inferred demand curve lies entirely above the true one. Therefore, the economist overestimates the extent to which consumers who buy the product care about it – that is, she overestimates consumer surplus in the market.

Second, when the shift in the supply curve is due to a tax change, a neoclassical economist also misestimates the excess burden from tax intervention. Given that she estimates the demand curve to be AC, she estimates the deadweight loss (the total value of trades that are not taking place because of the tax) to be ABC. In contrast, to identify the true deadweight loss, we first consider the consumer side. In equilibrium C, the consumer surplus is the triangle CHG. In equilibrium A, the gross consumer surplus (the consumer surplus gross of expenditure) is the trapezoid OHDJ. At the same time, consumers pay an amount equal to the rectangle OJAE. Therefore, the net consumer surplus becomes HEI minus ADI. Compared with the consumer surplus in equilibrium C, the loss in consumer surplus due to the tax is the trapezoid CIEG plus the triangle ADI. Because the producer in this market is neoclassical, the loss in producer surplus due to the tax is the trapezoid CBHG, just as in standard supply–demand analysis. Therefore, because the government collects taxes equal to the area ABHE, the excess burden is the triangle DBC, which is less than ABC.

Intuitively, a tax leads to a deadweight loss because it eliminates some socially beneficial trades. Consumers do not care as much about the good as their market reaction would seem to indicate (part of that reaction reflects their mistake in expressing preferences), so the trades that do not take place due to the tax are less valuable than the neoclassical estimate suggests. Consequently, based on our findings for any given reduction in the quantity traded, the total

value of lost trades tends to be smaller than the neoclassical estimate based on the assumption of preference stability.

A similar analysis reveals that a neoclassical methodology also leads to an overestimation of the excess burden from a subsidy. Intuitively, a subsidy leads to a deadweight burden because it facilitates extra trades that generate losses from a social point of view; the cost of producing the extra items is greater than the benefit consumers derive from them. Because consumers care more about the good than their market reaction indicates (part of the reduction in demand can be attributed to the lowered prices in the market), we anticipate that a neoclassical economist would overestimate the losses in the new trades generated by a subsidy.

*** Figure 8 ***

FINAL COMMENTS

In this work, we experimentally demonstrate that elicited preferences can be sensitive to a central characteristic of markets: the prices consumers expect to face. We investigate the source of this dependence in incentive compatible contexts without the use of an explicit anchor, and formally derive some possible implications of our findings. In Experiment 1, we find that in a market game with multiple deterministic trials expressed valuations for a product with a clear and certain value ($0.50 gift certificate for Amazon.com) can be higher when consumers experienced a high price distribution than when they experienced a low price distribution. In Experiment 2, we test three possible accounts for this influence: (1) uncertainty about the product's value, which leads to rational inferences about it from the price distribution; (2) true context-dependent preferences, such that higher prices generate a higher utility from consuming

the product; and (3) a bias (mistake) in the expression of price-independent true preferences. The results of Experiment 2 show that rational inferences cannot account for the results and suggest that the observed behavior is a mistake in the expression of preferences rather than a reflection of true context-dependent preferences. Finally, in experiment 3 we present further support for the bias-thesis by de-biasing participants' preferences.

Together, these results show that though most economic analysis assumes that exogenously given preferences determine the market price (along with supply), the influence also runs in the other direction: The price is an inherently powerful contextual variable that determines expressed preferences even in the absence of a clear, explicit anchor. Ironically, though the most important component of the preference-stability assumption is that it holds when applied to changes in the economic (as opposed to other) environment, the economic environment may in fact influence expressed preferences most, and do so simply because it is often the most salient factor in economic decisions. Finally, on the basis of a formalization of our results, we show that a neoclassical economist operating under the assumption of preference stability underestimates the elasticity of demand and overestimates the consumer surplus and excess burden from taxation. These findings have important implications for the interpretation of pricing models and subsequent marketing strategies.

It is also worthwhile to use these results to reflect on the measurement of utility in experimental economics and marketing research, which is often done with some variant of the BDM procedure (for a further discussion on the usefulness of the BDM procedure see also Wertenbroch and Skiera 2002). On the basis of the results presented herein, we argue that caution should be taken when interpreting absolute levels of preferences ("absolute valuation") from experimental data. However, our results do not necessarily undermine conclusions based on

directional comparisons between conditions ("relative valuations"). Consider, for example, a study that uses the standard BDM procedure to examine the valuation of a lottery ticket relative to its expected value compared with another study that contrasts valuations for the same lottery ticket under different conditions (e.g., buying and selling). Although the sensitivity of the BDM procedure likely is present in both studies, the resulting biased estimates pose a larger problem for the former one. In particular, when comparing elicited preferences with an external standard, the aspects of the price distribution can largely determine the conclusions of the study. In contrast, when the main focus is comparing the qualitative (directional) difference between two conditions, and as long as both conditions employ the same price distribution, the general conclusions should be unaffected by the mismeasurement of absolute preferences.

REFERENCES

Ariely, Dan, George Loewenstein, and Drazen Prelec (2003), "Coherent Arbitrariness: Stable Demand Curves without Stable Preferences," *Quarterly Journal of Economics*, 118 (1), 73-105.

____, and Thomas S. Wallsten (1995), "Seeking Subjective Dominance in Multidimensional Space: An Explanation of the Asymmetric Dominance Effect," *Organizational Behavior and Human Decision Processes*, 63 (3), 223-232.

Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak (1964), "Measuring Utility by a Single-Response Sequential Method," *Behavioral Science*, 9 (3), 226-232.

Bettman, James R., Mary Francis Luce, and John W. Payne (1998), "Constructive Consumer Choice Processes," *Journal of Consumer Research*, 25 (3), 187-217.

Bohm, Peter, Johan Linden, and Joakin Sonnegard (1997), "Eliciting Reservation Prices: Becker-DeGroot-Marschak Mechanisms vs. Markets," *Economic Journal*, 107 (443), 1079-1089.

Bostic, Raphael, Richard J. Herrnstein, R. Duncan Luce (1990), "The effect on the preference-reversal phenomenon of using choice indifferences," *Journal of Economic Behavior and Organization*, 13 (March), 193-212.

Budescu, David V., Shalva Weinberg, and Thomas S. Wallsten (1988), "Decisions based on numerically and verbally expressed uncertainties," *Journal of Experimental Psychology: Human Perception and Performance*, 14 (2), 281-94.

Chapman, Gretchen B., and Eric J. Johnson (2002), "Incorporating the Irrelevant: Anchors in Judgments of Belief and Value," in *Heuristics and Biases: The Psychology of Intuitive Judgments,* Thomas Gilovich, Dale Griffin, and Daniel Kahneman, eds. Cambridge, UK: Cambridge University Press, 120-138.

Fischer, Gregory W., and Scott A. Hawkins (1993), "Strategy Compatibility, Scale Compatibility, and the Prominence Effect," *Journal of Experimental Psychology: Human Perception and Performance*, 19 (3), 580-597.

Fox, Craig R., Brett A. Rogers, and Amos Tversky (1996), "Options traders exhibit subadditive decision weights," *Journal of Risk and Uncertainty*, 13 (1), 5-17.

Fox, Craig R. and Amos Tversky (1991), "Ambiguity aversion and comparative ignorance," *The Quarterly Journal of Economics*, 110 (August), 585-603.

Hsee, Christopher K., George Loewenstein, Sally Blount, and Max H. Bazerman (1999), "Preference reversals between joint and separate evaluations of options: a theoretical analysis," *Psychological Bulletin*, 125 (5), 576-590.

Huber, Joel, John W. Payne, and Christopher Puto (1982), "Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis," *Journal of Consumer Research*, 9 (1), 90-98.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler (1990), "Experimental tests of the endowment effect and the Coase theorem," *Journal of Political Economy*, 98 (6), 1325-1348.

Kamenica, Emir (2008), "Contextual Inference in Markets: On the Informational Content of Product Lines," *American Economic Review*, 98 (5), 2127-2149.

Lichtenstein, Sarah, and Paul Slovic (1971), "Reversal of Preferences between Bids and Choices in Gambling Decisions," *Journal of Experimental Psychology*, 89 (1), 46-55.

Monroe, Kent B. (1990), *Pricing: Making Profitable Decisions*, 2nd ed. New York: McGraw-Hill.

Nunes, Joseph C. and Peter Boatwright (2004), Incidental prices and their effect on willingness to pay, *Journal of Marketing Research*, 41 (1), 457-466.

Prelec, Drazen and Duncan Simester (2001), "Always leave home without it: a further investigation of the credit-card effect on willingness to pay," *Marketing Letters*, 12 (1), 5-12.

____, Birger Wernerfelt, and Florian Zettelmeyer (1997), "The Role of Inference in Context Effects: Inferring What you Want from What is Available," *Journal of Consumer Research*, 24 (1), 118-125.

Simonson, Itamar (1989), "Choice Based on Reasons: The Case of Attraction and Compromise Effects," *Journal of Consumer Research*, 16 (2), 158-174.

____ (2008), "Regarding Inherent Preferences," *Journal of Consumer Psychology*, 18 (3), 191-196.

____ and Aimee Drolet (2004), "Anchoring Effects on Consumers Willingness-to-

Pay and Willingness-to-Accept," *Journal of Consumer Research*, 31(4), 681-690.

Thaler, Richard H. (1985), "Mental Accounting and Consumer Choice," *Marketing Science*, 4 (3), 199-214.

Tversky, Amos, and Daniel Kahneman (1974), "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185 (4157), 1124-1131.

____, S. Sattath, and Paul Slovic (1988), "Contingent Weighting in Judgment and Choice," *Psychological Review*, 95 (3), 371-384.

Wernerfelt, Birger (1995), "A Rational Reconstruction of the Compromise Effect: Using Market Data to Infer Utilities," *Journal of Consumer Research*, 21 (4), 627-633.

Wertenbroch, Klaus and Bernd Skiera (2002), "Measuring consumers' willingness to pay at the point of purchase," *Journal of Marketing Research*, 39 (2), 228-41.

FOOTNOTES

1. In this new analysis, we excluded 4 participants that had bought gift certificates for $.1 but not for any of the other prices. If we do not exclude them but instead code their reservation prices as $0, our result gets stronger (left-skewed market: M = $0.27, SEM = 0.02; right-skewed market: M = $0.36, SEM = 0.02; $t(80) = 3.026$, $p = 0.003$).

2. In their experiment, Bohm and colleagues (1997) distributed a gas card with a value of SEK228 (Swedish Kronor) to each of the participants. Participants could then sell the card for money in the experiment. Those, who didn't sell it in the experiment could either use the card in the "real world" once the experiment was over (as a substitute for the money that they would have to pay for 30 liters of gas) or otherwise easily sell the card for its approximate current value. The distribution of the card to the participants was therefore equal to handing out house money. Thus, it seems only rational that participants in their experiment stated selling prices much higher than SEK228 when the upper bound of the BDM distribution (SEK 300) exceeded the market value of this card.

3. In particular, the method of preference elicitation influences revealed preferences (e.g., Fischer and Hawkins 1993; Tversky, Sattath, and Slovic 1988; Lichtenstein and Slovic 1971), "irrelevant alternatives" in a decision problem affect final choices (see Simonson 1989; Ariely and Wallsten 1995; Huber, Payne, and Puto 1982), and individuals have no fixed idea about how much they are willing to pay for experiences with which they are familiar (see Tversky and Kahneman 1974; Ariely, Loewenstein, and Prelec 2003; Chapman and Johnson 2002).

4. If the economy is originally in an equilibrium in which expressed consumer preferences are higher or lower than the true ones, the analysis below would have to account for

whether a tax or subsidy shifts the demand curve toward or away from the true one. For example, if consumers' expressed preferences are initially too high, a tax has the beneficial effect of lowering the volume of trade in the market. Since our methodology does not allow us to determine whether observed demand is above or below the true demand (the one determined by true preferences), we do not discuss such possibilities in detail.

5. We have drawn the shifted demand curve, D2, as parallel to the true one. This is not crucial for any of the results that follow. All we require is that relative consumer preferences do not "reverse" due to the price-induced bias: if one individual's true preferences value the product higher than another individual's true preferences, her expressed willingness to pay is also higher.

6. As is typical in economic analysis, we do not elaborate on how the new equilibrium is reached. Presumably, an initial increase in prices shifts the demand curve up, further increasing prices, shifting the demand curve further, and so on. Eventually, the process settles into a new equilibrium.

FIGURES

FIGURE 1: TWO DIFFERENT PRICE DISTRIBUTIONS UNDERLYING THE MARKET GAME EXPERIMENT 1.

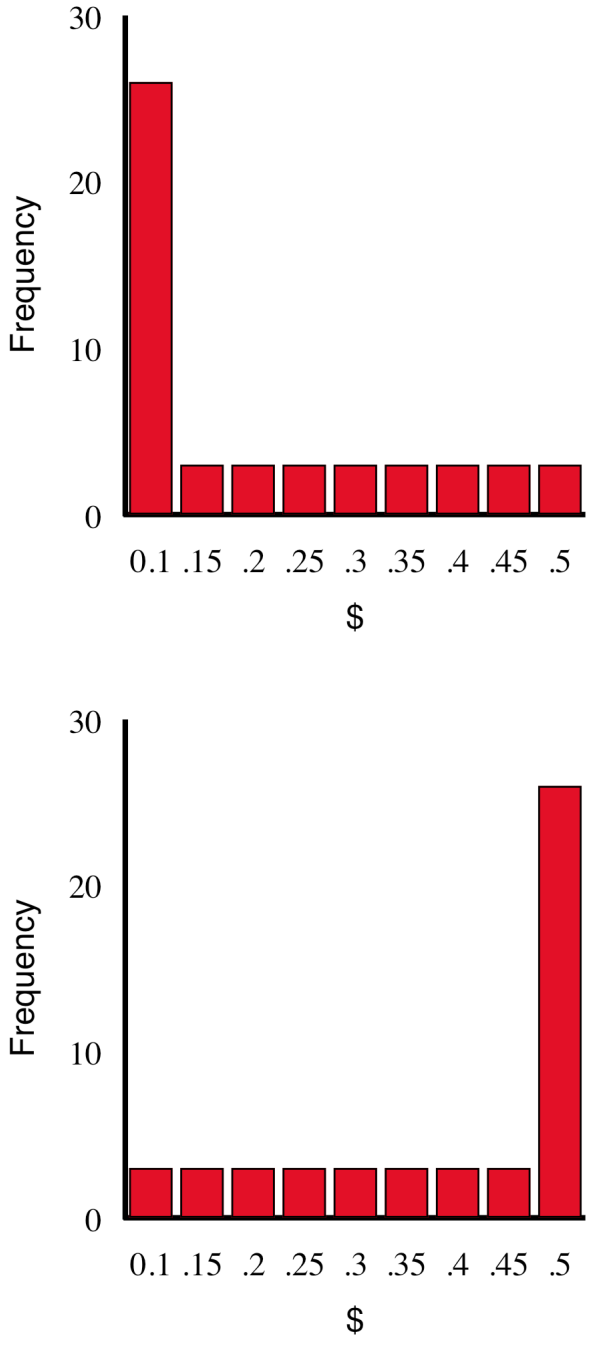

Figure 2: Cumulative frequency distribution OF RESERVATION PRICES across the TWO MARKET CONDITIONS in Experiment 1.

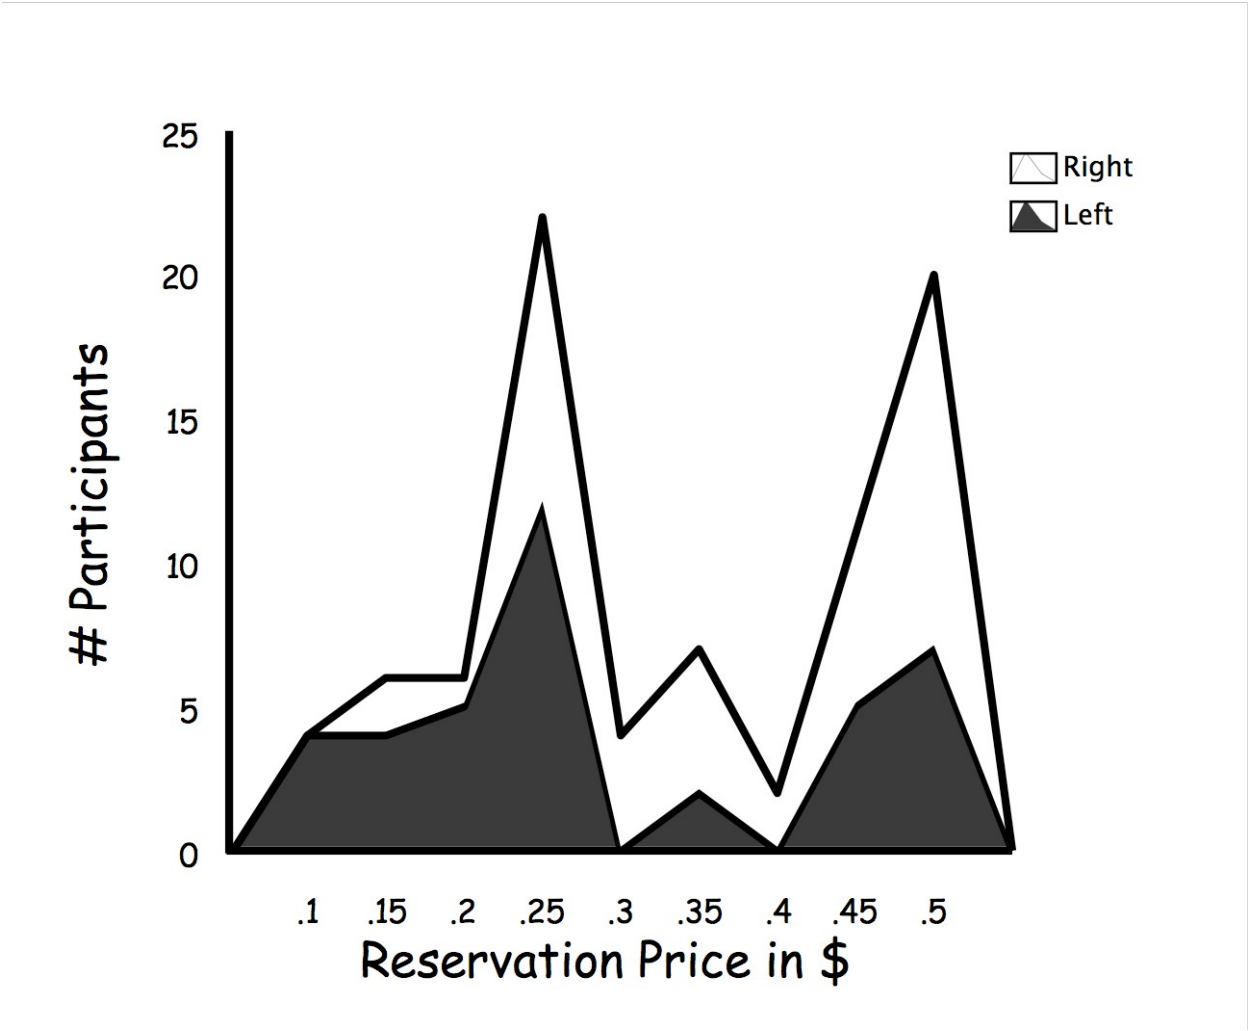

FIGURE 3: THE GENERAL BDM PROCEDURE IN EXPERIMENTS 2 AND 3.

FIGURE 4: TWO DIFFERENT BUYOUT DISTRIBUTIONS UNDERLYING THE BDM PROCEDURE IN EXPERIMENT 2.

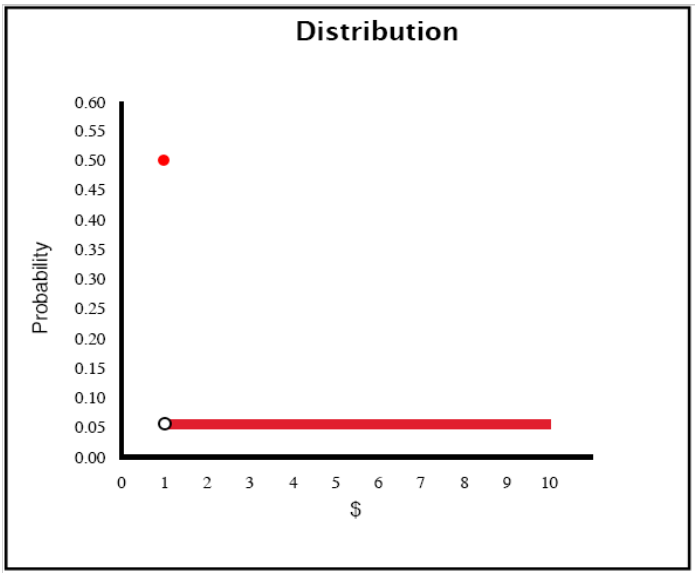

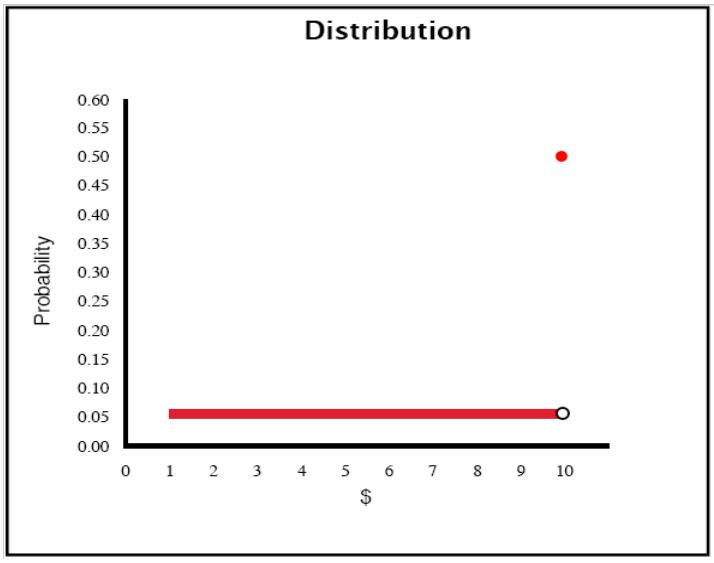
Distribution
Probability
0.60
0.55
0.50
0.45
0.40
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
0 1 2 3 4 5 6 7 8 9 10
$

FIGURE 5: MEAN RESERVATION PRICES FOR THE TWO DISTRIBUTIONS IN THE SEPARATE AND JOINT VALUATIONS CONDITIONS IN EXPERIMENT 2.

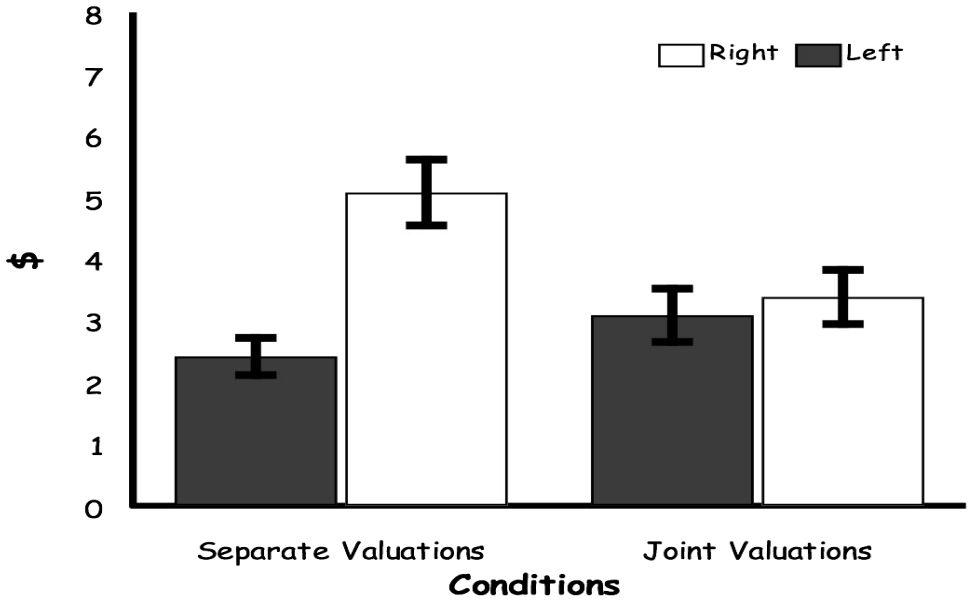

Notes: Error bars are based on standard error of the mean.

FIGURE 6: TWO DIFFERENT BUYOUT DISTRIBUTIONS UNDERLYING THE BDM PROCEDURE IN EXPERIMENT 3.

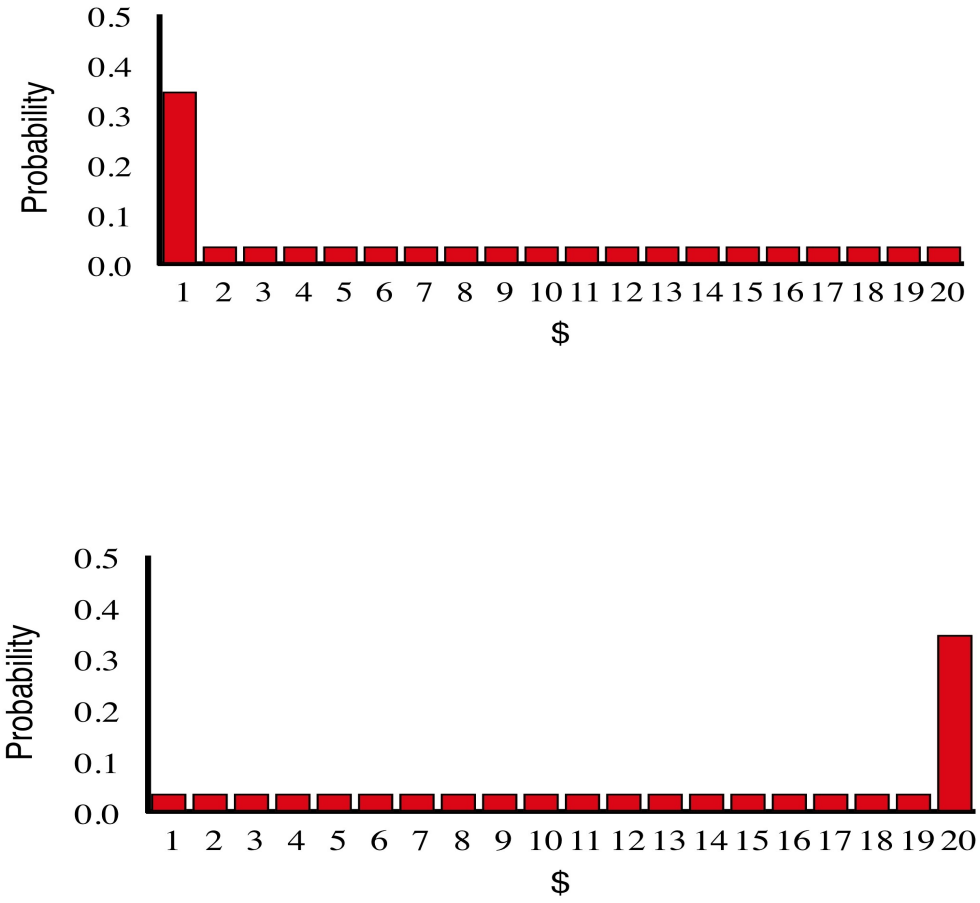

FIGURE 7: MEAN RESERVATION PRICES FOR THE TWO DISTRIBUTIONS ACROSS THREE ELICITION METHODS IN EXPERIMENT 3.

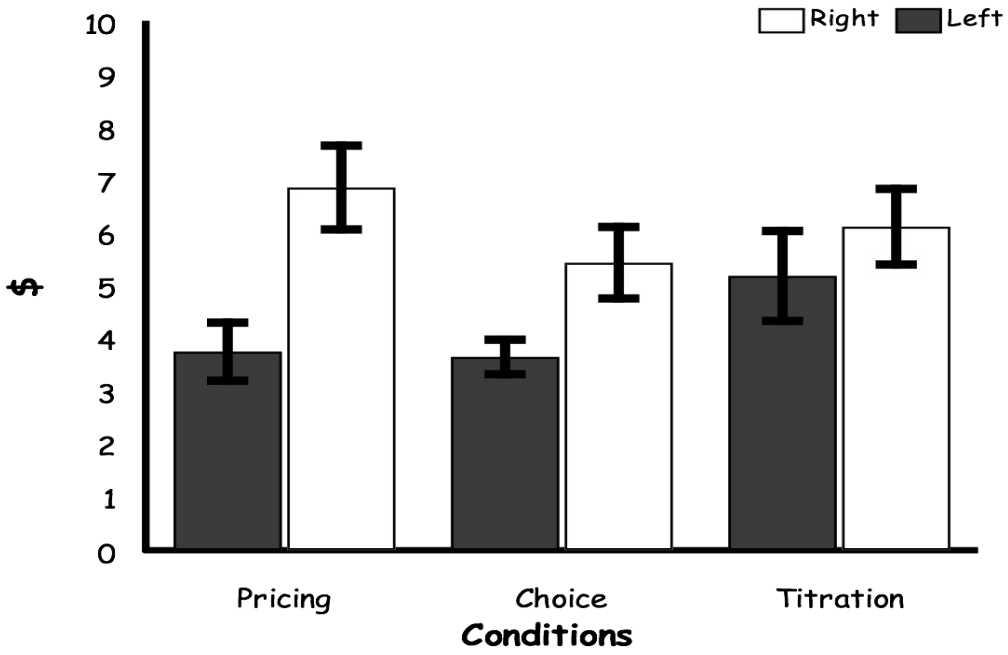

Notes: Error bars are based on standard error of the mean.

FIGURE 8: (MIS)MEASUREMENT OF THE ELASTICITY OF DEMAND, CONSUMER SURPLUS, AND EXCESS BURDEN.

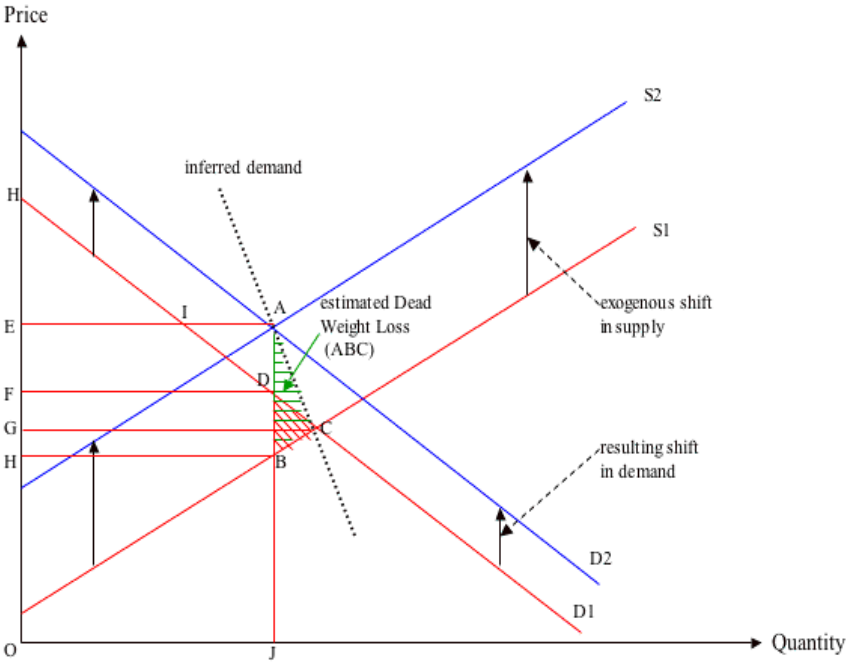

Estimated DWL= ABC, True DWL= DBC < ABC.

APPENDIX A

1. The separate-valuations instruction for the BDM procedure in Experiment 3.

In this study you have a chance to buy a travel mug.

To determine the price, we will use the BDM procedure – please pay careful attention to the instructions below:

There are two price distributions as can be seen below. These distributions are mirror images of each other. In this procedure you will bid against **one** of these distributions. You will have to flip a coin to determine the distribution. If it lands on head you get the distribution on the left and if it lands on tail you get the one on the right.

If the coin lands on head

[PICTURE OF DISTRIBUTION 1]

Please flip a coin now. The outcome of the coin is: _____

Please turn to the next page.

After you pick the price that you are willing to bid for the mug, a random number generator will produce a number (selling price) from the following distribution:

[PICTURE OF RANDOMLY ASSIGNED DISTRIBUTION 1 or 2]

If the price you bid is **higher** than the selling price generated from this distribution, **you will get the mug and pay the selling price selected by the computer!**

If the price you bid is **lower** than the selling price generated from this distribution, **you will not get the mug!**

Note that the MBD procedure is such that your best response is to write down the maximum you are willing to pay for this mug – not more and not less.

My bid is: $_____

Random selling price: _____ Outcome: Buy / not buy

2. The joint-valuations instruction for the BDM procedure in Experiment 3.

In this study you have a chance to buy a travel mug.

To determine the price, we will use the BDM procedure – please pay careful attention to the instructions below:

There are two price distributions as can be seen below. These distributions are mirror images of each other. In this procedure you will bid against **one** of these distributions. You will have to flip a coin to determine the distribution. If it lands on head you get the distribution on the left and if it lands on tail you get the one on the right.

If the coin lands on head

[PICTURE OF DISTRIBUTION 1]

Please turn to the next page.

After you pick the price that you are willing to bid for the mug, a random number generator will produce a number (selling price) from the distribution selected by your coin flip:

If the price you bid is **higher** than the selling price generated from this distribution, **you will get the mug and pay the selling price <u>selected by the computer</u>!**

If the price you bid is **lower** than the selling price generated from this distribution, **you will not get the mug!**

Note that the MBD procedure is such that your best response is to write down the maximum you are willing to pay for this mug – not more and not less.

If the coin lands on head	If the coin lands on tail
[PICTURE OF DISTRIBUTION 1]	[PICTURE OF DISTRIBUTION 2]
My bid is: $_____	My bid is: $_____

Please flip a coin now. The outcome of the coin is: _____

Random selling price: _____ Outcome: Buy / not buy

APPENDIX B

The instructions for the BDM with the direct elicitation method in Experiment 4.

Please complete this only if you are interested in buying **a box of Godiva Chocolate**.

You will now name the maximum price for which you are willing to buy the **box of Godiva Chocolate**. Then another price will be randomly chosen out of a set of prices form $1 to $20 with probabilities as shown in the chart bellow (named “random price”).

If the random price is less or equal to your price, then you will get the box and pay the random price. Otherwise you do not get the box and do not pay. For example if your price is x, and the random price is y, and $x \geq y$, then you get the box and pay y, but if $x < y$, then you do not get the box, and do not pay. This means that the optimal strategy for you is to state the maximum price for which you are willing to buy the box.

[PICTURE OF THE DISTRIBUTION]

Your price: $_____

APPENDIX C

The instructions for the BDM with the choice-based elicitation method in Experiment 4.

In this study you will have a chance to buy a **box of Godiva Chocolate**.

In each of the lines 1 through 20 on the next page you will be presented with a choice of buying a **box of Godiva Chocolate** or not for a certain price. In each of the lines, check the outcome you would choose. One of the lines will be picked at random with a probability that is indicated in the right column of each line. These probabilities are also shown in the chart bellow. The line that is picked will be carried out for real. If in that line, you checked "Get the box for $x", then you will get the box and pay $x. If in that line you checked "Not get the box and not pay", then you will not get the box and will not pay. The optimal strategy for you is to state your true preferences in each line.

[PICTURE OF THE DISTRIBUTION]

Please turn to the next page.

			Probability to be selected
1.	❑Get the box for $1	❑Not get the box and not pay	[10 or 1]/29
2.	❑Get the box for $2	❑Not get the box and not pay	1/29
3.	❑Get the box for $3	❑Not get the box and not pay	1/29
4.	❑Get the box for $4	❑Not get the box and not pay	1/29
5.	❑Get the box for $5	❑Not get the box and not pay	1/29
6.	❑Get the box for $6	❑Not get the box and not pay	1/29
7.	❑Get the box for $7	❑Not get the box and not pay	1/29
8.	❑Get the box for $8	❑Not get the box and not pay	1/29
9.	❑Get the box for $9	❑Not get the box and not pay	1/29

10.	❑Get the box for $10	❑Not get the box and not pay	1/29
11.	❑Get the box for $11	❑Not get the box and not pay	1/29
12.	❑Get the box for $12	❑Not get the box and not pay	1/29
13.	❑Get the box for $13	❑Not get the box and not pay	1/29
14.	❑Get the box for $14	❑Not get the box and not pay	1/29
15.	❑Get the box for $15	❑Not get the box and not pay	1/29
16.	❑Get the box for $16	❑Not get the box and not pay	1/29
17.	❑Get the box for $17	❑Not get the box and not pay	1/29
18.	❑Get the box for $18	❑Not get the box and not pay	1/29
19.	❑Get the box for $19	❑Not get the box and not pay	1/29
20.	❑Get the box for $20	❑Not get the box and not pay	[1 or 10]/29

APPENDIX D

The instructions for the BDM with the titration-based elicitation method in Experiment 4.

In this study you will have a chance to buy a **box of Godiva Chocolate**.

During the study you will now go through a set of iterations that will ultimately bring you to the maximum price at which you are willing to buy the **box of Godiva Chocolate** (this will be your answer in step #13). Then another price will be randomly chosen out of a set of prices form $1 to $20 with probabilities as shown in the chart bellow (named "random price").

If the random price is less or equal to your price in step #13, then you will get the box and pay the random price. Otherwise you do not get the box and do not pay. For example if your price in step #13 is x, and the random price is y, and $x \geq y$, then you get the box and pay y, but if $x < y$, then you do not get the box, and do not pay. This means that the optimal strategy for you in step #13 is to state the maximum price for which you are willing to buy the box. Your answers in all other steps will have no impact on the outcome and are designed to help you figure out the maximum price you are willing to pay for the box.

[PICTURE OF THE DISTRIBUTION]

Please turn to the next page.

Please follow the following steps:

1) Write down a number that is a reasonable estimate of how much you would be willing to pay for the box; $ _____
2) Ask yourself if this price is (mark with a √):
   a. Too high ____ (go to step #3)
   b. Too low ____ (go to step #3)
   c. This is the max I am willing to pay _____ (go to step #13)
3) Write down **a new** number that is a reasonable estimate of how much you would be willing to pay for the box; $ _____
4) Ask yourself if this price is (mark with a √):
   a. Too high ____ (go to step #5)

   b. Too low ____ (go to step #5)
   c. This is the max I am willing to pay _____ (go to step #13)
5) Write down **a new** number that is a reasonable estimate of how much you would be willing to pay for the box; $ _____
6) Ask yourself if this price is (mark with a √):
   a. Too high ____ (go to step #7)
   b. Too low ____ (go to step #7)
   c. This is the max I am willing to pay _____ (go to step #13)
7) Write down **a new** number that is a reasonable estimate of how much you would be willing to pay for the box; $ _____
8) Ask yourself if this price is (mark with a √):
   a. Too high ____ (go to step #9)
   b. Too low ____ (go to step #9)
   c. This is the max I am willing to pay _____ (go to step #13)
9) Write down **a new** number that is a reasonable estimate of how much you would be willing to pay for the box; $ _____
10) Ask yourself if this price is (mark with a √):
   a. Too high ____ (go to step #11)
   b. Too low ____ (go to step #11)
   c. This is the max I am willing to pay _____ (go to step #13)
11) Write down **a new** number that is a reasonable estimate of how much you would be willing to pay for the box; $ _____
12) Ask yourself if this price is (mark with a √):
   a. Too high ____ (go to step #13)
   b. Too low ____ (go to step #13)
   c. This is the max I am willing to pay _____ (go to step #13)
13) What is the maximum you are willing to pay for the box; $ _____