Journal of Neuroscience, Psychology, and Economics
2016, Vol. 9, No. 2, 65–77

1937-321X/16/$12.00 http://dx.doi.org/10.1037/npe0000057

Flexible Valuations for Consumer Goods as Measured by the Becker–DeGroot–Marschak Mechanism

Agnieszka Tymula
University of Sydney

Eva Woelbert
Maastricht University

Paul Glimcher
New York University

Economists, psychologists, and neuroscientists have long been interested in methods that elicit individuals' true valuations of goods. In this paper, we take 1 of the most popular of such mechanisms, the Becker–DeGroot–Marschak (BDM) procedure, and study the nature of the dependence of the valuations obtained using the BDM procedure on the distribution of prices presented to subjects when the method is implemented. In a within-subject design with products with a high market value, we show that this effect of price distribution occurs quite frequently, significantly impacts reported valuations, and is unlikely to be caused by misconceptions about the BDM procedure. This effect is largest when pricing distributions show a large peak close to an individual's average valuation of the good. A simple nonincentive-compatible subjective rating of the desirability of goods can be used to predict the likelihood that pricing distributions will influence BDM valuations; valuations for goods that subjects report to most want to purchase are most likely to be influenced by distributional structure. Our results challenge some of the dominant theoretical models of how BDM-like valuation procedures relate to standard notions of utility and shed light on how to interpret the data obtained using the BDM method.

Keywords: utility, reference, valuation

Supplemental materials: http://dx.doi.org/10.1037/npe0000057.supp

Knowing a person's true valuation for a good is important for research on human decision making and marketing and for policy decisions on the provision of public goods. Of course, simply asking someone to state their valuation, without giving them any incentive to tell the truth, has been found to sometimes lead to inflated value estimates (Harrison & Rutström, 2008; List & Gallet, 2001; Wertenbroch & Skiera, 2002). To increase the accuracy of value measurements, incentive-compatible value elicitation methods are thus often used, which guarantee that misstating valuation is costly for participants. In the Becker–DeGroot–Marschak (BDM) mechanism, which is one of the most widely used of such methods, subjects are asked to state the maximum amount they would be willing to pay for a given good. The actual price for the good is then determined randomly from a distribution of possible prices (known to the subject in advance). If the stated willingness to pay (WTP) is higher than the randomly determined price, the subject buys the good from the experimenter, paying the randomly determined price. If the maximum stated WTP is lower than the randomly determined price, the subject does not buy the good (Becker, DeGroot, & Marschak, 1964). This procedure is widely assumed to provide the right incentives for sub-

Agnieszka Tymula, School of Economics, University of Sydney; Eva Woelbert, Department of Cognitive Neuroscience and Department of Economics, Maastricht University; Paul Glimcher, Center for Neuroeconomics, New York University.

Agnieszka Tymula gratefully acknowledges financial support from the AXA Research Fund.

Correspondence concerning this article should be addressed to Agnieszka Tymula, School of Economics, University of Sydney, Room 370, Merewether Building (H04), New South Wales 2006, Australia. E-mail: agnieszka.tymula@sydney.edu.au

jects to truthfully reveal their valuation; by reporting truthfully, they both maximize their odds of getting the good at an acceptable price and avoid any risk of overpaying.

However, as Karni and Safra (1987) and Horowitz (2006) have shown theoretically, elicitation of the *true value* of both lotteries and goods using the BDM mechanism can be problematic. In the BDM mechanism, the subject faces uncertainty regarding the actual price, which also creates uncertainty about whether or not he or she will buy the good. If the utility function of the chooser depends on any of these attributes (as in, e.g., certain types of reference-dependent utility, anticipatory utility, or disappointment aversion models), then it may be optimal for the chooser to bid an amount different from his or her true value for the good (his or her maximum WTP when there is no uncertainty about the price). Hence, manipulations to the pricing distribution could well be expected to change individual bids if individuals are not simple expected utility maximizers.

In line with these theoretical concerns, existing experimental evidence suggests that the BDM mechanism does not elicit subjects' truthful valuations under some conditions (Kaas & Ruprecht, 2006; Noussair, Robin, & Ruffieux, 2004; Rutström, 1998; Shogren et al., 2001), and it has been suggested that bidders may, in line with Horowitz's (2006) theoretical contribution, be influenced by the price distribution (Bohm, Linden, & Sonnegard, 1997; Mazar, Kõszegi, & Ariely, 2014; Urbancic, 2011). While previous studies have generally found that such distributional dependence takes the form that price distributions with a *higher* expected value result in *higher* observed valuations, findings on the extent and the robustness of this effect have been conflicting, and studies have not been able to provide an intuition for the underlying mechanism.

In this empirical paper, we therefore set out to investigate the nature of distributional dependence in the BDM mechanism in detail. The distribution of possible prices is a central feature of the BDM mechanism that the researcher controls, and therefore it seems crucial to understand how this choice affects the data. Moreover, uncertainty in prices is not a feature that is limited to the BDM mechanism—it applies to ordinary purchase situations as well. Before observing the actual price for a good, consumers are thought to have an expectation of the distribution of possible prices, which may affect purchase decisions (Heidhues & Kõszegi, 2014). Thus, distributional dependence in the BDM mechanism may be an opportunity to study a more general pattern of the influence of expected prices on choices.

Reviewing the literature on distributional dependence, it becomes apparent that previously published papers manipulated the price distributions by changing the upper and lower boundaries of the distribution, either by changing the maximum price or by increasing the probability of the lowest or highest possible price: Bohm et al. (1997) elicited valuations for a petrol voucher using the BDM mechanism. Varying the upper bound of the price distribution between subjects, they observed higher bids when the upper bound was higher.

Also using a between-subjects comparison, Mazar et al. (2014) found that BDM bids for a travel mug were higher when the price distribution assigned a high probability to the highest possible price compared with when the lowest price was highly probable. Interestingly, this difference was much smaller in a within-subject comparison—that is, when the same subjects were asked to place two bids on the travel mug, one for each distribution. This is all the more surprising because all subjects were informed about the existence of the two price distributions, excluding the possibility that subjects used the price distribution to draw rational inferences about the quality of the good on offer. Because subjects reported similar valuations when forced to bid under both distributions, Mazar et al. (2014) concluded that people do not truly hold different WTP for different price distributions but instead, when considering only one price distribution, people make mistakes in reporting their true WTP. In contrast to the findings of Mazar et al. (2014), in an unpublished paper, Urbancic (2011), using a within-subject design, found that distributional dependence still holds when people repeatedly bid for a cookies voucher but did not provide an intuition of why his result differs from those by Mazar et al. (2014). More recently, Cason and Plott (2014) provided evidence suggesting that misconceptions play a big role in suboptimal bidding behavior in the BDM mechanism. In their experiment, they asked individuals to bid twice on a $2 voucher. Those who did not bid

$2 and were exposed to their mistake were more likely to choose the correct valuation of $2 in the second round.

To provide insights relevant to these seemingly contradictory results in the literature, we aimed to create an environment where pricing distributions were systematically manipulated while ensuring that mistakes and misconceptions about the BDM mechanism were minimized. We asked the subjects to *repeatedly* (rather than just once or twice as in Mazar et al. (2014), and Cason & Plott, 2014) value each of three goods for (on average) more than 50 rounds per good. Importantly, as in Cason and Plott (2014), after every round, we told subjects which price was realized in that round and whether they would purchase the good if this round was later counted for payment, thus reinforcing their understanding of the mechanism. We hypothesized that if subjects had any doubts about the procedure in the beginning, after tens of rounds with feedback they should have had a clear understanding about how the BDM mechanism works during the bulk of our empirical measurements. Further, the repeated nature of the task made it very transparent for subjects that the price distribution is not informative of the value of the good. Finally, we employed only goods with suggested retail prices above $50. We thus used relatively large stakes and ensured that subjects were motivated to buy the goods typically offered at prices far below $50 during the experiment.

The central methodological feature of our study—and one that allows us to reconcile our findings with the different results found in the literature—is the fact that we observed the same person repeatedly under a continuum of pricing distributions rather than just under two extreme price distributions. In the beginning of every round, one price from the fixed support of the price distribution (from $0 to $50) was randomly selected to be the *most likely price*; all other prices were equally likely. With this information available, subjects then stated their valuation. Subjects bid not on one but on three different goods, which enabled us to disentangle subject- and good-specific effects. Under these conditions, we found that subjects bid higher (lower) when the most likely price was higher (lower) than the subject's average bid for that good. This effect remains significant and does not weaken throughout many rounds of bidding.

Our novel design allowed us to make interesting observations that may shed light on the discrepancies among previous studies. Biasing of stated WTP was strongest when the most likely price was relatively close to the subject's average bid for the respective good, a finding compatible with previous studies finding a range of effect sizes. All in all, our results suggest that distributional dependence of the BDM mechanism is empirically robust and persists even after thorough training and experience. We discuss our findings in light of theories of reference-dependent preferences because it has been suggested that price distributions can act as reference points (Heidhues & Kõszegi, 2014; Wenner, 2015). Interestingly, one of the most prominent theories of reference-dependent preferences (Kõszegi & Rabin, 2006) makes predictions that run counter to our findings. Our findings appear to be more in line with a simple model proposed by Wenner (2015), suggesting that subjects seek to pay a price that compares favorably against the expected price. In both models (Kõszegi & Rabin, 2006; Wenner, 2015), reference prices are the key drivers of behavior. The key difference between these models is that in Kõszegi and Rabin (2006), prices at which individuals would not buy a good do not influence their utility (a feature of the personal equilibrium) and in Wenner (2015), they do. In the Discussion, we discuss our findings in relation to these two and other theoretical models.

Materials and Method

Twenty-seven paid volunteers (12 females) participated in this experiment. At the beginning of the experiment, participants were informed that they would be repeatedly asked to state their maximum WTP (bid) for three different goods: a backpack, an iPod Shuffle, and a pair of noise-canceling headphones. To ensure that subjects had enough information about the goods to understand what they were bidding on, they were given substantial time to inspect the products. They were informed that the suggested retail price of all three products, including tax, was higher than the maximum possible price in the experiment ($50) but were not told the exact market prices. At the time of the experiment, the actual market prices of the goods were $49 plus tax (iPod Shuffle), $49 plus tax (Sony noise-canceling headphones), and $59 plus tax (Case Logic laptop backpack with iPad pocket). Participants learned that, out of all rounds they completed, one round would be randomly selected and implemented at the conclusion of the experiment. Following the standard BDM procedure, they would purchase the good at the randomly selected actual price if their bid on that trial was higher than or equal to the randomly selected actual price. If they had bid lower than the actual price, they would not purchase the good. Each participant was given $50 that he or she could use toward his or her purchase. Participants could keep any money not spent during the experiment. Participants who made no purchases would keep the full endowment of $50 and receive no goods.

Each round followed the same structure: First, participants learned which good they were bidding on in that round and which price was most likely to be selected as the actual price (that price would be selected with a probability equal to 0.51; $41 in Figure 1A). We refer to the most likely price as the *revealed price*. All other prices between $1 and $50 (in steps of $1) had an equal probability of .01 to be selected as the actual price (represented as a question mark in Figure 1A). Second, after observing the information on the good and the revealed price, participants could enter a bid (constrained to be between $0 and $50) for the current good and round ($37 in Figure 1B). Third, the actual price for this round was then determined randomly by drawing from the specified distribution and revealed to the subject. Subjects were also told whether or not they would buy the good if this round was later selected for payment, so as to reinforce their understanding of the BDM procedure (Figure 1C).

Each participant worked through the experiment at his or her own pace and in private, completing, on average, 140 rounds (SD = 33.98).[1] The number of rounds experienced was determined by one of a small set of *termination rules*, varying by subject, that did not affect the results presented herein (see online supplemental material Appendix B for the description and analysis). Note that, irrespective of the termination rule used on that subject, for all of our subjects, standard theory predicts that the subject should be insensitive to the price distribution—constantly bidding the same amount for a given good.

Before the beginning of the experiment, these procedures were explained to the participants in detail using written instructions and extensive examples (see online supplemental material Appendix A). The instructions stressed that the actual price of a good was determined randomly and could not be influenced by the subject's bid. To test their understanding of the task, participants were given a set of three comprehension questions (see Part 1 in online supplemental material Appendix A). If they did not answer all of the questions correctly, additional explanations were given, and they were given a second set of three comprehension questions (Part 2). Participants knew in advance that if they failed to answer the second set of questions correctly, they would receive $5 as a showup fee and would not be allowed to participate in the study. There were no participants who failed to answer the second set of comprehension questions correctly; therefore, all invited participants were admitted to the experiment. In the debriefing questionnaire, all participants indicated that the instructions were either clear or very clear. Participants were also given five practice rounds, which were not relevant for their payoff, in order to become familiar with the experiment.

Data were collected at the Center for Experimental Social Science (CESS) at New York University (NYU). All procedures were approved by the NYU Institutional Review Board, and all participants gave informed consent. Sessions lasted approximately 90 min. The task was programmed using ePrime 2.0 software (Psychology Software Tools, Pittsburgh, PA).

Results

Subjects bid, on average, $19.57, indicating that they were interested in buying the goods on offer. Mean bids for each good and standard deviations are shown in Table 1. Subjects tended to bid multiples of 5 much more often than all other prices, which resulted in a multimodal distribution of bids (see Figure 2). Overall, 62% of the bids were placed on these focal points.

[1] For four subjects, the experiment was aborted early due to a technical failure. We include the data from these subjects that were recorded up to the failure (60–93 rounds). No self-report was collected from these subjects after the experiment.

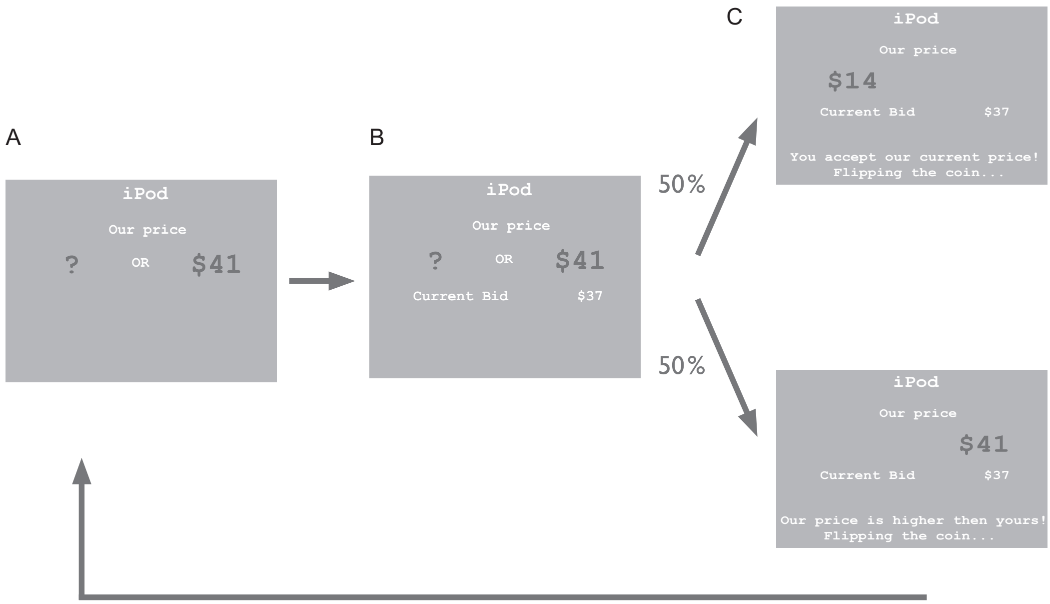

Figure 1. Structure of the experiment.

Importantly, contrary to the prediction under traditional expected utility theory (Becker et al., 1964), subjects did not state a constant maximum WTP for each good. The bid of an average subject for a given good varied substantially, with an average standard deviation of 3.35. Subjects differed with respect to the variability in the bids, one subject bidding constantly the same amount for a good and some subjects changing their WTP substantially from round to round. There was a decrease in variability over time, but bids still varied considerably in the second half of the experiment (see Table 2).

The Effect of the Revealed Price

We found that bids were significantly higher when the revealed price was high (equal to or higher than $25) compared with when it was low (below $25; $p < .01$, Wilcoxon signed-ranks test over subjects' mean bids), suggesting that the distribution of prices used in the BDM mechanism may affect the WTP in a systematic way. To verify if this is indeed true, we first plot the bid in each round as a function of the revealed price in that round. To account for between-subjects variability, we plot both the bid and the revealed price relative to the subject's mean bid for the respective good. A scatter plot of these normalized bids against normalized revealed price shows that the data points are separable into two very distinct patterns (see Figure 3). For some, there is a strong influence of the revealed price on the bid (with a slope near 1), whereas for others there is no influence of the revealed price at all.

Figure 3 also suggests that the revealed price affects bids—not across the whole range of possible values but mostly when it is relatively close to the mean bid that a subject places on a good. To verify whether this is indeed the case and to quantify the effect, we regressed the bid in each round on the revealed price on the full data set and on reduced data sets that include only rounds where the revealed price was relatively close to a subject's mean bid for a good

Table 1
Descriptive Statistics for Bids

Good	Mean bid	Standard deviation bid
Backpack	$17.16	15.62
Headphones	$20.60	14.90
iPod	$21.10	16.52
Total	$19.57	15.78

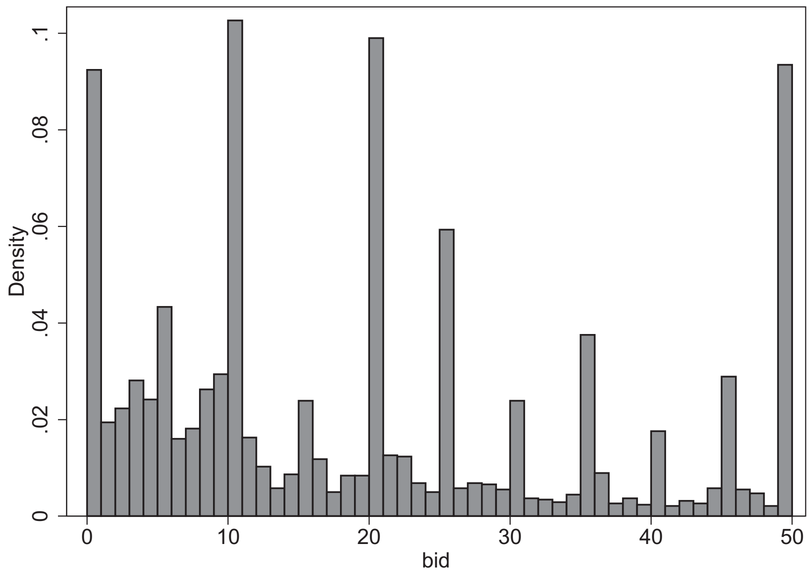

Figure 2. Histograms over the bids.

(see Table 3). The coefficient for the revealed price is larger when considering only rounds where the revealed price was relatively close to the subject's average bid. When the revealed price is in the range of ±\$3.35 (equivalent to 1 mean standard deviation) of the mean bid, the bid increases by \$0.41 on average for each \$1 increase in the revealed price. To further confirm whether the impact of revealed price weakens as the distance between the revealed bid and the individual's mean bid for the good increases, we included a *distance* variable equal to the absolute value of the difference between the revealed price and the individual's mean bid for the good in the model. It turns out that the coefficient on the interaction between distance and revealed price (Revealed Price × Distance in Model 5; Table 3) is significant and negative, meaning that revealed prices further away from the individual's mean bid have less of an effect on the individual's bid. The bid and the revealed price in the preceding round did not affect the bid in the current round, which is reassuring, as it suggests that despite the repeated nature of our task, there is no influence from one round to the next. In order to ascertain that this finding is not driven by the fact that subjects would typically bid on a different good from one round to the next, in Table C1 in online supplemental material Appendix C, we instead include the previous bid and the previous revealed price on the same good as a control variable. The pattern of results remains the same, and there is no indication that subjects were influenced by the previous round. There was also no systematic change in the bids or in the effect of the revealed price on the bids over the course of the experiment.

Table 2
Summary Statistics for Standard Deviations of the Bids an Individual Subject Placed for One Specific Good

Experimental half	Mean	Median	Minimum	Maximum
First half	3.46	2.30	0	16.38
Second half	2.42	1.71	0	13.96
All	3.35	2.30	0	15.88

Variability in the Effect of the Revealed Price

Figure 3 shows that sometimes there is a strong effect of the revealed price but sometimes no effect at all. In order to understand whether this effect occurs only for some people or some goods, we regressed the bid on the

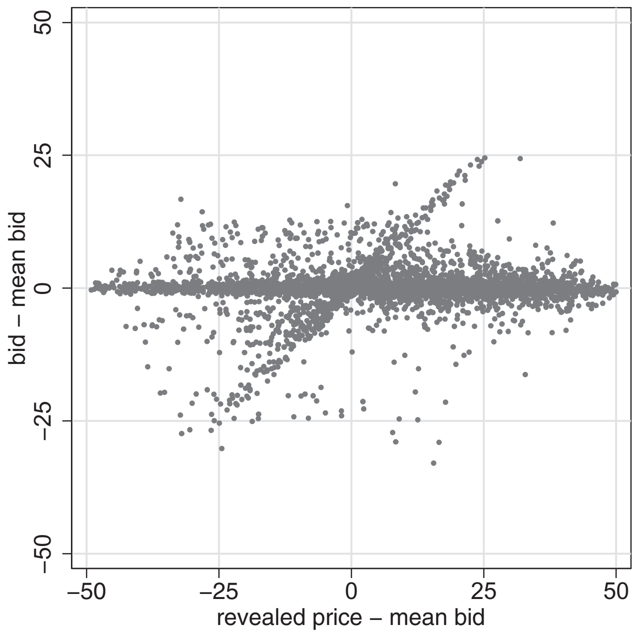

Figure 3. Scatter plot of bid against revealed price, both centered on the mean bid that the subject placed for the respective good over the course of the experiment.

revealed price separately for each good and each subject (this means running three regressions for most subjects and two or one regression[s] for those who finished the experiment before they got a chance to bid on all three goods). In these individual regressions, 48% (41%) of the subjects showed a significant positive effect of the revealed price for at least one good at the 10% (5%) level. Table 4 summarizes the results of these individual- and good-specific regressions. Interestingly, we find that out of the ones that do show an individually significant effect of the revealed price, 77% (82%) show it for one or two of the goods, but

Table 3
The Effect of the Revealed Price on the Bid

Independent variables	1 All data	2 ±2 *SD*	3 ±1.5 *SD*	4 ±1 *SD*	5 All data
Revealed price	.09 (.05)	.23** (.07)	.31*** (.07)	.41*** (.08)	.18* (.07)
Previous bid	.14 (.10)	.11 (.08)	.09 (.08)	.11 (.09)	.14 (.09)
Previous revealed price	.00 (.01)	.01 (.01)	.01 (.01)	.00 (.01)	.00 (.01)
Round	.02 (.01)	.01 (.01)	.01 (.01)	.01 (.01)	.02 (.01)
Revealed Price × Round	.00 (.00)	.00 (.00)	.00 (.00)	.00 (.00)	.00 (.00)
Revealed Price × Distance					−.00** (.00)
Distance					.03 (.02)
Constant	13.85*** (2.81)	13.59*** (1.83)	12.04*** (1.88)	9.25** (3.07)	12.84*** (2.92)
Adjusted R^2	.086	.172	.175	.216	.107
Number of observations	3,781	861	654	444	3,781

Note. Dependent variable is the bid in a given trial. Previous bid (revealed price) is the bid (revealed price) in the previous round. Distance = |revealed price − mean bid|. Fixed effects for each subject–good combination are included. Standard errors are clustered on subject level and given in parentheses.
$^{*}p < .05$. $^{**}p < .01$. $^{***}p < .001$.

Table 4
Frequency of the Revealed Price Bias on the Individual Level

Goods affected by revealed price	10% level		5% level	
	Frequency	Percentage	Frequency	Percentage
0	14	51.85%	16	59.26%
1	6	22.22%	4	14.81%
2	4	14.81%	5	18.52%
3	3	11.11%	2	7.41%
At least one good	13	48.15%	11	40.74%

Note. The bid was regressed on the revealed price in separate regressions for each subject–good combination. *Goods affected by revealed price* indicates for how many goods a subject showed a significant revealed price bias on the 10% and 5% levels in these individual regressions.

not for another good. This suggests that both individual- and good-specific factors play a role in determining how strongly a person responded to the revealed price and are driving the two distinct patterns shown in Figure 3.

Next, we attempted to determine whether we could identify a predictor that would correlate with the stronger or weaker distributional dependence of a subject's valuations of each good. To do that, we had each subject rate his or her *desire to buy* each of the three goods on a 5-point Likert scale (Likert, 1932), which ranged from *not at all* to *really wanted to buy*. As shown in the regression results presented in Table 5, subjects who report a high Likert number (a strong *wish to buy*) for a particular good are more likely to be influenced in their bidding behavior by the price distribution, even if the revealed price is relatively far away from their average bid (the interaction term of the revealed price and the self-reported wish to buy the good is significant in the full data set). When restricting the range of the revealed price to the vicinity of the mean bid, the interaction is insignificant and the revealed price itself is a strong predictor of the bid, suggesting that when people do not express a strong wish to buy the good, they will still react to the revealed price, but only if it is close to their mean bid.

Finally, we assessed whether there was any evidence that subjects are matching the revealed price exactly. To this end, we computed the proportion of bids, which was equal to the revealed price, and compared it to the expected frequency of the revealed price matching the subject's bid under the assumption that the bid is not influenced by the revealed price. We observed that, overall,

Table 5
The Effect of the Revealed Price Depends on How Strongly the Subject Wants to Buy the Good

Independent variables	−1 All data	−2 ±2 *SD*	−3 ±1.5 *SD*	−4 ±1 *SD*
Revealed price	.01 (.03)	.37*** (.09)	.48*** (.17)	.68 (.41)
Previous bid	.14 (.10)	.11 (.09)	.10 (.09)	.12 (.10)
Previous revealed price	−.00 (.01)	−.01 (.01)	−.01 (.01)	−.00 (.01)
Round	.01 (.01)	.01 (.01)	.01 (.01)	.02 (.01)
Wanting × Revealed Price	.02** (.01)	−.02 (.03)	−.04 (.05)	−.06 (.12)
Constant	15.02** (2.68)	12.43*** (2.67)	11.44*** (2.53)	8.41* (4.12)
Adjusted R^2	.095	.154	.148	.168
Number of observations	3,491	787	592	400

Note. Dependent variable is the bid in a given trial. Fixed effects for each subject–good combination are included. Standard errors are clustered on subject level and given in parentheses. Data on wanting were missing for four subjects due to technical failure.
* $p < .05$. ** $p < .01$. *** $p < .001$.

about 8.5% of bids are equal to the revealed price, which is significantly higher than the 2% of matches we would expect at random, $F(1, 26) = 14.33$, $p < .01$.

We further observed that subjects are significantly more likely to match the revealed price if they report a stronger wish to buy the particular good, $F(1, 22) = 5.45$, $p = .029$. Thus, it appears that subjects are indeed matching the revealed price, and they are particularly likely to do so if they are more interested in buying the good.

Discussion

We find that most subjects' bids in the BDM mechanism not only depend on the price distribution used but also, surprisingly, that distributional dependence persists and does not weaken over many repetitions. This is true even though subjects repeatedly bid on the same three goods and receive feedback regarding outcomes on every round. This, and the fact that we had careful instructions and comprehension questions, leads us to conclude that the effects that we observed in this study are not likely to be due to misconceptions about the BDM mechanism but instead reveal something about the valuation process and the underlying preferences of our subjects. The main novel insights from this study are that (a) the effect of the revealed price on individual bids is strongest when the revealed price is close to the individual mean WTP for a good, (b) individuals often exhibit distributional dependence in their bids for some but not all of the goods, and (c) subjects show the effect over a wider range of revealed prices for goods for which they report a strong desire to buy.

Some recent empirical studies have aimed to investigate the nature of distributional dependence on WTP (Cason & Plott, 2014; Mazar et al., 2014; Wenner, 2015). In line with our findings, these studies observed strong distributional dependencies in the direction that we observed. Interestingly, these studies found that distributional dependence was decreased under certain circumstances, such as when subjects reported their maximum WTPs simultaneously for two different distributions (Mazar et al., 2014) or when they were informed of the full range of possible prices (Wenner, 2015), which poses the question of whether distributional dependence is a robust phenomenon. Cason and Plott (2014) argue that misconceptions about the BDM mechanism rather than true distributional dependence of preferences explain the distributional dependence of WTP. In their study, 57% of individuals who observed the same price distribution the second time they made the decision changed their bid. Moreover, those who were exposed to their own mistake of not bidding \$2 for a \$2 money voucher were more likely to later bid \$2 than those who were not exposed to their mistake, suggesting that, with appropriate training and experience, the distributional dependence should disappear.

Our results, as well as those of an unpublished study (Urbancic, 2011), show that even if distributional dependence may be decreased in a within-subject design, it is not eliminated and is strongest when the manipulation occurs close to the subject's valuation for the good. In the study by Mazar et al. (2014), only either the maximum or minimum of the price distribution served as the most likely price, and the support of the distribution was selected such that the subject's valuation should be somewhere in the middle of the distribution. Similarly in Cason and Plott (2014), the possible prices for a \$2 voucher were drawn from a range of prices from \$0 to an upper limit that varied from \$4 to \$8. In light of our findings, it appears that in both Mazar et al. (2014) and Cason and Plott (2014), the distributional dependence was minimized by design. The study by Wenner (2015) used purchasing decisions rather than the BDM mechanism to investigate distributional dependence. In this case, distributional dependence is only detected if it changes the subjects' WTP from below to above the requested price (or vice versa) and therefore changes the buying decisions. Only a quarter of subjects seemed motivated to buy the good on offer which, according to our results, reduces the likelihood that they would observe distributional dependence even if it were present. Notably, when analyzing nonincentivized questionnaire data, the study found that even after controlling for a variety of factors, distributional dependence was still evident in stated WTP. It is likely, then, that the results reported by Cason and Plott (2014), Mazar et al.

(2014), and Wenner (2015) underestimate the true extent of distributional dependence.

Mazar et al. (2014) and Cason and Plott (2014) suggest that previous studies of the distributional dependence of bids in the BDM mechanism overestimated the effect because the subjects were confused about the link between their actions and outcomes. In line with such an interpretation, Hao and Houser (2012), studying a different elicitation mechanism, show that the way the mechanism is implemented plays a big role, especially for naïve decision makers. Indeed, in our study, the size of the effect is not as big as in previous studies and is isolated to changes in prize distribution around one's average valuation. Nevertheless, using large incentives and a continuous price distribution, we clearly show that distributional dependence persists after controlling for rational inferences and possible misconceptions.

Surveying the existing literature, we found two conceptually different types of explanations for this dependency. Neither of these ideas and models seems, however, to fully account for the effects that we observed. Explanations of the first type focus on the idea that differences in bids under different price distributions stem from imprecise knowledge about the subject's own valuation of the item. The second type assumes that subjects know their valuation of the good but that their utility directly depends on the changes to the distribution of prices in the BDM mechanism.

Assuming that subjects are unsure of their valuation for the product on offer, there are several ways that the price distribution could influence the subjects' bids. Price distribution could serve as a credible signal of market value and quality. Thus, higher price distributions may lead to higher bids because subjects believe that they are bidding on a more valuable product. Prices that are close to a person's initial guess for his or her valuation may also provide an incentive to think more about one's exact value for a good and thus could lead a subject to come to a more precise estimate of that value (Wathieu & Bertini, 2007). In a similar way, Lusk, Alexander, and Rousu (2007) and Kaas and Ruprecht (2006) have discussed imprecise preferences as a potential cause of distributional dependence. It is known that even bringing arbitrary numbers to subjects' attention can exert influence on WTP for consumer goods under some conditions, presumably because they interpret them as informative (Ariely, Loewenstein, & Prelec, 2003; Fudenberg, Levine, & Maniadis, 2012; Sugden, Zheng, & Zizzo, 2013). Because we implemented different price distributions by assigning a high probability to one specific price, subjects uncertain in their valuation may be anchoring on this price in the first rounds of the experiment. While an initial effect of the revealed price as an anchor is thinkable, it seems implausible that subjects would remain uncertain in their valuation and that they would continue to consider the constantly and obviously randomly changing revealed price as informative over the course of the experiment.

Here, we do observe a slight decline in the variability of individual bids over time, consistent with the idea that in the beginning of the experiment, subjects may be unsure how much to bid, and thus they change their mind more often early on. However, the effect of the revealed price on the bid remains equally strong throughout our experiment, suggesting that even if subjects are refining their knowledge of their own preferences, this is not reducing the effect of the price distribution on their bids. For this reason, the aforementioned explanations—each suggesting that the influence of the price distribution should be reduced by experience—seem to be insufficient to explain our results. Moreover, recent experiments on anchoring in purchasing decisions have found only very weak evidence for anchoring effects on WTP (Fudenberg et al., 2012; Sugden et al., 2013).

Assuming that subjects do know their fixed valuation for the good and respond optimally to the BDM mechanism, their bid could still be affected by the price distribution. This may be the case if bidders care about the buying probability, or the probability of certain prices. The probability distribution can, for example, define a reference point against which outcomes are compared. Outcomes that fall below expectation create disutility, possibly driven by feelings of regret, disappointment, or loss, and those that exceed expectations give positive utility. In our study, one could imagine subjects experiencing feelings of disappointment and regret when they expected to buy a product but end up not buying it, or else when they had reason to expect a very low price, but a higher price was selected as the actual price. Because most people dislike losses more than they like similarly sized gains (Kahneman & Tversky, 1979), this may create incentives to manage one's expectations in order to minimize the feeling of loss.

Interestingly, however, some of the most prominent models of reference-dependent preferences that incorporate this idea do not make predictions in line with our empirical findings. For example, if a subject had preferences such as those modeled in Kõszegi and Rabin (2006) and Heidhues and Kõszegi (2014), he or she could maximize his or her reference-dependent utility by adjusting the expectation regarding the probability of purchase to its optimal, equilibrium level. In our BDM experiment, these expectations are influenced by two factors: (a) the revealed price that is out of the subject's control (for a given bid, the subject is more likely to buy the good if the revealed price is less than or equal to her bid than if it is above her bid) and (b) the subject's own bid (for a given revealed price, the higher the bid, the more likely the subject is to buy). Notice that this implies that in order to maintain the optimal level of expectations about buying (and thus to maximize this kind of utility), the subject would need to adjust her bid in the following way: For higher (lower) revealed prices that imply a lower (higher) probability of buying, the subject should bid less (more; see Wenner, 2015, for formal proof). This prediction is completely at odds with our empirical findings.

To account for our finding that subjects are willing to bid more on the products when the revealed price is higher, an individual would have to have a utility function such that, where x is the value of the good, p is the price paid, and p_r is the revealed price that, for a given bid, affects the probability of winning and the expected price.

This observation of ours is in line with a simple model of reference dependence in purchasing decisions recently proposed by Wenner (2015). In his paper, Wenner challenges the personal equilibrium predictions of Kõszegi and Rabin (2006), who assert that individuals are influenced by their own expected behavior. Instead, according to this model, people are simply motivated to obtain the good at a price that looks like a good deal relative to the distribution of prices they expect to pay. In our setting, a high revealed price gives rise to price distribution with a high expected value; thus, subjects can bid higher without risking that the price they pay compares unfavorably to the expected price. However, this model cannot account for our observation that the influence of the revealed price is stronger when it is closer to the average bid of the subject.

In our experiment, as the revealed price exceeds a subject's bid, buying probability decreases sharply. Therefore, one way to think about our results is the following: When the revealed price is above but close to a subject's valuation, by placing a bid equal to the revealed price, the subject can significantly increase her chances of buying at a small cost. When the revealed price is far from the subject's valuation, the cost of matching the revealed price is much higher, and thus the subject prefers to stick to her "true valuation." In such a setting, it might be that subjects are willing to match higher revealed prices for goods that they wish to buy more, just as we observed (see Table 5). This behavior would also be in line with the observation that BDM bids are somewhat lower than valuations that are revealed in take-it-or-leave-it offers (Berry, Fischer, & Guiteras, 2015), where buying probability is 1 if agreeing to the offered price and 0 if refusing the offer. Although we do not have conclusive information about subjects' true valuations to test this proposition, it may well be that revealed prices that are just slightly too high do tempt subjects to accept them, even though the price paid is above the value obtained from the good.

Conclusions

We demonstrated that in repeated rounds of the BDM mechanism, subjects do not constantly bid the same amount for the same good but instead are surprisingly flexible in their bids. The variability in their bids is not random, however. Subjects tend to submit higher valuations when the price distribution assigns a high probability to a high price and bid lower when a low price is highly probable. Interestingly, distributional dependence is more frequent when the likely price is close to the average bid that a subject places on a given good. Subjects who report a strong wish to buy a particular good are more likely to show distributional dependence for a wide range of price distributions for that good. Such bidding behavior cannot be reconciled with the standard assumption that consumers have a fixed valuation for a good that is independent of the details of the BDM procedure. Distributional dependence in our experiment is unlikely to result from misunderstanding the BDM procedure because subjects repeatedly bid on the same products and received detailed instructions and feedback on the BDM mechanism. Our findings are broadly in line with the assumption that, within limits, subjects are motivated to make a good deal relative to the expected price. Additionally, they may be motivated to increase buying probability when this comes at a low cost. Taken together, these results show that distributional dependence in the BDM mechanism is a robust but complex phenomenon. Further understanding the driving factors of this phenomenon would contribute toward a better understanding of value construction.

References

Ariely, D., Loewenstein, G., & Prelec, D. (2003). "Coherent arbitrariness": Stable demand curves without stable preferences. *The Quarterly Journal of Economics, 118,* 73–106. http://dx.doi.org/10.1162/00335530360535153

Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science, 9,* 226–232. http://dx.doi.org/10.1002/bs.3830090304

Berry, J., Fischer, G., & Guiteras, R. P. (2015, July). *Eliciting and utilizing willingness to pay: Evidence from field trials in Northern Ghana* (CEPR Discussion Paper No. DP10703). Retrieved from http://ssrn.com/abstract=2630151

Bohm, P., Linden, J., & Sonnegard, J. (1997). Eliciting reservation prices: Becker–DeGroot–Marschak mechanisms vs. markets. *The Economic Journal, 107,* 1079–1089. http://dx.doi.org/10.1111/j.1468-0297.1997.tb00008.x

Cason, T., & Plott, C. (2014). Misconceptions and game form recognition: Challenges to theories of revealed preference and framing. *Journal of Political Economy, 122,* 1235–1270. http://dx.doi.org/10.1086/677254

Fudenberg, D., Levine, D. K., & Maniadis, Z. (2012). On the robustness of anchoring effects in WTP and WTA experiments. *American Economic Journal: Microeconomics, 4,* 131–145. http://dx.doi.org/10.1257/mic.4.2.131

Hao, L., & Houser, D. (2012). Belief elicitation in the presence of naïve respondents: An experimental study. *Journal of Risk and Uncertainty, 44,* 161–180. http://dx.doi.org/10.1007/s11166-011-9133-1

Harrison, G. W., & Rutström, E. E. (2008). Experimental evidence on the existence of hypothetical bias in value elicitation methods. In C. Plott & V. Smith (Eds.), *Handbook of experimental economics results* (pp. 752–767). Amsterdam, the Netherlands: Elsevier. http://dx.doi.org/10.1016/S1574-0722(07)00081-9

Heidhues, P., & Kõszegi, B. (2014). Regular prices and sales. *Theoretical Economics, 9,* 217–251. http://dx.doi.org/10.3982/TE1274

Horowitz, J. K. (2006). The Becker-DeGroot-Marschak mechanism is not necessarily incentive compatible, even for non-random goods. *Economics Letters, 93,* 6–11. http://dx.doi.org/10.1016/j.econlet.2006.03.033

Kaas, K. P., & Ruprecht, H. (2006). Are the Vickrey auction and the BDM mechanism really incentive compatible? Empirical results and optimal bidding strategies in cases of uncertain willingness-to-pay. *Schmalenbach Business Review, 58,* 37–55.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–292. http://dx.doi.org/10.2307/1914185

Karni, E., & Safra, Z. (1987). "Preference reversal" and the observability of preferences by experimental methods. *Econometrica, 55,* 675–685. http://dx.doi.org/10.2307/1913606

Kõszegi, B., & Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics, 121,* 1133–1165. http://dx.doi.org/10.1093/qje/121.4.1133

Likert, R. (1932). A technique for the measurement of attitudes. *Archives de Psychologie, 22,* 1–55.

List, J. A., & Gallet, C. A. (2001). What experimental protocol influence disparities between actual and hypothetical stated values? *Environmental and Resource Economics, 20,* 241–254. http://dx.doi.org/10.1023/A:1012791822804

Lusk, J. L., Alexander, C., & Rousu, M. C. (2007). Designing experimental auctions for marketing research: The effect of values, distributions, and mechanisms on incentives for truthful bidding. *Review of Marketing Science, 5.* http://dx.doi.org/10.2202/1546-5616.1059

Mazar, N., Kõszegi, B., & Ariely, D. (2014). True context-dependent preferences? The causes of market-dependent valuations. *Journal of Behavioral Decision Making, 27,* 200–208. http://dx.doi.org/10.1002/bdm.1794

Noussair, C., Robin, S., & Ruffieux, B. (2004). Revealing consumers' willingness-to-pay: A comparison of the BDM mechanism and the Vickrey auction. *Journal of Economic Psychology, 25,* 725–741. http://dx.doi.org/10.1016/j.joep.2003.06.004

Rutström, E. E. (1998). Home-grown values and incentive compatible auction design. *International*

Journal of Game Theory, 27, 427–441. http://dx.doi.org/10.1007/s001820050082

Shogren, J. F., Cho, S., Koo, C., List, J., Park, C., Polo, P., & Wilhelmi, R. (2001). Auction mechanisms and the measurement of WTP and WTA. *Resource and Energy Economics, 23,* 97–109. http://dx.doi.org/10.1016/S0928-7655(00)00038-5

Sugden, R., Zheng, J., & Zizzo, D. J. (2013). Not all anchors are created equal. *Journal of Economic Psychology, 39,* 21–31. http://dx.doi.org/10.1016/j.joep.2013.06.008

Urbancic, M. (2011). Testing distributional dependence in the Becker-DeGroot-Marschak mechanism. Retrieved from https://www.researchgate.net/publication/268025612_Testing_Distributional_Dependence_in_the_Becker-DeGroot-Marschak_Mechanism

Wathieu, L., & Bertini, M. (2007). Price as a stimulus to think: The case for willful overpricing. *Marketing Science, 26,* 118–129. http://dx.doi.org/10.1287/mksc.1060.0222

Wenner, L. M. (2015). Expected prices as reference points—Theory and experiments. *European Economic Review, 75,* 60–79. http://dx.doi.org/10.1016/j.euroecorev.2015.01.001

Wertenbroch, K., & Skiera, B. (2002). Measuring consumers' willingness to pay at the point of purchase. *Journal of Marketing Research, 39,* 228–241. http://dx.doi.org/10.1509/jmkr.39.2.228.19086

Received October 6, 2015
Revision received April 7, 2016
Accepted May 17, 2016 ■