

Knowledge and use of price distributions by populations and individuals

Timothy F. Lew (tflew@ucsd.edu) and Edward Vul (evul@ucsd.edu)

University of California, San Diego
Department of Psychology, 9500 Gilman Dr.
La Jolla, CA 92093 USA

Abstract

How much do individuals, compared to the population, know about the distribution of values in the world? Participants reported the prices of consumer goods such as watches and belts and we compared how accurately individuals vs. the overall population knew the mean and dispersion of prices. Although individuals and the population both knew objects' average prices and relative standard deviations, the population was more sensitive to the absolute standard deviation of prices. In a second experiment, we examined whether individuals' impoverished distribution knowledge impairs their ability to interpret advertisements. Consistent with people using Bayesian inference, the higher an object's actual price dispersion, the more participants relied on advertisements; however, this effect is considerably smaller than a simple proportional offset, suggesting again that individuals underestimate dispersion. Thus, despite having a sense of the distribution of real world quantities, individuals tend to know only a fraction of the world distribution.

Keywords: Probabilistic inference; decision-making; behavioral economics; prior knowledge

Introduction

How much does a television cost? How about a television endorsed by Quentin Tarantino? Chances are you can accurately guess the prices of these and many other everyday objects. People appear to use their knowledge of how values are distributed (e.g. the mean and variance of a value) to infer the prices of objects or other real-world values like people's life spans and cake baking times (Griffiths & Tenenbaum, 2006). This behavior can also arise from people only knowing a few examples of values (Mozer, Pashler & Homaei, 2008). For example, rather than know the average time it takes to bake a cake and how much baking times vary, each person may only know one or two cake baking times. Aggregating over multiple people in a population each with a few sample values can give the appearance of individuals representing complex distributions.

However, people only reporting a few car costs or baking times may actually reflect them using knowledge of distributions. Over multiple trials a single individual's responses resemble learned distributions (Lewandowsky, et al., 2009), suggesting that people do represent the distribution of real-world values. In fact, people appearing to rely on only a few samples may actually reflect people efficiently approximating complex distributions by taking a few samples (Vul, et al., 2014). Here, we examined how well individual people vs. populations know the distribution of real world values—the prices of everyday objects—and

how they use that prior distributional knowledge to update price estimates given new information (i.e. advertisements).

Although individuals appear to represent distributions of values, different sources of errors and biases may impede their ability to estimate the true mean and variance of values as accurately as populations. First, individuals are likely to experience an imperfectly representative fraction of the world, resulting in idiosyncratic biases in their expectations. Anchoring effects may bias individuals' judgments, such that hearing a high price causes people to overestimate the prices of objects (Tversky & Kahneman, 1974). People might estimate the price of an object by comparing it to other objects they have recently encountered (Ungemach, et al., 2011; Vlaev, et al., 2011). Additionally, rather than independently sample possible prices of objects, people may be biased by their memories of previous responses (Vul & Pashler, 2008; Hourihan & Benjamin, 2010). When aggregating the judgments of a population of independent individuals instead, many of these biases will wash out allowing the population to represent the distribution more accurately.

In turn, the limitations of people's distributional knowledge may determine how they infer the values of new objects and how they integrate new sources of information. For instance, how should a person infer the price of an object when it is presented with an advertisement? People may use Bayesian inference to determine how to use ads, similar to how listeners infer the meaning of utterances in the domain of pragmatic inference (Frank & Goodman, 2012). When referring to an object, listeners rely more heavily on speakers' utterances when no objects are particularly salient (have a low prior probability of being referred to). Similarly, people should rely on ads more heavily when they are more uncertain about the price of an object (i.e. when the price of a type of object varies greatly between instances). Consequently, if people do not know the dispersion of objects' prices, they will be unable to determine how much to rely on advertisements.

In the current study, we measured individuals' and populations' knowledge of the prices of everyday objects. In Experiment 1, we selected a set of object categories and asked participants to make multiple guesses about the price range of an object from each category. To measure participants' accuracy, we used Google Product search to obtain the true prices of objects from each category. In Experiment 2, participants performed the same price estimation task but we presented the object categories with positive or negative advertisement modifiers.

Participants accurately estimated the average prices of the objects. Aggregating all of our participants' responses together revealed that the population was very well calibrated to the dispersion of prices. However, while individual participants knew which objects' prices had greater or lower dispersion, they severely underestimated how much prices' dispersions varied overall. In Experiment 2, advertisements influenced participants' price estimates. Participants guessed objects were cheaper when presented with negative ads and more expensive when presented with positive ads and weighted ads in proportion to the dispersion of objects' prices. But, consistent with participants underestimating the dispersion of prices, the interaction between advertisements and price dispersion was relatively weak. Although people possess some distributional knowledge, that knowledge is highly limited compared to the knowledge of the population and falls far short of the true distribution of values in the world.

Experiment 1

We evaluated how much individuals vs. the population know about real-world price distributions by asking them to guess the prices of objects multiple times.

Methods




Figure 1. Example trial. Participants clicked to guess a price using the horizontal slider and then set a confidence range using the vertical slider. As participants adjusted the confidence slider, their current bet was displayed. Here, the participant has guessed the ring costs \$41.05 with a range from \$10.12 to \$166.00. If the object's price falls in this range, the participant will earn 4.02 points. The larger the range, the fewer points. Participants could alternate between setting the guess and range. Exact bets and ranges were displayed as participants made their choice.

Participants. People from the Cambridge, MA community participated as part of a paid, daylong behavioral economics battery (the other study included in this battery is irrelevant for the present purposes—it investigated delayed

discounting in behavior). Participants completed a varying number of blocks, depending on how quickly they completed each block. There were initially 29 participants; 23 returned for Session 2, 14 for Session 3 and 10 for Sessions 4 and 5. By spacing out responses across sessions, we hoped to decrease dependence between responses and ascertain participants' full distributional knowledge.

Stimuli. We selected 50 object categories (e.g. “Diamond solitaire ring”, “Printer”), attempting to cover a wide span of empirical log price means and variances while minimizing the correlation between the empirical means and variances. All of our subsequent analyses use the log prices. We manually obtained object prices from Google Products searches for the categories. Despite our efforts, variance in log prices remained weakly, but significantly, correlated with mean log price ($r=.28, p=.048$). Indeed, participants appeared to expect a strong correlation, such that the means and standard deviations of their responses were highly correlated ($r=.55, p<.001$).

Procedure. Each trial, participants saw the name of an object category and a generic picture of that category (Figure 1). We told participants that each trial we had selected the price of a specific object from that category using an Internet search; this discouraged subjects from making repeated guesses or focusing on values like the mean. Participants guessed the price of the object by clicking on a log scale and then selecting a confidence interval around their guessed price. Participants earned points if the correct price was in the confidence interval. As the size of the confidence interval increased, participants' potential reward decreased and vice versa. Thus, participants were incentivized to set a sufficiently narrow confidence interval so that getting the correct answer was rewarded, but a sufficiently broad confidence interval such that they were likely to get a correct answer. To keep participants motivated, every few trials we told participants their current score.

Each block, participants guessed the price of objects from a random subset containing 35 of the categories, block randomized.

Results

Did participants know the mean price of objects? We compared the empirical mean prices calculated from the Google Product search results to subjects' estimated mean prices (Figure 2). Reported mean price was strongly correlated with the actual mean price ($r=.85, p<.001$), demonstrating that subjects knew which products were more or less expensive. To evaluate how well participants knew the exact prices of goods, we found the best fitting linear regression. A slope of 1 would indicate that on average participants knew the exact prices of objects. The regression line was close to one ($m=.81, 95\% \text{ confidence interval}=.67-.96$), suggesting that together subjects knew the mean prices of objects fairly accurately.




Figure 2. The mean reported price of objects as a function of their true price. The grey line denotes equality. Participants possessed very accurate knowledge of the average prices of objects.

Did the population know the dispersion of prices? We compared the true dispersion of prices to the dispersion of prices across all participants' responses (Figure 3A). For each object, we calculated the standard deviation of all the guesses that all the participants made. The true and estimated variances were strongly correlated ($r=.68$, $p<.001$), demonstrating that the group knew which prices were more or less variable. The slope of the linear regression was $.48$ (95% confidence interval $=.33-.63$), suggesting that the group accounted for roughly half of the actual variability in prices.

Did individual participants know the dispersion of prices? We next compared the true dispersion of prices to the dispersion of individual participants' responses (Figure 3B). For each participant, we calculated the standard deviation of their responses for each object category. We then found the average standard deviation of responses.

Within-participant dispersion was significantly correlated with the true dispersion ($r=.51$, $p<.001$), demonstrating that participants knew which objects had more or less variable prices. However, the slope of the linear regression was much smaller ($.12$, 95% confidence interval $=.064-.19$) than the across-participant slope, indicating that despite knowing the relative dispersion of prices, individual participants were not as well calibrated to the absolute dispersion of prices as the overall population.

Did individual participants know the range of prices? Confidence interval width was strongly correlated with the true price variance ($r=.63$, $p<.001$) (Figure 3C) indicating that participants possessed explicit knowledge about the dispersion of prices. However, participants' reported confidence intervals and the dispersion of their responses explained comparable amounts of the true dispersion of prices. The 95% confidence interval of the confidence interval regression slope ($.11$, 95% confidence interval $=.068-.15$) overlapped with the within-participant dispersion regression slope, indicating that they captured similar levels of objects' price dispersions. This suggests that participants used the same impoverished knowledge of price distributions when generating confidence intervals and repeated guesses.

Experiment 2

Experiment 1 demonstrated that individuals represent the mean and dispersion of object prices, but are not as well calibrated on the dispersion of prices as the population. In Experiment 2, we examined whether participants' limited knowledge of prices' dispersions impaired their ability to use advertisements to infer the prices of new objects.




Figure 3. Across and within participant dispersion estimates as functions of the true price dispersions. A) The standard deviation of responses, aggregated across participants. B) The average standard deviation of each participant's responses. C) The average range of participants' confidence ranges. Grey lines indicate equality. Red lines with grey shading indicate linear regression fits with confidence intervals. Participants were generally sensitive to the dispersion of prices, but the population was much better calibrated than individuals.




Figure 5. The influence of advertisement modifiers on price estimates. The percentile ranks of objects when accompanied with each type of modifier. The centers of the red boxplots denote the average percentile ranks, the tops and bottoms of the boxes denote ± 1 SEM and the error bars indicate the 25th and 75th percentiles. The grey violin plots indicate the frequency of different percentile ranks for a given advertisement. The effectiveness of ads greatly varied: Very negative modifiers (left side) like “low quality” caused participants to think objects were cheaper, neutral modifiers (center) like “utilitarian” had minimal effect and very positive modifiers (right side) like “used by Michael Jordan” caused participants to report more expensive prices.

Methods

Participants. The participant pool was identical to Experiment 1. There were initially 25 participants; 24 returned for Session 2, 17 for Session 3 and 4 for Session 4 and 3 for more than 5 sessions (8, 14, and 24 sessions).

Stimuli. The stimuli were identical to Experiment 1. However, each stimulus was now accompanied by one of 55 ads. 28 of the ads were positive (e.g. “award winning”, “five-star”) and 27 were negative (e.g. “bargain”, “overstock”). Each trial, we presented a random modifier alongside the object category image, with the constraint that objects in the same block could not have the same modifier.

Procedure. The procedure was identical to Experiment 1.

Results

Did the valence of ads affect participants’ judgments? We first examined whether negative and positive advertisements caused participants to judge objects as cheaper and more expensive, respectively. For each object

category, we rank ordered participants’ responses and then converted the ranks to percentiles, such that the lowest and highest estimates were the 0th and 100th percentiles, respectively. We then found the average percentile for objects that had been accompanied by each modifier. Transforming responses into percentiles allowed us to compare people’s price estimates despite the wide variance in objects’ prices.

The efficacy of ads varied greatly (Figure 5). The most negative ad, “low quality”, resulted in a mean rank of the 28th percentile (SEM=2.5) whereas the most positive ad, “used by Michael Jordan” resulted in a mean rank of 70th (SEM=3.0). These patterns demonstrate that participants used advertisements to infer the price of objects.

Did uncertainty about the price of an object increase the efficacy of advertisements? Participants may have relied on not only the advertisements but also their own knowledge about the distribution of objects’ prices. In particular, we expected that participants would rely more heavily on advertisements when objects’ prices varied more, as a way to compensate for their uncertainty. Given that individuals underestimated the dispersion of prices, however, we were uncertain whether participants would be

able to use objects' true price dispersions to determine how much to rely on advertisements.

To test how people inferred the prices of objects, we used the *lme4* package for R to design a suite of mixed effects models that utilized different types price and advertisement information. First, participants may have solely relied on their knowledge of objects' average prices. We wrote an average price model that treats objects' true mean prices as a fixed effect and participants and objects as random effects. Second, to test whether participants used advertisements to infer the prices of objects, we designed an advertisement model that builds upon the average price model by treating advertisements as a random effect. Finally, we evaluated whether participants relied on ads more when the dispersion of an objects' prices was high. We designed an advertisement-dispersion model that also accounts for the standard deviation of prices by adding an interaction between advertisements and the true price standard deviation. Thus our full model is defined by:

$$\log_{10} estimate_{p,i,m} = \beta_0 + \beta_1 Mean[\log_{10} price_i] + \gamma_p + \delta_i + \zeta_m + \theta_m SD[\log_{10} price_i]$$

$\log_{10} estimate_{p,i,m}$ is participant p 's price estimate for object i when it is accompanied by advertisement modifier m . The second line corresponds to the nested average price model. β_0 is the intercept and β_1 is the coefficient for $Mean[\log_{10} price_i]$, the true mean price of the objects. γ_p is the offset for each participant and δ_i is the offset for each object. The third line corresponds to the additional advertisement modifier coefficients for the advertisement model. ζ_m is the offset for each ad. The fourth line corresponds to the interaction between true price variance and the advertisements in the advertisement-dispersion model. θ_m indicates how much each modifier affects estimates as the standard deviation of prices, $SD[\log_{10} price_i]$, changes.

Advertisements influenced participants' responses, such that the advertisement model fit ($AIC=3757$) participants' behavior much better than the average price model ($AIC=4336$, $\chi^2(1)=580$, $p<.001$). The advertisement-dispersion model had a lower but comparable AIC ($AIC=3753$) and provided a slightly better fit than the advertisement model ($\chi^2(2)=8.4$, $p=.015$). This pattern is consistent with participants using Bayesian inference to integrate prior knowledge about the prices of objects with new advertisements to determine value of novel objects. Nevertheless, the small improvement from the advertisement-dispersion interaction suggests participants' still failed to fully account for the variability of prices.

To more directly evaluate the effect of price dispersion on the use of ads, we used the full model to test whether the influence of ads increased with the dispersion of prices. We extracted the object-advertisement coefficient and then separately averaged the coefficients for positive valence and

negative valence ads for each object. This yielded an average shift for positive ads and an average shift for negative ads for each object. We then calculated the absolute difference between the average positive ads and average negative ads. Larger absolute differences indicate that advertisements had a greater influence on the object's price (that is, the price was more malleable to the influence of advertisements).

There was a significant positive correlation between objects' price dispersions and the difference between positive and negative price estimates ($r=.44$, $p=.0014$) (Figure 6), indicating that participants used price distribution knowledge and advertisements to infer the prices of novel objects. However, consistent with the small improvement from the advertisement model to the full model, the magnitude of the interaction between price dispersions and ads was small. We compared the magnitude of the ad random effects to the ad-dispersion random effects and found that the ad-dispersion interaction accounted for only 3% of the overall variance attributed to advertisements. Thus, although participants relied more heavily on advertisements when objects had highly variable prices, participants' impoverished knowledge limited their ability to effectively weigh the value of ads.

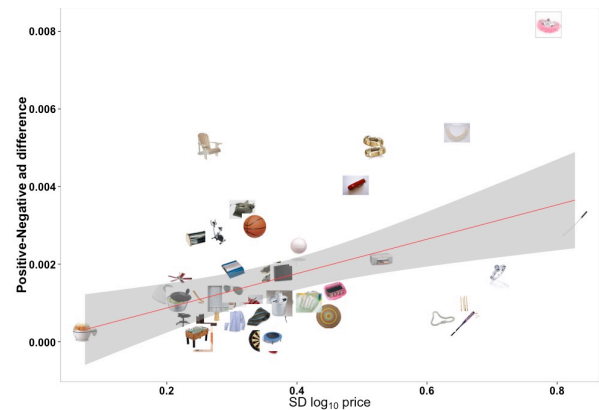


Figure 6. The influence of advertisements on price estimates given the dispersion of objects' prices. Advertisements had a larger impact on guesses when objects had more disperse prices.

Discussion

We asked people to guess the prices of objects from different categories and compared how well individuals vs. the population knew the distribution of values. Although individuals knew the mean price of objects well, they greatly underestimated the dispersion of prices compared to the population. Participants were also able to use their knowledge of price distributions to determine how much to rely on advertisements, but their behavior indicated that they still only knew a fraction of the overall dispersion of prices.

Both individuals and the overall population exhibited knowledge of the mean and the dispersion of prices in the world, though our population was much better calibrated to the absolute dispersion of prices. This deficiency may have

arisen from individuals being biased by their idiosyncratic experiences and contexts. Seeing more or less expensive objects could have resulted in participants anchoring their responses and respectively making higher or lower guesses (Tversky & Kahneman, 1974) or comparing our stimuli to objects they had recently encountered (Ungemach, et al., 2011; Vlaev, et al., 2011). In the future, accounting for people's histories, such as their socioeconomic backgrounds, may reveal differences in their price inferences.

Our experiments may have also been limited in their ability to reveal people's full knowledge of distributions. Although we tried to increase the independence of responses by spacing out sessions over hours, increasing delay intervals to weeks or months could reveal more extensive distribution knowledge. Vul & Pashler (2008), for example, found that responses made even after 3 weeks were relatively similar. Participants' responses may have also been influenced by decision biases. Participants, for instance, might have been risk averse and reported the mean to minimize their maximum loss or risk seeking and set small confidence intervals. These strategies could have limited participants' apparent distributional knowledge.

Additionally, people most likely possess more complex categorical knowledge about real-world values than just individual object categories (Hemmer & Steyvers, 2009; Hemmer & Persaud, 2014). The prices of cell phones, televisions and computers, for instance, may all fall under the category of electronics. People's judgments about the price of a new cell phone then might be constrained by both their knowledge of how much cell phones cost and also how much electronics in general cost. Furthermore, individual differences in the categorization of objects (is a smartwatch an electronic device or a fashion accessory?) may lead to distinct biases in how people estimate prices.

People used their prior knowledge of price distributions with advertisements to infer the prices of new objects, relying more heavily on advertisements when objects had highly variable prices. This behavior may reflect processes comparable to linguistic pragmatic inference (Frank & Goodman, 2012). In the case of pragmatic inference a speaker's utterance can help a listener select an object out of a crowd. Similarly, if a person buying a car is uncertain about the value of the car, an advertisement can help the buyer infer whether it's a steal or a lemon. Moving forward, using the framework of pragmatic inference to examine the role of advertisements in decision-making may help us learn why the effectiveness of our modifiers varied so much and how to craft more influential advertisements. More broadly, this approach may give insights into questions like how ads are interpreted in different contexts (Barner & Snedeker, 2008) and when people decide that ads are informative (Frank & Goodman, 2014).

People have some knowledge about the distribution of values in the world. Despite knowing the mean and the relative dispersion of prices, they have a poor idea of how much prices actually vary compared to the population.

Furthermore, their impoverished distribution knowledge impairs their ability to appropriately weigh new information like advertisements to infer the prices of novel objects. Our future work will examine the sources of the idiosyncratic limitations on people's distribution knowledge.

References

- Barner, D., & Snedeker, J. (2008). Compositionality and statistics in adjective acquisition: 4-year-olds interpret tall and short based on the size distributions of novel noun referents. *Child Development, 79*(3), 594–608.
- Frank, M.C., & Goodman, N.D. (2012). Predicting pragmatic reasoning in language games. *Science, 336*(6084), 998-998.
- Frank, M.C., & Goodman, N.D. (2014). Inferring word meanings by assuming that speakers are informative. *Cognitive Psychology, 75*, 80-96.
- Griffiths, T.L., & Tenenbaum, J.B. (2006). Optimal predictions in everyday cognition. *Psychological Science, 17*(9), 767-773.
- Hemmer, P., & Persaud, K. (2014). Interaction between categorical knowledge and episodic memory across domains. *Frontiers in psychology, 5*(584), 1-5.
- Hemmer, P., & Steyvers, M. (2009). A Bayesian account of reconstructive memory. *Topics in Cognitive Science, 1*(1), 189-202.
- Hourihaan, K.L., & Benjamin, A.S. (2010). Smaller is better (when sampling from the crowd within): Low memory-span individuals benefit more from multiple opportunities for estimation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 36*(4), 1068-1074.
- Lewandowsky, S., Griffiths, T.L., & Kalish, M.L. (2009). The wisdom of individuals: Exploring people's knowledge about everyday events using iterated learning. *Cognitive Science, 33*(6), 969-998.
- Mozer, M.C., Pashler, H., & Homaei, H. (2008). Optimal predictions in everyday cognition: The wisdom of individuals or crowds? *Cognitive Science, 32*(7), 1133-1147.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124-1131.
- Ungemach, C., Stewart, N., & Reimers, S. (2011). How incidental values from the environment affect decisions about money, risk, and delay. *Psychological Science, 22*(2), 253-260.
- Vlaev, I., Chater, N., Stewart, N., & Brown, G. D. (2011). Does the brain calculate value? *Trends in Cognitive Sciences, 15*(11), 546-554.
- Vul, E., Goodman, N., Griffiths, T.L., & Tenenbaum, J.B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science, 38*(4), 599-637.
- Vul, E. & Pashler, H. (2008). Measuring the crowd within: probabilistic representations within individuals. *Psychological Science, 19*(7), 645-647.