NEUROECONOMIC MODELS OF ECONOMIC DECISION-MAKING[†]

The Impact of Computation Time and Experience on Decision Values

By K. CARRIE ARMEL AND ANTONIO RANGEL*

Most economists and neuroeconomists believe that individuals make choices first by assigning values to objects and then by selecting the option with the highest value, perhaps with some noise (Rangel, Colin Camerer, and Read Montague 2007). This raises a question with important implications for economics: How does the brain compute the values that guide decisions (henceforth called decision values) and what are the properties of those processes? Important and more concrete examples of these questions include the following: Does the brain always assign values to objects that are commensurate with the benefits that they generate, or does it make mistakes sometimes? Does the amount of time spent computing the value matter? Are there incidental variables, such as the way an object is displayed at the time of sale, that affect the value that is assigned to it?

The properties of the brain's value computation processes should be of interest to economists since those properties determine the extent to which individuals are able to make quality choices. In addition, the properties of these processes could have important implications for the behavioral and welfare effects of practices such as in-store marketing.

In this paper we explore these questions theoretically and experimentally. We propose

a simple model of how computation time and experience can affect the value that is assigned to items. The model makes several stark predictions that we test using behavioral experiments.

The impact of computation time and experience on valuation is a new question in economics, but not in psychology or behavioral neuroscience. In particular, several models have been proposed of how the brain makes binary choices by aggregating information over time (see Rafael Bogacz (2007) and Rangel (forthcoming) for reviews). Interestingly, some of these models have received considerable empirical support in both human behavioral and monkey electrophysiology experiments. This paper makes several contributions to that literature. First, we develop a model of valuation for a single item. In contrast, all the existing models apply only to binary choice. Second, the model we propose is based on ideas from optimal Bayesian inference. In contrast, previous models often assume computational processes that are not linked to optimal inference. Third, although there is a large experimental literature on reaction times and *choice*, the role of computation time on *valuation* has not been studied. Fourth, our work is also related to the psychology literature on the Mere Exposure Effect (Robert F. Bornstein 1989; Robert B. Zajonc 1968, 2001) which studies the impact of passive exposure to a typically neutral and unfamiliar item, such as a Chinese ideogram, on subsequently reported liking-ratings. In contrast, we look at how experience computing the value of an item affects subsequent valuations. Finally, our work is more generally related to the literature on the construction of preference in behavioral economics (see Sarah Lichtenstein and Paul Slovic (2006) for a compendium of articles).

[†] *Discussants:* David Laibson, Harvard University; Drew Fudenberg, Harvard University; Alberto Bisin, Harvard University.

* Armel: Stanford Prevention Research Center, Stanford University School of Medicine, Hoover Pavilion, N229, 211 Quarry Road, Stanford, CA 94305 (e-mail: kcarmel@stanford.edu); Rangel: HSS and Computational and Neural Systems, California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91103 (e-mail: rangel@hss.caltech.edu).

I. A Model of the Computation of Decision Values

Consider the problem of an organism that needs to evaluate an item in order to make a choice. Items differ on the consumption value that they generate. Consumption values are distributed normally with mean 0 and variance τ^2 . The organism needs to forecast the value of items *before consuming them*. We call these forecasts decision values (DVs).

A central feature of the model is that the computed DVs are only noisy estimates of the actual value of consumption. In particular, we assume that when evaluating an item the brain can obtain only one noisy estimate of the true consumption value per unit of time. Let θ_t denote the estimate at time t . We assume that (a) the estimates are distributed normally with variance σ^2 and mean μ equal to the consumption value of the item being evaluated; and (b) the estimates are independently distributed. Let $d(\theta)$ denote the DV that is computed after t measurements $\theta = (\theta_1, \dots, \theta_t)$, and $m(\theta)$ denote the mean measurement.

The key assumption of the model is that the processes computing DVs are consistent with Bayesian updating. In particular, we assume that $d(\theta)$ is equal to the mean of the posterior distribution of Bayesian beliefs about the item's true consumption value. The posterior distribution is normal and its mean and variance are given, respectively, by

$$\frac{t\tau^2 m(\theta)}{\sigma^2 + t\tau^2} \quad \text{and} \quad \frac{\sigma^2 \tau^2}{\sigma^2 + t\tau^2}.$$

Note that $d(\theta)$ is a random variable since it depends on the sample of measurements, and that it incorporates the priors. Let $\phi(t|\mu)$ denote the expected value of $d(\theta)$ after t units of computation time for an item with consumption value μ . This function, which describes the average impact of computation time on DVs, is given by

$$\phi(t|\mu) = \frac{t\tau^2 \mu}{\sigma^2 + t\tau^2},$$

Figure 1 describes the dynamics of the DV signal for items with consumption values $\mu^{++} > \mu^+ > 0 > \mu^- > \mu^{--}$ and provides intuition for the predictions listed below. The solid lines depict the path of the expected DV signal

FIGURE 1. DYNAMICS OF THE DV SIGNAL FOR ITEMS WITH DIFFERENT CONSUMPTION VALUES

$\phi(t|\mu)$. The dashed lines depict potential realizations of $d(\theta)$. Note a few things about these paths. First, at $t = 0$ the DV is equal to the mean consumption value in the population. Second, the expected DV function $\phi(t|\mu)$ converges asymptotically to the consumption value of the item. Third, $d(\theta)$ fluctuates around the average DV, but the amount of noise goes to zero over time. Finally, the impact of computation time is increasing on the absolute value of μ , and is zero for items with $\mu = 0$.

So far we have assumed that the speed at which the organism can take measurements is fixed and independent of experience. However, cognitive psychologists have identified many instances in which cognition, perception, and memory improve with experience. We model this by assuming that previous experience might increase the speed at which measurements can be taken. More concretely, let $\gamma(t|e)$ denote the number of measurements taken in the first t units of time given the level of experience e . We say that the process exhibits positive experience effects if $\gamma(t|e) > \gamma(t|e')$ whenever $e > e'$.

The following predictions follow directly from the previous discussion and will be tested in the experiments:

- *Attention Duration Effect (ADE)*. The impact of computation time depends on the valence of the item: the effect is positive for positive items (i.e., items with positive true consumption values), negative for negative items, and constant for neutral items. Furthermore, the effect of computation time is stronger for more liked (or disliked) items.
- *Previous Experience Effect (PEE)*. If the DV computation process exhibits positive experi-

ence effects, then the computed DV increases with experience for positive items, decreases for negative items, and has no effect for neutral items. Furthermore, the effect of exposure is stronger for more liked (or disliked) items.

- *Amplification Effect (AE).* The difference between the DV of items with different underlying consumption values increases with computation time.

Note that the basic mechanism at work in the model is simple. The goal of the DV processes is to compute quality forecasts of the actual consumption value of items. The problem is that (a) it takes time to make the required measurements, and (b) the measurement processes are noisy. If the processes are consistent with Bayesian updating, computation time matters because, by increasing the amount of measurements that are taken, it decreases reliance on priors and allows the process to average out the noise. Experience matters because it speeds up the required computations.

It is important to emphasize that the model applies equally well to unconscious and conscious processes, as long as the computations that they implement are consistent with the model. In particular, the model does not assume that the organism is consciously and intentionally computing Bayesian updates. These computations could be performed by low-level unconscious processes. In fact, evidence for embedded Bayesian updating at low levels of processing has been found in the sensorimotor domain (Konrad P. Kording and Daniel M. Wolpert 2006).

Next, we describe the results of three experiments designed to test the three predictions listed above.

II. Experiment 1

In this experiment ($N = 77$) undergraduates purchased appetitive foods by placing bids on a Becker-DeGroot Mechanism. (See the Web Appendix, www.aeaweb.org/articles.php?doi=10.1257/aer.98.2.163, for a detailed description of the experiments and for figures describing the results.) All of the items were junk foods (e.g., candy bars) and were presented using high-resolution photographs in which both the food and packages were visible. There

were three parts to the experiment. First, subjects received an endowment of \$10 that they could use to purchase food by placing bids on 70 different food items. Second, they provided a liking-rating for each food. Finally, one of the bidding trials was selected by drawing a ball from an urn. That was the only trial that counted. At the end of the experiment, subjects stayed in the lab an additional 30 minutes; during this time they were allowed to eat as much as they wanted of the food they purchased from us, but no other food or drinks were allowed. Note that since only one of the trials counted, subjects treated each decision as if it were the only one, and thus they could bid up to \$10 every time.

A key feature of Part 1 of the experiment is that items were presented for a predetermined amount of time before subjects were allowed to enter their bids by clicking a mouse. Subjects bid twice on each food item. The bidding task was divided into two parts. In the first half of the trials, each food item was presented in random order for 500 ms. After subjects had placed bids on all foods at that duration, the items were presented in random order a second time for either 500, 2,000, 3,500, 5,000, or 6,500 ms. During the liking-rating trials, items were shown in random order for 3,000ms. Subjects then had unlimited time to type their liking-rating according to a scale of -50 ("worst food ever tasted") to 50 ("best food ever tasted"), where 0 denotes a neutral item.

Several aspects of the design are worth emphasizing. First, every item was presented twice: the first time at 500ms, and the second time at either 500, 2,000, 3,500, 5,000, or 6,500 ms. Presenting items twice allowed us to test for the presence of PEEs. Presenting the items at different lengths on the second presentation allowed us to test for the presence of ADEs. Second, we controlled computation time by manipulating the amount of time that items were displayed on the computer screen. We asked the subjects to look at the item the entire time that it was displayed, but we did not collect a measure of compliance. Third, we used the liking-rating measure collected at the end of the experiment as our measure of the true consumption value of the items. This is justified by the model: when the computation time is kept constant for all items, on average the resulting DVs provide an ordinaly correct ranking of the items.

The results described below are based on the estimation of the following mixed effects linear model:

$$\begin{aligned} \text{bid-diff}_{i,t} = & \text{constant} \\ & + \alpha(\text{exposure-time}_{i,t} - 0.5) \\ & + \sum_{k=0.5, 2, 3.5, 5, 6.5} \beta_k(I_{i,t}^k \cdot \text{liking-rating}_{i,t}), \end{aligned}$$

where i indexes subjects, t indexes items, exposure-time denotes the presentation time of the picture during the second part, I^k is an indicator function for the exposure time k , liking-rating is the rating for the item provided by the subject at the end of the experiment, and $\text{bid-diff}_{i,t}$ is the change in bid between the first and second parts of the experiment for item t and subject i . Note that this model allows the change in bid to depend both on the exposure time and on the liking-rating of the items. This last feature is important because, as we will see, the PEE and the ADE change significantly with the underlying value of the item. The estimates of this model provide tests for the predictions listed above. The PEE for an item with a liking-rating of l is given by $\text{constant} + \beta_{0.5}l$. The ADE generated by increasing exposure from t to t' seconds for an item with liking-rating l during the second presentation is given by $\alpha(t' - t) + l(\beta_{t'} - \beta_t)$. All the results described below are based on these three estimates.

The results of the experiment were as follows. First, the estimated PEE was positive and increased with the liking-rating; it was not significant for neutral items ($M = 7.3$ cents, $p = 0.28$) and low-valued items, but it was 17.72 cents ($p = 0.115$) for the most liked items. Figure S2(a) in the Web Appendix plots the estimated PEE and 95 percent confidence intervals for items with different liking-ratings. Second, at exposures of 2,000 and 3,500 ms, there was a positive and significant ADE for most items except for neutral ones. For example, the effect at 3,500 ms was 43 cents ($p < 0.001$) for the most liked items, which is close to the retail price for the median item used in the experiment. Figure S2(b) in the Web Appendix plots the estimated ADE for different liking-ratings. Third, since the ADE increased with the liking-ratings, there was an AE. In fact, the difference in computed DVs between items with the maximum and minimal liking-ratings, a variable that

we refer to as the spread, changed between the first exposure (at 500 ms) and a second exposure at 3,500 ms by 49.2 cents ($p < 0.000$).

Two additional results are shown in the Web Appendix. Figure S2(c) plots the net change in bids that resulted between the first presentation at 500 ms, and the second presentation at 3,500 ms. The net effect on bids was positive and sizable at all levels and increased with the liking-rating. For neutral items, the net effect was 11.5 cents ($p < 0.05$). For the most liked items, it was 60.7 cents ($p < 0.000$). Figure S2(d) plots the marginal attention duration effect that results from increasing exposure by 1.5s for an item at the top of the liking-rating scale. The first three seconds of exposure have a substantial impact on the bids, whereas the last three seconds have negligible effects. This is consistent with the model, which predicts an asymptotic effect of computation time on DVs.

III. Experiment 2

This experiment uses a variation of the previous design to test the predictions of the model for negative items. Subjects ($N = 59$) placed bids to avoid having to eat food items that they disliked. The key difference with Experiment 1 is that items were now aversive instead of appetitive. The stimuli included items such as Spam and spinach baby food. The bids served as our measure of the negative DVs. The procedures were similar to those for Experiment 1. Subjects were told at the beginning of the experiment that they would have to eat five spoonfuls of the item shown in a randomly selected trial. They were also endowed with \$10 that they could use to bid for *avoiding* having to eat the item.

The data were analyzed using the same method as in Experiment 1. First, the estimated PEE was negative and the more disliked the item, the higher the subject's willingness to pay to avoid eating it. As shown in Figure S3(a), the PEE was marginally significant for neutral items ($m = 29.8$ cents, $p = 0.08$), but it was 61.9 cents ($p < 0.01$) for the most disliked items. Second, for the 5,500 ms exposure, the ADE was negative and insignificant for highly disliked items, but positive and significant for mildly disliked items (i.e., mildly disliked items became less negative with additional exposure). For neutral items the ADE was -44 cents ($p < 0.05$). For the most disliked items the ADE was

27.2 cents ($p = 0.284$) (see Figure S3(b)). Third, consistent with the AE, the change in spread between the first exposure (at 500 ms) and a second exposure (at 5,500 ms) was 103.7 cents ($p < 0.000$). Similar results were found at the 2,500 ms exposures. Two additional results are shown in the Web Appendix. Figure S3(c) plots the net change in bids that results between the first presentation at 1,000 ms and the second presentation at 5,500 ms. The net effect is negative and significant for sufficiently disliked items, and it increases with the negativity of the item. For neutral items the net effect is -14.5 cents ($p = 0.271$), whereas for the most disliked items the net effect is 89.12 cents ($p < 0.000$). Figure S3(d) plots the marginal attention duration effect that results from increasing exposure by 1,500 ms for an item at the bottom of the liking-rating scale. The marginal effects fluctuate, and the maximum effect is found by increasing exposure from 1,000 ms to 5,500 ms.

IV. Experiment 3

This experiment extended the previous ones in two directions. First, we used posters instead of food to explore the extent to which the results also apply to non-primary, non-immediate consumption goods. Second, ($N = 56$) subjects placed liking-ratings instead of bids, which allowed us to study the effect of computation time on DVs for positive and negative items within the same experiment.

The data were analyzed as before. The only difference is that now we looked at the change in liking-ratings, instead of the change in bids. The results were as follows. First, the PEE was positive for positive items, negative for negative items, and approximately zero for neutral items ($m = -1.59$ rating units, $p = 0.12$). It was -11.2 rating units ($p < 0.000$) for the most disliked posters and 8 rating units ($p < 0.000$) for the most liked ones (Figure S4(a)). Second, the ADE was negative for negative posters, positive for positive posters, and approximately zero for neutral ones ($m = 0.28$ rating units, $p = 0.747$). For example, for an exposure of 5s, the estimated ADE was -5.1 rating units ($p < 0.05$) for items at the bottom of the scale, and 5.65 rating units ($p < 0.05$) for items at the top (Figure S4(b)). Third, there was an AE, since the estimated change in spread between the first exposure (at 0.5s) and a second exposure at 5s was 29.9

rating units ($p < 0.000$). Two additional results are shown in Figure S4. Panel C shows the net effect of the ADE and PEE for the case of a 5s exposure. Panel D plots the marginal attention duration effect that results from increasing exposure by 1.5s for an item at the bottom of the liking-rating scale. As can be seen from the graph, the largest marginal effect occurred during the first 1.5s, which again is consistent with the predicted asymptotic effect of computation time on DVs.

V. Discussion

The results of the three experiments are mostly consistent with the qualitative predictions of the model. In particular, they suggest that the DV assigned to an item depends on the amount of time spent computing it (i.e., there is an attention duration effect) and on the amount of previous experience (i.e., there is a previous experience effect). The experiments also suggest that these effects are approximately positive for positive items, negative for negative items, and negligible for neutral items, and that the size of the effects is proportional to the magnitude of the underlying consumption value of the item. Furthermore, the results are quantitatively significant. For example, increasing computation time from 0.5 to 3.5 seconds generates an increase of 43 cents in subjects' willingness-to-pay for highly liked junk foods, an amount close to their actual retail price.

The main disagreement between the model and the data was the presence of a positive attention duration effect for mildly disliked items. One potential explanation is that the liking-rating scale is imperfect at detecting negative items, and the items that might have been rated mildly aversive in reality were appetitive.

The results developed here have an immediate implication: it should be possible to manipulate binary choices by changing the amount of time a subject spends computing the DV of each of the items. Armel, Aurelie Beaumel, and Rangel (2007) show that this is indeed the case.

REFERENCES

- Armel, Carrie, Aurelie Beaumel, and Antonio Rangel. 2007. "Biasing Simple Choices by Manipulating Relative Visual Attention." Unpublished.

- Bogacz, Rafael.** 2007. "Optimal Decision-Making Theories: Linking Neurobiology and Behavior." *TRENDS in Cognitive Science*, 11(3): 118–25.
- Bornstein, Robert F.** 1989. "Exposure and Affect: Overview and Meta-Analysis of Research, 1968–1987." *Psychological Bulletin*, 106(2): 265–89.
- Kording, Konrad P., and Daniel M. Wolpert.** 2006. "Bayesian Decision Theory in Sensorimotor Control." *TRENDS in Cognitive Sciences*, 10(7): 319–26.
- Lichtenstein, Sarah, and Paul Slovic.** 2006. *The Construction of Preference*. New York: Cambridge University Press.
- Rangel, Antonio, Colin Camerer, and Read Montague.** 2007. "Neuroeconomics: The Neurobiology of Value-Based Decision-Making." Unpublished.
- Rangel, Antonio.** Forthcoming. "The Computation and Comparison of Value in Goal-Directed Choice." In *Neuroeconomics: Decision-making and the Brain*, ed. P. Glimcher, C. Camerer, E. Fehr, and R. Poldrack. New York: Elsevier.
- Zajonc, Robert B.** 1968. "Attitudinal Effects of Mere Exposure." *Journal of Personality and Social Psychology*, 9(2): 1–27.
- Zajonc, Robert B.** 2001. "Mere Exposure a Gateway to the Subliminal." *Current Directions in Psychological Science*, 10(6): 224–28.