

December 2017

**When and how can social scientists add value to data scientists?
A choice prediction competition for human decision making¹**

Ori Plonsky^{*}, Reut Apel[†], Ido Erev[†], Eyal Ert[‡], and Moshe Tennenholtz[†]

^{*}Duke university; [†]Technion – Israel Institute of Technology; [‡]The Hebrew University

Introduction

The importance of models that allow useful predictions of human decision making has increased in recent years. For example, the development of safe and effective autonomous vehicles requires accurate predictions of the choices pedestrians and human drivers make. In theory, predictions of this type could benefit from the knowledge accumulated in behavioral decision research. In practice, however, it is often easier to predict behavior with data-driven machine learning tools that do not rely on the behavioral literature.

This paper reviews and extends a research project that aims to improve our understanding of the ways by which behavioral decision research can be used to derive useful predictions. This project is motivated by the observation that behavioral models have been often developed to capture specific interesting behavioral phenomena, and different models have been proposed to capture different phenomena. The narrow scope of many behavioral models makes it difficult to predict behavior in a new task (*1*).

¹ This white paper gives an initial overview of CPC18, a choice prediction competition for decisions under risk, under ambiguity and from experience. The competition's website: <https://cpc18.wordpress.com>

A recent study (2) tried to address this difficulty by organizing a Choice Prediction Competition (hereinafter CPC15) that challenged researchers to develop a single model² that can capture 14 classical choice anomalies documented in behavioral decision research (including, e.g., the St. Petersburg, Allais', and Ellsberg's paradoxes, (3–5)). The competing models were compared based on their ability to predict behavior in a new set of choice tasks. The results suggest a proof of concept for the feasibility of developing behavioral models that both capture all 14 phenomena and have high predictive value. Interestingly, the best models do not share the assumptions made by popular theories of choice like expected utility theory (6) and prospect theory (7, 8). Rather, they assume sensitivity mainly to the expected values and to the probability of immediate regret. Moreover, the results revealed that theory-free statistical learning submissions did not predict well, and a subsequent test (9) of various off-the-shelf theory-free machine learning algorithms revealed they also had surprisingly poor predictive power for these data.

The current study presents a new choice prediction competition aiming to evaluate three possible interpretations of the apparent advantage of behavioral models over machine learning tools in predicting choice behavior (as evident in CPC15). The first interpretation rests on the distinction between predicting behavior in a familiar setting for which much data are already available, and predicting behavior in a new environment for which no specific data are available. The clearest evidence for the value of pure machine learning tools comes from predictions of familiar agents in familiar settings (like the reaction of a known user to a specific movie). It is possible that these tools are less useful, and are outperformed by behavioral models, in predicting

² In different fields, the term *model* can mean different things. Here, we use this term to refer to a computer program that provides quantitative predictions for each problem it is faced with (of course, given some constraints on the *type* of problems it is designed to capture), without additional training or fitting of parameters.

human behavior in a new setting (like the new choice tasks examined in CPC15). Specifically, in CPC15, relative to the size of the space of choice tasks models had to predict, the size of the training data was small. The sparsity of the data led to what is known in machine learning applications as the “cold start problem” (10). A second interpretation involves the possibility that the failure of the machine learning submissions in CPC15 reflects suboptimal choice of features (covariates or predictors) these submissions used. Partial support to this hypothesis comes from a recent analysis (9) showing that it is possible to outperform the best behavioral models on the CPC15 data using a machine learning tool that is provided with theoretically-driven engineered features. The third interpretation involves the possibility that the failure of machine learning tools tested on CPC15’s data reflects suboptimal choice of these tools. It is possible that certain sophisticated tools that ignore behavioral theories could have won CPC15. Partial support to this hypothesis comes from a study (11) showing that a deep neural network architecture can outperform the best available behavioral models in a (different) task of predicting human choices in strategic settings.

Predicting choice behavior

Background

Our focus is on prediction of human choice among monetary gambles over time. The research of choice among gambles underlies both the foundations of rational economic theory (3, 6, 12) and the analyses of robust deviations from rational choice (4, 5, 7). For example, it stands at the heart of the development of seminal decision theories, like expected utility theory and prospect theory. The large impact of this line of research reflects the assumption that choice and learning among gambles reveals basic features of human behavior in a wide set of situations. Thus, having good predictions of how people choose and learn between gambles can lead to

important theoretical as well as practical developments. Moreover, some of the best successes of behavioral decision research (including in CPC15) involve studies of choice between gambles. If in this domain behavioral decision research cannot contribute much to predictions, it is likely that in more complex and less theoretically structured domains, it would contribute even less.

Experimental Task

The data used to train and test the predictive models of choice comes from the same experimental paradigm used in CPC15 (2). In this paradigm, decision makers are faced with descriptions of two monetary prospects and are asked to choose between them repeatedly for 25 trials. In the first five trials, decision makers do not receive feedback on their choice. After each trial thereafter (i.e. starting Trial 6), decision makers get full feedback concerning the outcomes generated by each option in that trial (both the obtained payoff and the forgone payoff are revealed).

The use of the CPC15 paradigm has two advantages. First, a large dataset (334,500 unique consequential decisions) based on this paradigm already exists. This dataset can then be used to train and develop the predictive models that participate in the prediction competition. Second, many teams of researchers have already tried to develop predictive models for behavior in this paradigm. This not only reduces the cost associated with the development of new predictive models from scratch for these teams, but also provides strong benchmarks against which models' performance can be compared.

Although we use the same paradigm, we expand on the space of choice problems studied in CPC15 (2), in which only one of the two options could have more than two possible outcomes. In the space investigated here, each of the two options can have up to 10 possible

outcomes³. For succinct representation of the space, we use 10 dimensions to represent the payoff distributions of the options (five dimensions for each option). In addition, and similarly to the space studied in CPC15, one of the options could be ambiguous (i.e. its description does not include the exact probabilities of its possible outcomes), and the payoffs provided by the two options can be correlated. Thus, each problem in the space is uniquely defined by 12 dimensions. Appendix A provides full description of the space of problems investigated. Appendix B presents screenshots of the experimental task.

Experimental Data

As mentioned above, models competing in the current choice prediction competition can make use of the data collected as part of CPC15 (2). These data include Problems 1-150 given in Appendix C. In addition, to supplement and extend the data (and the space of choice problems), we ran Experiment 1, and plan to run Experiment 2 at a later date. Data from Experiment 2 will be used as the *competition set* for one of the tracks of the choice prediction competition (*Track I*, see below). The apparatus and design of these two new experiments are very similar to those used in CPC15 (2) and the reader is referred there for more details. Hereinafter, we provide the essentials.

Experiment 1 includes Problems 151-210 listed in Appendix C. These 60 choice problems were randomly selected according to the problem selection algorithm given in Appendix D. Two-hundred and forty participants (139 female, $M_{\text{Age}} = 24.5$) participated in the experiment, half at the Technion and half at the Hebrew University of Jerusalem. Each participant faced one set of 30 problems: half faced Problems 151-180 and the other half faced

³ Indeed, it has been suggested recently (18) that the number of possible outcomes in a choice problem can greatly affect the observed choice behavior.

Problems 181-210. The order of the problems was randomized among participants. Each problem was faced for 25 trials consecutively. The first five trials without feedback and the rest with full feedback.⁴ Participants were paid for one randomly selected choice they made, in addition to a show-up fee. The final payoff ranged between 10 and 136 shekels, with a mean of 40 (about 11 USD).

A small portion of Experiment 1's data will be used as the *competition set* for one of the competition's tracks (*Track II*, see below) and is therefore not publicly available. The rest of the raw data, as well as CPC15's raw data (510,750 observations) is publicly available (*I3*) at: <https://zenodo.org/record/845873#.WeDg9GhSw2x>. Each observation includes a unique decision maker identifier, basic demographic information for that decision maker, the 12 parameters defining the problem faced, the trial number, the choice made, and the payoffs each of the options provided in that trial.

Experiment 2 will be very similar to Experiment 1, and will be run later. To make sure that the competition set of Track I is similar to the *calibration set*, the 60 problems used in Experiment 2 will be selected pseudo-randomly using a stratified sampling procedure from a large pool of problems selected according to the problem selection algorithm in Appendix D. Specifically, these 60 new problems will include roughly the same numbers of problems of the types "each option up to 2 outcomes", "exactly one problem with more than 2 outcomes", and "both options with more than 2 outcomes" as their numbers in Experiment 1.

⁴ Unlike in CPC15, all participants in the new experiments have (and will) face each problem in full before moving on to the next problem. Thus, in CPC15 terms, all participants are run under Condition *ByProb*.

The Choice Prediction Competition

Two Prediction Tasks

The choice prediction competition includes two parallel tracks that differ in the level of generalization the predictive models are required to achieve.

Track I: Aggregate behavior, new problems.

The first competition track, *aggregate behavior, new problems*, is similar to CPC15 (2). In this track, the goal is to provide, for each of 60 novel choice problems (i.e. unbeknownst to the modeler at the time of model development), a prediction for the progression over time of the mean aggregate choice rate of one of the options (without loss of generality, Option B). Specifically, the 25 trials of each problem are pooled to five blocks of five trials each, and the goal is to predict the mean aggregate choice rates in each of the five time-blocks. The 60 novel problems will be randomly drawn according to the problems selection algorithm in Appendix D, and will not be revealed until after the model submission deadline. Thus, during the competition phase, a competing model would be required to get as input the values of the parameters defining the novel choice problems and provide as output a sequence of five predictions (each in the range $[0,1]$) for the mean choice rates in each problem. The mean aggregate choice rate would be determined by the average choice behavior of a new sample of decision makers (i.e. for whom no data is available).

The prediction of the “average user” behavior can be very important in many real world settings. For example, consider a company contemplating between two promotion schemes (e.g., frequent small discounts vs. infrequent large discounts) and wishes to predict the popular reaction to each promotion over time. The company has no prior information regarding the popular response to these specific promotions, but may use historical data from somewhat

similar promotions to construct a predictive model. Similarly, autonomous vehicles may get as input a current motion picture of unknown pedestrians nearing the road and from their movements and analysis of other factors (e.g., are they looking at the road or on their phones?) predict whether they are about to cross the road. Encountering these particular pedestrians in this particular road under this particular setting of factors is extremely unlikely and thus it seems reasonable to make a prediction based on the “average pedestrian behavior” expected in such a setting.

Task II: Individual behavior, familiar problems

In track *individual behavior, familiar problems*, the goal is to provide, for each of 30 target individual decision makers, the progression over time (in five time blocks of five trials each) of the mean choice rate of Option B in each of five target choice problems. The training data in this track includes full sequences, of 25 choices each, made by each of the 30 target individuals in 25 different (non-target) choice problems (taken from the same space of problems). In addition, data regarding behavior of other (non-target) decision makers in the five target problems is also available. Note that the target problems differ among target individuals. Thus, the goal of the model is to provide for each target individual, in each of that individual's target problems, a sequence of five predictions in the range $[0,1]$ (although the observed response in this track is in the set $\{0,0.2,0.4,0.6,0.8,1\}$).

This task is reminiscent of online recommender systems predicting favorability ratings, but the predicted ratings are given for a complete time series (of 5 time-blocks) in advance. That is, rather than recommending a particular movie to the user, the requirement may be the recommendation of a complete movie or TV series. Such recommendation should probably take into account not only how the user is likely to respond to each episode in isolation, but also the

temporal structure involved, for example because the user is unlikely to follow the recommendation to watch the fourth episode if she dislikes the first three. Similarly, the rise of the sharing economy highlights a need to identify the reaction over time of specific workers to specific incentive structures.

Competition Metric: MSE

The winning model in each competition track would be the one that achieves best (lowest) mean squared error (*MSE*) between the observed and the predicted choice rates. Specifically, the winner in the *aggregate behavior, new problems* track would be the model with lowest *MSE* computed over 300 choice rates (5 time blocks per problem, in each of 60 test problems), and the winner in the *individual behavior, new problems* track would be the model with the lowest *MSE* computed over 750 choice rates (5 time blocks per problem, in each of 5 target problems for each of the 30 target decision makers).

The *MSE* metric has the advantage of being a proper scoring rule (14), and is also easily translated to an intuitive and interpretable measure: *Equivalent Number of Observations* (ENO) (15). ENO is an estimation for the number of decision makers that has to be run until the mean choice rate across these decision makers in a problem provides a better prediction for the behavior of the next decision maker than the prediction of the model.

Note that the rank order of the models competing in Track I is likely to be dependent on the (random) selection of the competition set problems, as well as on the (random) sample of decision makers that would take part in Experiment 2. To account for the possibility that these random selections would lead to a random winning model out of a group of equally good models, we will also compute, using a bootstrap analysis, a confidence interval for the difference in prediction *MSEs* between each of the top submissions and the winner, for different samples of

decision makers facing different samples of problems. Analyses of the statistical differences between the winners and other submissions in Track II will also be performed. We will then explicitly note all models that obtain statistically similar performance as the winner in each track.

Baseline Models

Baseline models facilitate evaluation of models in the development stage, by providing information on the range of *MSE* scores achievable. Although it is possible that a single model may perform well in both tracks, we provide different baselines for each track. The baseline models' codes are available on the competition's website.

Track I baselines

Because Track I is similar to CPC15, both of the baselines we present for this track rely on the model BEAST (Best Estimate and Simulation Tools), which served as the baseline model for CPC15(2). BEAST has proven to be a good predictor for behavior in the CPC15 data. BEAST assumes that each prospect is evaluated as the sum of three terms: the best estimate for the prospect's expected value (EV), the average of a small sample mentally drawn based on the prospect's distribution, and estimation noise. Moreover, BEAST assumes that each element in the small sample is drawn using one of four simulation tools - unbiased (a draw from the objective distributions), equal weighting (a draw from a distribution in which all the prospects' outcomes are considered equally likely), sign (a draw from a distribution in which all the prospects' payoffs with the same sign have the same valence), and contingent pessimism (a bias towards the worst possible outcome) – and that the use of the unbiased tool becomes more likely with the accumulation of feedback.

In addition to these main assumptions, BEAST also assumes that in “trivial” problems, the estimation noise is reduced, where “trivial” problems are defined as problems in which one

option stochastically dominates the other. Thus, BEAST uses an objective definition for what constitutes a “trivial” problem.

The first baseline we present for Track I relaxes this definition. BEAST.sd (BEAST subjective dominance; Cohen, Plonsky, and Erev, under review) assumes that the estimation noise is reduced when the problems are *perceived* as “trivial”. Specifically, according to BEAST.sd, a problem is likely to be perceived as trivial if both the EV rule and the equal weighting rule favor the same prospect, and the choice of that prospect does not lead to immediate regret. In addition, BEAST.sd also assumes that whereas in trivial problems the estimation noise is reduced, in complex problems (defined as a problem in which one option has at least 2 possible outcomes and the other has at least 3 possible outcomes), the estimation noise is increased. Finally, BEAST.sd assumes faster learning from feedback in ambiguous problems. Appendix E presents more details of BEAST.sd. The MSE of the original BEAST (estimated only on the first 90 choice problems in Appendix C) for Experiment 1’s 60 problems is 0.0129. BEAST.sd, (estimated based on all 210 estimation problems) achieves MSE of 0.0071 on the same set of 60 problems.

The second baseline presented for Track I is Psychological Forest (9), a model integrating the theoretical insights of BEAST for what type of features are likely to be important for prediction with the power of a random forest algorithm (16). Full details on the model are given in the original paper. Importantly, Psychological Forest uses as one of its features the predictions of the original BEAST, as well as 13 hand-crafted features constructed to capture the logic underlying BEAST. Training Psychological Forest on the 150 choice problems from CPC15 (the first 150 problems in Appendix C) yields a purely out-of-sample prediction error MSE of 0.0080 on the 60 choice problems from Experiment 1.

Note that these results are obtained when training Psychological Forest on a dataset drawn from a subspace of the space of problems that it is meant to predict. Specifically, no problems with more than two outcomes in both options are provided in the training phase, but nearly 20% of the test problems are of this type. The full baseline does not suffer from this problem. It is the same model, trained on all 210 estimation problems, which are part of the larger space. Yet, in addition, we also suggest a method to handle similar challenges in which it appears that training data from some segment of the space to predict is particularly sparse. Assuming that a reasonable behavioral model for the full space exists, it is possible to generate artificial training examples and tag their responses using the predictions of the behavioral model. Specifically, we allowed BEAST to generate predictions for a choice problems randomly sampled from the large space, and then trained Psychological Forest on both the 150 problems from CPC15 and the artificially created data. The results suggest it is possible to reduce Psychological Forest’s prediction MSE on the 60 new problems to 0.0077 on average.

Track II baselines

The first naïve baseline for Track II predicts that each individual target decision maker in each block of its individual target problems would behave the same as the average decision maker behaves in the same block of that problem. The average decision maker’s behavior is estimated as the mean aggregate behavior of all decision makers for which training data exists (there are at least 90 decision makers for each such problem). To estimate the predictive performance expected on the competition set for this naïve baseline, we simulated (10 times) the competition data out of Experiment 1’s training data, used the rest of Experiment 1’s data for training, and performed a cross validation procedure. Specifically, in each iteration, we randomly selected 30 decision makers who participated in Experiment 1, and for each, we randomly

selected five problems to be used as the validation set. The average validation MSE for the naïve baseline is 0.1038. The average correlation between the naïve model's predictions and the observed choice rates is 0.63.

Surprisingly, we did not find it easy to defeat this very naïve baseline by a large margin. Using many statistical learning techniques, and employing knowledge extracted from the psychological literature (e.g. based on BEAST), the best baseline that we could find was the use of a Factorization Machine (17), a predictor based on Support Vector Machines and factorization models, which is employed in many collaborative filtering settings (i.e. settings in which the goal is to generate predictions regarding the tastes of particular users, for whom some data exists, using information on the tastes of many other users, as in the Netflix Challenge). Each observation supplied to the current factorization machine (FM) is composed of a long binary feature vector with only two non-zero elements that correspond to the active decision maker and the active block within an active problem. The response is the observed choice rate of the active decision maker in the active block of the active problem (first transformed to imply the maximization rate of the problem, and then after making the prediction transformed back to implying the choice rate of Option B⁵). Notice that this means the FM model does not directly use the knowledge that behavior across different blocks of the same problem is likely correlated. The cross-validated MSE score of the FM model (using the same procedure as that done for the naïve baseline) is 0.0976. The average correlation between the FM model's predictions and the observed choice rates is 0.66.

⁵ There is no clear relation between the B options of the different problems. Transformation to maximization is meant to link the response a decision maker makes in one problem with her responses in another. Of course, there could also be other ways to make this link, and competitors are encouraged to investigate if these can get lead to improved outcomes.

Competition Protocol

The exact schedule for the competition appears on the competition's website. In particular, interested participants are required to register for the competition in advance, indicating for which track they wish to submit a model. Each submission may have up to three co-authors. A team can register for only one or for both tracks in parallel (a single model competing in both tracks is also acceptable). However, each person may be registered as the (co-)author of no more than two submissions per track, and be the first author of no more than one submission per track. In addition, each person may make one additional early-bird submission, sent to the organizers by the end of January 2018.

By the *Track I Code Submission Deadline*, participants competing in Track I will be required to submit to the organizers a complete, functional, documented code of their submission. We will accept submissions written in Python, Matlab, R, or SAS. To count as an acceptable submission, the code will be required to read the parameters of the problem at hand (from an external source - like the one we would eventually provide as the competition data for the first track) and output a prediction for the choice rates. One day after the Track I Code Submission Deadline, we will publish the problems that will be used as competition data of the first track. This procedure helps ensure that submissions to the first track are blind to the exact problems on which they are tested. There is no code submission requirement for Track II.

Then, by the *Predictions Submission Deadline* all participants (regardless of track) will be required to submit to us their predictions for each of the competition choice rates (i.e. 300 values in Track I and 750 values in Track II). We will then publish the full competition data. Participants are required to evaluate their submissions on the competition data (on their own), reveal the MSE results, and send these to the organizers. We will independently verify the MSE

results of the 10 lowest (best) reported MSE scores in each track. In Track I this will be done by running the code submitted to us in advance. Submissions that include random components (i.e. include calls to a random number generator) will be run up to 20 times and the MSE results from these runs will be averaged. In case the MSE participants report differs from the MSE that we will find, we will contact the authors and try to jointly reveal the source of the discrepancy. In any case, only the code that is submitted to us before the code deadline day will be considered.

After results of all submissions are obtained, we will let participants know their relative ranking and announce the winner. Additionally, we will perform tests for the statistical significance of difference between the winner and other submissions and will announce those with performance that is statistically equivalent to the winner of each track.

References

1. I. Erev *et al.*, A choice prediction competition: Choices from experience and from description. *J. Behav. Decis. Mak.* **23**, 15–47 (2010).
2. I. Erev, E. Ert, O. Plonsky, D. Cohen, O. Cohen, From Anomalies to Forecasts: Toward a Descriptive Model of Decisions Under Risk, Under Ambiguity, and From Experience. *Psychol. Rev.* (2017), doi:10.1037/rev0000062.
3. D. Bernoulli, Exposition of a New Theory on the Measurement of Risk (original 1738). *Econometrica.* **22**, 22–36 (1954).
4. M. Allais, Le comportement de l’homme rationnel devant le risque: critique des postulats et axiomes de l’école américaine. *Econom. J. Econom. Soc.* **21**, 503–546 (1953).
5. D. Ellsberg, Risk, ambiguity, and the Savage axioms. *Q. J. Econ.* **75**, 643–669 (1961).
6. J. Von Neumann, O. Morgenstern, *Theory of Games and Economic Behavior* (Princeton university press, Princeton, NJ, 1944).
7. D. Kahneman, A. Tversky, Prospect theory: An analysis of decision under risk. *Econom. J. Econom. Soc.* **47**, 263–292 (1979).
8. A. Tversky, D. Kahneman, Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertain.* **5**, 297–323 (1992).
9. O. Plonsky, I. Erev, T. Hazan, M. Tennenholtz, in *The thirty-first AAAI conference on Artificial Intelligence* (2017).
10. A. I. Schein, A. Popescul, L. H. Ungar, D. M. Pennock, in *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval* (ACM, 2002), pp. 253–260.
11. J. S. Hartford, J. R. Wright, K. Leyton-Brown, in *Advances in Neural Information Processing Systems* (2016), pp. 2424–2432.
12. L. J. Savage, *The foundations of statistics* (John Wiley & Sons, Oxford, England, 1954).
13. O. Plonsky, I. Erev, E. Ert, Calibration Data For Choice Prediction Competition 2018

- (Cpc18) [Data set]. *Zenodo* (2017), , doi:10.5281/ZENODO.845873.
14. R. Selten, Axiomatic characterization of the quadratic scoring rule. *Exp. Econ.* **1**, 43–62 (1998).
 15. I. Erev, A. E. Roth, R. L. Slonim, G. Barron, Learning and equilibrium as useful approximations: Accuracy of prediction on randomly selected constant sum games. *Econ. Theory.* **33**, 29–51 (2007).
 16. L. Breiman, Random forests. *Mach. Learn.* **45**, 5–32 (2001).
 17. S. Rendle, in *Data Mining (ICDM), 2010 IEEE 10th International Conference on* (IEEE, 2010), pp. 995–1000.
 18. A. Glöckner, B. E. Hilbig, F. Henninger, S. Fiedler, The reversed description-experience gap: Disentangling sources of presentation format effects in risky choice. *J. Exp. Psychol. Gen.* **145**, 486 (2016).

Appendix A: Space of Problems

Each problem in the space is a binary choice problem in which either option is defined by a discrete distribution with up to 10 different possible outcomes. Ten dimensions (5 for each option) define these distributions: L_A , H_A , pH_A , $LotNum_A$, $LotShape_A$, L_B , H_B , pH_B , $LotNum_B$, $LotShape_B$. In particular, Option A provides a lottery, which has expected value of H_A , with probability pH_A and provides L_A otherwise (with probability $1 - pH_A$). Similarly, Option B provides a lottery, which has expected value of H_B , with probability pH_B , and provides L_B otherwise (with probability $1 - pH_B$). The distribution of the lottery of Option A (Option B) around its expected value H_A (H_B) is determined by the parameters $LotNum_A$ ($LotNum_B$) that defines the number of possible outcomes in the lottery, and $LotShape_A$ ($LotShape_B$) that defines whether the distribution is symmetric around its mean, right-skewed, left-skewed, or undefined (if $LotNum = 1$).

When a lottery is defined (i.e. $LotNum_A$ and/or $LotNum_B > 1$), its shape can be either “*Symm*”, “*R-skew*,” and “*L-skew*”. When the shape equals “*Symm*” the lottery’s possible outcomes are generated by adding the following terms to its EV (H_A or H_B): $-k/2$, $-k/2+1$, ..., $k/2-1$, and $k/2$, where $k = LotNum - 1$ (hence the lottery includes exactly $LotNum$ possible outcomes). The lottery’s distribution around its mean is binomial, with parameters k and $1/2$. In other words, the lottery’s distribution is a form of discretization of a normal distribution with mean H_A or H_B . Formally, if in a particular trial the lottery is drawn (which happens with probability pH_A or pH_B), the outcome generated is:

$$\left\{ \begin{array}{ll} \mathbf{H} - \frac{k}{2}, & \text{with probability } \binom{k}{0} \left(\frac{1}{2}\right)^k \\ \mathbf{H} - \frac{k}{2} + 1, & \text{with probability } \binom{k}{1} \left(\frac{1}{2}\right)^k \\ \vdots & \\ \mathbf{H} - \frac{k}{2} + k, & \text{with probability } \binom{k}{k} \left(\frac{1}{2}\right)^k \end{array} \right.$$

When the lottery's shape equals "*R-skew*," its possible outcomes are generated by adding the following terms to its EV: $C^+ + 2^1, C^+ + 2^2, \dots, C^+ + 2^n$, where $n = LotNum$ and $C^+ = -n - 1$. When the lottery's shape equals "*L-skew*," the possible outcomes are generated by adding the following terms to its EV: $C^- - 2^1, C^- - 2^2, \dots, C^- - 2^n$, where $C^- = n + 1$ (and $n = LotNum$). Note that C^+ and C^- are constants that keep the lottery's distribution at either H_A or H_B . In both cases (*R-skew* and *L-skew*), the lottery's distribution around its mean is a truncated geometric distribution with the parameter $\frac{1}{2}$ (with the last term's probability adjusted up such that the distribution is well-defined). That is, the distribution is skewed: very large outcomes in *R-skew* and very small outcomes in *L-skew* are obtained with small probabilities.

Each problem is further defined by two additional dimensions: *Corr* that determines whether there is a correlation (positive, negative, or zero) between the payoffs that the two options generate, and *Amb* that defines whether Option B is ambiguous (i.e. the probabilities of its possible outcomes are not revealed to the decision maker) or not. Therefore, each choice problem in the space is uniquely defined by 12 parameters. Notice that by defining $LotNum_A = 1$ (and $LotShape_A = "-"$), the current space of problems coincides with that investigated in Erev et al. (2017) thus this space is an expansion of that space.

Appendix B: Screenshots of the Experimental Paradigm




Figure B1. Example of a translated experimental screen in a choice problem with $Amb = 0$ and $Corr = -1$.




Figure B2. Example of a translated experimental screen in an ambiguous choice problem with $Amb = 1$ and $Corr = 0$.




Figure B3. Example of a translated experimental screen when full feedback is given (blocks 2–5). The participant here chose Option B.

Appendix C: The Estimation Set

Prob.	Option A					Option B					B-rate						
	H	pH	L	Lottery A		H	pH	L	Lottery B		Corr.	Amb	No-FB		With-FB		
				Num	Shape				Num	Shape			1	2	3	4	5
1	3	1	3	1	-	4	0.8	0	1	-	0	0	.42	.57	.57	.60	.65
2	3	0.25	0	1	-	4	0.2	0	1	-	0	0	.61	.62	.62	.64	.62
3	-1	1	-1	1	-	0	0.5	-2	1	-	0	0	.58	.60	.60	.58	.56
4	1	1	1	1	-	2	0.5	0	1	-	0	0	.35	.51	.54	.50	.54
5	-3	1	-3	1	-	0	0.2	-4	1	-	0	0	.49	.46	.42	.38	.36
6	0	0.75	-3	1	-	0	0.8	-4	1	-	0	0	.38	.40	.40	.42	.41
7	-1	1	-1	1	-	0	0.95	-20	1	-	0	0	.48	.63	.62	.62	.64
8	1	1	1	1	-	20	0.05	0	1	-	0	0	.39	.38	.33	.34	.29
9	1	1	1	1	-	100	0.01	0	1	-	0	0	.47	.40	.39	.39	.39
10	2	1	2	1	-	101	0.01	1	1	-	0	0	.55	.45	.43	.42	.42
11	19	1	19	1	-	20	0.9	-20	1	-	0	0	.13	.22	.21	.20	.21
12	0	1	0	1	-	50	0.5	-50	1	-	0	0	.34	.41	.43	.44	.38
13	0	1	0	1	-	50	0.5	-50	1	-	0	0	.36	.37	.40	.37	.36
14	0	1	0	1	-	1	0.5	-1	1	-	0	0	.49	.45	.42	.41	.38
15	7	1	7	1	-	50	0.5	1	1	-	0	0	.78	.84	.88	.83	.85
16	7	1	7	1	-	50	0.5	-1	1	-	0	0	.71	.79	.81	.83	.83
17	30	1	30	1	-	50	0.5	1	1	-	0	0	.24	.33	.33	.30	.29
18	30	1	30	1	-	50	0.5	-1	1	-	0	0	.23	.33	.40	.33	.33
19	9	1	9	1	-	9	1	9	8	R-skew	0	0	.37	.39	.36	.31	.30
20	9	1	9	1	-	9	1	9	8	R-skew	0	0	.38	.38	.39	.36	.36
21	10	0.5	0	1	-	10	0.5	0	1	-	0	1	.37	.42	.47	.48	.51
22	10	0.1	0	1	-	10	0.1	0	1	-	0	1	.82	.84	.78	.71	.66
23	10	0.9	0	1	-	10	0.9	0	1	-	0	1	.15	.16	.26	.33	.32
24	-2	1	-2	1	-	-1	0.5	-3	1	-	0	0	.48	.52	.48	.48	.45
25	2	1	2	1	-		0.5	1	1	-	0	0	.41	.50	.46	.46	.49
26	16	1	16	1	-	50	0.4	1	1	-	0	0	.50	.65	.61	.60	.55
27	16	1	16	1	-	48	0.4	1	3	L-skew	0	0	.50	.57	.60	.58	.57
28	6	0.5	0	1	-	9	0.5	0	1	-	-1	0	.91	.87	.83	.85	.84
29	2	1	2	1	-	3	1	3	1	-	0	0	.97	.98	.99	.99	1.0
30	6	0.5	0	1	-	8	0.5	0	1	-	1	0	.94	.97	.96	.98	.98
31	4	1	4	1	-	40	.6	-44	1	-	0	1	.23	.31	.34	.41	.37
32	24	.75	-4	1	-	82	.25	3	1	-	0	0	.68	.68	.67	.67	.69
33	-3	1	-3	1	-	14	.4	-22	1	-	0	0	.33	.28	.31	.25	.22
34	7	1	7	1	-	27	.1	4	3	Symm	0	0	.39	.45	.43	.41	.40

35	-5	1	-5	1	-	47	.01	-15	1	-	0	0	.18	.09	.04	.03	.05
36	28	1	28	1	-	88	.6	-46	4	R-skew	0	0	.39	.55	.56	.60	.57
37	23	.9	0	1	-	64	.4	-7	1	-	0	0	.38	.41	.42	.37	.37
38	24	1	24	1	-	34	.05	28	1	-	0	0	.91	.98	.99	1.0	1.0
39	29	1	29	1	-	33	.8	6	5	Symm	0	0	.50	.69	.70	.68	.66
40	3	.8	-37	1	-	79	.4	-46	7	L-skew	0	0	.49	.54	.61	.60	.56
41	29	1	29	1	-	44	.4	21	5	Symm	0	0	.68	.73	.74	.68	.68
42	-6	1	-6	1	-	54	.1	-21	1	-	0	1	.61	.48	.25	.23	.20
43	14	1	14	1	-	12	.9	9	1	-	0	0	.13	.04	.00	.00	.01
44	23	1	23	1	-	24	.99	-33	1	-	0	0	.27	.44	.47	.49	.47
45	13	1	13	1	-	13	1	13	9	Symm	0	0	.50	.56	.59	.53	.52
46	37	.01	9	1	-	30	.6	-37	1	-	0	0	.20	.27	.27	.31	.30
47	11	1	11	1	-	57	.2	-5	6	L-skew	0	0	.22	.20	.21	.16	.15
48	-2	1	-2	1	-	24	.5	-24	1	-	0	1	.39	.49	.52	.45	.42
49	23	1	23	1	-	23	1	23	3	Symm	0	0	.54	.50	.49	.48	.48
50	4	1	4	1	-	4	1	4	9	Symm	0	0	.44	.66	.60	.57	.57
51	42	.8	-18	1	-	68	.2	23	1	-	0	0	.79	.71	.74	.72	.70
52	46	.2	0	1	-	46	.25	-2	1	-	0	0	.36	.26	.25	.22	.22
53	28	1	28	1	-	42	.75	-22	1	-	0	0	.36	.43	.42	.42	.42
54	18	1	18	1	-	64	.5	-33	1	-	0	0	.32	.30	.31	.32	.29
55	43	.2	19	1	-	22	.25	17	9	Symm	-1	0	.21	.16	.09	.07	.07
56	-8	1	-8	1	-	-5	.99	-34	1	-	0	0	.76	.88	.91	.90	.89
57	49	.5	-3	1	-	33	.95	17	9	Symm	-1	0	.77	.71	.70	.73	.75
58	85	.4	-7	1	-	40	.25	24	1	-	0	0	.60	.51	.52	.54	.56
59	17	.25	16	1	-	43	.4	2	1	-	0	0	.51	.52	.52	.49	.49
60	51	.1	21	1	-	38	.6	1	1	-	0	0	.37	.38	.34	.29	.30
61	26	.25	25	1	-	29	.05	24	7	R-skew	0	0	.67	.62	.62	.60	.56
62	25	1	25	1	-	45	.2	17	1	-	0	0	.32	.32	.35	.34	.34
63	17	1	17	1	-	60	.1	15	5	Symm	0	0	.68	.70	.67	.66	.69
64	52	.1	-8	1	-	5	.9	-43	1	-	0	1	.35	.55	.70	.72	.68
65	12	.4	-16	1	-	-5	1	-5	1	-	0	0	.33	.40	.44	.45	.45
66	45	.6	2	1	-	54	.1	20	5	L-skew	0	0	.43	.35	.40	.44	.43
67	85	.25	4	1	-	54	.25	11	1	-	1	0	.45	.47	.46	.48	.43
68	12	1	12	1	-	102	.2	-14	1	-	0	0	.39	.27	.29	.32	.31
69	49	.5	11	1	-	31	.95	21	3	Symm	0	1	.39	.29	.37	.40	.45
70	18	1	18	1	-	35	.75	-19	1	-	0	0	.38	.55	.58	.60	.58
71	13	.6	-20	1	-	76	.2	-26	1	-	0	0	.38	.25	.29	.23	.28
72	-9	1	-9	1	-	13	.25	-8	1	-	0	0	.82	.96	1.0	1.0	1.0
73	2	1	2	1	-	51	.05	0	7	Symm	0	0	.37	.38	.39	.39	.41
74	44	.05	16	1	-	14	.9	10	3	Symm	0	1	.13	.05	.02	.00	.00
75	13	1	13	1	-	50	.6	-45	1	-	0	0	.35	.44	.42	.44	.50

76	35	.01	16	1	-	20	.5	13	5	Symm	0	1	.68	.71	.71	.68	.64
77	1	1	1	1	-	38	.4	-9	1	-	0	0	.65	.66	.65	.60	.63
78	19	1	19	1	-	44	.05	9	1	-	0	0	.11	.12	.11	.14	.12
79	32	.01	19	1	-	65	.01	9	1	-	0	0	.14	.07	.04	.02	.02
80	3	1	3	1	-	50	.4	-36	1	-	0	0	.47	.37	.41	.41	.43
81	10	.25	2	1	-	-1	.9	-32	1	-	0	1	.14	.04	.01	.01	.01
82	25	1	25	1	-	26	.01	25	7	Symm	0	1	.55	.72	.77	.81	.82
83	9	1	9	1	-	64	.01	9	1	-	0	0	.87	.96	.98	.98	.99
84	27	1	27	1	-	22	.99	-7	1	-	0	0	.08	.02	.00	.00	.00
85	20	1	20	1	-	70	.25	6	1	-	0	0	.43	.45	.49	.46	.44
86	71	.5	-11	1	-	61	.75	-49	1	-	0	1	.13	.23	.30	.32	.25
87	-2	1	-2	1	-	4	.99	-34	7	Symm	0	0	.81	.96	.98	.96	.98
88	17	.05	-7	1	-	13	.25	-15	1	-	0	1	.68	.57	.51	.44	.37
89	17	1	17	1	-	44	.1	17	1	-	0	0	.88	.96	.99	1.0	1.0
90	10	1	10	1	-	31	.75	-49	1	-	0	0	.42	.55	.53	.56	.55
91	7	1	7	1	-	16	.1	10	1	-	0	0	.92	.96	.99	.99	1.0
92	8	.8	-37	1	-	102	.2	-29	1	-	0	0	.39	.28	.35	.31	.32
93	5	1	5	1	-	103	.1	-9	4	L-skew	0	1	.56	.48	.44	.33	.30
94	7	1	7	1	-	6	.75	1	1	-	0	0	.10	.04	.02	.03	.02
95	-3	.05	-9	1	-	42	.4	-24	6	L-skew	1	0	.72	.60	.57	.57	.57
96	35	.5	-47	1	-	-10	.75	-15	1	-	0	0	.30	.34	.37	.34	.31
97	10	1	10	1	-	45	.2	-5	1	-	0	0	.22	.14	.27	.24	.21
98	94	.5	-40	1	-	36	.75	-21	7	Symm	0	0	.51	.45	.43	.44	.45
99	22	1	22	1	-	44	.4	15	5	Symm	0	0	.65	.74	.75	.71	.71
100	18	.6	-29	1	-	-1	1	-1	1	-	0	0	.54	.46	.50	.54	.53
101	28	1	28	1	-	73	.05	27	3	Symm	0	0	.83	.83	.79	.77	.76
102	11	1	11	1	-	25	.5	-3	3	Symm	0	1	.39	.51	.59	.62	.59
103	27	.8	-4	1	-	77	.1	22	6	R-skew	-1	0	.83	.82	.73	.78	.77
104	-6	1	-6	1	-	3	.99	-27	1	-	0	0	.85	.95	.98	.97	.98
105	30	1	30	1	-	90	.01	36	1	-	0	0	.90	.97	1.0	1.0	1.0
106	2	1	2	1	-	34	.05	-5	5	Symm	0	0	.20	.14	.15	.15	.12
107	25	1	25	1	-	65	.25	9	5	Symm	0	0	.39	.4	.40	.37	.32
108	16	1	16	1	-	91	.2	-11	1	-	0	0	.27	.17	.23	.19	.18
109	11	1	11	1	-	26	.5	-9	1	-	0	0	.33	.41	.47	.38	.40
110	12	1	12	1	-	29	.8	-35	2	L-skew	0	0	.56	.72	.71	.70	.77
111	28	1	28	1	-	47	.6	-13	1	-	0	0	.25	.41	.41	.40	.39
112	-7	1	-7	1	-	28	.2	-18	7	Symm	0	0	.51	.36	.31	.29	.26
113	9	.95	0	1	-	37	.25	-3	6	R-skew	0	0	.35	.37	.37	.38	.37
114	72	.01	-2	1	-	112	.25	-33	1	-	-1	0	.44	.45	.40	.42	.32
115	50	.4	5	1	-	20	.8	-17	7	Symm	0	0	.17	.19	.23	.20	.20
116	2	1	2	1	-	45	.05	3	5	Symm	0	0	.95	.99	1.0	1.0	1.0

117	-6	1	-6	1	-	7	.5	-30	1	-	0	0	.33	.39	.36	.35	.34
118	26	1	26	1	-	46	.5	10	6	L-skew	0	0	.47	.56	.62	.59	.57
119	19	.4	12	1	-	100	.25	-12	2	R-skew	0	0	.33	.32	.34	.34	.35
120	-9	.95	-26	1	-	-1	.1	-11	1	-	0	0	.57	.41	.43	.46	.42
121	-8	1	-8	1	-	21	.01	0	3	Symm	0	0	.79	.95	.99	1.0	.99
122	68	.05	-14	1	-	-11	.9	-36	1	-	0	0	.36	.39	.42	.46	.40
123	28	.75	-13	1	-	57	.1	16	1	-	0	0	.74	.64	.66	.70	.63
124	15	.95	7	1	-	42	.01	19	1	-	0	0	.85	.96	.97	.97	.98
125	28	1	28	1	-	41	.4	12	1	-	0	0	.29	.36	.33	.33	.36
126	-8	1	-8	1	-	80	.2	-18	7	Symm	0	0	.53	.52	.51	.49	.52
127	4	1	4	1	-	29	.6	-40	1	-	0	1	.26	.33	.44	.46	.44
128	-3	1	-3	1	-	32	.4	-16	1	-	0	0	.64	.57	.56	.57	.56
129	-2	1	-2	1	-	-2	1	-2	9	Symm	0	0	.46	.53	.51	.48	.53
130	72	.4	-41	1	-	16	.01	1	1	-	0	0	.61	.60	.56	.54	.54
131	18	1	18	1	-	45	.01	11	1	-	0	0	.19	.09	.08	.08	.06
132	11	1	11	1	-	20	.99	4	7	Symm	0	0	.81	.94	.97	.97	.98
133	3	1	3	1	-	8	.99	-17	9	Symm	0	0	.71	.91	.92	.92	.94
134	27	.05	24	1	-	31	.5	10	3	Symm	0	0	.34	.34	.38	.38	.37
135	6	1	6	1	-	8	.5	-1	1	-	0	0	.25	.32	.31	.29	.29
136	4	1	4	1	-	25	.01	-5	1	-	0	0	.16	.07	.07	.05	.05
137	3	1	3	1	-	4	.4	3	5	Symm	0	1	.73	.86	.90	.90	.90
138	23	1	23	1	-	21	.8	16	1	-	0	0	.13	.07	.01	.02	.02
139	14	1	14	1	-	35	.6	-9	7	Symm	0	0	.48	.67	.70	.64	.65
140	-2	1	-2	1	-	9	.25	8	1	-	0	0	.91	.98	.98	.99	.98
141	28	.8	-26	1	-	22	.75	2	1	-	0	0	.77	.70	.62	.60	.62
142	23	1	23	1	-	29	.8	-8	1	-	0	0	.30	.44	.43	.51	.54
143	67	.5	-39	1	-	93	.25	-15	1	-	0	0	.53	.58	.54	.60	.63
144	16	.8	12	1	-	15	1	15	9	Symm	0	0	.42	.50	.44	.42	.42
145	17	.5	-27	1	-	3	.75	-35	7	Symm	0	0	.42	.43	.42	.34	.34
146	45	.2	3	1	-	75	.05	13	5	Symm	0	0	.79	.82	.84	.87	.84
147	29	1	29	1	-	36	.1	32	7	Symm	0	0	.88	.96	.99	.98	.97
148	65	.01	1	1	-	12	.01	3	1	-	-1	1	.73	.81	.82	.84	.85
149	12	1	12	1	-	31	.1	12	3	Symm	0	0	.86	.90	.92	.91	.94
150	16	1	16	1	-	24	.05	12	3	L-skew	0	0	.35	.25	.22	.19	.17
151	27	0.8	-37	7	R-skew	28	0.01	20	3	Symm	0	0	.67	.50	.45	.45	.46
152	87	0.2	-26	5	Symm	64	0.05	0	9	Symm	0	0	.79	.80	.78	.79	.79
153	39	0.4	6	5	Symm	15	0.95	12	1	-	1	0	.37	.30	.36	.35	.38
154	46	0.1	16	1	-	65	0.2	2	5	Symm	0	1	.66	.58	.43	.40	.34
155	7	1	7	1	-	6	0.01	-4	5	Symm	0	0	.03	.01	.01	.00	.00
156	-1	1	-1	1	-	1	0.2	-1	7	Symm	0	0	.90	.94	.95	.97	.97
157	16	1	16	1	-	33	0.5	14	1	-	0	0	.91	.93	.90	.89	.90

158	3	0.8	-17	2	R-skew	6	0.6	-13	9	Symm	0	0	.52	.54	.61	.60	.60
159	14	1	14	1	-	20	0.95	16	1	-	0	0	.91	.96	.99	.99	.99
160	60	0.1	19	1	-	34	0.8	23	1	-	-1	0	.80	.85	.82	.84	.85
161	0	0.9	-4	5	R-skew	-2	0.99	-8	1	-	0	0	.18	.21	.22	.24	.28
162	39	0.6	14	1	-	48	0.25	24	1	-	0	0	.54	.51	.58	.56	.56
163	25	1	25	1	-	35	0.2	12	1	-	0	0	.09	.12	.12	.11	.11
164	1	0.5	1	1	-	40	0.01	-4	1	-	0	0	.11	.11	.12	.11	.11
165	35	0.75	-18	1	-	53	0.2	27	9	Symm	0	0	.90	.87	.84	.86	.82
166	-6	1	-6	1	-	15	0.6	-10	1	-	0	0	.91	.93	.89	.90	.93
167	13	1	13	1	-	12	0.8	8	1	-	0	0	.05	.02	.01	.01	.01
168	29	1	29	1	-	31	0.8	10	3	Symm	0	0	.25	.37	.41	.38	.36
169	58	0.1	0	1	-	26	0.2	8	3	R-skew	-1	1	.88	.91	.90	.87	.87
170	30	0.99	28	6	R-skew	30	0.6	19	6	L-skew	0	0	.07	.06	.11	.11	.10
171	-8	1	-8	1	-	7	0.25	-14	9	Symm	0	0	.58	.56	.53	.41	.38
172	21	1	21	1	-	26	0.8	-24	1	-	0	0	.19	.28	.31	.31	.34
173	83	0.5	-33	7	Symm	25	1	25	1	-	0	0	.61	.50	.49	.53	.50
174	74	0.6	-36	5	Symm	46	0.5	-11	3	Symm	0	0	.50	.37	.35	.33	.30
175	-4	1	-4	1	-	-2	0.8	-25	9	Symm	0	0	.31	.44	.48	.49	.43
176	-9	1	-9	1	-	12	0.5	-8	1	-	0	1	.86	.98	1.00	1.00	1.00
177	41	0.4	-4	5	Symm	42	0.4	-9	5	Symm	0	1	.56	.52	.48	.42	.39
178	32	0.6	-38	1	-	14	0.2	4	3	L-skew	0	0	.82	.74	.66	.64	.64
179	36	0.8	-18	5	Symm	52	0.25	21	3	L-skew	0	0	.83	.75	.68	.70	.71
180	50	0.6	-8	6	L-skew	52	0.4	-12	1	-	0	1	.29	.28	.22	.20	.19
181	-9	1	-9	1	-	71	0.2	-20	4	R-skew	0	0	.57	.53	.46	.43	.46
182	28	1	28	1	-	27	0.95	-11	1	-	0	0	.04	.02	.01	.01	.01
183	35	0.25	-14	3	Symm	42	0.25	-14	3	R-skew	0	0	.92	.92	.88	.87	.86
184	22	0.01	16	3	R-skew	84	0.25	-9	3	L-skew	-1	0	.26	.31	.33	.34	.31
185	7	1	7	1	-	18	0.99	-19	1	-	0	0	.78	.91	.94	.94	.96
186	19	1	19	1	-	97	0.1	5	1	-	0	0	.27	.24	.26	.27	.26
187	-3	1	-3	1	-	43	0.2	-20	1	-	0	0	.26	.22	.26	.25	.23
188	-7	1	-7	1	-	-4	0.01	-5	3	Symm	0	0	.83	.94	.98	.99	.99
189	2	1	2	1	-	4	0.75	-23	1	-	0	0	.17	.28	.27	.25	.26
190	10	1	10	1	-	28	0.2	-1	1	-	0	0	.14	.20	.21	.20	.21
191	16	1	16	1	-	15	0.8	-12	1	-	0	0	.05	.02	.01	.00	.01
192	10	1	10	1	-	59	0.01	4	1	-	0	0	.16	.13	.12	.13	.13
193	42	0.01	29	1	-	80	0.1	20	2	L-skew	0	0	.49	.41	.36	.35	.32
194	30	1	30	1	-	64	0.01	20	1	-	0	1	.46	.35	.21	.13	.11
195	10	1	10	1	-	92	0.01	9	9	Symm	0	0	.67	.62	.60	.56	.56
196	19	0.8	4	1	-	38	0.6	-27	1	-	0	1	.16	.23	.45	.43	.43
197	19	1	19	1	-	38	0.75	-33	4	R-skew	0	0	.37	.55	.62	.64	.64
198	31	0.8	-4	1	-	34	0.1	16	1	-	1	1	.68	.50	.40	.37	.35

199	76	0.5	-24	1	-	38	0.6	-20	3	L-skew	0	0	.44	.38	.35	.32	.35
200	16	1	16	1	-	40	0.6	-15	1	-	0	0	.37	.54	.56	.55	.56
201	14	1	14	1	-	14	1	14	2	R-skew	0	0	.44	.52	.52	.47	.40
202	53	0.1	27	1	-	32	0.95	9	1	-	0	0	.57	.63	.65	.61	.59
203	104	0.05	1	7	Symm	34	0.05	2	2	L-skew	0	0	.46	.47	.47	.45	.44
204	9	1	9	1	-	55	0.1	5	1	-	0	1	.64	.53	.45	.42	.37
205	36	0.05	16	6	L-skew	33	0.5	-1	1	-	0	0	.42	.49	.49	.49	.47
206	92	0.2	9	1	-	26	1	26	7	Symm	0	0	.68	.71	.69	.65	.67
207	22	0.9	-41	1	-	14	0.99	6	1	-	0	0	.67	.49	.41	.41	.42
208	6	1	6	4	R-skew	116	0.25	-31	9	Symm	0	0	.27	.31	.35	.33	.37
209	85	0.4	-22	7	L-skew	30	0.95	-2	1	-	0	0	.79	.77	.74	.71	.72
210	25	0.2	10	1	-	102	0.05	9	3	Symm	0	0	.64	.58	.49	.47	.48

Appendix D: Problem Selection Algorithm

The 60 problems in Experiment 1 were generated according to the following algorithm.

(This algorithm will also be used to determine the problems in the Experiment 2.)

1. Draw randomly $EV_A' \sim \text{Uni}(-10, 30)$ (a discrete uniform distribution)
2. Draw number of outcomes for Option A, N_A :
 - 2.1. With probability .4 ($N_A = 1$), set: $L_A = H_A = EV_A'$; $pH_A = 1$; $LotNum_A = 1$; and
 $LotShape_A = \text{"-"}$
 - 2.2. With probability .6 ($N_A > 1$), draw pH_A uniformly from the set $\{.01, .05, .1, .2, .25, .4, .5, .6, .75, .8, .9, .95, .99, 1\}$
 - 2.2.1. If $pH_A = 1$ then set $L_A = H_A = EV_A'$
 - 2.2.2. If $pH_A < 1$ then draw an outcome $temp \sim \text{Triangular}[-50, EV_A', 120]$
 - 2.2.2.1. If $\text{Round}(temp) > EV_A'$ then set $H_A = \text{Round}(temp)$;
 $L_A = \text{Round}[(EV_A' - H_A \cdot pH_A)/(1 - pH_A)]$
 - 2.2.2.2. If $\text{Round}(temp) < EV_A'$ then set $L_A = \text{Round}(temp)$;
 $H_A = \text{Round}\{[EV_A' - L_A(1 - pH_A)]/pH_A\}$
 - 2.2.2.3. If $\text{round}(temp) = EV_A'$ then set $L_A = H_A = EV_A'$
 - 2.2.3. Set lottery for Option A:
 - 2.2.3.1. With probability 0.6 the lottery is degenerate. Set $LotNum_A = 1$ and
 $LotShape_A = \text{"-"}$
 - 2.2.3.2. With probability 0.2 the lottery is skewed. Draw $temp$ uniformly from the
 set $\{-7, -6, \dots, -3, -2, 2, 3, \dots, 7, 8\}$
 - 2.2.3.2.1. If $temp > 0$ then set $LotNum_A = temp$ and $LotShape_A = \text{"R-skew"}$
 - 2.2.3.2.2. If $temp < 0$ then set $LotNum_A = -temp$ and $LotShape_A = \text{"L-skew"}$

2.2.3.3. With probability 0.2 the lottery is symmetric. Set $LotShape_A = \text{“Symm”}$
and draw $LotNum_A$ uniformly from the set $\{3, 5, 7, 9\}$

3. Draw difference in expected values between options, $DEV: DEV = \frac{1}{5} \sum_{i=1}^5 U_i$, where $U_i \sim \text{Uni}[-$

20, 20]

4. Set $EV_B' = EV_A + DEV$, where EV_A is the real expected value of Option A.

4.1. If $EV_B' < -50$ stop and start the process over

5. Draw pH_B uniformly from the set $\{.01, .05, .1, .2, .25, .4, .5, .6, .75, .8, .9, .95, .99, 1\}$

5.1. If $pH_B = 1$ then set $L_B = H_B = \text{Round}(EV_B')$

5.2. If $pH_B < 1$ then draw an outcome $temp \sim \text{Triangular}[-50, EV_B', 120]$

5.2.1. If $\text{Round}(temp) > EV_B'$ then set $H_B = \text{Round}(temp)$;

$$L_B = \text{Round}[(EV_B' - H_B \cdot pH_B)/(1 - pH_B)]$$

5.2.2. If $\text{Round}(temp) < EV_B'$ then set $L_B = \text{Round}(temp)$;

$$H_B = \text{Round}\{[EV_B' - L_B(1 - pH_B)]/pH_B\}$$

6. Set lottery for Option B:

6.1. With probability 0.5 the lottery is degenerate. Set $LotNum_B = 1$ and $LotShape_B = \text{“-”}$

6.2. With probability 0.25 the lottery is skewed. Draw $temp$ uniformly from the set

$$\{-7, -6, \dots, -3, -2, 2, 3, \dots, 7, 8\}$$

6.2.1. If $temp > 0$ then set $LotNum_B = temp$ and $LotShape_B = \text{“R-skew”}$

6.2.2. If $temp < 0$ then set $LotNum_B = -temp$ and $LotShape_B = \text{“L-skew”}$

6.3. With probability 0.25 the lottery is symmetric. Set $LotShape_B = \text{“Symm”}$ and draw

$LotNum_B$ uniformly from the set $\{3, 5, 7, 9\}$

7. Draw $Corr$: 0 with probability .8; 1 with probability .1; -1 with probability .1

8. Draw Amb : 0 with probability .8; 1 otherwise.

In addition, in the following cases the generated problem is not used for technical reasons: (a) there was a positive probability for an outcome larger than 256 or an outcome smaller than -50; (b) options were indistinguishable from participants' perspectives (i.e., had the same distributions and $Amb = 0$); (c) $Amb = 1$, but Option B had only one possible outcome; and (d) at least one option had no variance, but the options were correlated.

Appendix E: BEAST.sd

Like BEAST, BEAST.sd assumes that Option A is strictly preferred over Option B, after trial r , if and only if:

$$[BEV_A(r) - BEV_B(r)] + [ST_A(r) - ST_B(r)] + e(r) > 0 \quad (1)$$

where BEV_j is the best estimate of the expected value of option j , ST_j is the value of option j based on the use of sampling tools, and e is an error term.

Unlike BEAST, the error term is drawn from a normal distribution with mean 0, and standard deviation that equals $S_b(1 - subjDom(t))$, where S_b is a basic level of standard deviation that depends on the complexity of the problem, and $subjDom(t)$ is an estimate for the probability that one of the options is perceived as dominant, and t is the number of trials with feedback so far. In simple problems, $S_b = \sigma_i > 0$, whereas in complex problems $S_b = \lambda_i > \sigma_i$ (both λ_i and σ_i are properties of agent i). A problem is defined as complex if one of its options includes at least two outcomes, and the other includes at least three outcomes.

In non-ambiguous problems, the probability that there exists a subjectively dominant option is identical in all trials and equals:

$$subjDom(t) = \begin{cases} 1 & \text{if } MIN_A > MAX_B \text{ or } MAX_A < MIN_B \\ Pbet_A & \text{if } EV_A > EV_B \text{ and } EW_A \geq EW_B \\ Pbet_B & \text{if } EV_B > EV_A \text{ and } EW_B \geq EW_A \\ 0 & \text{otherwise} \end{cases}$$

where $Pbet_j$ is the probability that the payoff from choosing option j is not lower than the payoff from the alternative to j .

In ambiguous problems, the probability that there exists a subjectively dominant option, $subjDom(t)$, equals:

$$\begin{aligned}
& 1 && \text{if } MIN_A > MAX_B \text{ or } MAX_A < MIN_B \\
subjDom(t) = & pWin_A(t) && \text{if } EV_A > GM_B(t) \text{ and } EW_A \geq EW_B \text{ and } t > 0 \\
& pWin_B(t) && \text{if } GM_B(t) > EV_A \text{ and } EW_B \geq EW_A \\
& 0 && \text{otherwise}
\end{aligned}$$

Where $pWin_j(t)$ is the proportion of trials with feedback in which option j provided a higher payoff than the alternative so far, and $GM_B(t)$ is the mean payoff from all observed payoffs from Option B thus far.

Like in BEAST, $ST_j(r)$ equals the average of κ_i (a property of i) outcomes that are each generated by using one of four sampling tools. According to the sampling tool *unbiased*, draws are made from either the actual described distributions (before obtaining feedback) or the options' observed history of outcomes (after feedback). When draws are taken from the objective distributions, a *luck-level* procedure is used: The agent first draws a luck-level (uniformly between zero and one), and then uses it as a percentile in each of the prospects' cumulative distribution functions to draw the outcome that fits this percentile. When draws are made from the observed history, each past trial is equally likely to be selected, and the observed outcomes in both options in that trial are used.

The other three sampling tools are "biased", in the sense that they can be described as mental draws from distributions other than the objective distributions. The probability of choosing one of the biased tools, $PBias$, decreases when the agent receives feedback. Specifically, $PBias(t) = \beta_i / (\beta_i + 1 + t^{\theta_i})$, where $\beta_i > 0$ captures the magnitude of the agent's initial tendency to use one of the biased tools, t is the number of trials with feedback, and $\theta_i > 0$ captures agent i 's sensitivity to feedback.

Each biased tool is used with the same probability, $PBias(t)/3$. The sampling tool *uniform* yields each of the possible outcomes with equal probability using the luck-level procedure described above (but the draws are made from the uniform cumulative distribution function even after feedback is obtained). The sampling tool *sign* is identical to the tool *unbiased*, with one important exception: positive drawn values are replaced by R , and negative outcomes are replaced by $-R$, where R is the payoff range (the difference between the best and worst possible payoffs in the current problem). The sampling tool *contingent pessimism* yields the worst possible payoffs for each option (MIN_A and MIN_B), but only if $SignMax > 0$ and $RatioMin \leq \gamma_i$ ($0 < \gamma_i < 1$ is a property of i), where $SignMax$ is the sign of the best possible payoff, and $RatioMin$ is:

$$RatioMin = \begin{cases} 1, & \text{if } MIN_A = MIN_B \\ \frac{\text{Min}(|MIN_A|, |MIN_B|)}{\text{Max}(|MIN_A|, |MIN_B|)}, & \text{if } MIN_A \neq MIN_B \text{ and } \text{sign}(MIN_A) = \text{sign}(MIN_B) \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

when one of the two conditions is not met, *contingent pessimism* is identical to tool *uniform*.

Again, like BEAST, when the payoff distribution of option j is known (i.e., non-ambiguous), the best estimation for its expected value, BEV_j , equals to its true expected value EV_j . When the option is ambiguous, its initial expected value is estimated as a weighted average EV_A , MIN_B , and UEV_B , the latter being the estimated EV from Option B, under the assumption that all the possible outcomes are equally likely. BEAST.sd also assumes the same weighting for EV_A and UEV_B and captures the weighting of MIN_B with $0 \leq \varphi_i \leq 1$, an ambiguity aversion trait of agent i . That is,

$$BEV_B(0) = (1 - \varphi_i)(UEV_B + EV_A)/2 + \varphi_i \cdot MIN_B, \quad (5)$$

Each trial with feedback in the ambiguous problems moves $BEV_B(t)$ toward EV_B . Specifically,

$$BEV_B(t+1) = (1 - w_i) \cdot BEV_B(t) + w_i \cdot O_B(r) \quad (6)$$

where w_i is agent i 's sensitivity to feedback in ambiguous problems, and $O_B(r)$ is the observed payoff generated from the ambiguous Option B at trial r . Note that unlike in BEAST, in BEAST.sd the speed of convergence in ambiguous problems is a free parameter.

Finally, and again like BEAST, when unknown probabilities of the m possible outcomes need to be estimated, they are estimated under the assumption that the probability of the worst outcome, SP_{MINB} , is higher than $1/m$, and each of the other $m - 1$ probabilities equals $(1 - SP_{MINB})/(m - 1)$. Specifically, SP_{MINB} is computed as the value that minimizes the difference between $BEV_B(0)$ and the estimated expected value from Option B based on the estimated probabilities: $SP_{MINB} \cdot MIN_B + (1 - SP_{MINB}) \cdot U_{Bh}$, where $U_{Bh} = (m \cdot UEV_B - MIN_B)/(m - 1)$ denotes the average of the best $m - 1$ outcomes. This assumption implies that

$$SP_{MINB} = \begin{cases} 0, & \text{if } BEV_B(0) > U_{Bh} \\ 1, & \text{if } BEV_B(0) < MIN_B \\ \frac{U_{Bh} - BEV_B(0)}{U_{Bh} - MIN_B}, & \text{otherwise} \end{cases} \quad (7)$$

The eight properties of each agent are assumed to be drawn from uniform distributions between 0 and the model's parameters: $\sigma_i \sim U(0, \sigma)$, $\lambda_i \sim U(0, \lambda)$, $\kappa_i \sim (1, 2, 3, \dots, \kappa)$, $\beta_i \sim U(0, \beta)$, $\theta_i \sim U(0, \theta)$, $\gamma_i \sim U(0, \gamma)$, $\varphi_i \sim U(0, \varphi)$, and $w_i \sim U(0, W_{amb})$. That is, BEAST.sd has eight free parameters: σ , λ , κ , β , γ , φ , θ , and W_{amb} . Best fit of these parameters to the calibration set was obtained with the parameters $\sigma = 13$, $\lambda = 35$, $\kappa = 3$, $\beta = 1.4$, $\gamma = 1$, $\varphi = .25$, and $\theta = .7$, and $W_{amb} = .25$.