

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Non-parametric estimation of the individual's utility map

Permalink

<https://escholarship.org/uc/item/11h9r2df>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 35(35)

ISSN

1069-7977

Authors

Noguchi, Takao
Sanborn, Adam
Stewart, Neil

Publication Date

2013

Peer reviewed

Non-parametric estimation of the individual's utility map

Takao Noguchi (t.noguchi@warwick.ac.uk)
Adam Sanborn (a.n.sanborn@warwick.ac.uk)
Neil Stewart (neil.stewart@warwick.ac.uk)
Department of Psychology, University of Warwick

Abstract

Models of risky choice have attracted much attention in behavioural economics. Previous research has repeatedly demonstrated that individuals' choices are not well explained by expected utility theory, and a number of alternative models have been examined using carefully selected sets of choice alternatives. The model performance however, can depend on which choice alternatives are being tested. Here we develop a non-parametric method for estimating the utility map over the wide range of choice alternatives. The estimated maps are compared against the three of the most well-known models of risky choice: expected utility theory, cumulative prospect theory, and the transfer of attention exchange model. Model comparison indicates that cumulative prospect theory provides a better prediction of individuals' choices, but the estimated maps show that the overall shape of utility map is different from what the model predicts.

Keywords: decision making; risky choice; utility; MCMC with People; expected utility; cumulative prospect theory; transfer of attention exchange

Background

Understanding how people trade off risk and reward is a fundamental goal of behavioural economics. The most common approach to modelling how people make decisions between risky alternatives is based on the idea of utility: individuals integrate the probability of reward with the utility of the reward to produce an expected utility that describes how well the alternative is preferred. The alternative with the highest utility is most often chosen.

The normative calculation of utility that maximizes long-term gain is to multiply the probability with the utility of the associated outcome and to derive the expected utility. For an illustration, suppose an individual is considering a choice alternative with three possible outcomes: £20, £10, and £0. This particular alternative has a 20% probability for £20, 40% for £10, and 40% for £0. Then, the expected utility is $20\% \times u(\text{£}20) + 40\% \times u(\text{£}10) + 40\% \times u(\text{£}0)$, where u is the function to map the monetary value to the utility.

However, previous research has demonstrated that an individual's choice frequently deviates from the predictions of expected utility theory (for review, Schoemaker, 1982). To explain the deviations, descriptive models of how risk and reward are integrated have been developed (for review, Starmer, 2000). A common and useful way to visualize the predictions of these models is to look at the indifference lines, which connect choice alternatives of equal utility, over a Marschak–Machina probability triangle (Marschak, 1950; Machina, 1982). The probability triangle is a two-dimensional space which maps alternatives with varying probabilities for the same set of three potential outcomes. Throughout this pa-

per, we use £20, £10, and £0 as the potential outcomes from a choice alternative.

Figure 1 displays the predicted utility maps from three of the most well-known models of risky choice: expected utility theory, cumulative prospect theory (Tversky & Kahneman, 1992) and transfer of attention exchange (TAX) model (Birnbaum, 2008). In the probability triangle, the probability of attaining the best outcome (£20) is represented in the vertical axis, and the probability of the worst outcome (£0) is represented in the horizontal axis. The probability for the other outcome (£10) is represented as the distance from the diagonal boundary along the horizontal axis. The diagonal boundary ensures that the sum of the probabilities for £20, £10 and £0 does not exceed 1.

The red area in the triangles indicates the area of high utility, and the blue area is the area of low utility. Also, the coloured lines connect the alternatives of equal utility. These indifference lines highlight the differences between expected utility theory and the two descriptive models. Expected utility theory predicts indifference lines that are parallel and straight. Both cumulative prospect theory and the TAX model predict concave lines in the top corner of the triangle but convex lines in the lower right corner.

The usual experimental practice is to investigate choices in the regions of the triangle where models most differ from each other (e.g., Wu & Gonzalez, 1998). When the models are tested in this way, the “best” model may not predict choices away from the diagnostic regions well. For instance, Harless (1992) suggests that cumulative prospect theory outperforms expected utility theory only at the edges of the triangle. Thus, the model comparison could benefit from being tested on the whole area of triangle. One way is to estimate the utility map over the whole triangle and compare the estimated map against the model prediction. However to the best of our knowledge, the available estimation methods impose an assumption on how subjective value and probabilities are integrated (e.g., Abdellaoui, 2000), which could favour the model with the identical assumption.

To this end, we develop a non-parametric method to estimate entire utility maps, an extension of Markov chain Monte Carlo (MCMC) with People (Sanborn, Griffiths, & Shiffrin, 2010). We have modified MCMC with People to investigate the regions of the probability triangle where the choice alternatives are less preferred. The new method is tested in a simulation to show that it can deliver useful results within a reasonable number of trials. We then estimate utility maps from human. Finally, we discuss the results and future applications for this approach.

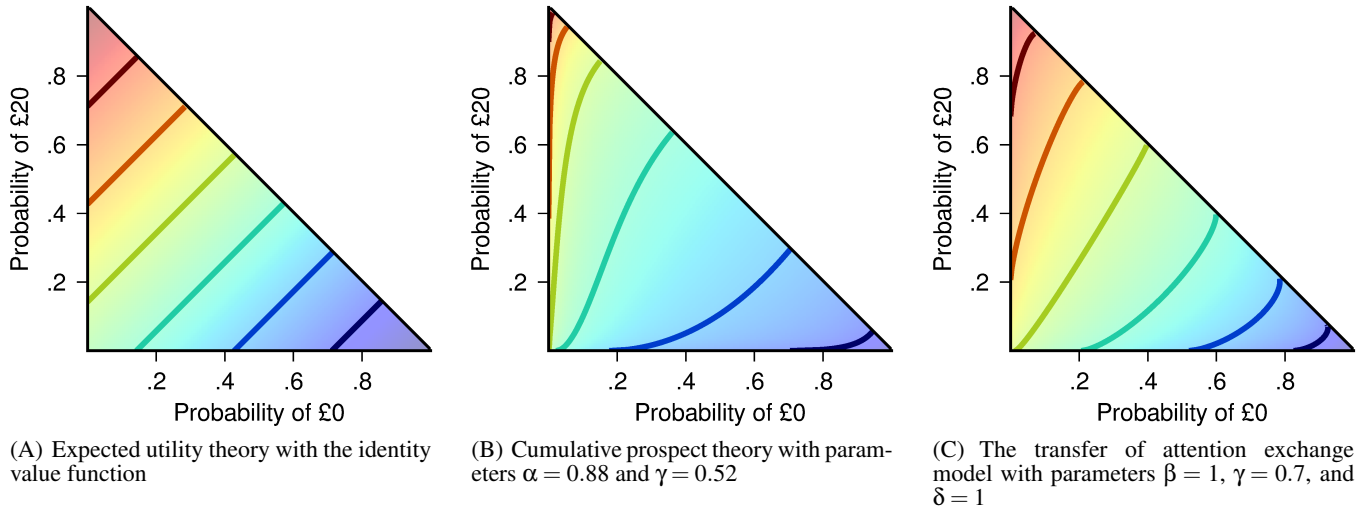


Figure 1: Theoretical predictions

Markov chain Monte Carlo with People

Markov chain Monte Carlo (MCMC) is a common method for drawing samples from a distribution. It has been widely used to perform probabilistic inference especially when solving the exact function of interest is computationally difficult (Neal, 1993).

MCMC begins in a start state z . A sample z' is first drawn from the proposal distribution q , and then z' is evaluated with the function of interest, π , to determine whether to accept z' as a new state or discard it and retain the current state z . The sequence of accepted samples forms a Markov chain, and after this Markov chain converges, accepted samples can be regarded as samples from the π distribution. To ensure that the Markov chain converges to π , it is sufficient to satisfy detailed balance (as well as ergodicity):

$$\pi(z)q(z'|z)A(z',z) = \pi(z')q(z|z')A(z,z'), \quad (1)$$

where $q(z'|z)$ is the probability of drawing z' when the current state is z and $A(z',z)$ is the probability of accepting proposal z' over the current state z .

Throughout the paper, we assume a symmetric distribution for q , $q(z'|z) = q(z|z')$, so Equation 1 becomes

$$\pi(z)A(z',z) = \pi(z')A(z,z'). \quad (2)$$

Detailed balance can be satisfied by carefully designing the acceptance function A . The most commonly used function is the Metropolis acceptance function (Metropolis, Rosenbluth, Rosenbluth, Teller, & Teller, 1953), but the Boltzmann acceptance function (Flinn & McManus, 1961) is of interest here:

$$A(z',z) = \frac{\pi(z')}{\pi(z) + \pi(z')}.$$

If an individual is asked to make a choice between alternatives z' and z , then the Boltzmann acceptance function can

model that individual's choice. This is because the Boltzmann function is equivalent to Luce's choice rule (Luce, 1959), which has been frequently used to model risky choice (e.g., Blavatsky & Pogrebn, 2010). As a result, by sequentially presenting pairs of choice alternatives to an individual (where the new alternative z' is selected by the computer), the collection of choice alternatives chosen by the individual can be treated as samples from the probability distribution whose density is proportional to the individual's utility (Sanborn et al., 2010).

Extending MCMC with People

However, sampling from the individual's utility distribution does not necessarily serve to estimate the shape of the utility map: pilot work confirms that all of the samples will be concentrated around the most favourable alternative (100% probability of £20 in the triangle), and that it would take a very large number of trials to explore the rest of the utility map. To enable the reasonable estimation of the utility map, the stationary distribution needs to be more diffused, so that the Markov chain travels better around the triangular space.

For this purpose, we implement a latent agent in the experimental program. This agent makes an independent choice between the same alternatives as the participant, and only when the agent and the participant both select the new choice alternative does the new alternative become the new state. Otherwise, the current state remains the same and another alternative is generated from the proposal distribution.

When the agent is implemented in this way, the acceptance function becomes a joint function of the participant's and the agent's choices. Specifically, the acceptance function is defined as

$$A^*(z',z) = \frac{f(z')}{f(z) + f(z')} \frac{g(z')}{g(z) + g(z')},$$

where f is the utility function for the participant and g is the

agent's utility function. Here, both the participant and the agent follow the Boltzmann acceptance function. Then Equation 2 becomes

$$f(z)g(z)A^*(z',z) = f(z')g(z')A^*(z,z').$$

With the implementation of the agent, the trajectory of the Markov states depends on both the participant's and the agent's choices. If the agent's utility is the lowest at the top corner of the triangle, the Markov chain would be pushed away from that region. With this extended method, the stationary distribution of the Markov chain is the joint utility function of the participant and the agent, fg . The participant's utility map can subsequently be recovered by dividing the joint utility by the latent agent's known utility, g .

Simulation

To demonstrate that the developed method can estimate a participant's utility map within a reasonable number of trials, we conducted a simulation.

Method

The simulation used two of the utility functions in Figure 1: the latent agent's utility function, g , was set to the inverse of expected utility theory, and the simulated participant's function, f , was cumulative prospect theory:

$$g = \frac{1}{20 \times p(\pounds20) + 10 \times p(\pounds10)},$$

and

$$f = 20^\alpha \times w(p(\pounds20)) + 10^\alpha \times (w(p(\pounds20) + p(\pounds10)) - w(\pounds20)),$$

where $p(\pounds20)$ is the probability of attaining $\pounds20$, and $w(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{1/\gamma}}$. The parameter values for α and γ were 0.88 and 0.52, respectively. The proposal distribution, q , was uniform over the triangular space. The possible outcomes were fixed to be $\pounds20$, $\pounds10$ and $\pounds0$, and hence, the agent and the simulated participant repeatedly made choices between two alternatives with varying probabilities for fixed outcomes: e.g., a choice between an alternative with a 30% probability for $\pounds20$, 40% for $\pounds10$ and 30% for $\pounds0$, and another alternative with a 10% probability for $\pounds20$, 60% for $\pounds10$ and 30% for $\pounds0$.

With the above specifications, a choice trial was simulated as follows. First, the agent used the g function to evaluate each alternative and used the Boltzmann acceptance function to select between the current state and the proposed alternative. If the agent preferred the current state over the proposed alternative, another alternative was sampled from the proposal distribution. If the agent chose the new alternative over the current state, the simulated participant used the f function to make a choice between the same two alternatives.

Although the agent and the simulated participant could have made a choice at the same time over the same two alternatives, we had the agent decide first: if the agent does not select the new alternative, the previous state remains the state

regardless of the choice the participant makes. This reduces the number of choices the participant must make.

Each simulation consisted of three chains: one chain started with the Markov state of 60% of $\pounds20$, 20% of $\pounds10$ and 20% of $\pounds0$. Another chain started with the state of 20% of $\pounds20$, 60% of $\pounds10$ and 20% of $\pounds0$. The final chain started with 20% of $\pounds20$, 20% of $\pounds10$ and 60% of $\pounds0$.

Results and Discussion

The first 100 trials were considered to be trials before convergence of the Markov chain (burn-in period) and were discarded from each chain. The remaining samples from the three chains were pooled and smoothed by kernel density estimation. Because of the triangular boundary of the estimation space, it is actually quite difficult to produce unbiased indifference lines. We chose to use a Dirichlet kernel, an extension of the Beta kernel (Chen, 1999) to the triangular space, because it produced less bias than the other alternatives we investigated. The Dirichlet kernel is defined as

$$\hat{f}(x)g(x) = \sum_i Dir(z_i|\alpha_1, \alpha_2, \alpha_3),$$

where z_i is the i th state in the Markov chain, x is a vector of probabilities for three outcomes, and α_j is $x_j / \min(h, x_j, 1 - x_j)$. The kernel width h was set to 0.09. This smoothed joint distribution is then divided by g to derive the estimation \hat{f} .

To assess the similarity between f and \hat{f} , we computed Kullback–Leibler (KL; denoted as $KL(f||\hat{f})$) divergence (Kullback & Leibler, 1951), which measures how much information is lost in the estimation process.

The KL divergences for different sample sizes are plotted in the left panel of Figure 3. This figure illustrates that the estimation shows the increasingly smaller divergence within the first few hundred trials. The estimation becomes reasonably accurate on average after 700–800 trials.

The middle and right panels of Figure 3 display the estimations after 1,000 trials. The estimation with the smallest KL divergence among the 10 simulation runs is in the middle panel, and the right panel show the estimation with the largest KL divergence. Both panels show the key property of cumulative prospect theory: the indifference lines show fanning-out property from the lower left corner toward the diagonal boundary.

Thus, the simulation demonstrated that the proposed method with the Dirichlet kernel density estimation can recover the key characteristic of the utility map using a reasonable number of samples.

Experiment

Method

Participant Ten participants were recruited through the subject panel at the University of Warwick. One participant did not complete the experiment, leaving nine (five male and four female) participants. Their age ranged from 19 to 30 with a mean of 22.9.


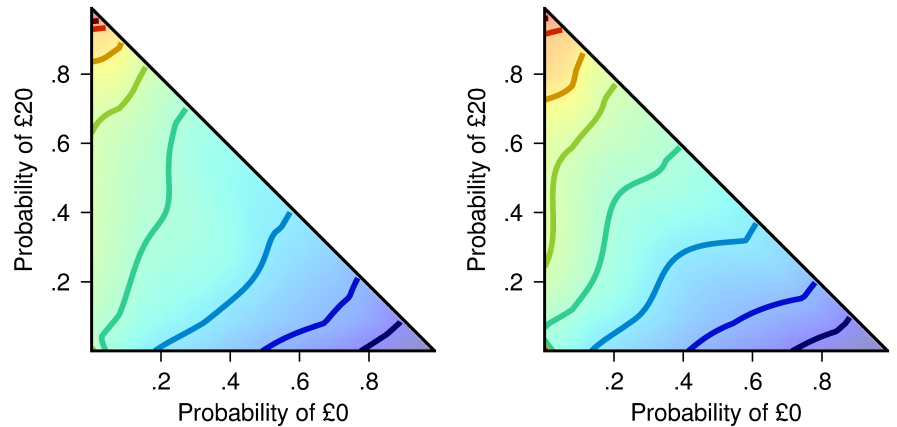


Figure 2: The KL divergences between f and \hat{f} for various numbers of trials. The solid line represents the mean measurement of the 10 simulation runs, and the dotted lines are maximum and minimum values archived in the simulations.



(A) The estimation with the smallest KL divergence ($KL(f||\hat{f}) = 0.002$) (B) The estimation with the largest KL divergence ($KL(f||\hat{f}) = 0.007$)

Figure 3: Estimation of cumulative prospect theory with 1,000 trials

Procedure The experimental procedure closely followed that of the simulation. The agent’s utility function, g , was set to the inverse of expected utility theory raised to the power of 8, and the proposal distribution, q , was uniform over the triangular space. The possible outcomes were fixed to be £20, £10 and £0, and hence, the agent and the participant repeatedly made choices between two alternatives with varying probabilities for fixed outcomes.

In each trial, the agent made a decision first, and a new alternative was drawn until the agent chose the new alternative. Three chains with the same start states as the simulation were run interleaved until participants had made 1,000 choices per chain. In addition, 50 catch trials were inserted into the experiment, so that we could assess whether participants were engaged in the task. In each catch trial, one alternative had larger probabilities for both £20 and £10. If a participant was not engaged with the task and randomly making choices, it is expected that he or she would occasionally not select the non-dominant alternative.

The experiment presented a choice alternative as a pie chart with three slices. Each slice represented one possible outcome, and the size of the slice was proportional to the probability of the outcome. Participants were forced to log out from the online experiment and take a break after spending one hour on it. After the minimum break of three hours, participants were allowed to log in again and resume the experiment.

The choices participant made were incentivized: we invited participants to the lab when participants completed the experiment. At the lab, we randomly selected one trial from the experiment and played the selected alternative for real. Participants were paid what they earned from the play.

Results and Discussion

All the nine participants selected the dominant alternative in all of the catch trials, which was evidence that all participants understood and were engaged in the task.

Utility maps were estimated as in the simulation study. All participants show a sharp peak at the top corner of the triangle in the estimated maps. The sharp peak makes it difficult to see the shape of the map, and thus for illustration purposes, we spaced out the indifference lines by taking the natural logarithm of the estimation. As a result, differences in small utilities are exaggerated, but the shapes of the indifference lines are not affected. The resulting maps are displayed in Figure 4. Each panel in the figure corresponds to one participant’s map.

The estimated maps show the steep indifference lines, especially where the probability of £0 is small. The steep lines indicate aversion to the worst outcome (c.f., Tversky & Kahneman, 1992; Birnbaum, 2008), where the increment in probability for the worst outcome needs to be compensated with a larger increment in probability for the most desirable outcome. The steepness tends to be lessened near the lower right corner of the triangle. As a result, for Participants A, D and H in particular, the indifference lines show the fanning-out property. The fanning-out suggests that participants more willingly accept an increment in probability for the worst outcome when the probability is already large. The fanning-out is consistent with the prediction from cumulative prospect theory and the transfer of attention exchange (TAX) model.

The estimated maps also show the convex indifference lines throughout the triangle. The convexity makes the estimated maps appear rather different from the predicted utility maps from cumulative prospect theory and the TAX model, which expect the concavity toward the top corner of the triangle (Figure 1).

To quantitatively assess the model performance, we fit the models to the individuals’ choices by maximizing the likeli-

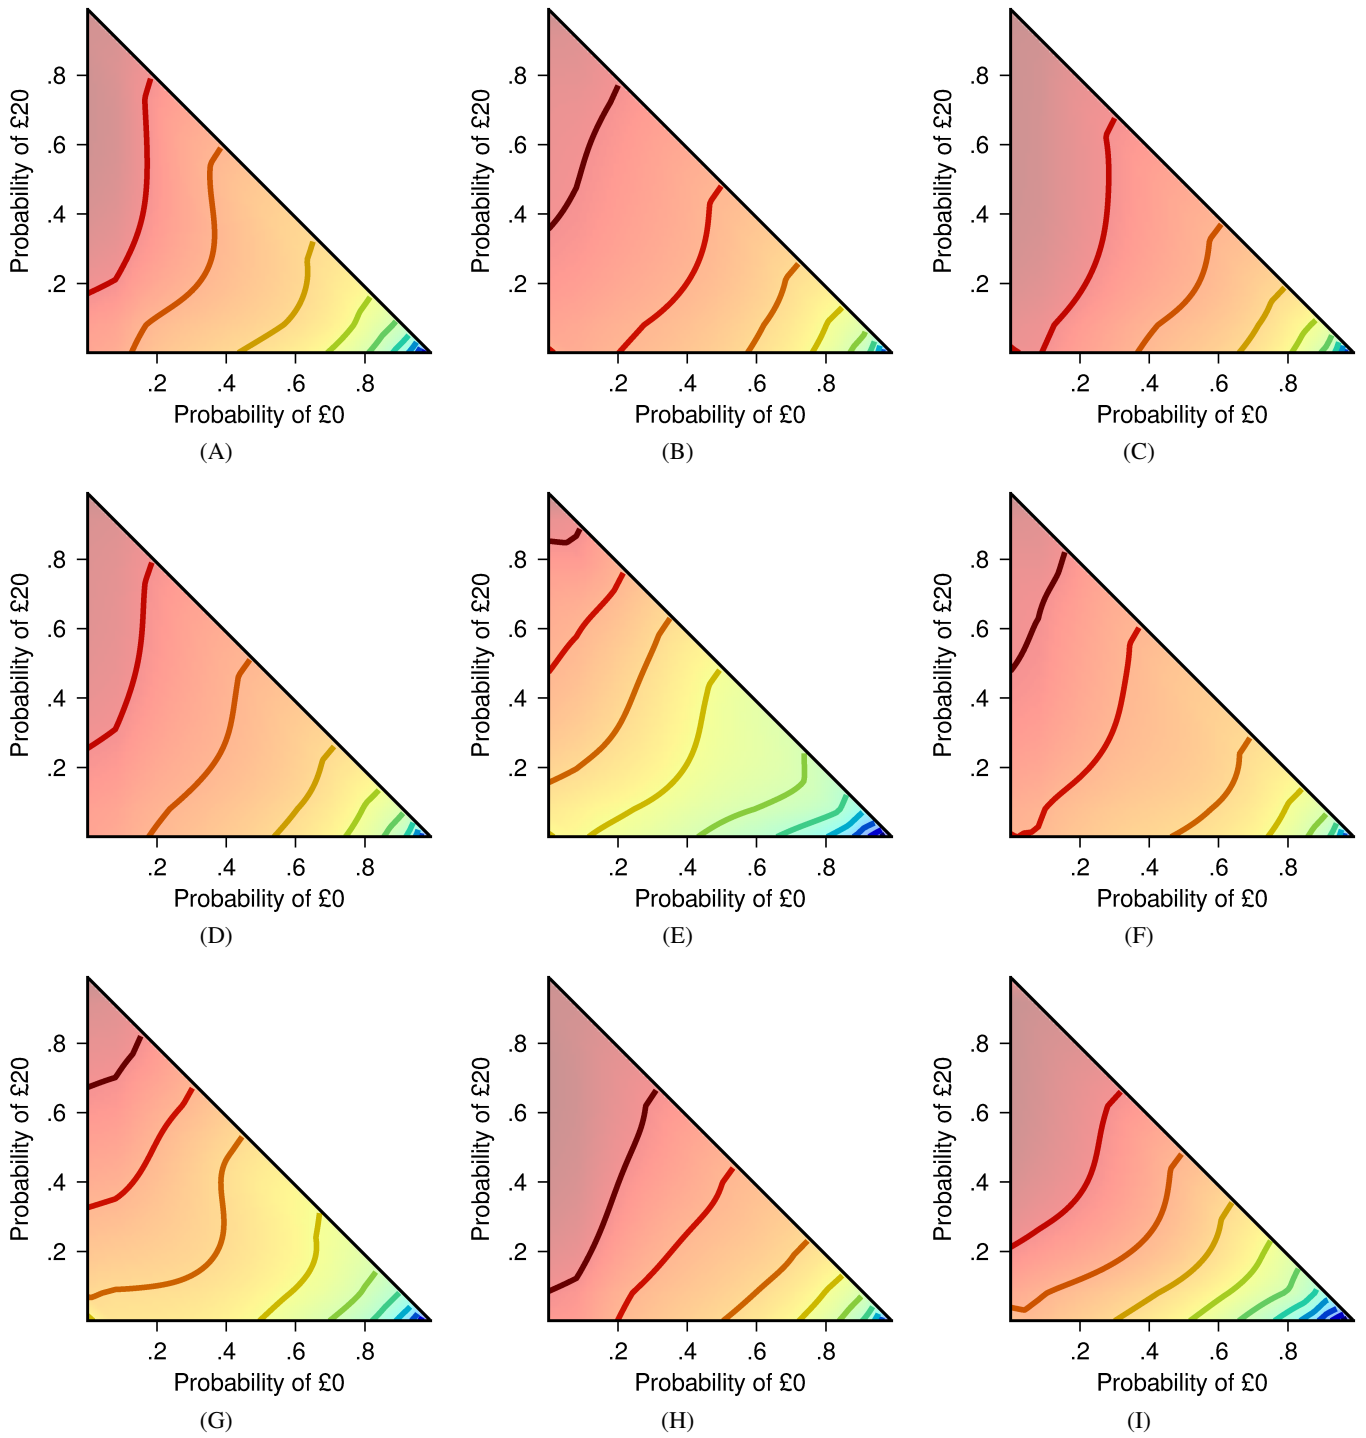


Figure 4: $\ln(\hat{f})$

hoods. When fitting the model, we used the power law utility function for expected utility theory: $u(s) = s^\alpha$. The range of parameter values are restricted to be between 0 and 1 for all the parameters. Also, each model included one additional parameter to raise the predicted utility. This exponent controls how steep the peak is toward the most favourable alternative.

The value for this exponent parameter is restricted to be non-negative.

Bayesian information criteria (BIC) indicates that the choices are best predicted by cumulative prospect theory for seven out of nine participants (Panels A through G). The TAX model achieves smallest BIC for one of the remaining partic-

ipants (Panel H), and the expected utility theory has smallest BIC for the other participant (Panel I).

General Discussion

Previous research has demonstrated that individuals' choices deviate predictions from expected utility theory, and variety of descriptive models have been proposed. However, the deviation from expected utility theory has often been studied with relatively limited range of choice alternatives. The present study developed the non-parametric method to estimate utility maps over the whole probability triangle.

The curvature of the indifference lines in the estimated maps implies differences to the predictions from the existing models: The lines tend to be convex where concavity is expected. Even though cumulative prospect theory (CPT) does not predict this curvature, CPT provides a better fit to the choice data than expected utility theory or the attention exchange model does for the majority of participants. Thus, a new model could explain the choices better than CPT, if the new model produces a utility map similar to the estimated maps.

In developing such a model, it is useful to identify choice alternatives where the CPT prediction differs from the individuals' choice behaviour. To this end, the estimation method that we have developed can be further extended. As the developed method lets the MCMC chain converge to the joint distribution of the individual's and the agent's utility, manipulation of agent's utility function can reveal interesting joint distributions. For instance, by setting the latent agent's utility to the inverse of the CPT prediction, the MCMC chain converges to the distribution whose density is proportional to the individual's utility divided by the CPT prediction. The condensed area in this joint utility distribution is where the CPT prediction is smaller than the individual's utility (i.e., the area where CPT underpredicts the utility), and the thin area is where the CPT prediction is larger than the individual's utility (i.e., the area where CPT overpredicts the utility).

To conclude, we have developed the method for estimating the utility map. The developed method can be further leveraged in future study.

References

- Abdellaoui, M. (2000). Parameter-free elicitation of utility and probability weighting functions. *Management Science*, *46*, 1497–1512.
- Birnbaum, M. (2008). New paradoxes of risky decision making. *Psychological Review*, *115*, 463–501.
- Blavatsky, P. R., & Pogrebna, G. (2010). Models of stochastic choice and decision theories: why both are important for analyzing decisions. *Journal of Applied Econometrics*, *986*, 963–986.
- Chen, S. X. (1999). Beta kernel estimators for density functions. *Computational Statistics & Data Analysis*, *31*, 131–145.
- Flinn, P. A., & McManus, G. M. (1961). Monte Carlo calculation of the order-disorder transformation in the body-centered cubic lattice. *Physical Review*, *124*, 54–59.
- Harless, D. (1992). Predictions about indifference curves inside the unit triangle. *Journal of Economic Behavior and Organization*, *18*, 391–414.
- Kullback, S., & Leibler, P. C. (1951). On information and sufficiency. *Annals of Mathematical Statistics*, *22*, 79–86.
- Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: John Wiley & Sons, Inc.
- Machina, M. J. (1982). "Expected Utility" analysis without the independence axiom. *Econometrica*, *50*, 277–323.
- Marschak, J. (1950). Rational behavior, uncertain prospects, and measurable utility. *Econometrica*, *18*, 111–141.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, *21*, 1087–1092.
- Neal, R. (1993). *Probabilistic inference using Markov chain Monte Carlo methods* (Tech. Rep. No. CRG-TR-93-1). Department of Compute Science, University of Toronto.
- Sanborn, A. N., Griffiths, T. L., & Shiffrin, R. M. (2010). Uncovering mental representations with Markov chain Monte Carlo. *Cognitive Psychology*, *60*, 63–106.
- Schoemaker, P. J. H. (1982). The expected utility model: Its variants, purposes, evidence and limitations. *Journal of Economic Literature*, *20*, 529 - 563.
- Starmer, C. (2000). Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature*, *38*, 332–382.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.
- Wu, G., & Gonzalez, R. (1998). Common consequence conditions in decision making under risk. *Journal of Risk and Uncertainty*, *16*, 115–139.