Rationalizing subjective probability distortions

Nisheeth Srivastava (nsrivast@cse.iitk.ac.in)
Department of CSE, IIT Kanpur
Kalyanpur, UP 208016 India

Edward Vul (evul@ucsd.edu)
Department of Psychology, 9500 Gilman Drive
La Jolla, CA 92093 USA

Abstract

You cannot know the contents of a memory until after you have actually retrieved it. This paper considers the implications of this straightforward observation upon the psychological process of preference construction. We show that this constraint renders observers with random access memory susceptible to tail risks. We show that this difficulty can be rectified by permitting observers to weight memory retrieval for such observations, that outcome utility cannot be used for this purpose, but information-theoretic surprise can serve as a useful proxy for it. Using two novel experiments, we present evidence in support of our account. With the first, we show that humans find surprising experiences easier to remember. With the second, we show that surprising experiences in the past have a greater influence on future decisions than is statistically warranted. This twofold demonstration substantiates a psychologically plausible account for the origin of subjective probability distortions.

Keywords: prospect theory; rational analysis; decision-making; extreme events

Probabilistic sophistication is impossible

Offer an economist a choice between a risky option that yields $ 10 with a probability of 0.2 versus a safe option that always yields $ 1, and you will barely have time to blink before they give you both the 'correct' answer, and their rational basis for choosing it (expected value maximization). Offer the same economist a choice between a risky option with known payouts but unknown probabilities, and the same economist will try to estimate the probability from previous experience with related choices, and then complete the expected value calculation using this estimated probability.

What happens when we ask a cognitive scientist these questions? To the first,they will respond in much the same way as the economist: given payoffs and probabilities, calculate expected value and choose the better option by this measure. To the second, however, their response cannot be as straightforward. The cognitive scientist would realize that estimating implicit probabiliies from past experience is a poorly constrained problem absent knowledge about the sampling distribution inside their head. Thus, the cognitive scientist's response to the second question would have to be conditioned on an assumption about the sampling distribution of the memory retrieval process from within the set of all relevant past experiences.

The economist and the cognitive scientist's responses to the second question would diverge most significantly in situations that deal with small probabilities. When confronted with high (dis)utility, rare events, an economic observer can appropriately modify its behavior by taking their multiplicative impact into account. For example, in evaluating the decision to insure against an implicit 1% chance of losing $1000 for a $ 5 premium, they will easily calculate that the expected value of the uninsured loss is lower and buy insurance. They can do this calculation because, even though utility and probability are definitionally separate, they are available to the observer *simultaneously* and so, can inform their behavior adequately.

A psychological observer, in contrast, cannot simply assume that utility and probability information is available to it - it must explicitly model the process by which this information becomes available. A psychological observer that assumed a uniform sampling distribution on past memories would retrieved past experiences in the relative proportions it experiences them. Assuming veridical access to empirical frequencies, the probability this observer will be able to sample at least one memory corresponding to experienced uninsured loss in N tries would be $1-0.99^N \sim 0.01N$ when drawing between 5-50 samples. The sequential nature of obtaining frequency and utility information for such an observer would render it insensitive to rare, extreme events and, even after sampling 20 experiences from memory before deciding, for instance, it would decide 80% of the time not to buy insurance, because the memory corresponding to drastic disutility was never sampled.

An animal that discounts the possibility of rare, extreme events occurring is unlikely to survive for very long in the wild. So, a uniform prior on past experiences cannot be an adequate representation of how memory retrieval informs future behavior. While assuming that observers explore the entire set of memories available at the time of each decision could solve this problem, by ensuring that even the rarest of memories are sampled, such a proposal is unattractive on mechanistic grounds. It would, for example, predict that observers would get slower at making decisions for similar problems over time. Additional objections to such a proposal can be found in previous studies supporting both the existence of limited memory sampling (Ashby, Tein, & Balakrishnan, 1993) and its value on both ecological (Todd & Gigerenzer, 2000) and information-theoretic (Vul, Goodman, Griffiths, & Tenenbaum, 2014) grounds. Likewise problematic are reinforcement learning accounts that assume that observers maintain a weighted average of expected value, rather than compute it on the fly (Montague, Dayan, & Sejnowski, 1996). Such accounts are difficult to generalize to peoples' responses to novel settings (Dayan & Niv, 2008), and to the high variability in behavior responses within subjects across trials in choice experiments (Beck, Ma, Pitkow, Latham, & Pouget, 2012).

So we have a problem. Preferences have to be constructed by retrieving memory particles, taking the impact of extreme events (characterized by extreme utility) into account. But this construction has to respect the epistemic constraint imposed by the nature of memory retrieval - utility information contained in a memory particle cannot be observed until the memory particle has already been retrieved, rendering direct utility-weighting of probabilities at the time of retrieval, as proposed in (Bordalo, Gennaioli, & Shleifer, 2012; Lieder, Hsu, & Griffiths, 2014) psychologically unfeasible.

The logical solution is to determine how the utility of prior events might affect the probability with which memory particles are retrieved *at the time of encoding*. We show that the impact of extreme events enters observers' psychological calculations via *surprise* - memories of surprising events are easier to retrieve. We assume that memory encoding is sensitive to extreme prediction error, likely through devoting greater attention to events causing such large errors, thereby facilitating their subsequent recall. Compilation of prior history now becomes a biased sample, privileging retrieval of surprising events. The ecological correlation between extreme (dis)utility and surprise makes the latter a reliable proxy for the former.

In this paper, we develop a formal model for preference construction via such surprise-biased memory sampling, and present results from two experiments that justify the two crucial assumptions in our model: surprise-sensitivity in memory retrieval, and disproportionate impact of low probability events in the past on future choices.

Preference construction via biased memory sampling

The idea that behavior is determined by averaging over a subset of previous experiences sampled from long-term memory is shared by several classic cognitive theories (Anderson, 1996). Whereas classically such accounts have been detailed within specific cognitive architectures such as ACT-R, we focus on developing our model using only the abstract generality of preference construction via sample averaging. Such a preference formation model can be described computationally as,

$$p(u,x) = \sum_{m \in \mathcal{M}} p(u,x|m)p(m), \tag{1}$$

where $x \in \mathcal{X}$ are options available to the agent, $u : u \rightarrow \mathcal{R}_{+}$ is a utility indicator, and $m \in \mathcal{M}$ are memory *particles* corresponding to past choice selections. We use the term memory *particle* to specify a particular assumption we make about the nature of memory representation, i.e. that observers encode utility information alongside option labels encountered at a point in time as one memory encoding event $m = \{u,x,t\}$. Once encoded in memory, the time label loses meaning, and is ignored. The probability distribution $p(m)$ - which we call the *memory distribution* - encodes the probability of recalling the memory particle m, while the distribution $p(u,x|m)$ encodes the utility information learned during the experience corresponding to the memory particle m. The marginal distribution $p(u,x)$ represents the agent's revealed preference for the option x, constructed by summarizing information retrieved from memory.

Notice that state probability estimation is already built into this basic model. Options that occur more frequently in the world will have more memory particles associated with them. A flat memory distribution on the set of memory particles would privilege retrieval of memories related to options following the relative frequencies of their observations in the world. To introduce a surprise bias, we have to sensitize the memory encoding process to the presence of large prediction errors. We use a standard specification of prediction error using an information divergence,

$$R(p(u,x), p(u,x|m)) = \sum_{x \in \mathcal{X}} p(u,x) \log \frac{p(u,x)}{p(u,x|m)}, \tag{2}$$

where, $p(u,x)$ is the observer's preference right before encountering an event, and $p(u,x|m)$ is the preference information encoded by the observer corresponding to that event. We further instantiate the memory prior as a softmax function of this prediction error,

$$p(m) = \frac{\exp(A(m))}{\sum_{m \in \mathcal{M}} \exp(A(m))}, \tag{3}$$

where the weight $A(m)$ is computed as deviation of surprise corresponding to that memory particle from the average surprise $\overline{R}$ experienced by the agent in the sequential choice selection process,

$$A(m) = \max(0, R(p(x), p(x|m)) - \overline{R}), \tag{4}$$

In conjunction with Equation 1, this formal specification of the memory process constitutes a surprise-sensitive model of preference construction via memory retrieval, formalizing an account wherein observers sample a limited set of available information when constructing new preferences, and do so in a way that is statistically biased, but ecologically justified by the necessity of accounting for extreme (dis)utility events under the procedural constraint that (dis)utility information itself has to be retrieved from memory.

Experiments

Memory retrieval is sensitive to surprise

Our model relies heavily on the assumption that memories of surprising outcomes are easier to recall. We directly tested this using a behavioral experiment[1]. 30 UCSD undergraduate students participated in the experiment for course credit. The experiment was designed such that participants were shown 24 visual stimuli (drawn randomly from a set of 72 flags of cities from around the world) in sets of 4 (randomized across participants, fixed within participants at the beginning of the game), and were asked to identify and subsequently remember, from each set of 4, one special flag. Since we did not

[1] A working version of this experiment is available online at http://experiments.evullab.org/memory_expt/game.html

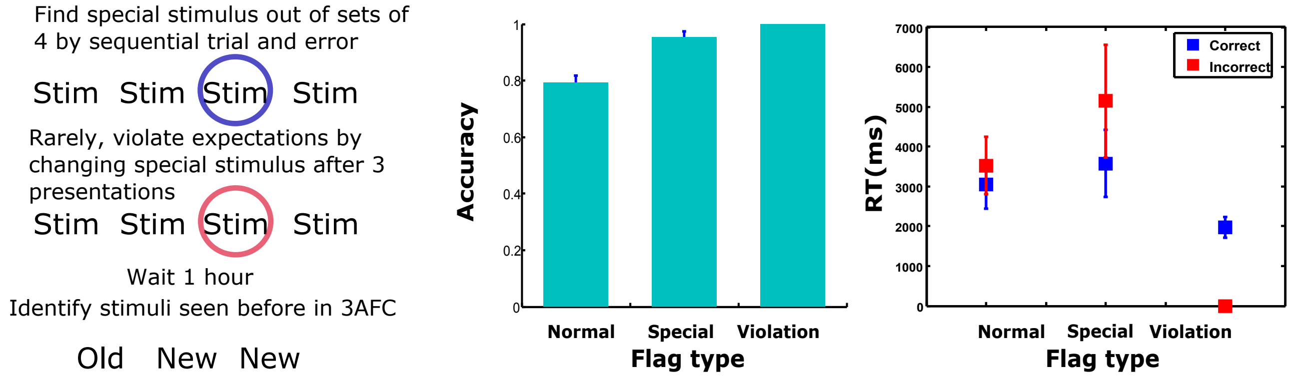

Figure 1: *(Left)*: Subjects played a game identifying special images from sets of 4, then had to recognize all images from this set as 'seen' in a 3AFC task 1 hour later. We surprised subjects within the first leg of the experiment by occasionally changing the identity of special images after the participant had identified them.*(Middle)* Recognition accuracy was high overall, but was significantly higher for images corresponding to the surprising manipulation in the first leg of the experiment. *(Right)* Such images were recognized both with perfect accuracy and with very short response times, indicating that they were easierto retrieve. All error bars represent s.e.m.

specify any criteria for recognizing the *special* flag, participants were expected to use trial-and-error to identify the correct flag. However, we imposed a restriction that they could only sample one flag per visual presentation of a stimulus set. Thus, the experiment cycled through all 6 image sets repeatedly (and in random order within the 6) until the participants could identify and recall the correct flag for each of the 6 image sets.

The spatial location of the four images was randomized on each presentation, so participants had to pay attention to the content of the image, not just the location. A rational strategy would be to remember which images one has sampled before in a particular set, and trying different ones each time, until the correct response is identified. After that, only the correct image need be remembered per image set. Into this setup, we introduced a manipulation - on rare occasions (once out of every 10 times participants picked a correct response), the game flipped the correct response for an image set after the participant had already picked it out correctly more than 3 times (including the first instance when they discovered it to be special). We expected this violation of the game's mechanics to surprise participants, and wanted to see how this surprise corresponded with subsequent recognition memory.

To this end, after a 1 hour delay and without forewarning, participants completed a second phase of the experiment, wherein they were asked to recognize the 24 images they had seen in the first phase in a series of 3AFC presentations, with each target image paired with two foil images from the same flag dataset. The surprise-sensitivity hypothesis predicts that memories of surprising events would be easier to remember. Thus, we expected that participants would make fewer mistakes recognizing images associated with induced expectation violations, and that they would recognize them faster, than other 'special' images in the game.

We calculated the performance on the recognition task (proportion correct and mean RT) for each type of stimuli: (1) normal stimuli unattended during the game, (2) special stimuli that were supposed to be remembered during the game, and (3) violations – viz. stimuli that violated participants' expectations about the game mechanics by appearing to be the correct answer, and then switching. Figure 1 summarizes our results. As the central panel shows, participants exhibited very good recognition accuracy for all images, in line with previous observations on the surprisingly good fidelity of visual recognition memory (Brady, Konkle, Alvarez, & Oliva, 2008). Even so, there were clear differences: special images were recognized more accurately, and special images that coincided with our surprise manipulation were recognized with perfect accuracy, significantly better than special images seen without expectation violations (one-tailed T-test $p < 0.01$). The right panel shows that such images were also recognized quicker than other special images ((one-tailed T-test $p < 0.001$)). In conjunction, these results present compelling evidence that stimuli associated with surprising events are easier to remember.

Surprising events disproportionately influence future choices

Demonstrating that surprising events are more accessible in memory justifies the particular definition of salience we adopt for our memory model. To demonstrate that the memory prior influences preference formation in the manner our model predicts, it is necessary to test whether surprising prior experience influences subsequent preference construction in an experiment set up such that the relationship between prior experience and future outcomes is empirically observable, without being trivial.

60 UCSD undergraduates participated in this experiment for course credit. This experiment took the form of a betting game, where participants bet on the outcome of simulated

soccer matches[2]. As an interesting sidenote, only three out of these 60 professed to follow soccer as a sport in general, reducing the potential for prior knowledge of the game significantly influencing aggregate measurements of participant behavior.

The game, schematized in the left panel of Figure 2, involved two levels: a junior level where participants bet on club soccer teams for small stakes, but with known team quality ratings (presented on a 1-100 scale), and a senior level where they bet on national teams for large stakes, but with unknown team quality ratings. We used 32 real club teams with notional quality ratings, 4 from each of 8 prominent soccer-playing countries. In both cases, the actual probability of either team winning was derived as a softmax of the quality ratings, but was not visually displayed. The true betting odds, likewise, were calculated as the inverse of the win probability, but were not displayed. Participants played the betting game beginning with a preset lump sum of \$1000 and were expected to continue playing until they had made \$10000. For each bet placed, we collected \$20 at the club level and \$200 at the country level as a betting free, independent of the outcome of the bet. After placing their bet, they were shown the outcome of the game, and either won or lost money based on the direction of their bet. Participants were instructed that they would not receive course credit if they ran out of their allocated betting budget, but that it was always possible to build up a depleted budget by switching back to the club level.

The link between the two levels lay in the mechanism of team construction, explicitly conveyed to the participants: countries form teams using players from the best national clubs. In practice, we calculated country quality ratings as the (max+avg)/2 of the clubs for each country. Participants were explicitly advised to learn country team ratings using experience gained while betting at the junior level, and were permitted to switch back and forth between levels. Participants were incentivized to bet at the country level by the size of bets they could place (\$1000 vs \$100 at the club level). Since the objective in the game was to win a fixed large amount of money (\$ 10000), betting correctly on the country level would minimize the total time taken in the task. Participants were disincentivized from playing at the country level without actually having learned the country quality ratings by the large betting fee at that level, and the possibility of not receiving course credit through running out of betting budget.

To account for participants who did not understand the game mechanism, we first identified players who had learned the mapping between club and country level team quality, which we operationalized as observing a correlation above 0.2 between country preference at the club level and country preference at the international level of the game.This procedure eliminated six participants. For the remaining participants (n=54), we postulated a generative process by which club level experiences x, measured as the absolute counts of the number of times club teams belonging to a country won minus the times they lost in the participant's experience, combined to construct country preferences for international level preferences y,measured as the absolute count of the number of times participants bet on that country minus the times they bet against it in the international matches. Specifically, we assumed that participants used linear combinations of previous experiences weighted by memory salience cues θ. Given this generative model, we use a Bayesian analysis to estimate θ^* that favor a linear relationship between x and y. Formally, we estimated

$$p(\theta)=\sum_{x,y} p(\theta,x,y)=\sum_{x,y} p(\theta|x,y)p(x,y)=\sum_{x,y} p(y|x,\theta)p(x)p(\theta), \tag{5}$$

where x and θ are independent by construction. Of the three empirical distributions needed to obtain the posterior θ estimate, we assume the prior distribution on θ to be uniform on the unit interval, $p(x)$ is the best fit student's t-distribution to the observed x values, and $p(y|x,\theta) \sim N(mx,\sigma^2)$, where $\{m,\sigma\}$ are obtained from the best fitting linear regressor of $w(x)$ and y, and $w(x)$ is obtained by weighting club experiences with different underlying probabilities, parameterized by θ as $w(x_p) = \theta(p)$,where x_p denotes a betting trial conducted with an underlying true probability of p. In particular, we stratify the probability domain into 10 partitions, with smaller partitions at the extremes of the range, and larger ones in the middle, using partition boundaries {0.05, 0.1, 0.2, 0.3, 0.5, 0.7, 0.8, 0.9, 0.95}, and permit each of the partitions to take a natural number weight $w = \{1,\cdots,6\}$. Next,we identified which weightings yield better linear relationships between the country selections and their weighted club level predictors, which we estimated via a brute-force grid search over all possible weighting functions,seeking to identify the weighting that minimized the residual error in the regression of x on y. If our generative model is correct, this analysis whould give us the memory weights on the probabilities of different club-level bets experienced in the past.

As the middle panel in Figure 2 illustrates, this memory weight function shows a striking U-shape, heavily privileging extremely rare events, also overweighting high probability events, and essentially ignoring events with intermediate probabilities. This bias pattern in memory retrieval indicates frugal sampling of memory in service of preference formation - observers selectively encode events with high information value, ignoring a very large majority of experienced events (Srivastava & Schrater, 2014). The bias in favor of rare, surprising events (an underdog beating a favorite) is irrational under value-of-information considerations in the context of the game, but is rationalized by the greater accessibility of such surprising memories, consistent with our earlier finding. The right panel in Figure 2 plots the probability distortion function from (a) our sample, determined by drawing 1000 samples of probability p using $w(p)$ normalized such that $w(1) = 1$ (clamping rightmost points to the true prob-

[2] A working version of this experiment is available online at http://experiments.evullab.org/fifa_betting/soccergame.html

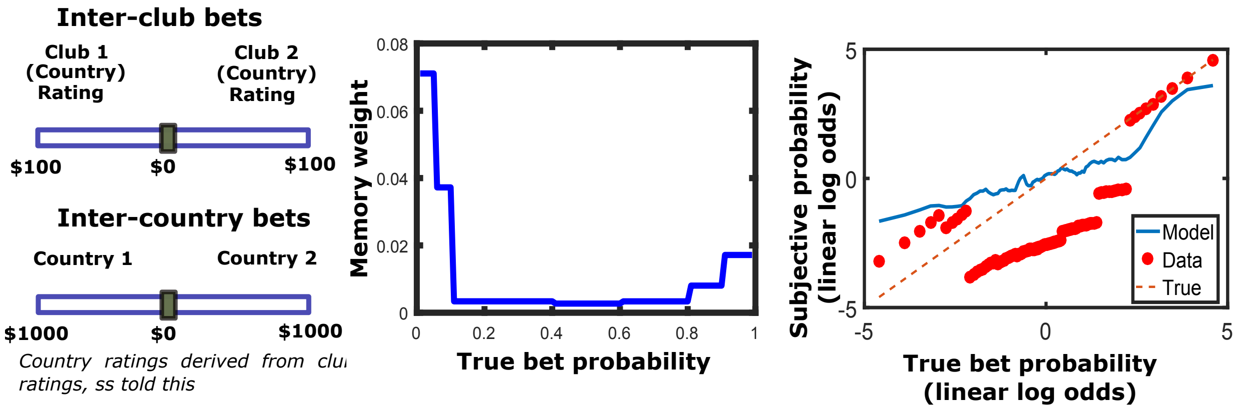

Figure 2: *(Left)* Setup for experiment 2. Participants bet on soccer matches, with small stakes and known team strengths at one level, and then with large stakes and unknown team strengths at a higher level. Their goal was to learn the strengths of the higher level teams, knowing that the higher level teams were made up of players from lower level teams. *(Middle)* Best fit memory weights for subject performance on this task, assuming that higher-level bets were constructed using a weighted linear combination of lower-level betting experiences.*(Right)* Aggregate subjective probability distortion implied by this weighting function from our experimental data compared with predicted subjective probability distortion from our surprise-based weighting model.

ability line), and (b) from our surprise-based overweighting model (model parameter $|\mathcal{M}| = 7$) for a generic binary choice task given a random 100 sample choice histories for each possible probability outcome, on a linear log odds scale (Zhang & Maloney, 2012). The lack of fit for intermediate probabilities suggests that our experiment participants essentially ignored these experiences, while making model-congruent predictions when using the outcomes they actually remembered.

Discussion

This paper originated from a simple observation: psychologically realistic observers cannot know the utility of a past experience until they have actually retrieved its corresponding memory. A random access model of memory would indicate that experiences that are more frequent become more likely to be retrieved. Observers with random access memory would retrieve memories of experiences in ways that reflect the statistics of their environment. Such a procedure, we show, would make the observer ignore the impact of statistically rare, but important experiences, since these would be seldom retrieved. Since extreme events are likely to affect the biological fitness of observers, we predicted that nature has to have found a way to incorporate the impact of such extreme events. We suggested one such mechanism: using surprise at the time of memory encoding as a retrieval-facilitator. Using this assumption, we constructed a model of memory encoding and retrieval that privileges surprising experiences at the time of encoding and showed that it reproduces the familiar inverse-S shaped distortion of subjective probability seen in multiple cognitive and perceptual tasks. Results from two novel experiments substantiate the two core assumptions of our proposal (i) that human observers find it easier to recall surprising memories and (ii) such surprising memories inordinately influence subsequent preferences.

Our results shed light on the psychological mechanims underpinning a large set of economic phenomena typically explained using prospect theory (Kahneman & Tversky, 1979). For instance, investors' overvaluation of stocks with right-skewed return distributions (indicating a small chance of large payoffs), aligns well with both the predictions of prospect theory as well as our account (Barberis, 2013). A similar explanation holds for the propensity of consumers in US health insurance markets to prefer policies with high premiums and low deductibles, likely overweighting the probability of making claims with respect to the average claim probability of the risk pool (Sydnor, 2010). Whereas the descriptive stylized fact of probability overweighting is sufficient to explain such phenomena, the psychological mechanism proposed in our account makes it possible to further test how such phenomena are influenced by changes in the evidence integration process. For instance, since memories of rare experiences remain rare even accounting for salience-based overweighting, artificially curtailing the integration process by adding time pressure or cognitive load should result in less risk-averse insurance decisions. Our account also predicts that interventions in risky decision-making will be effective either if they are very persistent or very rare and surprising. Testing these predictions presents an exciting opportunity for future work.

In these examples and in our experiments, the implicit probability of the underlying stake is unknown - people have to infer it from previous experience. These are, thus, decisions from experience, and yet show small probability overweighting, as predicted by prospect theory. Overweighting of rare events in such settings is apparently in conflict with work identifying an opposite pattern of distortions in decisions from experience - underweighting of rare events (Hertwig, Barron, Weber, & Erev, 2004). Our work supports the view of Fox and Hadar that the decisions-from-experience results are a function of sampling error - observers cannot overweight the probabilities of rare events they haven't yet experienced (Fox & Hadar, 2006).

Probabilistic sophistication - the separation of value and probability remains a foundational assumption in economic choice models (Machina & Schmeidler, 1992). And yet, for all its centrality in the design of economic observers, it has proved difficult to establish that humans are, in reality, probabilistically sophisticated. The possible localization of probability (Gold & Shadlen, 2002) and value (Schultz, Dayan, & Montague, 1997) estimation in the brain suggests such a separation to some (Gershman & Daw, 2012). Our proposal interacts with the assumption of probabilistic sophistication in a couple of interesting ways. On one hand, it falsifies some strong forms of the hypothesis. Peoples' subjective probability estimates are biased by salience cues, and extreme utility values can serve as a source of this salience, violating the assumption of epistemic separation of value and probability. On the other hand, our account supports a novel *procedural* version of this hypothesis. People *cannot* know the utility of an experience unless they encounter it, in reality, or in memory. In either case, knowledge of utility *succeeds* knowledge of probability. Thus a separation between probabilities and utilities is justified on purely procedural grounds.

In fact, the fundamental conclusion from this work is that separation between utility and probability is so profound that nature has had to work around it by biasing the process of subjective probability estimation using information-theoretic surprise. Since utility is not the only possible source of surprise, nature's development of this particular solution to handling tail risks has left humans incredibly susceptible to statistically unjustifiable decisions based on chasing shiny, bright objects (Folkes, 1988), commonly attributed to the use of an availability heuristic (Tversky & Kahneman, 1973).

Acknowledgments

Paul Schrater helped us design the football game experiment.

References

Anderson, J. F. (1996). ACT: A simple theory of complex cognition. *American Psychologist*, *51*(4), 355-365. doi: 10.1037/0003-066X.51.4.355

Ashby, F. G., Tein, J.-Y., & Balakrishnan, J. (1993). Response time distributions in memory scanning. *Journal of Mathematical Psychology*, *37*(4), 526–555.

Barberis, N. C. (2013). Thirty years of prospect theory in economics: A review and assessment. *The Journal of Economic Perspectives*, *27*(1), 173–195.

Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron*, *74*(1), 30–39.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2012). Salience theory of choice under risk. *The Quarterly journal of economics*, qjs018.

Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, *105*(38), 14325–14329.

Dayan, P., & Niv, Y. (2008). Reinforcement learning: the good, the bad and the ugly. *Current opinion in neurobiology*, *18*(2), 185–196.

Folkes, V. S. (1988). The availability heuristic and perceived risk. *Journal of Consumer Research*, *15*(1), 13–23.

Fox, C. R., & Hadar, L. (2006). " decisions from experience"= sampling error+ prospect theory: Reconsidering hertwig, barron, weber & erev (2004). *Judgment and Decision Making*, *1*(2), 159.

Gershman, S., & Daw, N. (2012). Perception, action and utility: the tangled skein. In M. Rabinovich, K. Friston, & P. Varona (Eds.), *Principles of brain dynamics: Global state interactions.* MIT Pres.

Gold, J. I., & Shadlen, M. N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, *36*(2), 299–308.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological science*, *15*(8), 534–539.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263-291.

Lieder, F., Hsu, M., & Griffiths, T. L. (2014). The high availability of extreme events serves resource-rational decision-making. In *Cogsci.*

Machina, M. J., & Schmeidler, D. (1992). A more robust definition of subjective probability. *Econometrica: Journal of the Econometric Society*, 745–780.

Montague, P., Dayan, P., & Sejnowski, T. (1996). A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of Neuroscience*, *16*, 1936-47.

Schultz, W., Dayan, P., & Montague, P. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593-1599.

Srivastava, N., & Schrater, P. R. (2014). Frugal preference formation. In *Cogsci.*

Sydnor, J. (2010). (over) insuring modest risks. *American Economic Journal: Applied Economics*, *2*(4), 177–199.

Todd, P. M., & Gigerenzer, G. (2000). Précis of simple heuristics that make us smart. *Behavioral and brain sciences*, *23*(05), 727–741.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, *5*(2), 207–232.

Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? optimal decisions from very few samples. *Cognitive science*, *38*(4), 599–637.

Zhang, H., & Maloney, L. T. (2012). Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition. *Frontiers in Neuroscience*, *6*, 1.