



THE UNIVERSITY OF  
**SYDNEY**

Economics Working Paper Series

2016 - 08

**Expected Subjective Value Theory  
(ESVT): A Representation of  
Decision Under Risk and Certainty**

Paul W. Glimcher & Agnieszka A. Tymula

April 28, 2016

# Expected Subjective Value Theory (ESVT): A Representation of Decision Under Risk and Certainty

Paul W. Glimcher\*      Agnieszka A. Tymula†

April 28, 2016

## Abstract

We present a descriptive model of choice that incorporates neurobiological constraints, representational structures and costs into a traditional economic framework. An individual's behavior is fully described by two, in principle observable, primitives: an individual's neural/mental capacity and an endogenous rational expectation. The model captures the phenomena captured by Prospect Theory: reflection in risk attitudes and loss aversion, but unlike Prospect Theory accounts for individual heterogeneity in each and employs fewer parameters. Our theory provides an alternative explanation for endowment effect and makes a series of novel predictions amenable to future testing.

JEL: D03, D81, D87

Keywords: utility, decision-making, reference point, neuroeconomics

## 1 Introduction

There is now growing evidence that during the last five years neurobiologists have succeeded in identifying the representational forms used by the human brain to encode decision-variables for causally producing choice. This class of representational forms has attractive near-normative features (Wainwright et al., 2002) and can predict a number of interesting human and animal choice behaviors that have so far escaped description in economics literature (Louie et al., 2013). In this paper we compare and contrast one of the simplest members of this representational class with the Prospect Theory of Kahneman and Tversky (1979).

---

\*Institute for the Interdisciplinary Study of Decision Making, New York University, 300 Cadman Plaza West, 7th floor, Brooklyn, NY 11201, USA

†School of Economics, University of Sydney, H04 - Merewether, room 370, Sydney, NSW 2006, Australia. Phone: +61 2 9351 2840, Fax: +61 2 9351 4341. Tymula acknowledges the support from the AXA Research Fund and Australian Research Council (ARC-DE150101032) while working on this project. Send correspondence to: agnieszka.tymula@sydney.edu.au.

The model we present assumes only two primitives: individual-specific neural mental constraints on capacity and a rational expectation about the environment. Unlike in most economic models of choice, attitudes towards risk and losses are not the primitives of the model. Instead they emerge from, and vary with, changes in individual capacity constraints and rational expectations. We find that our neurobiologically and theoretically identified representational form can account for many of the phenomena described by Prospect Theory as well extending its descriptive domain beyond some of Prospect Theory's limitations. Importantly, it also intrinsically captures heterogeneity in individual behavior, which neither Prospect Theory nor Expected Utility Theory capture without extension. Finally, it makes a number of novel testable predictions which differentiate it from Prospect Theory in a positive sense.

## **1.1 Model primitives**

There exists extensive evidence from the literature in economics and psychology that the primitives of our model - rational expectations and individual capacity constraints are key determinants of choice.

### **1.1.1 Rational expectation**

Abundant evidence shows that individuals evaluate options not in absolute terms but rather relative to some reference point, at least under some conditions. This reference point was initially conceptualized as closely related to the status quo (Kahneman and Tversky, 1979). More recently, plentiful evidence from behavioral economics has suggested that the reference point may be better described as a rational outcome expectation.

Rational expectations appear to affect individual behavior in a wide range of settings. In labor markets, for example, workers have been shown to adjust their effort according to how their actual wage compares to the wage that they expected to receive (Abeler et al., 2011, Mas, 2006, Crawford and Meng, 2011). In tournaments, workers adjust their effort to a rival's effort as if guided by an expectations-based reference point (Gill and Prowse, 2012). People's willingness to pay has also been shown to be affected by rational expectations in financial markets, auctions and in everyday consumption decisions. The history of stock prices (e.g. Baucells et al. (2011)) and the history of personal gains and losses in the market (e.g. Gneezy (2005)) affect trading behavior. An individual's bids in an auction has been shown to be best explained by an expectations-based reference-dependent preference (Lange and Ratan, 2010, Banerji and Gupta, 2014, Rosato and Tymula, 2016). Even daily decisions about whether to buy everyday consumption goods depends on price histories (Putler, 1992, Jacobson and Obermiller, 1990, Mazumdar et al., 2005); consumers are believed to evaluate the price of the good against a history-dependent reference price when deciding whether to make a

purchase. The strength of the famous endowment effect has also been shown to be affected by the probability with which people expect to keep their endowment (Marzilli Ericson and Fuster, 2011). This finding highlights the fact that it is not the physical ownership of the good, but rather it is a rational expectation about ownership that matters for valuation, a point to which we return later. Finally, Bartling et al. (2015) show that even professional soccer players and their coaches exhibit behavior during matches consistent with rational expectations. And their fans are affected by rational expectations as well. When sports fans' expectations regarding their favorite team are not met, they become more violent (Card and Dahl, 2011).

The importance of the reference point is today well established, but one of the biggest challenges that economists face in formalizing the theory of reference-dependent behavior stems from the fact that the reference point is unobservable using traditional techniques. Three different approaches have been taken to address this problem. Scholars have estimated reference points from observed choice (Post et al., 2008, Bartling et al., 2015, Card and Dahl, 2011, Gneezy, 2005, Baucells et al., 2011), hypothesized about how the reference point could be integrated into choice (Kőszegi and Rabin, 2006), and derived whether and how rational expectations should affect valuation in normative sense (Frederick and Loewenstein, 1999, Robson, 2001, Rayo and Becker, 2007, Netzer, 2009, Woodford, 2012). The overwhelming conclusion of the majority of these inferential studies is that the utility function is, in some sense, dynamically adjusted in response to rational expectations. Of course, without observing the reference point directly, one cannot know exactly how it enters, or modifies, individual's utility function.

However in neuroscience, which relies on a different set of observables, a tremendous amount of directly observed data is available on rational expectations and reference points - particularly in the domain of sensory perception. Detailed differential equation-based models of the neuronal networks that compute and represent rational expectations over upcoming sensory events like the brightness of an environment or the speed and direction of a moving object are now available. A description of these mechanisms can be found for example in Carandini and Heeger (2012). Measuring neuronal responses in a cat's visual cortex to different levels of a visual stimulus contrast Ohzawa et al. (1985) showed that the activity of neurons in the visual cortex is conditional on a rational expectation of future contrast. Interestingly, these physiological responses of neurons show a reference dependence that precisely matches the manner in which normative models from economics predict that rational expectations should modulate the utility function (Rayo and Becker, 2007, Netzer, 2009, Glimcher, 2010, Woodford, 2012) - they adjust response magnitudes dynamically to maximize the precision of their capacity-limited representation around the rational expectation of future contrast.

Of course, rational expectations in decision-making may be more complicated than in

sensory systems (Tymula and Plassmann, 2016). In this paper, we do not focus on how the rational expectation is formed. Instead, our focus is on a functional form that is novel in economics, but grounded in years of study by neuroscientists on the way in which rational expectations enter response functions as applied to the domain of economic choice.

### 1.1.2 Neural constraints

The first mention of the other primitive in our model, capacity constraints, goes back to Herbert Simon. In the 1950s, he famously hypothesized that many of the idiosyncratic features of human choice behavior that Prospect Theory would later capture might well arise from the limits of human cognition (Simon, 1955). He proposed that our evolutionary history had placed limits on our cognitive abilities. He suggested that apparent failures to maximize might represent an implicit tradeoff between the high biological cost of having a brain large enough for perfect maximization, and the diminishing marginal returns that ideal brain would yield as it grew (Glimcher, 2010).

Over the course of the last decade, there has been renewed interest in the notion that our growing biological understanding of the neurobiological costs of maximization might be tractably incorporated into a new class of economic models. Neuroeconomists in particular have embraced this idea, exploring the possibility that both inspiration and insight from neuroscience might allow us to define new models that increase both the positive and the normative reach of economics (Sims, 2003, Glimcher, 2010, Woodford, 2012, 2014).

More specifically, a variety of studies in economics and neuroeconomics point to the important role that neural resources play in decision-making. Theoretical studies in economics have begun to employ neural constraints as a modeling assumption to derive optimal properties of utility for a neurally-constrained individual (Rayo and Becker, 2007, Netzer, 2009, Woodford, 2012). Interestingly, all of these models theoretically derive S-shaped functions that adapt to changes in the environment as optimal. Empirical studies have now directly demonstrated the link between neural constraints and decision making. For example, evidence from neuroeconomics suggests that in principle, individual decision-making properties such as attitudes to risk can be inferred from neuroanatomy (Gilaie-Dotan et al., 2014, Kable and Levy, 2015). Differences in brain structure are believed to be the main cause of excessive risk taking in adolescence (Steinberg, 2007). What is even more interesting is that the gray matter volume (the number of brain cells in a particular functional module, a proxy for neural capacity) is more predictive of individual risk attitudes than standard predictors traditionally used in economics such as age (Grubb et al., in preparation). This of course makes sense if one assumes that age serves as a proxy for neural capacity which generally declines as we get older but at an individual-specific pace.

In our model, the neurobiological constraint on capacity serves as the second primitive that determines behavior. Just as Herbert Simon hypothesized, individual attitudes towards

risk and losses originate in our model from biological constraints in representing value. One advantage of such approach is that it allows us to start exploring the causal link between heterogeneity in neurobiology and behaviour.

## 1.2 Why a neuroeconomics approach?

Many theories have, over the last half-century, attempted to extend Expected Utility in a positive way to capture the widely observed violations of that Theory, but citation rates suggest that one of these descriptive theories dominates the academic discourse: the Prospect Theory of Kahneman and Tversky (1979). Kahneman and Tversky (1979) captured these violations by proposing a novel variant of the utility function they called the value function and an extension of Savage's subjective probability they called the probability weighting function. The value function captured, descriptively, many of these violations with a reference point, a reflection of utility curvature around this reference point (the reflection effect) and scaled asymmetry around the reference point (loss aversion). The probability weighting function captured residual violations of the independence axiom with an additional one or two parameter function. These three key elements (the reference point, the value function and the probability weighting function) have become so deeply embedded in academic discourse that they are often taken as ground truth at a mechanistic level rather than as the useful descriptive functions Kahneman and Tversky meant them to be.

Despite the obvious power of Prospect Theory, the possibility remains that the variance Prospect Theory fails to capture may someday be explainable with other approaches. One obvious limitation of Prospect Theory is that it does not explain heterogeneity in individual decision-making but rather captures individual heterogeneity by variation in each of its fit parameters. In general, in economic models heterogeneity in choice does not arise directly from the model, even though empirical work documents a range of reliable factors that correlate with observed tolerance for risk or asymmetry in valuation of gains and losses. In our model, in contrast, the key variables are not parameters per se, but rather individual-specific features that are, at least in principle, directly observable.

Further, Prospect Theory does not always successfully capture variation in choice over different contexts, even with its many parameters. Diminishing sensitivity with respect to the reference point (which leads to risk aversion in gains and risk seeking in losses) and loss aversion are assumptions of the model and thus they appear in the model unconditionally and independent of context. The generality of these properties, although often taken for granted, has been challenged in the empirical and theoretical literature. Indeed, loss aversion as traditionally parameterized, though widely observed, is notoriously variable across studies and contexts (Ert and Erev, 2013). The phenomenon of risk seeking in losses has eluded robust description and prediction at the single agent level (Cohen et al., 1987, Tymula et al.,

2013). And even the endowment effect is predictable only at the most qualitative level under most circumstances (Plott and Zeiler, 2005, 2011, List, 2003, Heffetz and List, 2014).

The instability of parameters across studies, failures to replicate, theoretical and empirical controversies, and the existence of hundreds of competing theories all suggest that there is still room to develop new positive descriptions of human choice behavior. And of course, the lack of a normative foundation in Prospect Theory, the implication that choosers are fundamentally inconsistent, is also something that many find troubling. We are able to address all these issues with our approach in a novel way that may be of general interest.

We find that normalization models from neuroscience can account for all of the phenomena captured by Prospect Theory. Further, these functional forms account for chooser heterogeneity as arising from endogenous differences in neural capacity and endogenous rational outcome expectations. Even more importantly, the functional form makes novel predictions about behavior that could, in the future, be empirically tested but which has not yet been examined.

In the next section, we develop in detail the neurobiological foundations of this functional form. (A reader uninterested in the foundations of our functional form, or uncomfortable with the engineering-like approach of computational neurobiology, might choose to skip this section and to proceed directly to section 3.) We begin with early theoretical work on the normative features of efficient representation in a capacity-limited brain. We then review evidence that the primate brain (both human and monkey) actually employs these classes of representations in the brain structures that make our decisions. Next we review behavioral evidence from humans and animals that supports the hypothesis that normalization-based models built on our two primitives can be used to make novel behavioral predictions that are fulfilled. We conclude with a presentation of the more detailed functional forms used by neurobiologists.

In section 3 we present the model and explain the simplifications employed for our analysis. Then we present the results of our mathematical analysis on preferences. Here, we set ourselves the task of asking a simple question: If human choosers represented, in their brains, the values of prizes using the same class of transformation that they are known to use in their sensory systems, and then performed an argmax operation on these subjective values, what kinds of preference structures would one expect to see in human choice behavior? Would one see phenomena like those Kahneman and Tversky sought to explain with Prospect Theory? Could such a representational theory capture the idiosyncrasies of human choice behavior captured by the Kahneman and Tversky value function, but in a completely different way? Further, could such a representation outperform Prospect Theory as a descriptive tool while opening a novel normative approach?

In section 4 we discuss the novel aspects of our modelling approach and suggest future tests and extensions of the model.

## 2 Normalization Model-Based Representations of Value

### 2.1 Well-Studied Neurobiological Theories of Representation: Statics

The study of how information is represented in the brain largely begins with the study of how sensory systems (the neurobiological mechanisms underlying sight, touch, hearing, taste and smell) encode and represent properties of the outside world (Ratliff and Hartline, 1959, Barlow, 1961, Attneave, 1954). Consider, to begin, how the brain might represent an image of a sunset presented on a black and white computer monitor's 1000 x 1000 grid of 1,000,000 pixels.

One obvious possibility might be that a single nerve cell in the brain's visual system would be devoted to 'conveying' a value for each of the 1,000,000 pixels, thus requiring 1,000,000 nerve cells to convey to other brain areas the content of the picture. Nerve cells convey information by transient electrochemical impulses called action potentials, and it is the rate at which these action potentials are generated that is the variable by which information is conveyed in the brain. To make this example concrete, we might imagine that each nerve cell would represent that its pixel was black by producing (or firing) one action potential per second and that its pixel was white by firing at a biophysically maximum rate of 100 action potentials per second. Intensities ( $I$ ) between these two extremes would be represented by a continuous monotone function ( $F$ ). Formally,  $I \mapsto F \in [0, 100]$ , where  $F$  is the firing rate (the number of action potentials produced per second).

In such a situation a perfectly white computer screen would cause all 1,000,000 neurons to fire 100 action potentials per second. In this way each pixel is associated with an independent neuron that has an independent firing rate and no information about the pixels in the display is lost by this transformation to firing rates.

What neurobiologists realized in the 1960s and formalized in the 1990s, however, is that adjacent pixels in real world images are non-independent. To get an intuition of why this is so, consider any pixel representing part of a typical real-world image. If you look closely at any real image you will empirically observe that adjacent pixels are in fact highly correlated (see Figure 1). To put that in more formal terms we can describe each image as a set of 1,000,000 pixels with the intensity of each pixel  $i$  labelled as  $I_i \in \mathbb{R}_+$ . The set of all possible images forms a set  $A$ . What a study of natural images reveals is that the set of all observable images,  $B$ , is much smaller than  $A$  ( $B \subset A$ ) and is marked by high adjacent pixel correlations (as well as a number of other features: see Olshausen and Field (1997)). Empirical measurements of huge photographic databases estimate adjacent pixel correlations to be approximately 0.7. This non-independence is an incredibly important point. Consider

the nine adjacent pixels shown in Figure 1. If one knew in advance that the eight outer pixels were white and the adjacent pixel correlation was 0.7, one would know with better than 99.9% certainty that the central pixel was also white. And of course that means that under these conditions the activity of the neuron encoding that central pixel is almost entirely redundant.

Figure 1

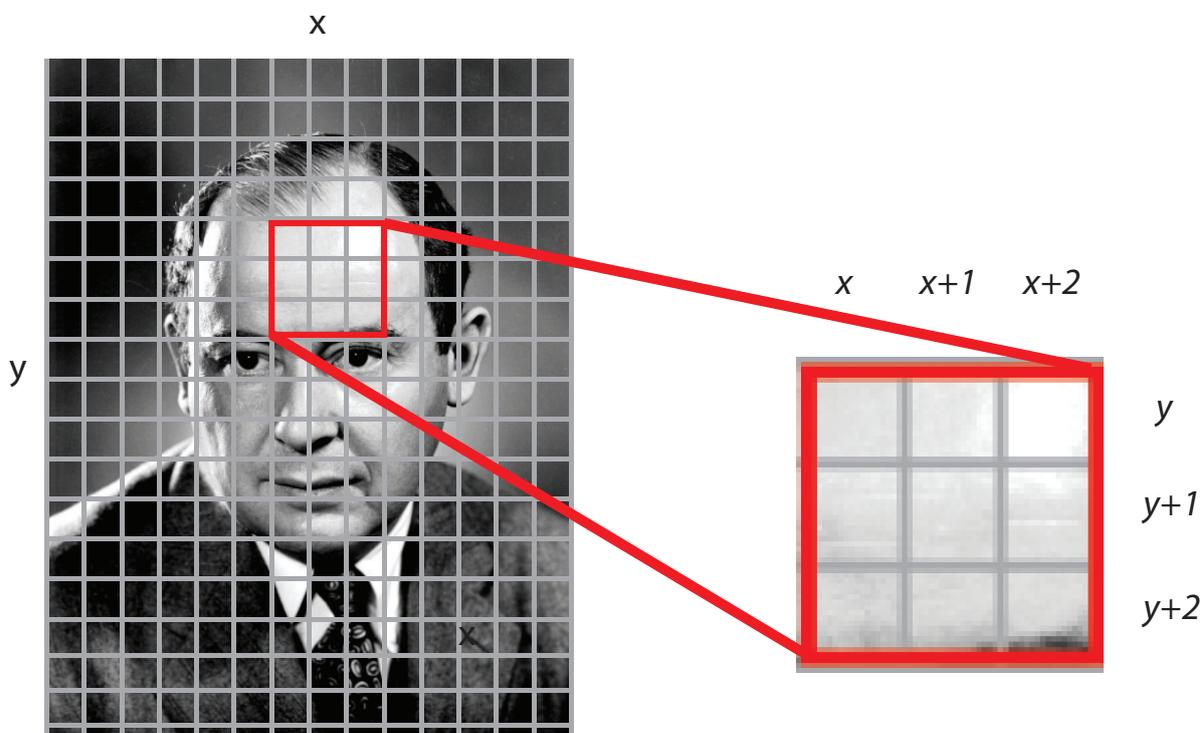


Figure 1: An array of local pixels from a photograph.

The reason that this is so important is that action potentials are incredibly costly in terms of calorie consumption. The cost of neural activity (measured in action potentials per second) is in fact monotone in calories, and almost certainly near linear (Lennie, 2003). While a calorie may not seem much at first glance, these costs are surprisingly significant. Nearly 20% of the calories we consume go to support our brains even though it accounts for only about 3% of our body weight. Brain tissue is on average about 7 times as expensive as the average tissue in our bodies. Put another way, were we to increase total representational capacity of our brains by a factor of ten (either by increasing the number of neurons by 10 or expanding the range of possible firing rates to  $[0,1000]$ , total caloric consumption would have to increase by roughly 200%. Consider the approximate marginal cost structure that imposes. For a factor of 10 scaleup in precision to be efficient in the economic sense, it would have to yield a tripling of total consumption wealth. The costs of a factor of 100 scaleup are almost unimaginable.

Faced with these high costs of representation, a representational schema that uses fully independent neuronal firing rates to encode non-independent pixels intensities is absurd. Recognizing this fact, neurobiologists working in the later 1990s and early 2000s began to ask what forms of representation would be maximally efficient: What representational forms would maximize Shannon Information (Shannon and Weaver, 1949) in the aggregate of action potential rates conditional on the correlational structure of the observed sets  $B$  (Heeger, 1992, Heeger et al., 1996, Reynolds and Heeger, 2009, Schwartz and Simoncelli, 2001). The general answer to that question is that neurons employ a representation of the general form:

$$F_i = \frac{I_i^\alpha}{M^\alpha + \sum_j w_j I_j^\alpha}$$

where  $F_i$  is the firing rate of the neuron encoding the objective intensity of pixel  $i$ ,  $I_i$  is the intensity of pixel  $i$ ,  $M$  is an expectation over the intensity levels in  $B$ ,  $w_j$  is a weight which captures the empirically measured ex-ante correlations between the pixel at location  $j$  and other pixels across  $B$ , and  $\alpha$  is an adjustable set-wide capacity that governs the total number of action potentials the set of all neurons is allowed to produce.<sup>1</sup> In essence, what is happening here is that the denominator is serving to remove from the firing rate of pixel  $i$  information already carried by other neurons, thus maximizing information per action potential.<sup>2</sup>

In 2000s a number of papers examining this general formulation demonstrated that for any observable set  $B$ , there exists an  $M$  and a set of weights  $w_i$  such that Shannon Information

---

<sup>1</sup>The notion that the  $\alpha$  term captures a total capacity across the set of neurons is really quite simple to see as was first developed by Webb et al. (2014). Rearranging a bit for clarity and following the derivation in Webb et al. (2014) we can say for a two neuron case:

$$F_1 = \frac{I_1}{(I_1^\alpha + I_2^\alpha)^{1/\alpha}}$$

$$F_2 = \frac{I_2}{(I_1^\alpha + I_2^\alpha)^{1/\alpha}}$$

When  $\alpha = 1$ , then it always holds that  $F_1 + F_2 = 1$ . So an increase in  $F_1$  must be linearly offset by a decrease in  $F_2$  and vice versa – a tightly binding constraint. At  $\alpha = \infty$  aggregate firing rates are only limited by the bound  $F_i \in [0, 100]$ . At intermediate levels of  $\alpha$ , capacity is constrained in intermediate ways. With  $\alpha = 1$ , the set wide limit on number of action potentials per second is effectively 1. This reveals how intense capacity restrictions can be and how large their impact is on the information-carrying capacity of the system. For more detailed discussion see Webb et al. (2014).

<sup>2</sup>Following the conventions in this field, we neglect here the fact that firing rates are biophysically restricted to be weakly positive, though we have drawn attention to the bounded range. This will turn out to be an important point later as we begin to cardinalize our representation around a unique zero point, but one to which we will devote little attention here. The full model that incorporates this unique zero into firing rates is:

$$F_i = \frac{I_i^\alpha + \beta}{M^\alpha + \sum_j w_j I_j^\alpha}$$

where  $\beta$  serves as the unique zero (baseline firing rate) and is an observable that has been widely studied and modeled empirically at the level of neurobiology and biophysics (Reynolds and Heeger, 2009, Louie et al., 2011).

across the firing rates of the neuronal set is maximized in the neuronal rate (Wainwright et al., 2002). Put another way, these proofs indicate that for any system that has monotone costs over number of action potentials, a representational schema of this general form is maximally efficient at any level of capacity.

## 2.2 A Note About Neuronal Stochasticity

Throughout the preceding discussion, neuronal firing rates have been treated as deterministic. Some readers may have found this slightly troubling. If the normalization equations map intensity to firing rate along a segment of  $\mathbb{R}_+$  between 0 and 100, in what sense can we think of the information-bearing capacity of these neurons as limited? Given infinite precision, the restriction that at  $\alpha = 1$  the total aggregate firing rate of the neuronal set is capped seems trivial.

Neuronal firing rates are, however, non-deterministic and this is a critical point to which much study has been devoted. In fact, firing rates have been known for decades to be roughly Poisson-distributed.<sup>3</sup> To a first approximation this means that mean firing rate and variance are correlated at roughly 1.0. If we say that the mean observed value of  $F_1$  is equal to 10 then the observed variance of  $F_1$  is approximately 10 as well. In a more rigorous treatment of the firing rates in the foregoing section, the equations would have been presented as descriptions of means with specified variances – a feature which has received tremendous attention in the neurobiological and neuroeconomic literatures and which is also related to the costs of producing each action potential (see Glimcher (2005)).

For our purposes this is an important point because it makes the limited capacity of the system hugely problematic in finite time. If the entire set of neurons encoding an image has only an average aggregate action potential rate of 100 to spend and each element (neuron) has a variance proportional to its mean, then the information carried by the set can indeed be quite limited. For an economist it is probably useful to then think about these kinds of representations through the lens of random utility models where each neuron encodes a mean value plus a variance term. More on this approach to neural theories of representation can be found in Webb et al. (2014). For the purposes of this summary we now neglect this point.

## 2.3 Derivative Representational Theories for Decision-Variables: Statics

Now consider the implications of the normative finding that, for any observable set  $B$  there exists  $M$  and a set of  $w_i$  such that Shannon Information across the firing rates of the neuronal

---

<sup>3</sup>The actual distribution has been very well studied and is an extreme type II distribution with interesting properties, but is beyond the scope of this discussion. See Kiani et al. (2008) for more.

set is maximized in the neuronal rate (Wainwright et al., 2002). Let us begin by imagining that we were directly observing physical neurobiological system designed to represent the cardinalized and unique form of utility (usually called subjective value, (Glimcher, 2011)) for  $n$  elements of a choice set  $X$ , each denoted with  $x_i$ ,  $i \in [1, n]$ . We might imagine (for simplicity) devoting a single neuron to representing the subjective value of each element,  $x_i$ , in the choice set. Let us begin by assuming a function that maps some (or any) concept of ‘value’ to the firing rates of these  $n$  neurons. Recall that neuronal firing rates are bounded between 0 and 100. We could imagine a simple monotone (or even linear) function that maps all possible subjective values to the finite range of realizable firing rates in each of the neurons devoted to representing the elements of this choice set. In this representation, neuronal firing rates simply represent monotonically, or even linearly, whatever theory of utility or subjective value one proposes.

To ask whether the theory one proposes is efficient, given the finite representational capacity of these systems, we would also have to know something about the structure of the set of observable choice sets, just as we had to know something about  $B$  in the previous example. Consider the set of all possible prizes which in the real-world includes all possible prizes from candy bars to cars. From this we can construct the set of all possible choice sets of all possible sizes,  $X$ . For simplicity, let us consider for now only a subset of  $X$ , the subset that defines all possible 4-element choice sets which we redesignate  $X_4$ .

What should be immediately clear from the example above about vision, however, is that there may be something very unrealistic about  $X_4$  for an empiricist.  $X_4$  contains choice sets like: {snickers, lamborghini, volkswagen, hershey} as often as it contains sets like {snickers, hershey, bounty, twix} and {lamborghini, volkswagen, porsche, toyota}. What if, at an empirical level, the choice sets consumers actually encounter with finite probability, which we might call  $Y_4$ , were a very small subset of  $X_4$  that had some interesting structure to it? Put more crudely: what if  $Y_4 \subset X_4$  showed a high degree of internal correlational structure just as  $B \subset A$  did in our visual example? If you knew that 3 of the elements of a given real-world 4-element choice set offered to you by a colleague were candy bars, would this reveal to you anything at all about the most-likely properties of the fourth element of the choice set? Before you rush to say “no”, what if Reinhart Selten really reached into the roller-bag he always has with him to pull out 4 things from which you could choose one, and he had pulled out snickers, hershey, bounty? Everyone in the real-world knows that the likelihood that Herr Selten is about to pull a Lamborghini out of his little roller-bag to complete this choice set is as close to zero as one can ever get. Real choice sets have structure. That has to be true. Although it is also true that we do not know what that structure is. Of course the point we are trying to make here is that every intuition we have is that just as in the case of vision, the set of all observable choice sets must have a high degree of internal correlational structure, a fact to which economists have paid little attention in

the past. Thus for a stochastic system with limited capacity, one would want to distribute firing rates across the values encountered in a choice set so as to minimize choice errors - a distribution accomplished optimally by normalization-based forms (Wainwright et al., 2002).

## 2.4 Direct Measurements of Subjective Value

For neurobiologists studying how we see, tools exist for assessing the independence of adjacent pixels in huge sets of existing images. In one study, for example, Simoncelli and colleagues aggregated a database of roughly 50,000 photographs of natural scenes and used these images to exhaustively compute the pixel-by-pixel correlations in natural images (Simoncelli and Olshausen, 2001, Simoncelli, 2003). Unfortunately, using such an approach to the set of all observed consumer choice sets, is problematic. Fortunately, however, another approach to this problem is possible.

Recall that for a neuroeconomist, subjective values are by definition (unlike the utilities to which they are related) both causally responsible for choice and directly observable. Using a variety of neurobiological tools it is now possible to measure directly, or to infer quite precisely, the activity levels of neurons encoding the subjective values of options in single choice sets actually being presented experimentally to choosers. Given that this is possible, rather than asking whether choice sets have correlational structure, one can ask whether the representation of option value in the brain follows a normalization-based representation.

The first effort to do that was undertaken by Louie and colleagues in 2011 (Louie et al., 2011). They studied the activity of single neurons in monkey subjects making choices over two and three element choice sets while systematically varying the structure of those choice sets in a way designed to reveal normalization in value representation if it occurred.<sup>4</sup> *They found that the firing rates of the neurons encoding each element in a choice set were in fact non-independent and were extremely well described with normalization models.* From these findings, they concluded that subjective value representations appear to be normalized in exactly the way one would expect if choice sets they sought to represent efficiently showed deep correlational structure.

Once it was clear that the neural instantiation of subjective value appeared to follow a normalized representation, a number of scholars began to ask whether these normalized representations made any unique behavioral predictions. Speaking broadly, these theoretical investigations have suggested that normalized representations should lead to very specific violations of the axiom of regularity (Becker et al., 1963), a conclusion that has been tested successfully in both monkeys and humans in choice sets of varying size and type (Louie

---

<sup>4</sup>They performed these initial measurements on monkeys because it allowed the use of invasive technologies that yield measurements of subjective value virtually without measurement error. Similar studies have subsequently been made in humans (Hunt et al., 2014, Webb et al., 2012) using techniques with much lower signal to noise ratios and similar results have been obtained.

et al., 2013, 2015, Webb et al., 2012). The effects of choice set size on stochasticity in choice have also been examined theoretically and empirically (Louie et al., 2013) and once again these studies have found evidence that normalization class representations can go far in explaining many classes of choice behavior. To summarize this literature, one can draw both the conclusion that neural representations of subjective value follow the predictions of the normalization models and the conclusion that the behavioral predictions made about normalization-class representations are observed. And while one must acknowledge that this cannot be taken as proof of anything, it does raise the possibility that these functional forms may be of interest to economic scholars of choice.

## 2.5 Dynamic Studies of Normalization: Theory and Empirics

It may be important to note, in support of the notion that these functional forms are important to choice theorists, that how networks of neurons biophysically generate the transformation specified by the normalization equations has also received significant attention. This is an important point because studies of the dynamic computation that gives rise to the efficient representation have provided important insights into the normalization equations themselves at a positive and at a normative level.

Current evidence suggests that each prize in a choice set is represented neurobiologically by what can be schematized as two dynamically interacting neurons: an excitatory and an inhibitory neuron. For simplicity, we now focus on the representation of a single prize in a degenerate one prize choice set – a situation which will be our primary focus in this paper. In the existing literature (e.g. LoFaro et al. (2014)) the function that represents the action potential production in response to objective input  $x$  is presented in the excitatory neuron as the function  $R$  and in the inhibitory neuron as the function  $G$ . These functions are known to be described by the following first-order differential equations:

$$\tau \frac{dR}{dt} = -R + \frac{x}{1+G}$$

$$\tau \frac{dG}{dt} = -G + R$$

where  $\tau$  is the time constant of neuronal information integration and  $x$  is the objective value of the prize. LoFaro et al. (2014) showed that in networks of this type, for any given  $x$  the network always has a unique equilibrium state:

$$R = \frac{x}{1+G}$$

There are two interesting things to note about this equation. The first is that the unique equilibrium state of the neuronal computation is the standard normalization equation we

have been presenting. The second is that the normalization emerges spontaneously at equilibrium from the formal integration of the value inputs over time by the system. Applying some additional simplification (LoFaro et al., 2014) we can represent the action potential rate at time  $t$  as:

$$R(t) = \frac{x_t}{\sum_0^{t-1} D(t)x_t + x_t + 1}$$

where  $D(t)$  is a weighting function such that the term on the left in the denominator ( $\sum_0^{t-1} D(t)x_t$ ) is a time discounted average of the prizes previously encountered by the system. In the language of reference-dependent theories of choice, *it is the endogenous rational expectation*. In the next section we show that its function is to adjust the curvature of the firing rate function in accordance with a rational expectation.

Empirically, these normalization functions have been used to model the time varying firing rates of subjective value encoding neurons in a way that goes beyond the statics described above (Louie et al., 2014, LoFaro et al., 2014). Perhaps unsurprisingly, neurons in a number of subjective value-representing areas (Louie et al., 2014, Padoa-Schioppa, 2009) have been found to be well described by these dynamic equations. This suggests that these equations might also be used to model dynamic effects of expectation changes on human choice behavior. Those studies are ongoing, but initial work suggests that the dynamics of the equations can, in fact, make novel testable predictions about the dynamics of choice.

## 2.6 Summary

In this brief literature review we have provided some insight into the functional form we examine in the following section. Early work showed that this class of functional forms can have normative qualities at steady-state; it maximizes Shannon Information per action potential under some conditions (Wainwright et al., 2002). Studies of how the brain actually represents subjective value in the circuits causally responsible for choice, show that these equations do an excellent job of predicting steady-state neuronal firing rates (Louie et al. (2011, 2015)). Econometric studies of representations employing these equations and of behavior suggest a number of unique preference structures that would be the product of representations having these features and many of those properties have now been observed in humans and animals (Louie et al., 2013, Webb et al., 2012, Louie and Glimcher, 2012). Finally, dynamic versions of these equations have also been developed (LoFaro et al., 2014), used to predict dynamic neuronal firing rates, used to better explain the normalization constants and used to predict novel choice dynamics that are now being studied (Louie et al., 2015).

One interesting feature of this approach, then, is that it raises the possibility that many behaviors which violate Expected Utility may not be failures to maximize as previously

thought. Expected Utility, of course, does not incorporate the costs of precision. Once we assume a strictly monotone cost for precision, it becomes obvious that no rational actor would employ infinite precision. This observation raises the possibility that some violations of EU might in fact be formally rational, trading the marginal benefits of additional precision against the marginal costs of additional precision.

In the next section we present our novel analysis of the normalization function with regard to risk preferences, and other related preference phenomena.

### 3 Model

Consider a set of prizes  $x \in \mathbb{R}_+$ . The subjective value of a prize  $x$  at time  $t$  is given by:

$$V(x) = \frac{x^\alpha}{x^\alpha + M_t^\alpha} \quad (1)$$

where  $\alpha > 0$  is the neural capacity available to valuation and  $M_t$  is the rational expectation, or reference point, at time  $t$ .  $\alpha$ , as noted in section 2, is a neurobiological variable that approximates how many neural resources the individual can allocate to valuation. A higher  $\alpha$  indicates a higher capacity, meaning that the individual has more resources that can be devoted to valuation.<sup>5</sup>

The reference point,  $M_t > 0$ , is the prize that the individual expects to encounter at period  $t$  based on the time-weighted average of past prizes up until period  $t$ , the biophysical and biological properties of which are described in the preceding section:

$$M_t = \sum_0^{t-1} D(t)x_t$$

$M_t$  can be interpreted as an experience-based expectation (a reference point) about average rewards in the environment during the recent past. Notice that the subjective value of the same prize  $x$  may be different at different points in time if the expectation ( $M_t$ ) changes.<sup>6</sup> Neurobiological data suggests that the weighting function  $D(t)$  is an exponentially weighted

---

<sup>5</sup>Here we focus on a simplified steady state form of value function that captures most of the features described in section 2. This form represents a single option,  $x$ , as in section 2.5 but with the standard exponential capacity term presented in section 2.1. We adopt the single option representation as canonical for ease of analysis. By examining how a single option is represented, we avoid the intractability of partial differential equations for the subsequent analyses. We retain the exponential term related to capacity limits because it turns out to allow tremendous generalization. In this way, we are able to focus on how the expectation term  $M$  and the capacity term  $\alpha$  influence preference structures in a way that captures much of the behavioral phenomena engaged by Prospect Theory. It also allows us to identify several novel behavioral predictions about preferences and agent heterogeneity that lie outside the predictive scope of Prospect Theory.

<sup>6</sup>Another source of different valuations of the same  $x$  over time is the naturally occurring stochasticity in the neuronal firing rates, a point also raised in the preceding section. For a detailed analysis on how neural noise affects choice under certainty see Louie et al. (2015), Webb et al. (2014) and Louie et al. (2013). A normally distributed additive noise term would not alter our results. Skewed error distributions, which are of interest, could generate new predictions which we leave for future research.

sum of previous offers as a function of time since that offer. To preserve generality, we do not specify the exact weighting function  $D(t)$ , but simply make the uncontroversial assumption that:

**Assumption 1.** *For all  $t$ ,  $D(t - 1) < D(t)$*

This assumption guarantees that when forming the reference point more recent prizes are given more weight than prizes in more remote past as is indicated unambiguously by the neurobiological data (e.g., Ohzawa et al. (1985)).

Both capacity and expectation are in principle measurable, although methods for measuring capacity are less well-developed today than for measuring expectation (although see Gilaie-Dotan et al. (2014) and Grubb (in preparation) for an example). The subjective value function takes as its argument  $x \in \mathbb{R}_+$ . The model thus allows for nominal losses after they are transformed to  $\mathbb{R}_+$ . The subjective value function takes values between 0 and 1 ( $V \in [0, 1]$ ) consistent with the idea that the decision-makers are biologically constrained in the range of values that they can assign to the rewards. For direct measurement purposes the neuronal firing rates are assumed to map linearly to subjective value.

The expected subjective value of a lottery  $(x, p)$  with possible prizes  $x_1, x_2, \dots, x_n$  that occur with corresponding probabilities  $p_1, p_2, \dots, p_n$  is calculated by taking a standard expectation:

$$EV(x, p) = \sum_{i=1}^n p_i V(x_i) \quad (2)$$

Figure 2 plots examples of subjective value functions for a range of capacity levels and reward histories. The model can produce a variety of familiar value functions: a value function concave for all  $x$  that looks just like the utility functions that economists have been using since the XVIII century, when these functions were introduced by Bernoulli, as well as more modern S-shapes used in Prospect Theory and other reference-dependent models.

To give the reader a quick intuition for the role of the capacity and expectation parameters in the subjective value function, in Figure 2 A we manipulate capacity while keeping the expectation constant and in Figure 2 B we manipulate expectation while keeping the capacity constant.

In Figure 2 A we see that capacity fundamentally affects the curvature of the value function. When the decision maker has only very few resources to devote to decision-making ( $\alpha = 1$ , blue solid line), the subjective value function is concave for all  $x$  and thus the decision-maker is always risk-averse. Under these conditions the subjective value function looks very much like the utility functions that dominated economics before the development of reference-dependent utility theories. As  $\alpha$  increases (and the capacity increases) the subjective value function begins to take on the familiar S-shape from Prospect Theory - it starts off as a convex function and then at some point changes to concave.

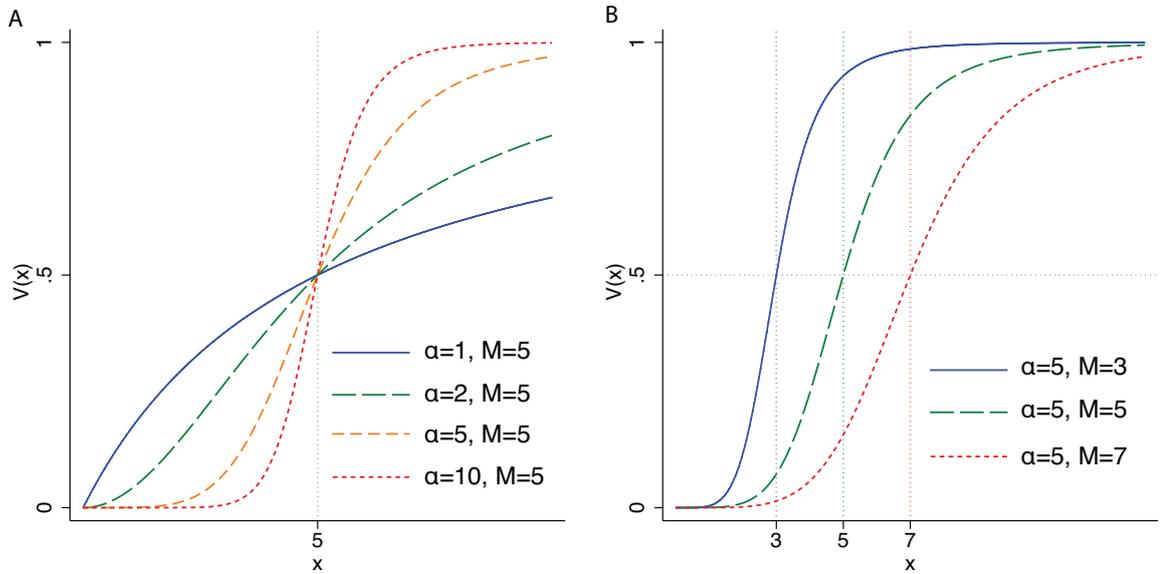


Figure 2: Examples of typical subjective value functions. In each graph, the vertical dashed lines are drawn at  $x = M_t$ . A: The effect of changes in the neural capacity on the subjective value function. B: The effect of changes in the prize expectation ( $M$ ) on the subjective value function.

The other insight from Figure 2 A is that as capacity increases, the function becomes steeper around the reference point. This allows decision makers with higher capacities to maximize choice precision (assuming non-zero noise as in section 2) at expectation. Any efficient organism, would of course spend more energy to increase precision at likely rather than unlikely valuations. Therefore  $V(M_t)$  should be the point at which precision is always maximal. The degree to which this precision is achieved here depends on the allocated capacity. In our model, individuals with higher capacity have more neural resources to spend and thus are able to encode such departures from expectation with weakly higher  $V$  difference for all  $x$ . As illustrated in Figure 2 A, for all  $x$  and for all  $M_t$ ,  $|V(x, M_t, \alpha) - V(M_t, M_t, \alpha)| > |V(x, M_t, \alpha') - V(M_t, M_t, \alpha')|$  iff  $\alpha > \alpha'$ .

To illustrate the role of the reference point, Figure 2 B plots three different value functions with different reference points and the same neural capacity. As the reward expectation increases, the subjective value function shifts to the right as in many reference dependent models. Importantly, the shift is such that at the reference point the function always takes the same value,  $V(M_t) = 0.5$ . This can be easily checked by substituting  $M_t$  in place of  $x$  in equation 1. For every  $\alpha$ , whenever  $x = M_t$ , then  $V(x) = 0.5$ . This property is consistent with vast behavioral and neural evidence that people evaluate the prizes as departures from the expectation rather than in absolute terms (Köszegi and Rabin, 2006, Abeler et al., 2011, Marzilli Ericson and Fuster, 2011).

### 3.1 Risk taking

To fully understand the role that reference point and capacity play in risk taking, we use a standard measure of individual risk attitude, the Arrow-Pratt index of relative risk aversion ( $RRA$ ):

$$RRA_t = -\frac{xV''}{V'} = \frac{(1 + \alpha)x^\alpha - (\alpha - 1)M_t^\alpha}{x^\alpha + M_t^\alpha} \quad (3)$$

The index can take both positive (risk aversion) and negative (risk seeking) values. In the traditional reference-dependent models, individuals are assumed to be risk-averse (seeking) whenever the reward is larger (smaller) than the expectation. In our model this condition is slightly more complex as the inflection point, that determines the switch from risk seeking to risk aversion, is not only determined by the expectation, but also by neural capacity.

**Proposition 1.** *The individual is risk averse if*

$$RRA_t > 0 \Leftrightarrow x > \sqrt[\alpha]{\frac{\alpha - 1}{\alpha + 1}} M_t \quad (4)$$

*An individual is risk neutral (seeking) whenever  $x = \sqrt[\alpha]{\frac{\alpha - 1}{\alpha + 1}} M_t$  ( $x < \sqrt[\alpha]{\frac{\alpha - 1}{\alpha + 1}} M_t$ )*

*Proof.* Follows from the definition of Arrow Pratt relative risk attitude index (equation 3). ■

#### 3.1.1 Neural capacity and risk attitude

Proposition 1 implies that an extremely constrained ( $\alpha = 1$ ) decision maker is risk averse for all  $x$ , just as in the traditional utility model. At the other extreme, when the individual approaches unlimited resources to spend on valuation ( $\alpha \rightarrow \infty$ ), the model predicts risk aversion for gains ( $RRA_t > 0$  iff  $x > M_t$ ) and risk seeking for losses ( $RRA_t < 0$  iff  $x < M_t$ ) in line with Prospect Theory. The model can thus accommodate the two most popular and competing models of choice as its extreme cases without additional parameters or complexity. Of course, human decision makers typically lie somewhere in between these two extremes: their subjective value function starts off as convex and as  $x$  increases switches to concave. The inflection in  $V$  occurs for some  $x < M_t$  because  $\sqrt[\alpha]{\frac{\alpha - 1}{\alpha + 1}} \leq 1$ . This leads us to our first two novel predictions:

**Proposition 2.** *For small losses with respect to expectation, an individual is risk averse. As the capacity increases, the range of rewards for which the individual is risk averse shrinks.*

*Proof.*  $\sqrt[\alpha]{\frac{\alpha - 1}{\alpha + 1}} \leq 1$  implies that for small losses relative to  $M_t$ , individual is risk averse. Since  $\frac{\partial \sqrt[\alpha]{\frac{\alpha - 1}{\alpha + 1}}}{\partial \alpha} > 0$  and  $\lim_{\alpha \rightarrow \infty} \sqrt[\alpha]{\frac{\alpha - 1}{\alpha + 1}} = 1$ , as  $\alpha$  increases the inflection point approaches  $M_t$ , meaning that the range of losses for which individual is risk averse shrinks. ■

In Figure 2 A we can see that changes in capacity differentially affect individual risk attitudes for gains and for losses. Formally:

**Proposition 3.** *As  $\alpha$  increases the decision maker becomes more risk averse for gains and more risk seeking for losses.*

*Proof.* To show:  $\frac{\partial RRA}{\partial \alpha} > 0$  for  $x > M_t$  and  $\frac{\partial RRA}{\partial \alpha} < 0$  for  $x < M_t$

$$\frac{\partial RRA}{\partial \alpha} = \frac{-M_t^{2\alpha} + x^{2\alpha} - 2\alpha M_t^\alpha x^\alpha \text{Log}(M_t) + 2\alpha M_t^\alpha x^\alpha \text{Log}(x)}{(M_t^\alpha + x^\alpha)^2}$$

The denominator of this expression is always positive so we focus on numerator only.

Hence,  $\frac{\partial RRA}{\partial \alpha} > 0 \Leftrightarrow -M_t^{2\alpha} + x^{2\alpha} - 2\alpha M_t^\alpha x^\alpha \text{Log}(M_t) + 2\alpha M_t^\alpha x^\alpha \text{Log}(x) > 0$

which we can rewrite

$$x^{2\alpha} - M_t^{2\alpha} + 2\alpha M_t^\alpha x^\alpha (\text{Log}(x) - \text{Log}(M_t)) > 0$$

which holds whenever  $x > M_t$  and completes the proof. ■

This is graphically illustrated in figure 2 A. This change in risk attitude is a result of improved discriminability around expectation for higher  $\alpha$ 's in a system that has limited resources to encode value.

### 3.1.2 Adaptive reference point and (in)stability of preferences

Since the expectation,  $M_t$  varies in our model from period to period, the individual risk attitude for a given  $x$ , as determined by equation 3 and Proposition 1, can fundamentally change over time. Our model thus adds to the existing discussion on the stability of preferences. Economists and psychologists have traditionally taken different views on the issue of preference stability, with psychologists largely arguing against stability (Slovic, 1995) and economists largely wishing for stability (Stigler and Becker, 1977). In our model, the fact that the expectation changes over time as the individual experiences different prizes implies that an individual's willingness to take on a given risky option will also change over time. This change is however fully predictable from an individual's reward history. Our model thus walks a line between traditional psychological and economic approaches. It yields a varying risk tolerance due to what could be construed as a contextual effect without resorting to discontinuities in the preference structure.

As the expectation increases, the Arrow Pratt index of relative risk attitude declines which gives rise to the following proposition:

**Proposition 4.** *For any given gamble, an individual is more risk tolerant the higher is her reference point.*

*Proof.* To show:  $\frac{\partial RRA_t}{\partial M_t} < 0$

$$\frac{\partial RRA_t}{\partial M_t} = -\frac{2\alpha^2 M_t^{\alpha-1} x^\alpha}{(x^\alpha + M_t^\alpha)^2} < 0 \text{ which completes the proof. } \quad \blacksquare$$

### Increasing versus decreasing income streams

The role of income history in explaining risk attitudes in our model goes beyond the pure wealth effects that economists have been modeling for centuries. Not only the total income earned, but also the timing when each reward was received, determines an individual's willingness to take risks. Imagine two individuals in period  $t$ , Oskar and John, who have the same neural capacity  $\alpha$  and have earned the same total income until period  $t$ . Oskar initially earned a lower income than John, but then his income continued increasing over time. The opposite happened to John. Since their total accumulated wealth is the same in period  $t$ , Expected Utility Theory and Prospect Theory would both predict that Oskar and John are going to have the same risk preference. Our model, however, makes a different prediction. Since the more recent events have a larger impact on the expectation (see Assumption 1), Oskar will have higher reference point. We can now use proposition 4 to predict that in period  $t$  Oskar will be more risk tolerant. Formally,

**Proposition 5.** *Keeping the total income experienced in the past and the neural capacity constant, increasing reward streams lead to more risk tolerance than decreasing reward streams.*

*Proof.* See proof of proposition 4. ■

### 3.1.3 Reflection affects

One of the key elements of Prospect Theory was inspired by the observation that people, on average, tend to be risk averse in gains and risk seeking in losses. This *reflection effect* in risk attitude has now been replicated in many studies on the aggregate level. The finding is, however, far less reliable when we look at the level of individuals (Tymula et al., 2013, Cohen et al., 1987, Schoemaker, 1990, Baucells and Villasís, 2010). Prospect Theory does not provide us with any explanation for why this may be the case. Our model, by incorporating individual heterogeneity in neural capacities and reward expectations, allows us to predict when reflection effects will and will not be observed at the individual level.

Figure 3 provides a graphical intuition about what types of samples are more or less likely to show reflection effects at the individual level. To produce the graphs A-E in Figure 3 we simulated one thousand individuals who varied in their neural capacities and expectations. In these simulations, the mean neural capacity was set equal to 5 (except for panel D where neural capacity varies) and the mean reference point was set equal to 10 (except for panel E where the reference point varies). To create samples of individuals who are more or less similar to one another, between graphs A-E we manipulated the variability of neural capacity and expectation. We then calculated for each of our simulated individuals their Arrow Pratt indexes of relative risk aversion (RRA) for a loss of size 5 and a gain of size 5 relative to the mean reference point of 10. Thus here RRA-gain is individual's Arrow Pratt index

calculated at  $x = 15$  and RRA-loss is the index calculated at  $x = 5$ . To visually assess whether individual's risk attitudes reflect, we then plotted RRA-gains against RRA-losses for all 1,000 simulated individuals. In each of the graphs, horizontal (vertical) green lines indicate risk neutrality in gains (losses). The black lines and the gray quadrant represent the areas consistent with Prospect Theory. If we take a very strict definition of the reflection effect, all the dots should fall on the black diagonal lines of slope -1, because people who are more risk averse in gains should similarly be more risk seeking in losses. Taking a less strict definition of reflection effects, Prospect Theory would predict that all the dots should fall in the gray quadrant where people are broadly risk averse in gains and risk seeking in losses. Clearly, while some graphs reproduce reflection effects almost perfectly (Figure 3 B, C, and part of E) others do not (Figure 3 A, D, and part of E).

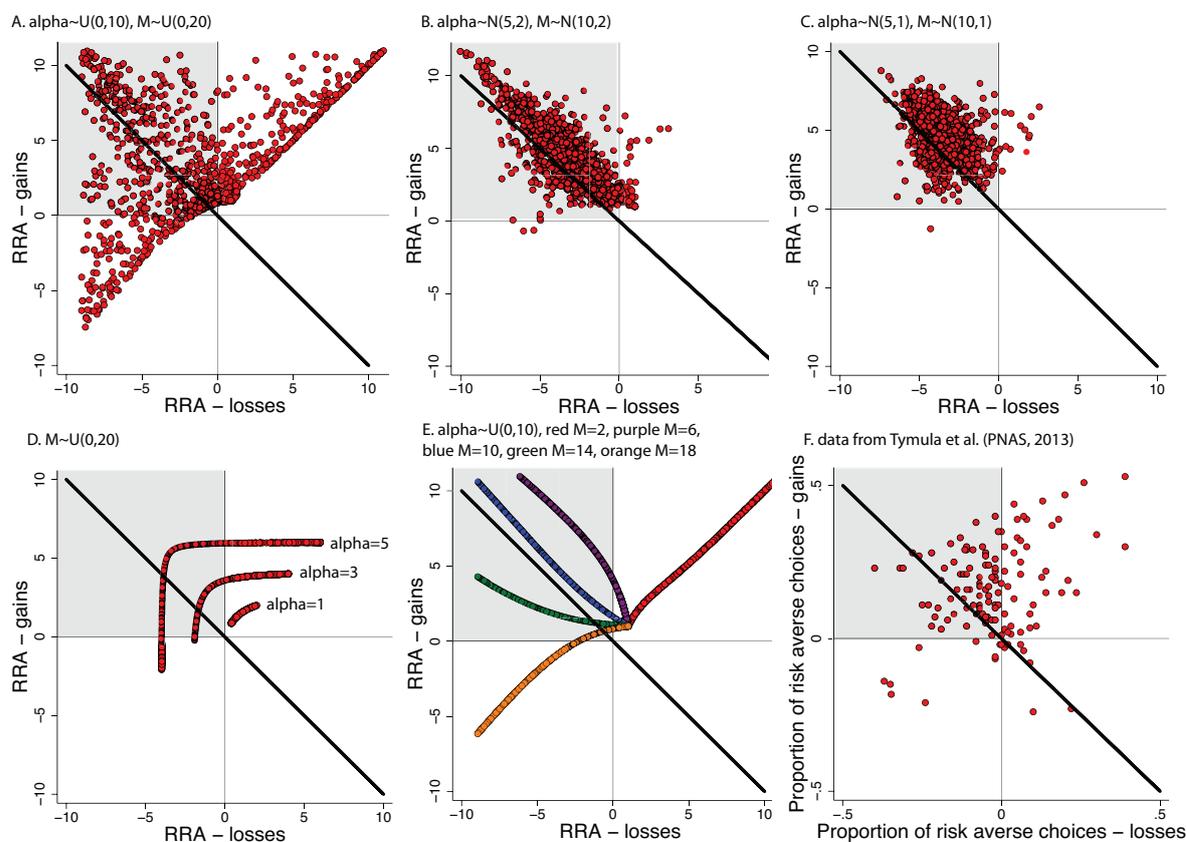


Figure 3: A-E plots of individual Arrow Pratt relative risk aversion index in gains against the index in losses for 1,000 individuals simulated from different distributions of  $\alpha$  and  $M$ . E - individual risk attitudes in gains against risk attitudes in losses, adapted from Tymula et al. (2013). Positive values in all graphs indicate risk aversion, negative risk seeking.

In all of the simulations, the dots fall into the upper left (triangular) half of the graphs. This is because in our model the higher the reward is, the more risk averse the individuals are. This model feature is consistent with a quite substantial experimental literature on the

effect of stakes on risk attitudes in gains and losses summarised in Vieider et al. (2012) and this observation brings us to our next proposition.

**Proposition 6.** *For all  $\alpha$  and for all  $M_t$ , individuals are less risk tolerant for higher  $x$ .*

*Proof.* To show: For all  $\alpha$  and for all  $M$ ,  $\frac{\partial RRA}{\partial x} > 0$

$$\frac{\partial RRA}{\partial x} = \frac{\alpha(1+\alpha)x^{\alpha-1}M_t^\alpha}{(x^\alpha + M_t^\alpha)^2} > 0 \quad \blacksquare$$

The above proposition has a flavor of the reflection effect, but is not sufficient to guarantee that it exists. For example, in Figure 3 A where we simulated a quite diverse sample of individuals drawn from a *uniform* distribution,  $\alpha \sim U(0, 10)$  and  $M \sim U(0, 20)$ , we did not see reflection effects. If anything, the correlation between risk attitude in gains and in losses is positive; regressing RRA-gains on RRA-losses, we obtain a coefficient of 0.287 on the RRA-losses ( $p < 0.001$ ). This is driven by large diversity in individual expectations and neural capacities. Individuals with lower capacities or reference points are more likely to be risk averse for both losses and gains relative to the population's mean reference point. Similarly, individuals with high expectations and sufficiently high capacity will be risk seeking at both 5 and 15. Both of these groups contribute to the overall positive correlation between risk attitude in gains and in losses. Graphically, the quadrant below the gray one consists of individuals with very high expectations and capacities. The quadrant to the right of the gray one is a combination of individuals who have low capacity and low expectations.

When we make the population more similar and draw the individual capacities and expectations from a *normally* distributed population with  $\alpha \sim N(5, 2)$  and  $M \sim N(10, 2)$  we begin to see strong reflection effects (OLS regression coefficient = -0.765,  $p < 0.001$ , Figure 3 B) that get stronger if we further reduce the variation in  $\alpha$  and  $M$  (OLS regression coefficient: -0.913,  $p < 0.001$ , Figure 3 C). Notice however, that for these figures we picked a distribution of capacities and expectations in the population to favor the existence of reflection effects. At the average  $\alpha = 5$  and  $M = 10$  in these figures, an average individual has an S-shaped value function (see Figure 2) with an inflection point close to 10 that produces the reflection effects. Had we assumed that the average capacity is lower in the population, more of the observations would fall into the white quadrant to the right of the gray quadrant, and we would have been less likely to see reflection effects. Trivially, when  $\alpha$  is very small, all individuals would be risk averse for both gains and losses (see Proposition 1 where we established that for a very low values of  $\alpha$ , the utility function is always concave). Similarly, if we picked higher capacities and expectations as the means for our simulated populations, more of the observations would have fallen into the white quadrant below the gray quadrant, making us less likely to observe reflection effects.

To illustrate these interactions, in Figure 3 D, we fixed the capacity,  $\alpha$ , at three different levels and simulated the model, drawing the expectation from a uniform distribution over 0 to 20 for each  $\alpha$ . In all the simulations there is a positive correlation between the risk

attitude in gains and in losses, inconsistent with reflection in utility curvature. Regressing RRA-gains on RRA-losses we get a significant coefficient of 0.444, 0.553 and 0.792 on RRA-losses for  $\alpha = 5$ ,  $\alpha = 3$ , and  $\alpha = 1$ , respectively ( $p < 0.001$  in all regressions). Nevertheless, looking at the aggregate level, more than half of the simulated individuals (563 out of 1000) fall into the grey region, consistent with a Prospect Theory-like reflection effect if  $\alpha = 5$ , 626 out of 1,000 when  $\alpha = 3$  and none when  $\alpha = 1$ .

Next, to investigate the role of expectation in reflection effects, we kept the expectation constant at five different levels (2, 6, 10, 14 or 18) in the population and assumed that capacity is uniformly distributed between 0 and 10. Figure 3 E plots the results of this exercise. Clearly, depending on the assumed level of expectation our simulations can either produce perfect reflection effects or a complete lack of thereof. When either the expectation is so low (red) or so high (orange) that 5 and 15 are both considered to be a gain or a loss, an individual’s risk attitudes in gains and losses are positively correlated. The closer the expectation is to 10, the more our data looks consistent with reflection effects.

Finally, to give our readers an idea of what kind of relationship one can expect between risk attitudes in gains and in losses, in Figure 3 F we reproduce experimental findings from an earlier paper of ours, Tymula et al. (2013), where we estimated the risk attitudes of 128 individuals in symmetric gain and loss gambles. Our sample in that paper was quite diverse, including individuals between 12 and 90 years old from New York, NY and New Haven, CT who likely differed significantly in their expectations and neural capacities. Consistent with the simulations of our model, in this diverse sample we found evidence for reflection effects on the aggregate level: on average people were risk averse in gains and risk seeking in losses with the majority (79 out of 128) individuals in the gray quadrant. However, on the individual level we found evidence against the reflection effect as risk attitudes in gains are weakly positively correlated with risk attitudes in losses (correlation coefficient = 0.254,  $p=0.003$ ).

### 3.2 Loss aversion

*Loss aversion* is perhaps the key component of most reference-dependent models of choice. This concept derives from the observation that in many situations the behavior of choosers is consistent with losses looming larger than equally sized gains. It has been usually modelled as a “kink” in the utility function that makes utility steeper for losses than it is for gains. The most commonly used utility specification for a loss averse individual is:

$$U(x) = \begin{cases} u(x) & \text{if } x \geq 0 \\ \lambda u(x) & \text{if } x < 0 \end{cases}$$

where  $\lambda > 1$ . Interestingly, a kink in the utility function ( $\lambda$ ) is not necessary to produce

loss aversion. In fact any utility function with diminishing marginal utility is behaviorally consistent with loss aversion (Blavatsky and Pogrebna, 2009, Ert and Erev, 2013). In a similar way our value function produces loss aversion even though there is no kink.

To make this statement we use a traditional definition of loss aversion where  $(\lambda)$  is equal to the ratio of the subjective value of loss of size  $k$  and gain of size  $k$ , each relative to the subjective value at the reference point  $M_t$ :

$$\lambda_t = \frac{V(M_t) - V(M_t - k)}{V(M_t + k) - V(M_t)} \quad (5)$$

Plugging in the formula for the subjective value function (equation 1), and simplifying, we get the expression for loss aversion in our model:

$$\lambda_t = \frac{(M_t^\alpha - (M_t - k)^\alpha)(M_t^\alpha + (M_t + k)^\alpha)}{((M_t - k)^\alpha + M_t^\alpha)((M_t + k)^\alpha - M_t^\alpha)} \quad (6)$$

Whenever  $\lambda > 1$ , an individual is loss averse by definition. If  $\lambda = 1$ , an individual treats losses and gains equally and is defined as not loss averse.

**Proposition 7.** *An individual with preferences represented by value function  $V(x)$  is always loss averse, though not necessarily to a significant degree.*

*Proof.* To show: For all  $k$ , for all  $M_t$  such that  $0 < k < M_t$  and for all  $\alpha > 0$ , individual is loss averse ( $\lambda > 1$ ).

$$\lambda > 1 \Leftrightarrow \frac{(M_t^\alpha - (M_t - k)^\alpha)(M_t^\alpha + (M_t + k)^\alpha)}{((M_t - k)^\alpha + M_t^\alpha)((M_t + k)^\alpha - M_t^\alpha)} > 1$$

Since the denominator is always positive ( $((M_t - k)^\alpha + M_t^\alpha)((M_t + k)^\alpha - M_t^\alpha) > 0$ ), we can multiply both sides of the inequality by it without changing the sign and rewrite the condition for loss aversion,  $\lambda > 1$ , as:

$$(M_t^\alpha - (M_t - k)^\alpha)(M_t^\alpha + (M_t + k)^\alpha) > (M_t^\alpha + (M_t - k)^\alpha)(-M_t^\alpha + (M_t + k)^\alpha)$$

Performing the algebra, we get that  $\lambda > 1 \Leftrightarrow 2k^2 > 0$ , which holds for all  $k$  and completes the proof. ■

In contrast to other reference-dependent theories that assume that  $\lambda$  is a fixed parameter, in our model loss aversion is a function of neural capacity and how far the offer under consideration would place you above or below the reference point  $M_t$  (captured by the stake size  $k$ ). This property allows us to make specific predictions about the precise degree of loss aversion that depend on both of these attributes.

Figure 4 plots loss aversion ( $\lambda$ ) as a function of  $k$  (size of gain and loss from the reference point) at different capacities. It is clear that both capacity and  $k$  play important roles in determining the strength of loss aversion. In general, the model predicts that people with higher capacity show less loss aversion. This could be seen as consistent with previous experimental findings suggesting that older people and people with lower levels of education

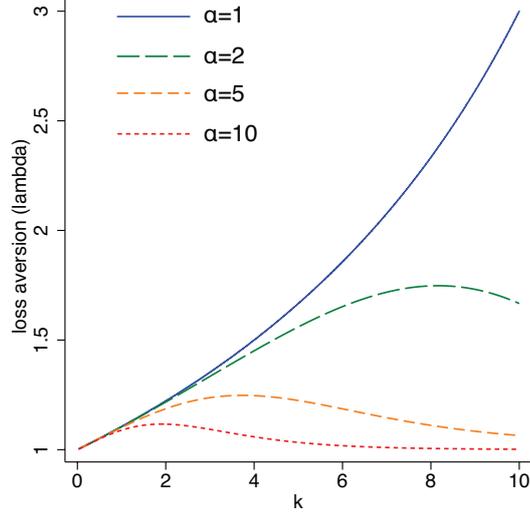


Figure 4: Loss aversion ( $\lambda$ ) as a function of the gain (loss) size ( $k$ ) at different neural capacities ( $\alpha$ ).

are more loss averse (Gaechter et al., 2006). The results can potentially explain the mixed findings in the literature on the effect of stake size on loss aversion (Pope and Schweitzer, 2011, Blavatsky and Pogrebna, 2007). For extremely constrained individuals, loss aversion is always increasing in stake size, however as capacity increases, increases in stake first lead to more loss aversion and then to less aversion. The higher the capacity, the lower is the stake at which loss aversion peaks.

**Proposition 8.** *For any stake, an individual with lower capacity is more loss averse. Loss aversion first increases and then decreases with stake size. Loss aversion peaks at lower stakes for individuals with higher capacity.*

*Proof.* See Figure 4. ■

It is important to notice that in our model loss aversion is not linked at all to an inflection in utility curvature as it is in Prospect Theory. For example, an individual who is extremely capacity constrained ( $\alpha = 1$ ) will be risk averse for both gains and losses and will at the same time exhibit loss aversion for all values of  $k$ . She will be loss averse, and in any experimental set up would be classified as loss averse, even though there is no kink in the utility function and no switch from risk seeking to risk aversion. Such a separation between the kink in the utility function and loss aversion has been discussed before (Blavatsky and Pogrebna, 2009). Our theory is the first to provide neurobiological foundations for the lack of relationship between the inflection point, the kink and loss aversion in the utility function.

### 3.2.1 An endowment effect without a “kink”

The traditional explanation for the endowment effect is *loss aversion*: People endowed with a mug demand more money to sell it than people who do not have the mug are willing to pay for it because losses loom larger than gains. Whether this is indeed the underlying cause of the observed differences between willingness to accept and willingness to pay has been the subject of a heated debate recently (List (2002), List (2003), Plott and Zeiler (2005), Engelmann and Hollard (2010), Isoni et al. (2011), Plott and Zeiler (2011)). Our model contributes to this debate by providing an alternative explanation for why the endowment effect occurs and when we can expect it. In our model, the endowment effect occurs because the individual is sluggish in updating the expectation term. Instead of the current status quo, the expectation is a time-weighted average of current and previous ownership (or offer) levels. Our model thus predicts the relationship between the existence and strength of the endowment effect and ownership history as documented by Strahilevitz and Loewenstein (1998), Marzilli Ericson and Fuster (2011), and Heffetz and List (2014) in experimental studies.

**Proposition 9.** *Let  $x_1$  be the individual’s initial ownership level of the good under consideration. Let  $x_2 > x_1$  be the ownership level after endowment. Let  $M_1$  be the reference point before the endowment and  $M_2$  be the reference point after endowment. The endowment effect occurs if  $M_2 < \frac{x_1 x_2}{M_1}$ . The endowment effect does not occur if  $(M_1^\alpha - M_2^\alpha)(M_1^\alpha M_2^\alpha - x_1^\alpha x_2^\alpha) = 0$ .*

*Proof.* The endowment effect occurs whenever:

$$WTA = \frac{x_2^\alpha}{(M_2^\alpha + x_2^\alpha)} - \frac{x_1^\alpha}{(M_2^\alpha + x_1^\alpha)} > \frac{x_2^\alpha}{(M_1^\alpha + x_2^\alpha)} - \frac{x_1^\alpha}{(M_1^\alpha + x_1^\alpha)} = WTP$$

Which is equivalent to:

$$WTA = \frac{(x_2^\alpha - x_1^\alpha)M_2^\alpha}{(M_2^\alpha + x_2^\alpha)(M_2^\alpha + x_1^\alpha)} > \frac{(x_2^\alpha - x_1^\alpha)M_1^\alpha}{(M_1^\alpha + x_2^\alpha)(M_1^\alpha + x_1^\alpha)} = WTP$$

After simplifying, we get that  $WTA > WTP$  if and only if

$$(M_1^\alpha - M_2^\alpha)(M_1^\alpha M_2^\alpha - x_1^\alpha x_2^\alpha) > 0$$

Since by assumption,  $M_1^\alpha - M_2^\alpha < 0$ ,  $WTA > WTP$  if and only if

$$M_1^\alpha M_2^\alpha - x_1^\alpha x_2^\alpha < 0 \Leftrightarrow M_2 M_1 < x_1 x_2.$$

The endowment effect does not occur whenever

$$WTA = WTP \Leftrightarrow (M_1^\alpha - M_2^\alpha)(M_1^\alpha M_2^\alpha - x_1^\alpha x_2^\alpha) = 0$$

■

Let us intuitively understand what this proposition implies. There are two interesting situations when the endowment effect does not occur. First when there is no change in expectation after endowment  $M_1 = M_2$ . Second, when the expectation is equal to the current level of ownership, so before the endowment  $M_1 = x_1$  and after the endowment  $M_2 = x_2$ . The second case is likely to occur when a sufficiently long time (or perhaps number of offers depending on whether one models times as continuous or discrete around

offers) has passed between the endowment and elicitation of willingness to accept, a period of time sufficient for the individual to fully adjust her reference point to the new ownership level. Graphically, this case is represented in Figure 5 B.

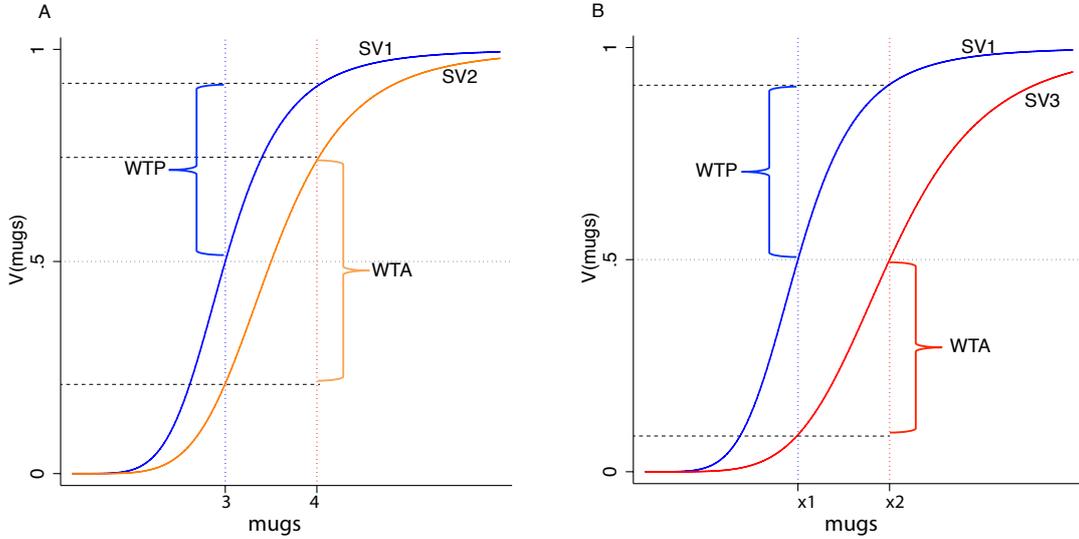


Figure 5: Endowment effect with A. not full B. full adjustment of the reference point.

To intuitively understand when the endowment effect occurs in our model, let us concentrate on two cases: First, let us assume that before the endowment occurs, the individual's expectation coincides with the ownership level  $M_1 = x_1$ . Then the endowment effect will be observed only if the reference point after the endowment does not fully update to the new ownership level  $M_2 < x_2$ . This is illustrated in Figure 5 A. the blue curve illustrates an individual's pre-endowment subjective value function. After the endowment, the individual's subjective value function shifts to the right, but does not fully adapt to the new ownership level. Instead the new reference point is a time weighted average somewhere between the new and old ownership level. Since the value function is, roughly speaking, steepest around the expectation, the individual demands more to part with the good after the endowment than she is willing to pay for it in the first place.

As a second case, let us assume that after the endowment sufficient time of offers pass that the reference point updates fully to the new ownership level, such that  $M_2 = x_2$ . Then the endowment effect will occur only if the reference point before endowment effect was below the ownership level  $M_1 < x_1$ . This could happen if the individual only recently gained  $x$  up to the level of  $x_1$  prior to the initial event and her reference point was not be fully adapted to that ownership level ( $M_1 < x_1$ ) when she is asked how much she is willing to pay for another unit of  $x$ .

Our framework can be used to derive predictions on the differences between WTA and WTP that depend on the reference point before and after an endowment. The timing and size

of ownership changes determine, in our model, whether the endowment effect occurs or not. One could even imagine situations where the endowment effect is reversed,  $WTA < WTP$  in our model. This could happen when WTP's are elicited soon after recent losses in ownership such that the individual does not fully adjust downwards to the new ownership level. This is a situation equivalent to the orange curve in Figure 5 A, but in the figure this is the subjective value function for an individual who used to own 4 mugs but very recently lost one of them.

Finally, notice that any reference dependent model based on taking a difference between the prize and the reference point makes an unrealistic prediction that after being given a mug the individual is going to have the same willingness to pay for an additional mug as he had for the first mug in the first place. Since after an endowment the utility function shifts in a parallel fashion after the endowment (as illustrated in Figure 6 A), if we kept endowing an individual with additional mugs, one at a time, she should (in Prospect Theory) be willing to pay the same amount for money for the third, fourth and hundredth mug! This is, of course, inconsistent with the notion of diminishing marginal utility. Our model does not, however, make this odd prediction (see Figure 6 B) as individual's willingness to pay for an extra mug always declines after an endowment.

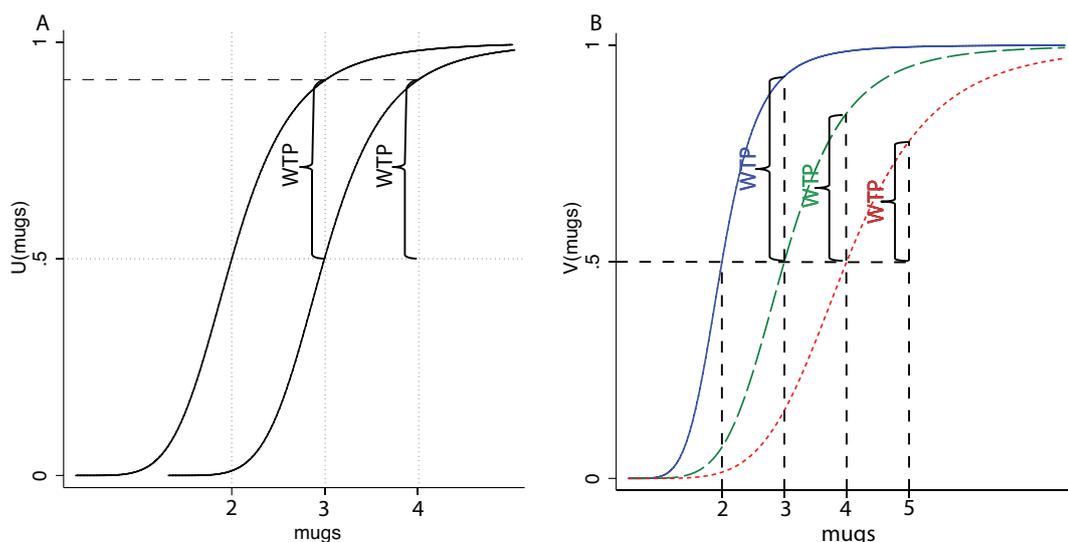


Figure 6: Endowment effect in A. difference-based reference point model B. divisive normalization model.

**Proposition 10.** *Suppose individual's reference point pre-endowment is set at the current ownership level. The individual's willingness to pay for another unit of  $x$  always drops after an endowment.*

*Proof.* To show: willingness to pay before the endowment is larger than willingness to pay after endowment.

The willingness to pay to increase ownership from  $x_1$  to  $x_2$  is given by  $SV_1(x_2) - SV_1(x_1) = \frac{(x_2^\alpha - x_1^\alpha)M_1^\alpha}{(M_1^\alpha + x_2^\alpha)(M_1^\alpha + x_1^\alpha)}$

To show:  $SV_1(x_2 - x_1) > SV_2(x_3 - x_2)$  (diminishing marginal utility condition), where  $x_1 < x_2 < x_3$  and  $x_2 = \frac{x_1 + x_3}{2}$

Given that  $\forall x > M_t$ , the marginal value of the additional unit of  $x$  decreases, the willingness to pay for another unit of  $x$  is the highest when the reference point coincides with the current level of ownership. Therefore  $SV_2(x_3 - x_2)$  is maximized and the diminishing marginal utility condition is hardest to satisfy when  $M_2 = x_2$ . If we show that this condition is satisfied for  $M_2 = x_2$ , then it holds for all other  $M_2$ .

Given  $M_1 = x_1$  and  $M_2 = x_2$  we rewrite the diminishing marginal utility condition as

$$SV_1(x_2) - SV_1(x_1) = \frac{(x_2^\alpha - x_1^\alpha)x_1^\alpha}{(x_1^\alpha + x_2^\alpha)2x_1^\alpha} > \frac{(x_3^\alpha - x_2^\alpha)x_2^\alpha}{(x_2^\alpha + x_3^\alpha)2x_2^\alpha} = SV_2(x_3) - SV_2(x_2)$$

This condition simplifies to:

$$x_2^{2\alpha} > x_1^\alpha x_3^\alpha \Leftrightarrow$$

$$x_2^2 > x_1 x_3 \Leftrightarrow (\text{since } x_2 = \frac{x_1 + x_3}{2}, \text{ we can rewrite as:})$$

$$\left(\frac{x_1 + x_3}{2}\right)^2 > x_1 x_3 \Leftrightarrow$$

$$0.25(x_1^2 + 2x_1 x_3 + x_3^2) > x_1 x_3 \Leftrightarrow$$

$$0.25(x_1^2 - 2x_1 x_3 + x_3^2) > 0 \Leftrightarrow$$

$$0.25(x_1 - x_3)^2 > 0 \text{ which always is true and completes the proof}$$

■

## 4 Discussion

In this paper we have presented an analysis of the risk and preference structures produced by Expected Subjective Value Theory, presenting it as a descriptive model of human choice behavior. Expected Subjective Value Theory rests on the normative *Normalization Models* of neuroscience which have been gaining popularity in Neuroeconomics. The analysis presented in Section 3 of this paper suggests that nearly all of the key behavioral phenomena that Prospect Theory was designed to explain, can also be explained by Expected Subjective Value Theory. We also find that Expected Subjective Value Theory makes a series of novel predictions that distinguish it from Prospect Theory behaviorally, predictions which are amenable to future testing.

Although Expected Subjective Value Theory has many similarities to Prospect Theory it differs from its predecessor in two very important ways that we wish to highlight. First, it rests on an entirely different, and smaller, set of logical primitives than does Prospect Theory. Second, it descends from a normative foundation in neuroscience, a fact which opens the possibility that in the future one could relate this descriptive theory to a stable normative foundation (Sims, 2003, Glimcher, 2010, Woodford, 2012, 2014). In the sections of this paper that follow we present: 1) an overview of the primitives of the theory and

an analysis of how they differ from the primitives of both Prospect Theory and Expected Utility Theory, 2) An assessment of the pros and cons of our theory relative to Prospect Theory from the point of view of practitioners, 3) A brief discussion of the implications of the normative foundation on which the theory was built, 4) A brief discussion of the relationship between the current model, the independence axiom and the probability weighting function of Prospect Theory and 5) Some tentative future directions for ongoing analyses of the theory.

## 4.1 Primitives of the Theory

It is important to note that ESVT differs fundamentally from nearly all other descriptive choice theories in the basic nature of its primitives – though in a way driven by recent insights into the neurobiology of decision-making. Expected Utility, at least when typically used as a predictive tool, relies on the specification of a functional form for utility to define a chooser’s preferences. And of course this means that when one fits EU to observed behavior to formulate a quantitative prediction about future choices one parameterizes the utility function one has selected – typically with a curvature specification and a wealth level. To restate this in a more philosophical vein; EU represents choosers as being exogenously provided with a utility function, and as having an endogenous wealth level, nothing more.

Prospect Theory (Kahneman and Tversky, 1979) broadened the parameter set for the utility function by introducing the *value function*. This was accomplished by introducing two curvature terms (one for gains and one for losses), a loss aversion term, and a reference point. (Practitioners often simplify these to three parameters by assuming that the negative going curvature is a pure reflection of the positive going curvature.) Effectively, Kahneman and Tversky broadened the primitives of their theory (relative to Expected Utility) to include not just the notion of utility function curvature but also the notions of a reference point and loss aversion. These three ideas (utility curvature, a reference point and loss aversion) form the core primitives of value in their theory and are the three (or four) exogenous elements with which choosers are provisioned by the theory, although the reference point is sometimes viewed as endogenous. (We neglect for now the Probability Weighting Function to which we will return in a later section.)

In contrast, ESVT has no fundamental primitive for utility function curvature. Instead curvature of the ESVT function *emerges* from the two core primitives of the theory: *capacity* and *expectation*. This is a fact that we acknowledge some economists will find disturbing. In our theory, choosers experience (or are offered) wealth shocks in time and from these derive an expectation about the future. A quantitative form of the expectation was derived originally by neurobiological measurement (LoFaro et al., 2014), and is in fact quite closely related to literature in economics. Choosers are thus exogenously endowed in our theory with only a single primitive, their so-called capacity, which interacts with the endogenously

specified expectation.

When one says that a chooser employing Prospect Theory shows *loss aversion*, we mean that loss aversion is both a logical primitive and an endowed property of the chooser. In some sense we mean that choosers behave *as if* they had a specific mechanism inside their brains which differentially weighted losses and gains in a multiplicative sense. In contrast, when we say that a chooser employing ESVT shows *loss aversion* we mean *neither of these things*. Instead we mean that were one to use standard econometric tools to measure the observational variable called *loss aversion*, one would observe values greater than 1 in this parameter. But *loss aversion* per se would not be a feature of the chooser, rather it would simply be the name of a measurement.

A key point is thus the fact that our theory does not really have any clear analogue to the static utility function of Expected Utility, a function which maps inputs,  $x$ , to *utility*, or the static value function of Prospect Theory which maps  $x$  to *value*. One can of course see the relationship between  $x$  and *expected subjective value* in our theory, but ESVT is clearly different in the sense that this mapping is conditional on the capacity of the chooser and the dynamic expectation of the chooser – the only two real variables in the theory.

## 4.2 Assessing the Pros and Cons of ESVT

Given the unusual nature of our approach, it seems reasonable to ask what advantages it offers over existing theories at both practical and theoretical levels. At a practical level the answer to this question seems quite straightforward: ESVT can capture all of the basic phenomena that the value function of Prospect Theory can capture, but with fewer parameters. At its simplest, Prospect Theory can capture the phenomena it addresses with 2-3 parameters. ESVT relies only on a single key parameter, capacity.

To make this clearer, when one fits Prospect Theory’s value function to an empirical dataset one typically fits a curvature term (or two), a loss aversion term, and then either assumes a reference point or fits the reference point with an additional parameter. When modeling the same dataset with ESVT, one relies on fitting a capacity term and then either assuming an expectation function (LoFaro et al., 2014) or fitting what is essentially a “reverse discount rate” to the stream of shocks experienced by the chooser to derive empirically the expectation term of the theory. The theory thus uses fewer parameters than Prospect Theory but captures many of the same behaviors – an obvious advantage.

The real test of the theory, however, is not whether it can replicate the performance of Prospect Theory with fewer parameters, but whether it can exceed the descriptive performance of Prospect Theory in new ways. We take it as already established that ESVT significantly outperforms Prospect Theory with regard to neurobiological observables (Glimcher (2015), Louie et al. (2014), LoFaro et al. (2014), and Louie et al. (2013)), but we

acknowledge that this will be only marginally interesting to pure economists. As a result, one focus of this paper has been to show how and where one can compare the performance of Prospect Theory and ESVT in behavioral data. Section 3 serves as a guide to that process. We believe it suggests that ESVT compares favorably with Prospect Theory in many regards, though we acknowledge that much work remains to be done.

Another practical area of differentiation between the theories lies with the nature of our preference-like parameter, the capacity term  $\alpha$ , and the method of its instantiation in the human brain. We are, of course, keen to render the capacity term,  $\alpha$ , directly observable through neurobiological or cognitive methods in a way that would be hard to imagine for the curvature and loss aversion terms of Prospect Theory. We suspect that it might be fruitful to hypothesize that the capacity term has two kinds of variation; a between-subjects component and a within-subjects component. Between subjects variation might well be easily restricted in interesting ways by assuming that aging (Tymula et al., 2013), IQ, educational level and structural features of the brain (Gilaie-Dotan et al., 2014, Kable and Levy, 2015) all map in some manner to capacity. This is an important future direction to which we plan to devote significant attention. Were a significant portion of the variance in the capacity constraint to become directly observable, ESVT would be reduced to a nearly parameter-free model of choice at the representative agent level, which we would take to be a significant advance. At the within-subject level, the notion of variation in  $\alpha$  immediately suggests ideas of self-control and commitment, another exciting area requiring additional future work.

### 4.3 Implications of the Normative Foundation

The core idea underlying ESVT is that the costs of precision are monotone and that these costs have been absent from much but not all (for example Simon (1976), Robson (2001), Rayo and Becker (2007), Netzer (2009), Glimcher (2010), Woodford (2012)) economic thinking. Neurobiological evidence available today clearly indicates that increasing the precision of any neural representation is without a doubt costly. And while we rush to acknowledge that ESVT in its current form is not normative, our goal in formulating it has been to incorporate that insight from neuroscience into our descriptive model. When EU was formulated, the core notion it captured was that choosers should behave in a consistent manner with regard to their wealth levels (Samuelson, 1947, von Neumann and Morgenstern, 1944, Savage, 1954). The entire notion of the neoclassical utility function emerges from that restriction. When Kahneman and Tversky (1979) developed Prospect Theory, they were reacting to growing evidence that people were not consistent with regard to wealth levels, that they were *irrational*. What we hope that ESVT raises as a possibility, is that the problem with Expected Utility might not have been that people were being inconsistent, but that the things over which they were being consistent were misspecified in that theory. von Neumann

and Morgenstern (1944) assumed that transitivity over *prizes* was the only factor relevant to thinking about consistency. ESVT proposes the possibility that people are much more consistent than previously thought, *but over prizes and costs combined* rather than just over *prizes*.

While future axiomatic work will hopefully clarify the relationship between these kinds of representations and consistency, it is clear that ESVT invokes a slightly different notion of consistency than do traditional models. Here we draw from neurobiology the notion that costs are both knowable and important – a point we take as first developed deeply by Herb Simon in the 1950s.

#### 4.4 The Probability Weighting Function

To this point, we have focused entirely on understanding how the *utility function*, the *value function* and the *expected subjective value function* relate. We have devoted none of our space to a discussion of subjective probabilities (Savage, 1954). Allais (1953) famously noted that human decision makers violate the Independence Axiom of Expected Utility, a point that has been made again and again (see Rieskamp et al. (2006) for review). Prospect Theory addressed that violation in two ways. First by constructing a more complex value function and second by defining a new form of subjective probability (Savage, 1954) they called the *probability weighting function* (Kahneman and Tversky, 1979). Often overlooked is the fact that both of these functions contribute to risk attitudes in Prospect Theory. Risk attitudes of choosers in low-probability lotteries reflect both of these mechanisms and fits of these two functions to data highlight that the parameters for the two functions interact (Fox and Poldrack, 2014).

We do not yet know the degree to which ESVT leads to violations of the Independence Axiom and that is a point that we are working to address. We also do not know enough today about the neural representation of symbolic probabilities. Without that understanding we cannot speak meaningfully about the efficient representation of probabilities we might desire. But it does seem clear that future work on ESVT will need both to clarify ESVT’s relationship to the Independence Axiom and will require a better understanding of the costs of representing probabilities. We are currently investigating both of these areas.

#### 4.5 Future Directions

Just like our brains, this paper too has limited capacity, and some of the novel predictions of our theory remain to be tested and verified in future work. Experimental studies that manipulate the history of rewards and capacity would help us understand whether our predictions are correct and will guide refinements of the model. The divisive normalisation functional form that we chose to use was on purpose simplified and there are interesting ways we would

like to see it enriched. In particular, we did not elaborate much on the endogenous rational expectation term,  $M$ . We leave it up to future extensions of the model to relate more precisely the past statistics of the environment to behavior. We suspect that changes to the skew and kurtosis of the distribution of past rewards will have particularly interesting implications for behaviour and these are not explicitly modelled in this paper. Nevertheless, the model provides the framework where skew and kurtosis can easily be incorporated and studied both theoretically and empirically.

## 5 References

### References

- Abeler, J., A. Falk, L. Goette, and D. Huffman (2011, apr). Reference Points and Effort Provision. *American Economic Review* 101(2), 470–492.
- Allais, M. (1953). Le comportement de l’homme rationnel devant le risque: critique des postulats et axiomes de l’école Américaine. *Econometrica* 21(4), 503–546.
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review* 61, 183–193.
- Banerji, A. and N. Gupta (2014, feb). Detection, Identification, and Estimation of Loss Aversion: Evidence from an Auction Experiment †. *American Economic Journal: Microeconomics* 6(1), 91–133.
- Barlow, H. A. (1961). Possible principles underlying the transformations of sensory messages. In W. A. Rosenblith (Ed.), *Sensory Communication*. Cambridge, MA: MIT Press.
- Bartling, B., L. Brandes, and D. Schunk (2015, nov). Expectations as Reference Points: Field Evidence from Professional Soccer. *Management Science* 61(11), 2646–2661.
- Baucells, M. and A. Villasís (2010, may). Stability of risk preferences and the reflection effect of prospect theory. *Theory and Decision* 68(1-2), 193–211.
- Baucells, M., M. Weber, and F. Welfens (2011, mar). Reference-Point Formation and Updating. *Management Science* 57(3), 506–519.
- Becker, G. M., M. H. DeGroot, and J. Marschak (1963). Stochastic models of choice behavior. *Behavioral Science* 8(1), 41–55.
- Blavatsky, P. and G. Pogrebna (2007, jun). Loss Aversion? Not With Half-a-Million on the Table! *SSRN Electronic Journal*.

- Blavatsky, P. and G. Pogrebna (2009, may). Endowment effects? “Even” with half a million on the table! *Theory and Decision* 68(1-2), 173–192.
- Carandini, M. and D. J. Heeger (2012, jan). Normalization as a canonical neural computation. *Nature reviews. Neuroscience* 13(1), 51–62.
- Card, D. and G. B. Dahl (2011, mar). Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior. *The Quarterly Journal of Economics* 126(1), 103–143.
- Cohen, M., J.-Y. Jaffray, and T. Said (1987). Experimental comparison of individual behavior under risk and under uncertainty for gains and for losses. *Organizational Behavior and Human Decision Processes* 39(1), 1 – 22.
- Crawford, V. P. and J. Meng (2011, aug). New York City Cab Drivers’ Labor Supply Revisited: Reference-Dependent Preferences with Rational-Expectations Targets for Hours and Income. *American Economic Review* 101(5), 1912–1932.
- Engelmann, D. and G. Hollard (2010). Reconsidering the Effect of Market Experience on the “Endowment Effect”. *Econometrica* 78(6), 2005–2019.
- Ert, E. and I. Erev (2013). On the descriptive value of loss aversion in decisions under risk: Six clarifications. *Judgment and Decision Making* 8(3), 214–235.
- Fox, C. R. and R. A. Poldrack (2014). Prospect Theory and the Brain. In P. W. Glimcher and E. Fehr (Eds.), *Neuroeconomics: Decision Making and the Brain* (Second Edition), pp. 533–563. Academic Press.
- Frederick, S. and G. F. Loewenstein (1999). Hedonic Adaptation. In D. Diener, N. Schwarz, and D. Kahneman (Eds.), *Well-being: The foundations of hedonic psychology*, pp. 302–329. New York, NY: Russell Sage Foundation.
- Gaechter, S., E. J. Johnson, and A. Herrmann (2006). Individual-level loss aversion in riskless and risky choices.
- Gilaie-Dotan, S., A. Tymula, N. Cooper, J. Kable, P. Glimcher, and I. Levy (2014). Neuroanatomy predicts individual risk attitudes. *Journal of Neuroscience* 34(37).
- Gill, D. and V. Prowse (2012, feb). A Structural Analysis of Disappointment Aversion in a Real Effort Competition. *American Economic Review* 102(1), 469–503.
- Glimcher, P. (2010). *Foundations of Neuroeconomic Analysis*. Oxford University Press.

- Glimcher, P. (2015, jan). Understanding the Hows and Whys of Decision-Making: From Expected Utility to Divisive Normalization. *Cold Spring Harbor symposia on quantitative biology* 79, 169–76.
- Glimcher, P. W. (2005, sep). Indeterminacy in Brain and Behavior. *Annual Review of Psychology* 56, 25–56.
- Glimcher, P. W. (2011, sep). Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences of the United States of America* 108 Suppl(Supplement\_3), 15647–54.
- Gneezy, U. (2005). Experimental Business Research. Chapter Updating t. Boston, MA: Springer US.
- Heeger, D. J. (1992, aug). Normalization of cell responses in cat striate cortex. *Visual neuroscience* 9(2), 181–197.
- Heeger, D. J., E. P. Simoncelli, and J. A. Movshon (1996, jan). Computational models of cortical visual processing. *Proceedings of the National Academy of Sciences of the United States of America* 93(2), 623–7.
- Heffetz, O. and J. A. List (2014, oct). Is the endowment effect an expectations effect? *Journal of the European Economic Association* 12(5), 1396–1422.
- Hunt, L. T., R. J. Dolan, and T. E. J. Behrens (2014, nov). Hierarchical competitions subserving multi-attribute choice. *Nature neuroscience* 17(11), 1613–22.
- Isoni, A., G. Loomes, and R. Sugden (2011, apr). The Willingness to Pay—Willingness to Accept Gap, the “Endowment Effect,” Subject Misconceptions, and Experimental Procedures for Eliciting Valuations: Comment. *American Economic Review* 101(2), 991–1011.
- Jacobson, R. and C. Obermiller (1990, mar). The Formation of Expected Future Price: A Reference Price for Forward- Looking Consumers. *Journal of Consumer Research* 16(4), 420.
- Kable, J. W. and I. Levy (2015, oct). Neural markers of individual differences in decision-making. *Current Opinion in Behavioral Sciences* 5, 100–107.
- Kahneman, D. and A. Tversky (1979). Prospect Theory. *Econometrica* 47(2), 311.
- Kiani, R., T. D. Hanks, and M. N. Shadlen (2008, mar). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 28(12), 3017–29.

- Kőszegi, B. and M. Rabin (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics* 121(4), 1133–1165.
- Lange, A. and A. Ratan (2010, mar). Multi-dimensional reference-dependent preferences in sealed-bid auctions – How (most) laboratory experiments differ from the field. *Games and Economic Behavior* 68(2), 634–645.
- Lennie, P. (2003, mar). The Cost of Cortical Computation. *Current Biology* 13(6), 493–497.
- List, J. A. (2002). Preference Reversals of a Different Kind: The "More Is Less" Phenomenon on JSTOR. *The American Economic Review* 92(5), 1636–1643.
- List, J. A. (2003). Does Market Experience Eliminate Market Anomalies? *The Quarterly Journal of Economics* 118(1), 41–71.
- LoFaro, T., K. Louie, R. Webb, and P. W. Glimcher (2014, mar). The Temporal Dynamics of Cortical Normalization Models of Decision-making. *Letters in Biomathematics* 1(2), 209–220.
- Louie, K. and P. W. Glimcher (2012). Set-size effects and the neural representation of value. In R. J. Dolan and T. Sharot (Eds.), *Neuroscience of Preference and choice: Cognitive and neural mechanisms*, pp. 143–169. London: Academic Press.
- Louie, K., P. W. Glimcher, and R. Webb (2015, oct). Adaptive neural coding: From biological to behavioral decision-making. *Current opinion in behavioral sciences* 5, 91–99.
- Louie, K., L. E. Grattan, and P. W. Glimcher (2011, jul). Reward value-based gain control: Divisive normalization in parietal cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 31(29), 10627–10639.
- Louie, K., M. W. Khaw, and P. W. Glimcher (2013, apr). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences of the United States of America* 110(15), 6139–6144.
- Louie, K., T. LoFaro, R. Webb, and P. W. Glimcher (2014, nov). Dynamic divisive normalization predicts time-varying value coding in decision-related circuits. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 34(48), 16046–57.
- Marzilli Ericson, K. M. and A. Fuster (2011, oct). Expectations as Endowments: Evidence on Reference-Dependent Preferences from Exchange and Valuation Experiments. *The Quarterly Journal of Economics* 126(4), 1879–1907.
- Mas, A. (2006, aug). Pay, Reference Points, and Police Performance\*. *Quarterly Journal of Economics* 121(3), 783–821.

- Mazumdar, T., S. Raj, and I. Sinha (2005, oct). Reference Price Research: Review and Propositions. *Journal of Marketing* 69(4), 84–102.
- Netzer, N. (2009, may). Evolution of Time Preferences and Attitudes toward Risk. *American Economic Review* 99(3), 937–955.
- Ohzawa, I., G. Sclar, and R. D. Freeman (1985). Contrast gain control in the cat’s visual system. *Journal of neurophysiology* 54(3), 651–667.
- Olshausen, B. A. and D. J. Field (1997, dec). Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vision research* 37(23), 3311–25.
- Padoa-Schioppa, C. (2009, nov). Range-adapting representation of economic value in the orbitofrontal cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 29(44), 14004–14.
- Plott, C. R. and K. Zeiler (2005). The Willingness to Pay–Willingness to Accept Gap, the “Endowment Effect,” Subject Misconceptions, and Experimental Procedures for Eliciting Valuations. *The American Economic Review* 95(3), 530–545.
- Plott, C. R. and K. Zeiler (2011, apr). The Willingness to Pay—Willingness to Accept Gap, the “Endowment Effect,” Subject Misconceptions, and Experimental Procedures for Eliciting Valuations: Reply. *American Economic Review* 101(2), 1012–1028.
- Pope, D. G. and M. E. Schweitzer (2011, feb). Is Tiger Woods Loss Averse? Persistent Bias in the Face of Experience, Competition, and High Stakes. *American Economic Review* 101(1), 129–157.
- Post, T., M. J. van den Assem, G. Baltussen, and R. H. Thaler (2008, feb). Deal or No Deal? Decision Making under Risk in a Large-Payoff Game Show. *American Economic Review* 98(1), 38–71.
- Putler, D. S. (1992, aug). Incorporating Reference Price Effects into a Theory of Consumer Choice. *Marketing Science* 11(3), 287–309.
- Ratliff, F. and H. K. Hartline (1959, jul). The responses of Limulus optic nerve fibers to patterns of illumination on the receptor mosaic. *The Journal of general physiology* 42(6), 1241–55.
- Rayo, L. and G. S. Becker (2007). Evolutionary Efficiency and Happiness. *Journal of Political Economy* 115(2), 302–337.
- Reynolds, J. H. and D. J. Heeger (2009, jan). The normalization model of attention. *Neuron* 61(2), 168–85.

- Rieskamp, J., J. R. Busemeyer, and B. A. Mellers (2006, sep). Extending the Bounds of Rationality: Evidence and Theories of Preferential Choice. *Journal of Economic Literature* 44(3), 631–661.
- Robson, A. J. (2001). The Biological Basis of Economic Behavior. *Journal of Economic Literature* 39(1), 11–33.
- Rosato, A. and A. Tymula (2016). No Loss Aversion and Competition in Vickrey Auctions: Money Ain't No Good.
- Samuelson, P. (1947). *Foundations of economic analysis*. Cambridge, MA: Harvard University Press.
- Savage, L. (1954). *The Foundations of Statistics*. New York: Wiley.
- Schoemaker, P. J. H. (1990, dec). Are Risk-Attitudes Related Across Domains and Response Modes? *Management Science* 36(12), 1451–1463.
- Schwartz, O. and E. P. Simoncelli (2001, aug). Natural signal statistics and sensory gain control. *Nature neuroscience* 4(8), 819–25.
- Shannon, C. E. and W. Weaver (1949). *The mathematical theory of communication*. Urbana, IL: University of Illinois Press.
- Simon, H. (1976). *Administrative Behavior*. New York: The Free Press.
- Simon, H. A. (1955, feb). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics* 69(1), 99.
- Simoncelli, E. P. (2003, apr). Vision and the statistics of the visual environment. *Current Opinion in Neurobiology* 13(2), 144–149.
- Simoncelli, E. P. and B. A. Olshausen (2001, jan). Natural image statistics and neural representation. *Annual review of neuroscience* 24, 1193–1216.
- Sims, C. A. (2003, apr). Implications of rational inattention. *Journal of Monetary Economics* 50(3), 665–690.
- Slovic, P. (1995). The construction of preference. *American Psychologist* 50(5), 346–371.
- Steinberg, L. (2007, apr). Risk Taking in Adolescence: New Perspectives From Brain and Behavioral Science. *Current Directions in Psychological Science* 16(2), 55–59.
- Stigler, G. J. and G. S. Becker (1977). De Gustibus Non Est Disputandum. *The American Economic Review* 67(2), 76–90.

- Strahilevitz, M. A. and G. Loewenstein (1998). The Effect of Ownership History on the Valuation of Objects on JSTOR. *Journal of Consumer Research* 25(3), 276–289.
- Tymula, A. and H. Plassmann (2016). Context-dependency in valuation.
- Tymula, A., L. a. Rosenberg Belmaker, L. Ruderman, P. W. Glimcher, and I. Levy (2013). Like cognitive function, decision making across the life span shows profound age-related changes. *Proceedings of the National Academy of Sciences of the United States of America* 110(42), 17143–8.
- Vieider, F. M., P. Martinsson, and H. Medhin (2012). Stake effects on ambiguity attitudes for gains and losses.
- von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Wainwright, M. J., O. Schwartz, and E. P. Simoncelli (2002). Natural image statistics and divisive normalization: Modeling nonlinearity and adaptation in cortical neurons. In R. P. Rao, B. A. Olshausen, and M. S. Lewicki (Eds.), *Probabilistic models of the brain: Perception and neural function*, pp. 203–222. Cambridge, MA: MIT Press.
- Webb, R., P. W. Glimcher, I. Levy, S. C. Lazzaro, and R. B. Rutledge (2012, oct). Neural Random Utility. *SSRN Electronic Journal*.
- Webb, R., P. W. Glimcher, and K. Louie (2014, jul). Rationalizing Context-Dependent Preferences: Divisive Normalization and Neurobiological Constraints on Choice. *SSRN Electronic Journal*.
- Woodford, M. (2012, may). Prospect Theory as efficient perceptual distortion. *American Economic Review* 102(3), 41–46.
- Woodford, M. (2014, may). Stochastic Choice: An Optimizing Neuroeconomic Model †. *American Economic Review* 104(5), 495–500.