THE UNIVERSITY OF
SYDNEY

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

Paul W. Glimcher & Agnieszka A. Tymula

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

Paul W. Glimcher[*]   Agnieszka A. Tymula[†]

January 10, 2017

Abstract

We present a descriptive model of choice with normative foundations based on how the brain is thought to represent value. An individual's behavior is fully described by two primitives: an individual's expectation and one free parameter we call "predisposition". The model captures the same apparent preference phenomena illustrated by Prospect Theory but unlike Prospect Theory accounts for individual heterogeneity in parameters, employs far fewer parameters than full prospect theory, and retains neurobiological plausibility as a causal model of the choice process. Additionally, our theory makes a series of novel predictions amenable to future testing and includes an alternative explanation for endowment effect.

JEL: D03, D81, D87
Keywords: utility, decision-making, reference point, neuroeconomics

1 Introduction

There is growing evidence that during the last five years neurobiologists have succeeded in identifying the basic representational forms used by the human brain to encode decision variables for causally producing choice. This class of representational forms can predict a number of interesting human and animal choice behaviors that have so far escaped description in the economics literature (Louie et al., 2013). Moreover, this class of models has been shown to have strong normative foundations — it ensures that the information embedded in

[*]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 representation is minimally redundant (Wainwright et al., 2002) and in terms of choice theory it has been shown to maximize utility when precision (a reduction in stochasticity) is costly, as is the case in the human brain (Steverson et al., 2016). In this paper we compare and contrast behavioral predictions of one of the simplest members of this representational class with the Prospect Theory of Kahneman and Tversky (1979).

A central feature of this class of models is the reference point. There is now abundant evidence from behavioral economics that individuals evaluate options not in absolute terms but rather relative to some reference point. This reference point was initially conceptualized as related to *status quo* (Kahneman and Tversky, 1979). More recently, extensive evidence from behavioral economics has suggested that the reference point may be better described as an outcome expectation. Today the importance of the reference point is 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 (Köszegi and Rabin, 2006), and derived whether and how expectations should affect valuation in a normative sense (Frederick and Loewenstein, 1999, Robson, 2001, Rayo and Becker, 2007, Netzer, 2009, Woodford, 2012, Robson and Whitehead, 2016). The overwhelming conclusion of the majority of these inferential studies is that the utility function is, in some sense, dynamically adjusted in response to expectations. Of course, without observing the reference point directly, one cannot know exactly how it enters, or modifies, an 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 expectations and reference points — particularly in the domain of sensory perception. (Studies of sensory perception have historically been very successful at identifying the primitives from which all neurobiological representations are built, a point to which we return later.) Detailed differential equation-based models of the neuronal networks that compute and represent 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 Wilson and Humanski (1993), Carandini and Heeger (2012). Measuring neuronal responses in a cat's visual cortex to different levels of contrast, Ohzawa et al. (1985) showed that the activity of neurons in the visual cortex is conditional on the 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 expectations should modulate the utility function (Rayo and Becker, 2007, Netzer, 2009, Glimcher, 2010, Woodford, 2012, Robson and Whitehead, 2016) — they adjust response

magnitudes dynamically to maximize the precision of their capacity-limited representation around the expectation of future contrast. An economist unfamiliar with biology might object to the application of these ideas to choice, but extensive evidence now makes it clear that these studies of sensory perception have identified the fundamental primitives of many, if not all, neurobiological representations (Carandini and Heeger, 2012, Glimcher, 2015).

Our initial focus here is the novel-to-economists functional form via which expectations enter the value function. We stress that this uncontroversial functional form is grounded in years of study by neuroscientists (see section 2). For the first time we apply this common approach to the domain of economic choice to generate predictions about how expectations affect behaviors traditionally explained through the lens of risk attitudes, loss aversion, and probability weighting. We organize the predictions of our model around influential concepts that originate from Prospect Theory. Many theories have, over the last half-century, attempted to extend Expected Utility in a positive way to capture its widely observed violations. Citation rates suggest that one descriptive theory dominates the academic discourse: the Prospect Theory of Kahneman and Tversky (1979). Kahneman and Tversky (1979) captured the violations of Expected Utility Theory 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 accepted as ground truth at a mechanistic level rather than as the useful descriptive functions Kahneman and Tversky intended them to be.

Despite the obvious power of Prospect Theory, the possibility remains that the variance Prospect Theory fails to capture may someday be explained by other approaches.[1] Prospect Theory does not always successfully capture variation in choice between individuals and over different contexts, even with its many parameters (up to 5 plus the exogenous reference point). Diminishing sensitivity with respect to the reference point (which leads to risk aversion in gains and risk seeking in losses), loss aversion and probability weighting 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

---

[1]We acknowledge the widely held belief among some economists and some reviewers that the days of unified theories of choice are forever behind us, but do not feel that this belief should be used as an *a priori* reason to dismiss any model.

parameterized, and though widely observed, is notoriously variable across studies and contexts (Ert and Erev, 2013). 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 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). The evidence on probability weighting clearly shows that on an individual level many patterns occur, often far from the cannonical inverse S-shape (Bruhin et al., 2010, Harbaugh et al., 2002, Harrison and Rutström, 2009, Wilcox, 2015).

The instability of parameters across studies, failures to replicate, theoretical and empirical controversies, and the existence of hundreds of competing theories all suggest there is still room to develop new 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 we believe will be of general interest.

Our model assumes only two primitives: an individual-specific parameter exogenous to the model that we refer to as a "predisposition" and an expectation about the environment which can be treated as either exogenous or endogenous to the model.[2] Critically, unlike in most economic models of choice, attitudes toward risk and losses are *not* the primitives of the model. Instead risk attitudes emerge from, and vary with, changes in individual expectations. We find that our model can account for all key phenomena described by Prospect Theory. It predicts reflection in risk attitudes, aversion to losses, and illusory distortions in probability weighting for certain predisposition and expectation levels, thus nesting the key predictions of Prospect Theory in a richer framework. However, unlike Prospect Theory and Expected Utility Theory, our approach intrinsically captures heterogeneity in individual behavior, arising from differences in expectations. Furthermore, our model makes novel predictions about behavior that can provide explanations for what currently seem to be contradictory pieces of evidence on behavior. For example, it predicts under what circumstances reflection effects will and will not occur, proposes that observed loss aversion will vary with the stake size, and generates predictions regarding the sources of heterogeneity in the estimated probability weighting functions. In the pages that follow we note that these and other novel predictions could be empirically tested in the future.

In the next section, we develop the neurobiological foundations of our functional form.

---

[2]For our initial treatment, we handle the expectation as exogenous. However, because our ultimate goal is to provide a useful positive tool we also present an endogenized form of the model which includes a precise formula for how expectations are calculated by the nervous system, which is based on standard neurobiological and dynamic programming approaches explained in Appendix B. The expectation formation literature from which we draw this endogenized form is largely based on studies in sensory systems and we acknowledge that expectations in decision-making may be more complicated than in sensory systems (Tymula and Plassmann, 2016). The investigation of these other determinants and how they may enhance the basic model presented here in both endogenous and exogenous forms lies beyond the scope of this paper.

Additional information about these foundations can be found in the appendix and in Glimcher (2011). (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, which reflect generalizations of the form we present here.

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 like ours capture the idiosyncrasies of human choice behavior captured by the Kahneman and Tversky value function, but in a completely different way? Furthermore, 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 modeling approach and suggest future tests and extensions of the model.

There are three appendixes to our paper. In the first appendix we present an empirical use of our model with a dataset consisting of the risky choices of 2,530 individuals. There we also compare our model fit to that of Prospect Theory and Expected Utility. We discuss the advantages of our approach in terms of policy implications using the gender gap in risk attitudes as an example. It is in this section that we make the case for our theory as a practical tool for positive use. The goal of the second and third appendixes is to provide more information from neuroscience on the primitives of our model — expectation in Appendix B and predisposition in Appendix C.

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). To begin, consider 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 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 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 (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. Under these conditions, therefore, the activity of the neuron encoding that central pixel is almost entirely redundant.
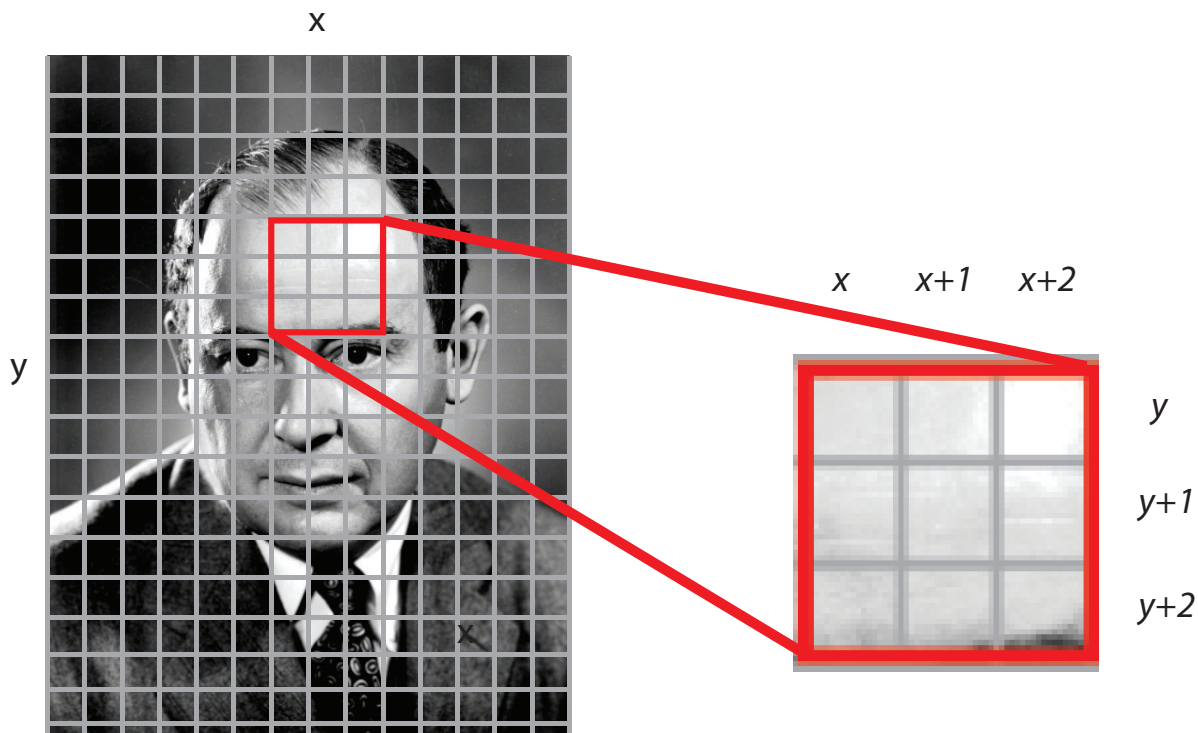
Figure 1

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

The reason 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 near linear (Lennie, 2003). While a calorie may not seem like much at first glance, these costs are surprisingly significant. Nearly 20% of the calories we consume go to support our brain even though it accounts for only about 3% of our body weight. Brain tissue is on average roughly 7 times as expensive as the average tissue in our bodies. Put another way, were we to increase the total representational capacity of our brains by a factor of 10 (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 marginal cost structure that imposes. For a tenfold increase in precision to be efficient in the economic sense, it would have to yield a tripling of total consumption wealth. The costs of a scaling up by a factor of 100 are almost unimaginable.

Faced with these high costs, 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 late 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:

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

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. α is a free parameter. 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.[3]

In the 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 across the firing rates of the neuronal set is maximized in the neuronal rate (Wainwright et al., 2002). Put another way, these results indicate that for any system that has monotone costs over a number of action potentials, a representational schema of this general form is efficient. Although less than a proof, this observation has motivated axiomatic work discussed below.

2.2 A Note About Neuronal Stochasticity

Throughout the preceding discussion, neuronal firing rates have been treated as deterministic. Some readers may find 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 the 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.[4] 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

---

[3]We neglect here the fact that firing rates are biophysically restricted to be weakly positive. For a more complete treatment of this issue see Reynolds and Heeger (2009), Louie et al. (2011), Glimcher (2011).

[4]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.

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).

2.3 Derivative Representational Theories for Decision-Variables: Statics

Now consider the implications of the 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 set is maximized in the neuronal rate (Wainwright et al., 2002). Let us begin by imagining that we are directly observing a 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 by x_i, $i \in [1, n]$. For simplicity, we might imagine 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, are a very small subset of X_4 that had some interesting structure to them? 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 often 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. The point 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 to minimize choice errors — a distribution accomplished optimally by normalization-based forms (Wainwright et al., 2002, Steverson et al., 2016).

2.4 Direct Measurements of Subjective Value

For neurobiologists studying how we see, there are tools 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 both causally responsible for choice and directly observable (unlike the utilities to which they are related). 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 Kenway 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.[5] *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 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

In support of the notion that these functional forms are important to choice theorists, 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

---

[5]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.

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 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 described by the following first-order differential equations:

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

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

where τ 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 notable aspects 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:

$$R(t) = \frac{x_t}{\sum\limits_{0}^{t-1} D(t)x_t + x_t + 1}$$

where $D(t)$ is a weighting function such that the term on the left in the denominator ($\sum\limits_{0}^{t-1} D(t)x_t$) 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 expectation*. In the next section we show that its function is to adjust the curvature of the firing rate function in accordance with 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 the 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.

Recently, an axiomatic analysis of these classes of functions has taken this argument a

step further (Steverson et al., 2016). In that paper it was shown that the models described above are, in fact, optimal maximization tools whenever representational precision is costly.

2.6 Summary

In this brief literature review we have provided some insight into the origins of the functional form we examine in the following section. Existing work has shown that this class of functional forms has normative qualities and is equivalent to reducing Shannon entropy in terms of implied behavior (Steverson et al., 2016). Studies of how the brain actually represents subjective value in the circuits causally responsible for choice show that these equations do a better job of predicting steady-state neuronal firing rates than alternatives (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 (Steverson et al., 2016). This observation suggests the possibility that some violations of Expected Utility might in fact be formally rational, trading the marginal benefits of additional precision against the marginal costs.

In the next section we present our novel analysis of the normalization function with regard to risk preferences, loss aversion, and probability weighting.

3 Model

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

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

where M_t is the expectation, or reference point, at time t and $\alpha > 0$ is a subject-specific parameter we call *predisposition* to differentiate it from preferences.[6]

---

[6]Here we focus on a simplified steady state form of the 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

The reference point, $M_t > 0$, is the prize that the individual expects to encounter in period t. For most of this paper, we do not devote attention to how this expectation is formed. Instead, for tractability, we assume that it is exogenously given and derive many behavioral predictions of this functional form that is novel to economists but established and uncontroversial in neuroscience. However, because our ultimate goal is to provide a testable model, in Appendix B we present a precise formulation of how expectations are calculated, based on standard neurobiological and dynamic programming approaches. To gain some additional theoretical insights and preserve generality at the same time, we constrain the expectation to be the time-weighted average of past prizes up until period t:

$$M_t = \sum_{0}^{t-1} D(t) x_t$$

We assume that for all t, $D(t-1) < D(t)$ which guarantees that when forming expectation, more recent prizes are given more weight than prizes in the more remote past as is indicated unambiguously by the neurobiological data, e.g. Ohzawa et al. (1985) (also see Appendix B). M_t can thus be interpreted as an experience-based expectation (a reference point) about average rewards in the environment based on 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.[7]

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 $x_1, x_2, ..., x_n$ that occur with corresponding probabilities $p_1, p_2, ..., p_n$ is calculated by taking a standard expectation:

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

exponential term α 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 because it allows tremendous generalization. In this way, we are able to focus on how the expectation term M and parameter α 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.

[7]Another source of different valuations of the same x over time is the naturally occurring stochasticity in neuronal firing rates. For a detailed analysis of 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 the results presented here. Skewed error distributions, which are of interest, could generate new predictions, which we leave for future research.

Figure 2 plots examples of subjective value functions for a range of predisposition values 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.
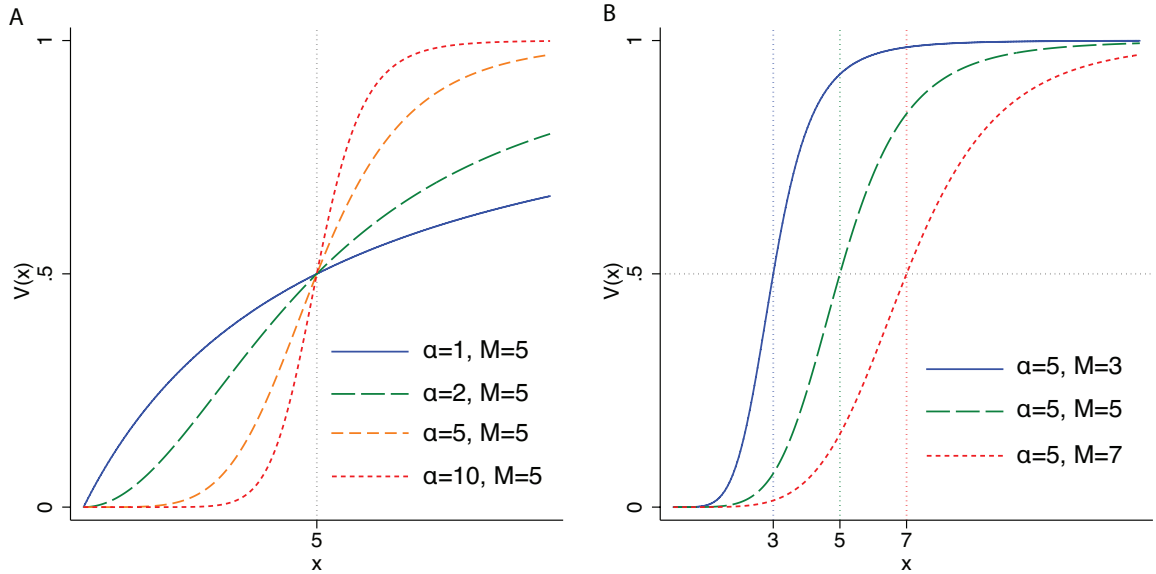

Figure 2: Examples of typical subjective value functions. In each graph, the vertical dashed lines are drawn at $x = M_t$. A: The effect of changes in the predisposition on the subjective value function. B: The effect of changes in the prize expectation (M) on the subjective value function.

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

In Figure 2 A we see that predisposition fundamentally affects the curvature of the value function. When predisposition is low ($\alpha = 1$, 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 α 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.

The other insight from Figure 2 A is that for higher predisposition levels, the function becomes steeper around the reference point. This allows decision makers with higher predisposition to maximize choice precision (assuming non-zero noise as in section 2) at expectation. In the main body of the paper we treat the predisposition, α, as a free parameter

and think of it as an individual-specific trait that is stable in the short term. In Appendix C we discuss properties of predisposition in more detail, and suggest ways to measure and manipulate it in future research so that eventually it may be rendered directly observable.

To illustrate the role of the reference point, Figure 2 B plots three different value functions with different reference points and the same predisposition. 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 α, 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.

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):

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

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 predisposition.

Proposition 1. *The individual is risk averse if*

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

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   Predisposition and Risk Attitude

Proposition 1 implies that an individual with really low-valued predisposition ($\alpha \leq 1$) is risk averse for all x, just as in the traditional utility model. At the other extreme, when the individual's predisposition has really high values ($\alpha \to \infty$), 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 much complexity. Decision makers likely are somewhere in between these two extremes: their subjective value function starts off as convex and as x increses 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 predisposition 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, an individual is risk averse. Since $\frac{\partial \sqrt[\alpha]{\frac{\alpha-1}{\alpha+1}}}{\partial \alpha} > 0$ and $\lim_{\alpha \to \infty} \sqrt[\alpha]{\frac{\alpha-1}{\alpha+1}} = 1$, as α increases the inflection point approaches M_t, meaning that the range of losses for which an individual is risk averse shrinks. ∎

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

Proposition 3. *As α 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 $x < M_t$
$\frac{\partial RRA}{\partial \alpha} = \frac{-M_t^{2\alpha}+x^{2\alpha}-2\alpha M_t^\alpha x^\alpha Log(M_t)+2\alpha M_t^\alpha x^\alpha Log(x)}{(M_t^\alpha+x^\alpha)^2}$
The denominator of this expression is always positive so we focus on numerator only. Hence, $\frac{\partial RRA}{\partial \alpha} > 0 \Leftrightarrow -M_t^{2\alpha} + x^{2\alpha} - 2\alpha M_t^\alpha x^\alpha Log(M_t) + 2\alpha M_t^\alpha x^\alpha Log(x) > 0$

which we can rewrite
$x^{2\alpha} - M_t^{2\alpha} + 2\alpha M_t^\alpha x^\alpha (Log(x) - Log(M_t)) > 0$
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 α 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, predictable from an individual's reward history. Our model therefore 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: $\frac{\partial RRA_t}{\partial M_t} < 0$
$\frac{\partial RRA_t}{\partial M_t} = -\frac{2\alpha^2 M_t^{\alpha-1} x^\alpha}{(x^\alpha + M_t^\alpha)^2} < 0$ which completes the proof.   ∎

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. It is not only the total income earned, but also the timing when each reward was received, that determines an individual's willingness to take risks. Imagine two individuals in period t, Oskar and John, who have the same predisposition, α, 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, Oskar will have a 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 leads to more risk tolerance than decreasing reward streams.*

Proof. See proof of Proposition 4.   ∎

3.1.3 Reflection Effects

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. By incorporating individual heterogeneity in predisposition and reward expectations, our model allows us to predict when reflection effects will and will not be observed at the individual level.

Figure 3 provides a graphical demonstration of 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 1,000 individuals who varied in their predisposition and expectations. In these simulations, the mean predisposition was set equal to 5 (except for panel D where predisposition 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 similar to one another, between graphs A-E we manipulated the variability of predisposition and expectation. We then calculated for each of our simulated individuals their Arrow-Pratt indexes of relative risk aversion (RRA) for a loss of 5 and a gain of 5 relative to the mean reference point of 10. Thus the RRA-gain is an individual's Arrow-Pratt index calculated at $x = 15$ and RRA-loss is the index calculated at $x = 5$. To visually assess whether individuals' 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).

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 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 α and for all M_t, individuals are less risk tolerant for higher x.*

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

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

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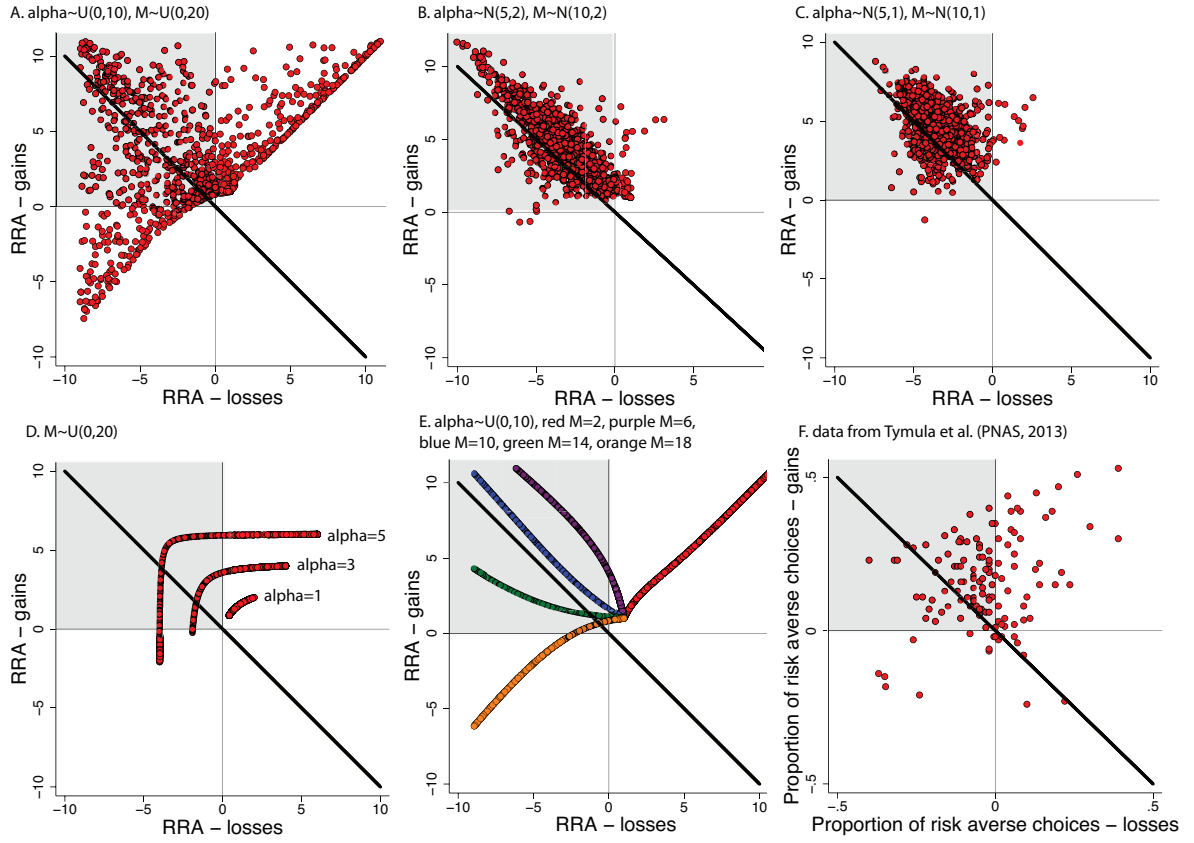
Figure 3: A-E plots of individual Arrow-Pratt relative risk aversion index in gains against the index in losses for 1,000 individuals simulated from different distributions of α and M. E: individual risk attitudes in gains against risk attitudes in losses, adapted from Tymula et al. (2013). Positive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

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 a large diversity in individual expectations and predisposition. Individuals with lower predisposition 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 predisposition 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 predisposition. The quadrant to the right of the gray one is a combination of individuals who have low predisposition and low expectations.

When we draw individual predispositions 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 α and M (OLS regression coefficient: -0.913, $p < 0.001$, Figure 3

C). Notice, however, that for these figures we picked a distribution of predispositions 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 predisposition is lower in the population, more of the observations would have fallen into the white quadrant to the right of the gray quadrant, and we would have been less likely to see reflection effects. When α is very small, all individuals would be risk averse for both gains and losses (see Proposition 1 where we established that for very low values of α, the utility function is always concave). Similarly, if we picked higher predispositions 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 predisposition, α, at three different levels and simulated the model, drawing the expectation from a uniform distribution over 0 to 20 for each α. 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 1,000) fall into the gray 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 predisposition 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 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 demonstrate 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 predispositions. 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

20

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 $\lambda > 1$. Notably, a kink in the utility function (λ) is not necessary to produce loss aversion. In fact, any utility function with diminishing marginal utility is behaviorally consistent with loss aversion (Blavatskyy and Pogrebna, 2009, Ert and Erev, 2013). In a similar way, our value function produces loss aversion even in the absence of a kink.

To make this statement we use a traditional definition of loss aversion where (λ) 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:

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

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

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

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 $\lambda > 1 \Leftrightarrow 2k^2 > 0$, which holds for all k and completes the proof. ∎

In contrast to other reference-dependent theories that assume λ is a fixed parameter, in our model loss aversion is a function of both predisposition 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 depends on both of these attributes.
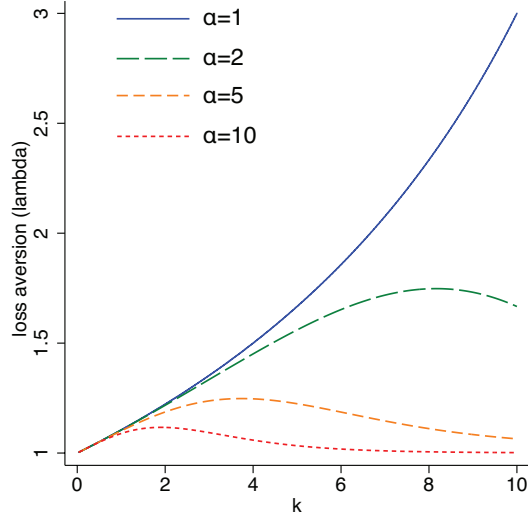

Figure 4: Loss aversion (λ) as a function of the gain (loss) size (k) at different values of predisposition (α).

Figure 4 plots loss aversion (λ) as a function of k (size of gain and loss from the reference point) at different predispositions. It is clear that both predisposition and k play important roles in determining the strength of loss aversion. In general, the model predicts that people with higher predisposition show less loss aversion. The results can potentially explain the mixed findings in the literature on the effect of stake size on loss aversion (Pope and Schweitzer, 2011, Blavatskyy and Pogrebna, 2007). For individuals with extremely low-valued predisposition, loss aversion is always increasing in stake size, however as predisposition increases, increases in stake first lead to more loss aversion and then to less aversion. The higher the predisposition, the lower the stake at which loss aversion peaks.

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

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 has very low predisposition ($\alpha = 1$) 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 (Blavatskyy and Pogrebna, 2009). Our theory is the first to provide neurobiological foundations for the lack of a 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 (List, 2002, 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 framework thus predicts the relationship between the existence and strength of the endowment effect and ownership history as documented in the experimental study by Strahilevitz and Loewenstein (1998).

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 intriguing 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 time 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.
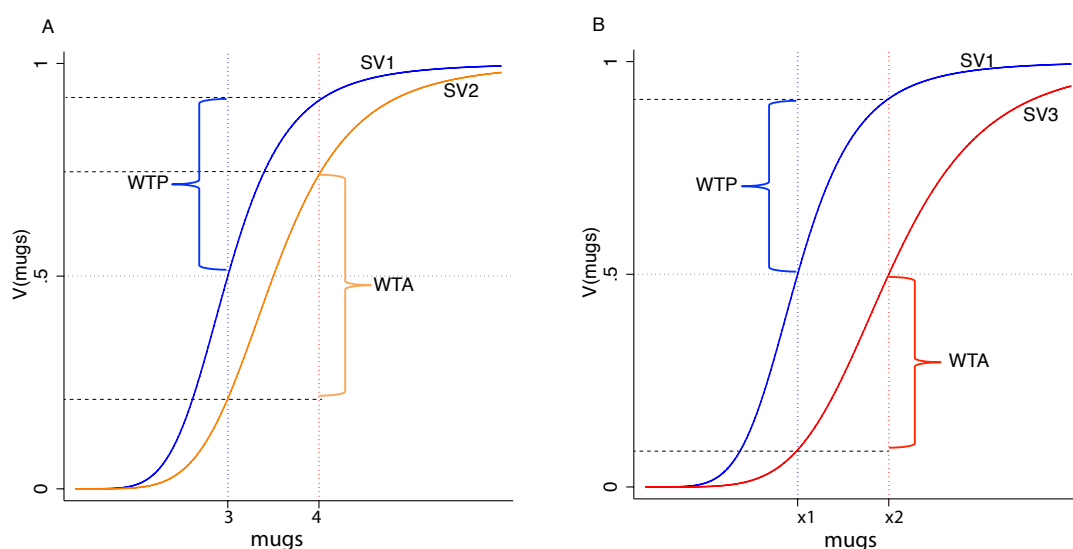

Figure 5: Endowment effect with not full (in A) and full (in B) adjustment of the reference point.

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 where 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 passes and 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 WTPs are elicited soon after recent losses in ownership such that the individual does not fully adjust downward to the new ownership level. This is a situation equivalent to the orange curve in Figure 5 A, but the figure depicts 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 she had for the first mug in the first place. Since after an endowment the utility function shifts in a parallel fashion (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 of money for an additional mug no matter whether it is her third, fourth or hundredth mug! This is, of course, inconsistent with basic consumer theory and empirical evidence. Our model does not, however, make this odd prediction (see Figure 6 B) as an individual's willingness to pay for an extra mug always declines after an endowment.

Proposition 10. *Suppose an 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
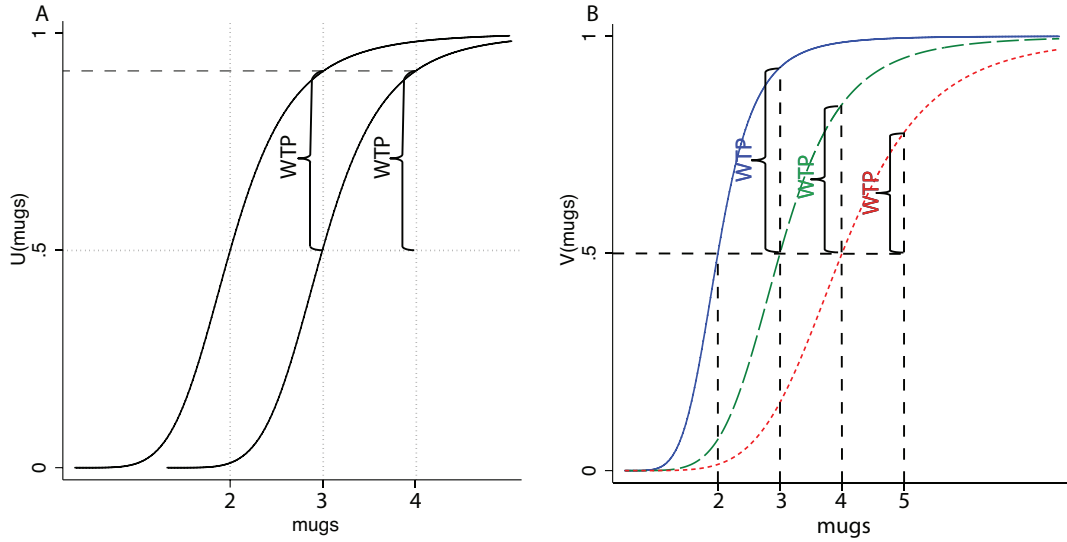
Figure 6: Endowment effect in difference-based reference point model (A) and in divisive normalization model (B).

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:

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

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

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

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

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

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

∎

3.3 Probability Weighting

To this point, we have devoted none of our paper to a discussion of subjective probabilities. Driven by experimental evidence, Savage's insight (Savage, 1954) and Allais paradox (Allais, 1953), Kahneman and Tversky (1979) proposed that people do not perceive probabilities objectively. Instead, they conjectured that people overweight small probabilities and underweight large probabilities and that this violation of the independence axiom accounts for some of the apparent irrationality in human choice behavior. While the proposed probability weighting function is assumed to be inverse S shaped, it is worth noting that empirically observed probability weighting functions estimated at the individual subject level take shapes

that are often far from the canonical inverse S (Abdellaoui, 2000, Bruhin et al., 2010, Harbaugh et al., 2002, Harrison and Rutström, 2009, Wilcox, 2015, Hsu et al., 2009, Tobler et al., 2008). One problem with the current use of the probability weighting function is that the existing models do not capture the heterogeneity in the estimated probability weighting functions on the individual level. For example, factors such as emotions (Fehr-Duda et al., 2011) have been found to affect an individual's subjective probabilities, but a coherent economic theory that would explain individual differences in probability weighting functions does not exist.

Perhaps even more importantly, much evidence on the probability weighting functions comes from joint elicitation procedures where a two- or three-parameter value function and a one- or two-parameter probability weighting function are estimated at the same time. Often overlooked is the fact that in Prospect Theory both probability weighting and utility curvature contribute to the observed risk attitudes. Fits of these two functions (value and probability weighting) to data highlight that the parameters for the two functions interact (Fox and Poldrack, 2014). *In Prospect Theory value function curvature, loss aversion and the parameters contributing to probability weighting function curvature are all collinear and are all used to account for observed risk attitudes in choice behavior.* We believe that this formal non-identifiability is very important.

Consider a theory like ours in which individuals normalize the value of the rewards but then obey the independence axiom. One could use such a model generatively to produce a series of choices in response to a set of lotteries and then could fit that choice data with Prospect Theory. If, for example, one was to use traditional CRRA preference functions and a probability weighting function to fit the data actually produced generatively by our model, one might well expect to observe a biased probability weighting function that depends on the individual-specific predisposition and expectation. What we stress here is that the violations of the independence axiom one would observe would simply result from a misspecification of the preference structure (the value function), not from any true violation of independence.

To illustrate this point, we simulated nine individuals with different predispositions ($\alpha_{low} = 1$, $\alpha_{med} = 2$, and $\alpha_{high} = 5$) and different expectations ($M_{low} = 40$, $M_{med} = 60$, and $M_{high} = 80$). Given these predispositions and expectations and assuming ESVT preferences and no distortion in probability (no violations of the independence axiom), for each of the individuals we generated their choices in the well-known decision task used by Gonzalez and Wu (1999). We then fit these choices using the standard joint estimation approach with the classic parametrization of Prospect Theory, assuming CRRA utility function and two-parameter probability weighting function ($w(p) = \frac{\delta p^\gamma}{\delta p^\gamma + (1-p)^\gamma}$). Figure 7 presents the results of this analysis.

What is striking is that our fit of Prospect Theory to the normalization-produced data replicates previously observed probability weighting function patterns, even though our in-
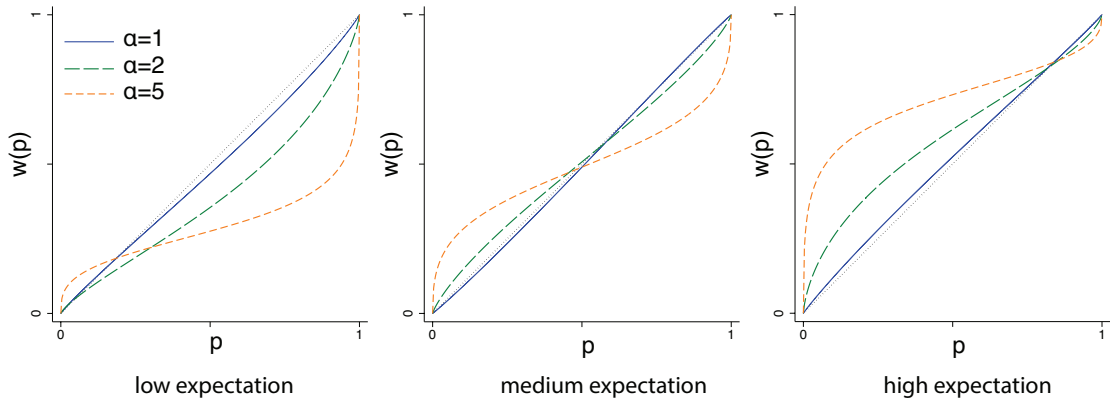
Figure 7: Estimated probability weighting functions

dividuals do not distort probabilities. Instead they employ a radically different kind of value function the one proposed by Kahneman and Tversky (1979). Individuals with the highest-valued predisposition show the most noticeable "probability distortions," which depend on expectations. Individuals with medium expectations show the traditional Prospect Theory S-shape. Individuals with low (high) expectations underweight (overweight) probability at almost all levels, *when viewed through the lens of Prospect Theory*. The smallest probability distortions occur for individuals with low-valued predisposition parameters. This is expected because their value function resembles CRRA the most and thus no distortions in probability are incorrectly inferred by the Prospect Theoretic analysis.

This simulation suggests that the estimated probability weighting functions widely inferred from studies in the human population may be an artifact of the assumed functional utility form. Observed heterogeneity in *estimated* probability weighting might in fact be due to differences in individuals' predisposition and expectation. *If choosers generated their choice behavior using the ESVT model but that choice behavior was fit by economists using Prospect Theory, those economists might incorrectly conclude that violations of the independence axiom were more widespread than is actually the case.*

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 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 and its normative properties have been recently characterized in economics and choice theory (Steverson et al., 2016). 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 practioners; 3) A brief discussion of the implications of the normative foundation on which the theory was built; and 4) 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 broadened the parameter set for the utility function by introducing the *value function* (Kahneman and Tversky, 1979). 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 curvature in the loss domain is a pure reflection of the curvature in the gain domain.) 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.

In contrast, ESVT's function *emerges* from the two core primitives of its theory: *predisposition* and *expectation.* We acknowledge that some economists will find this disturbing. In

our theory, choosers experience (or are offered) wealth shocks in time and from these derive (endogenously) 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 the existing literature in economics. Choosers are thus exogenously endowed in our theory with only a single primitive, their so-called predisposition, 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.

Thus, a key point is 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 predisposition 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 3-5 parameters (neglecting the reference point which we share with Prospect Theory). ESVT relies only on a single key parameter, predisposition, which we note in Appendix C may be related to mental 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. Finally, one then selects a functional form for the probability weighting function and employs one or two parameters for this element of the theory. When modeling the same dataset with ESVT, one relies on fitting a predisposition term and then either assuming an expectation function

(LoFaro et al., 2014) or fitting what is essentially a "forgetting rate" to the stream of shocks experienced by the chooser to empirically derive 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 in whether it can replicate the performance of Prospect Theory with fewer parameters (in Appendix A we show that in a typical dataset it performs better than Prospect Theory), but in 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, 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.

Finally, our model may be attractive to those interested in policy. Expectations and perhaps predisposition as well (see Appendix C) can be to some extent altered by policy interventions which would in turn lead to changes in behavior. Expectations can most certainly be influenced by changing people's reward history. In Appendix A we use the example of gender differences in risk attitudes, which have been suggested as one of the causes of the gender wage gap to illustrate the policy relevance of our theory. While studies using both Prospect Theory and Expected Utility Theory approaches find that men are more risk tolerant than women, these theories provide no insights into why this is the case. With ESVT, we find that these differences are driven by lower reward expectations among women. This finding clearly suggests that a policy intervention that increases women's expectations should make women willing to take as many risks as men.

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), Robson and Whitehead (2016)) economic thinking. Neurobiological evidence available today clearly indicates that increasing the precision of any neural representation is without a doubt costly. 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*.

It is clear that ESVT invokes a slightly different notion of consistency than do traditional models (see Steverson et al. (2016) for more details). 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   Future Directions

Just like our brains, this paper too has limited capacity. The novel predictions of our theory remain to be tested and verified in future work. Experimental studies that manipulate the history of rewards would help us understand whether our predictions are correct and will guide refinements of the model. The divisive normalization function we chose to use is purposefully simplified and there are important ways we would like to see it enriched. In particular, we did not elaborate much on the expectation term, M, (although see Appendix B for a description of how expectations are computed in the brain and how they can be handled econometrically). We leave it up to future extensions of the model to relate more precisely 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 relevant implications for behavior. Nevertheless, the model provides the framework where skew and kurtosis can easily be incorporated and studied both theoretically and empirically.

Despite some progress (Abler et al., 2006, Preuschoff et al., 2006, Berns et al., 2008, Hsu et al., 2009, Tobler et al., 2008) we 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 that would be desirable from a normative standpoint. 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.

References

Abdellaoui, M. (2000). Parameter-Free Elicitation of Utility and Probability Weighting Functions. *Management Science 46*(11), 1497–1512.

Abler, B., H. Walter, S. Erk, H. Kammerer, and M. Spitzer (2006, jun). Prediction error as a linear function of reward probability is coded in human nucleus accumbens. *NeuroImage 31*(2), 790–795.

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.

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.

Berns, G. S., C. M. Capra, J. Chappelow, S. Moore, and C. Noussair (2008, feb). Nonlinear neurobiological probability weighting functions for aversive outcomes. *NeuroImage 39*(4), 2047–2057.

Blavatskyy, P. and G. Pogrebna (2007, jun). Loss aversion? Not with half-a-million on the table! *SSRN Electronic Journal*.

Blavatskyy, P. and G. Pogrebna (2009, may). Endowment effects? "Even" with half a million on the table! *Theory and Decision 68*(1-2), 173–192.

Bruhin, A., H. Fehr-Duda, and T. Epper (2010). Risk and Rationality: Uncovering Heterogeneity in Probability Distortion. *Econometrica 78*(4), 1375–1412.

Caplin, A. and M. Dean (2008, may). Dopamine, Reward Prediction Error, and Economics. *Quarterly Journal of Economics 123*(2), 663–701.

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.

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.

Fehr-Duda, H., T. Epper, A. Bruhin, and R. Schubert (2011). Risk and rationality: The effects of mood and decision rules on probability weighting. *Journal of Economic Behavior and Organization 78*(1-2), 14–24.

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 Edi ed.)., 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.

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).

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.

Gonzalez, R. and G. Wu (1999, feb). On the shape of the probability weighting function. *Cognitive psychology 38*(1), 129–66.

Grubb, M., A. Tymula, S. Gilaie-Dotan, P. W. Glimcher, and I. Levy (2016). Neuroanatomy accounts for age-related changes in risk preferences. *Nature Communications*.

Harbaugh, W., K. Krause, and L. Vesterlund (2002). Risk attitudes of children and adults: Choices over small and large probability gains and losses. *Experimental Economics 5*, 53–84.

Harrison, G. W. (2008). Maximum likelihood estimation of utility functions using Stata.

Harrison, G. W. and E. E. Rutström (2009, jun). Expected utility theory and prospect theory: One wedding and a decent funeral. *Experimental Economics 12*(2), 133–158.

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.

Hsu, M., I. Krajbich, C. Zhao, and C. F. Camerer (2009, feb). Neural Response to Reward Anticipation under Risk Is Nonlinear in Probabilities. *Journal of Neuroscience 29*(7), 2231–2237.

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.

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 28*(12), 3017–29.

Köszegi, B. and M. Rabin (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics 121*(4), 1133–1165.

Lattimore, P. K., J. R. Baker, and A. D. Witte (1992). The influence of probability on risky choice: A parametric examination. *Journal of Economic Behavior & Organization 17*(3), 377–400.

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. *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 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 34*(48), 16046–57.

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 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.

Preuschoff, K., P. Bossaerts, and S. R. Quartz (2006, aug). Neural Differentiation of Expected Reward and Risk in Human Subcortical Structures. *Neuron 51*(3), 381–390.

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.

Robson, A. and L. A. Whitehead (2016). Rapidly adaptive hedonic utility.

Robson, A. J. (2001). The Biological Basis of Economic Behavior. *Journal of Economic Literature 39*(1), 11–33.

Rutledge, R. B., M. Dean, A. Caplin, and P. W. Glimcher (2010). Testing the Reward Prediction Error Hypothesis with an Axiomatic Model. *Journal of Neuroscience 30*(40).

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.

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.

Slovic, P. (1995). The construction of preference. *American Psychologist 50*(5), 346–371.

Steverson, K., A. Brandenburger, and P. Glimcher (2016). Rational imprecision: information processing, neural, and choice-rule perspectives.

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.

Sutton, R. S. and A. G. Barto (1998). *Reinforcement learning: an introduction*. Cambridge, MA: MIT Press.

Tobler, P. N., G. I. Christopoulos, J. P. O'Doherty, R. J. Dolan, and W. Schultz (2008). Neuronal distortions of reward probability without choice. *The Journal of Neuroscience 28*(45), 11703–11.

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*.

Weber, E. U. and C. Hsee (1998). Cross-Cultural Differences in Risk Perception, but Cross-Cultural Similarities in Attitudes Towards Perceived Risk. *Management Science 44*(9), 1205–1217.

Wilcox, N. (2015). Unusual Estimates of Probability Weighting Functions.

Wilson, H. R. and R. Humanski (1993). Spatial frequency adaptation and contrast gain control.

Woodford, M. (2012, may). Prospect Theory as efficient perceptual distortion. *American Economic Review 102*(3), 41–46.

Appendix A   Empirical Fit of the Model

We fit our model to a dataset that includes the risky decisions of 2,530 individuals. These individuals visited the National Academy of Sciences Museum in Washington, DC where as a part of the LifeLab exhibit they had an opportunity to assess their own risk attitudes. Risk attitudes were measured using a standard economic task in which participants are asked to choose between a certain option ($5 for sure) and a lottery with changing reward ($5, $8, $20, $50, and $125) and changing probability of winning this reward (13%, 25%, 38%, 50%, and 75%). The task involved a total of 40 decisions per individual — 25 in the risky condition and 15 in the ambiguous condition. Here we analyze data from the risky condition only. One randomly selected participant per month was paid based on one of his/her own choices that was randomly selected as well. Payments were delivered as gift cards. For more details on the procedure, see Glimcher (2015).

In this dataset, as in most, we do not have any information on people's expectations and we thus treat them as free parameters. To begin we estimate ESVT parameters and compare our model fit to Prospect Theory and Expected Utility. In our estimation procedures we follow Harrison (2008). For ESVT, we assume the value function as specified in our model in Equation 1 and *no distortion in the perception of probabilities*. For Prospect Theory, we assume a CRRA utility function (since we did not examine losses) $(u(x) = x^r)$ and a probability weighting function as in Lattimore et al. (1992) $(w(p) = \frac{\delta p^\gamma}{\delta p^\gamma + (1-p)^\gamma})$. For Expected Utility, we assume CRRA utility function $(u(x) = x^r)$ and no distortion in probability weighting.

We then created an index that says that an individual selects the risky option whenever the expected value of the risky option (v_r) is larger than the value of the safe option (v_s) subject to some additive stochastic error $\epsilon \sim N(0, \sigma^2)$:

$$v_r - v_s + \epsilon > 0$$

We then related this latent index to our choice data using a logistic choice function where the probability of choosing the risky lottery is given by:

$$P(x, p) = \frac{1}{1 + exp(-(v_r - v_s)/\sigma)}$$

We fit the data using a maximum likelihood procedure, clustering standard errors on the subject level. Table C1 presents the results. Out of all the models, ESVT performs the best according to Akaike and Bayesian information criteria. Prospect Theory comes second.

In the next step we demonstrate an example of how ESVT can generate additional insights, that go beyond insights gained from either Prospect Theory or Expected Utility Theory, into why decision-making varies with demographic variables such as gender and age. 1,287 out of our 2,530 participants in our sample are male. The average age in the sample

Table A1: Paramter estimates of ESVT, Prospect Theory, and Expected Utility. Standard errors clustered on participant are in parenthesis. Akaike and Bayesian information criterion is in the last two rows.

A. ESVT		B. Prospect Theory		C. Expected Utility	
α	1.153	r	0.185	r	0.456
	(0.019)		(0.014)		(0.004)
M	12.248	δ	3.332	$noise$	0.739
	(0.351)		(0.300)		(0.014)
$noise$	0.084	γ	0.661	AIC	57632
	(0.002)		(0.019)	BIC	576511
AIC	55427	$noise$	0.198		
BIC	55456		(0.020)		
		AIC	56813		
		BIC	56849		

is 37.4 years with a standard deviation of 14.99 years. To estimate the relationship between age and gender variables and the parameters of the models, we replace the parameters in the models that we so far treated as constants (r in the Expected Utility model and α and M in the ESVT model) with a linear combination of the constant, age and gender (Harrison, 2008).[8] Using a standard CRRA utility function we find that men in our sample are more risk tolerant than women (see Table C2 A). Such a gender difference in risk attitudes is often observed in the lab and has been suggested to be one of the reasons behind the gender wage gap. Using this standard approach to estimating risk attitudes we learn that the difference between genders exists but such approach does not suggest ways in which policy makers could make this difference smaller.

Fitting our participants' choices with ESVT, we gain what appears to be an additional insight. We find that the gender variable does not significantly affect predisposition but that the difference in risk taking is explained by differences in expectations. Men have significantly higher expectations and therefore behave in a more risk averse manner in line with Proposition 4. This result could have practical implications for policy, suggesting that if we increase women's expectations, they will be willing to take more risks, an observation compatible with Weber's empirical studies suggesting that carefully shaping expectations can fundamentally alter risk taking behavior (Weber and Hsee, 1998).

Interestingly in our sample we do not see any effect of age on risk attitudes using the standard CRRA utility function. One could infer from this that utility function does not change as people age. Fitting our participants' choices with ESVT, however, we find evidence

---

[8]$r = r_{constant} + r_{male}male + r_{age}age$, $\alpha = \alpha_{constant} + \alpha_{male}male + \alpha_{age}age$ and $M = M_{constant} + M_{male}male + M_{age}age$

Table A2: Gender and age effects through the lens of the Expected Utility and the Expected Subjective Value Theory. Standard errors, clustered on the level of individual, are reported in parentheses. $^*p < .05$, $^{**}p < .01$, $^{***}p < .001$.

A. Expected Utility			B. ESVT	
α			α	
male	0.056***		male	0.050
	(0.007)			(0.041)
age	0.000		age	−0.005***
	(0.000)			(−0.001)
constant	0.429***		constant	1.302***
	(0.010)			(0.064)
σ			M	
constant	0.736***		male	1.687**
	(0.015)			(0.581)
No. of obs	62928		age	0.146***
				(0.025)
			constant	6.478***
				(0.852)
			σ	
			constant	0.082***
				(0.002)
			No. of obs	62928

that suggests that such a conclusion might be wrong. Consistent with lifetime wealth effects we find that expectations increase. Predisposition, on the other hand, decreases. These two changes push risk attitudes in opposite directions. The decline in predisposition is pushing older people toward more risk averse behavior while an increase in expectations is pushing toward more risk tolerant behavior. As these effects counteract one another behavior remains the same, but the value function underlying these choices changes.

Appendix B   Computation of the Expectation Term

In the main body of this report we place minimal restrictions on how the reference point, or expectation (M_t) is computed. In this we seek to follow in the footsteps of previous work like Kahneman and Tversky (1979), which also imposed minimal structure on the reference point. We do, however, have strong neurobiological insights into how expectations are functionally computed in the human brain and this data guides our selection of a functional form for the term M_t, when one is required.

A large body of evidence has documented the existence of a well-studied brain system for computing an experience-based expectation of future reward (for a review see Glimcher (2011)). This system, sometimes called the dopaminergic reward prediction error system, employs a simple recursive computation that computes a recency-weighted average of past rewards. Computation in this brain system is typically modeled using discrete time and at each time step t, the brain computes the reward expectation P_t as:

$$P_t = P_{t-1} + \gamma(x_t - P_{t-1}) \tag{7}$$

Where x_t is the subjective (or in some applications of the model the objective) value of the reward delivered to the subject at time t, in cardinal units. γ is a number between 0 and 1 that effectively sets the forgetting rate of the averaging system.

This iterative computation performed in known brain circuits is equivalent to the following function (although it imposes simpler requirements for data storage):

$$P_t = (1 - \gamma)^t P_0 + \sum_{i=1}^{t} \gamma(1 - \gamma)^{t-i} x_i \tag{8}$$

Following the reinforcement learning literature (Sutton and Barto, 1998) we assume that initial expectations are equal to zero. For the ESVT model we present here we thus write the expectation (M_t) as:

$$M_t = P_{t-1} = \sum_{i=0}^{t-1} \gamma(1 - \gamma)^{t-1-i} x_i \tag{9}$$

As can be seen, this mechanism computes a time-weighted average of previous reward magnitudes and can be fit with a single free parameter γ.

It is important to note that the fact that a known brain system performs exactly this computation is not controversial in neuroscience. Thousands of papers have been devoted to this topic (for review see Glimcher (2011)). Each element of the computation has been studied at the level of individual brain cells, at the level of cell assemblies, and at the level of brain regions. Techniques ranging from brain imaging to techniques that monitor brain activity at subcellular levels have been employed, and species ranging from humans to monkeys to rodents have also been examined. We stress that a host of evidence now clearly indicates that this computation is being actively performed by a known brain system in the service of learning the conditional expected utilities of environmental stimuli and personal actions. While we do not yet know for certain whether these circuits communicate this expectation to brain areas that compute normalization-based representations, there is good reason to believe that they do (Caplin and Dean, 2008, Rutledge et al., 2010). For the purposes of this paper that may not be critical, however. What is critical is that when ESVT is implemented as an alternative to Prospect Theory, it does provide a simple functional form for computing expectation with a single free parameter $\gamma \in (0, 1)$.

Appendix C   Predisposition

At this point we do not know with certainty how to relate predisposition (α) to the nervous system, as it has been mostly treated as a free parameter in the research on sensory systems. In one notable exception, Webb et al. (2014) refer to the exponent of the normalization function (in a related functional form) as *capacity*, and show that individuals with higher capacity have a bigger budget of action potentials to represent the value of options in a choice set. In terms of interpretation, recalling that each individual neuron can produce only a fixed maximal number of action potentials per unit of time, we infer that an increased maximum on the total number of action potentials is effectively synonymous with the statement that a larger number of total neurons is committed to the representation as capacity increases. Although the argument in Webb et al. (2014) does not easily extend to our model, our intuition is to look for the representation of predisposition in the gray matter volumes (number of cells) in areas of the human brain.

Indeed, we have recently shown that risk attitude itself relies on the number of neurons in a highly specific brain region (posterior parietal cortex) believed to play a role in risky decision-making (Gilaie-Dotan et al., 2014, Grubb et al., 2016). We measured the number of nerve cells using brain scan data in each part of the human cortex. We found that the number of nerve cells in this one brain area accounts for about half of the inter-individual variance in risk attitude and it does so better than chronological age. Lower cell densities (and hence

lower total numbers of allocable action potentials) are associated with higher levels of risk aversion. This evidence is only correlational. Studies that temporarily reduce the maximum rate at which action potentials can be produced, either via transcranial magnetic stimulation or pharmacology interventions, would be needed to test for the causal relationship between gray matter volume and behavior.

Although this is just speculation, the interpretation of the predisposition parameter as representing gray matter volume is quite intuitive at this stage. It *is* consistent with the empirical finding in Appendix A that predisposition declines with age. It is a well-known fact that as people age, they loose gray matter volume. It is also consistent with how our functional form is affected by predisposition. In particular, we can show that the upper bound on our value function is lower for individuals with lower predisposition consistent with having less neurons to represent value. This is established in the following proposition:

Proposition 11. *Suppose an individual uses the value function, V, to represent value of rewards drawn from distribution with mean M and some variance. Individuals with higher-valued capacity parameters α have a higher upper bound on the subjective value function for all $x \neq \infty$.*

Proof. We need to prove that $\frac{\partial V(x)}{\partial \alpha} > 0$.

$\frac{\partial V(x)}{\partial \alpha} = \frac{M^\alpha x^\alpha (log(x) - log(M))}{(M^\alpha + x^\alpha)^2}$

It is straightforward to see that $\frac{\partial V(x)}{\partial \alpha} > 0 \iff x > M$.

∎

In other words, the bounds on the value function are more relaxed for individuals with a higher-valued predisposition. The demonstration of this result can be seen in the Figure 2A that shows the subjective value function for different values of predisposition. Indeed, there we see that V_{max} (imposed to be at $x = 10$ in this figure) is higher for higher predispositions. One may now see the similarity between predisposition, α, in our model and "bounded happiness" condition in Rayo and Becker (2007) that constrains the value function to fall within some bounded range ($V(x) \in [V_{min}, V_{max}]$).

Another key property of our model is that a higher predisposition also corresponds to a steeper value function around the expectation. This is a feature that many have employed in theoretical models of choice (for example see Rayo and Becker (2007), Netzer (2009), Woodford (2012), Robson and Whitehead (2016)) arguing that the optimal utility function that minimizes choice errors will be the steepest around the expectation where the chooser is expecting to be making choices. Although such steepening is intuitive, at this point we do not have a good sense of how to directly relate it to biology. We suppose that this will require an extension to the model that incorporates how variance in experienced rewards affects the value function. This is something that we find very interesting but leave outside the scope of this paper.