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

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

September 27, 2018

Abstract

We present an overview of a novel 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 reference point and a free parameter we call "predisposition". We demonstrate the breadth of the model, showing that it captures many of the behavioral phenomena described by Prospect Theory, but unlike Prospect Theory i) accounts for individual heterogeneity in parameters, ii) employs fewer parameters than prospect theory, and iii) 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 alternative explanations for the endowment effect and the probability weighting function.

JEL: D03, D81, D87

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

1 Introduction

There is growing evidence that during the last seven 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 now predict a number of interesting human and animal choice behaviors that had previously escaped description in the economics literature (e.g. Louie et al. (2013), Khaw et al. (2017)). Importantly, these forms have normative foundations. Steverson et al. (2016) have shown that

---

[*] Neuroscience Institute, New York University School of Medicine, 435 East 30th St, 11th Floor, New York, NY 10016, USA

[†] School of Economics, University of Sydney, A02, 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 specific class of the utility-like functions used in these models, allows agents to optimally balance the costs of reducing stochasticity (or equivalently increasing choice precision) against the gains from more accurately representing the utility of options, as has been observed biologically in the mammalian brain. This normative work also demonstrates that these utility-like functions impose a uniqueness on the represenation making it amenable to neurobiological study. Neuroscientists have long argued that these models also ensure that information encoded by the nervous system is minimally redundant (Wainwright et al., 2002) — a closely related conclusion to Steverson et al. (2016). In this paper we explore one of the simplest members of this class of models to determine whether or not (as neurobiologist have often claimed) it has interesting choice theoretic features. We find that the model has many interesting homologies to Prospect Theory, capturing a wide array of behaviors that Prospect Theory (Kahneman and Tversky, 1979) was designed to describe. We find that it also appears to capture some features of behavior that have traditionally lain outside the domain of Prospect Theory.

In this paper we reveal three central features of this class of models that have previously gone undescribed: 1) that the utility-like construct is cardinal (aligning well with Steverson et al. (2016) observation of uniqueness) and measured in the number of *action potentials* produced by a group of neurons per second ("firing rate"); hence to distinguish it from traditional utility we refer to it as *subjective value* in the paper; 2) the subjective value function is both bounded and has finite precision (aligning with Steverson et al. (2016) observation of efficiency); and 3) the subjective value function adjusts dynamically to changes in the reference point via a divisive functional form (an observation also relating to informational efficiency).

Although it is not a widely accepted fact in economics, it has long been argued in the biological literature that biologically feasible value functions are necessarily bounded and show finite precision due to the biological limitations of the nervous system: Information in the brain is conveyed by the rate with which neurons produce electrochemical impulses, called action potentials. The mechanism for generating these action potentials is imperfect and hence introduces noise into the representation. The number of neurons that are used to evaluate rewards is limited (see Bartra et al. (2013) for locations of reward-coding brain regions) as is the maximum rate at which each neuron can produce action potentials. As a result of these constraints, there is a meaningful limit on the information carried in any neural representation, a limit which takes the form of a maximum value (as measured in action potentials per second) that a brain can associate with any given choice object. The fact that the value function is biophysically bounded and imprecise is likely the reason why it is optimal for individuals to evaluate options relative to the reference point and not in absolute terms using representation of this type (Steverson et al., 2016, Rayo and Becker, 2007, Netzer, 2009, Glimcher, 2010, Woodford, 2012, Robson and Whitehead, 2016), a point which we develop further below.

Our initial focus here, however, is to explore the novel-to-economists functional form via which the reference point enters utility-like constructs (the subjective value function) in these widely used models. We stress that this functional form is uncontroversial in neuroscience, is normative (Steverson et al., 2016, Wainwright et al., 2002), and is grounded in years of study (see Appendix C, Louie et al. (2015), Glimcher (2015)). Our goal in this paper is not to innovate, or even so much to test, this functional form. It is rather to reveal its choice theoretic features from an economist's point of view. Here, for the first time, we explicitly apply this normative form to the domain of economic choice to generate predictions about behaviors traditionally explained through the lenses of risk attitudes, loss aversion, and probability weighting. The reference point, loss aversion, 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. Nevertheless, the generality of these properties has been challenged in the empirical and theoretical (e.g. Bordalo et al. (2012a,b, 2013), Kőszegi and Rabin (2007)) literature. Indeed, loss aversion as traditionally parameterized and 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 reflection in risk attitudes eludes 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 canonical inverse S-shape (Bruhin et al., 2010, Harbaugh et al., 2002, Harrison and Rutström, 2009, Wilcox, 2015). We find that these neurobiological models, in the specific functional form we examine, are able to address all these issues in a novel way that we believe will be of general interest to some scholars in economics.It is the elucidation of these properties of these models that is the subject of this paper.

The model we present assumes only two primitives: a free parameter exogenous to the model that we refer to as a "*predisposition*" and a reference point which has previously been treated as either exogenous or endogenous to the model depending on the study. Critically, unlike in most economic models of choice, attitudes toward risk and losses are *not* the primitives of the models in this class or the exemplar we describe. Instead risk attitudes emerge from, and vary with, changes in individual reference points as a function of predisposition. Our primary finding is, perhaps surprisingly, that the model can account for nearly all of the key behavioral phenomena described by Prospect Theory. It predicts reflection in risk attitudes, aversion to losses, and illusory distortions in probability weighting for certain predisposition and reference point levels. Furthermore, the model makes novel predictions about behavior that might at some future date provide explanations for what currently seem to be contradictory pieces of evidence on behavior. For example, it predicts the generally observed reflection effects on the population level, but not necessarily on the individual level. It 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. Additionally, it presents a novel explanation for the endowment effect that is based on how the reference point is updated rather than on loss aversion. Finally, the lack of a discontinuity in the utility-like function at the reference point (that is the "kink" of the Prospect Theory Value Function) means that, if needed, the reference point can be estimated, making the model attractive for empirical and policy analysis (we return to this point in Appendix A when we discuss a possible reason behind gender differences in risk attitudes implied by ESVT). In the pages that follow we note several novel predictions that could be empirically tested in the future if the model continues to be of deep interest to neuroscientists, psychologists, and economists. In this paper our goal is, however, quite limited. We seek to examine the classes of behaviors that this model can, and cannot, capture and to establish the usefulness of this normative functional form as a model of choice. We therefore present our findings at a relatively high level and due to space constraints cannot give the due credit to all important work that has been done on each of the issues that we touch on. We hope authors in these many areas will indulge our goal of presenting the model in this manner.

In the next section, we present the model and 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 by neurobiologists to use in the sensory systems and decision-making systems of the human brain, and then performed an argmax operation on these *subjective values*, what kinds of preference structures would one expect to see in choice behavior? Would one see phenomena like those Kahneman and Tversky sought to explain with Prospect Theory? Could such a representational theory capture the idiosyncrasies of human choice behavior captured by the Kahneman and Tversky value function, but in a different way? Furthermore, could such a representation outperform Prospect Theory as a descriptive tool under at least some conditions while opening a novel normative approach?

In section 3 we present the novel aspects of this neurobiological modeling approach and suggest future tests and extensions of the model.

There are three appendixes to our paper. In Appendix A we present an empirical use of the model with a dataset consisting of the risky choices of 2,530 individuals. There we compare the model to that of Prospect Theory and Expected Utility. We discuss the possible advantages of this neurobiologically derived approach in terms of policy analysis using the gender gap in risk attitudes as an example. It is in this section that we explore the possible use of the theory as a practical tool for positive use. In Appendix B, we provide more information from neuroscience on the calculation of the reference point. Here we detail operationally how the reference point can (though imperfectly) be endogenized into the theory based on recent neurobiological and psychological experiments (e.g. Khaw et al. (2017)). The issue of whether or not the reference point should be treated as exogenous to the model has not yet been resolved in the neurobiological literature so we attempt to remain fairly agnostic on this point, presenting the model with the reference point both endogenized and exogenized so that the reader may draw her own conclusions. In Appendix C, we develop the neurobiological foundations of our functional form for readers interested in the model's early origins and recent developments. Additional information about these foundations can be found in Glimcher (2011). 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 in the main body of the paper. A reader uninterested in the foundations of our specific functional form, or uncomfortable with the engineering-like approach of computational neurobiology, might choose to skip this last appendix.

2 Model

Consider a set of prizes $\mathbf{X} \in \mathbb{R}_+$. Following the literature summarized in Appendix C, we assume that the subjective value (the "decision value"(Kahneman and Tversky, 1984, Hare et al., 2008)) of a prize i at time t, $x_{i,t} \in \mathbf{X}$, is given by:

$$S_t(x_{i,t}) = \frac{x_{i,t}^{\alpha}}{x_{i,t}^{\alpha} + M_t^{\alpha}} \tag{1}$$

where M_t is the reference point at time t and $\alpha > 1$ is a free parameter we call *predisposition* to differentiate it from traditional preference parameters. Appendix C describes the origins and extensions of this formula widely employed in the neurobiological literature.

The model allows for nominal losses after they are projected onto $\mathbb{R}_+$. The subjective value function takes values between 0 and 1 ($S \in [0, 1]$) consistent with the idea that the decision makers are bounded in the range of values that they can biophysically assign to the rewards. Unlike traditional utility (but following Steverson et al. (2016)), the subjective value function is cardinal. Its unit is firing rate, that is the number of action potentials per second generated by neurons representing the subjective value of an alternative, here projected onto the unit line segment for simplicity. To transform S into true action potential units, it is typically multiplied by 120, the true maximal action potential rate of most neurons. We forego this transformation here as trivial. For direct measurement purposes by biologists, the neuronal firing rates are assumed to map linearly to subjective value. The subjective value of $x_{i,t}$ is divisively normalized by the reference point and itself. The fact that it is normalized by itself guarantees that the subjective value function always adjusts to the problem at hand.

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:

$$ES_t(x_{i,t}, p) = \sum_{i=1}^{n} p_i S_t(x_{i,t}) \tag{2}$$

2.1 Reference point

In economics, the reference point was initially conceptualized as *status quo* (Kahneman and Tversky, 1979), but recent extensive evidence has suggested that it may be better described as an expectation (Ericson and Fuster, 2011, Abeler et al., 2011, Gill and Prowse, 2012, Gneezy et al., 2017). Although the importance of the reference point is well established, 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.[1]

However, in the neuroscientific models of this class, which rely on a different set of observables, a tremendous amount of directly observed data is available on reference points and expectations — particularly in the domain of sensory perception (for a good example see Ohzawa et al. (1985)). 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 (Wilson and Humanski, 1993, Carandini and Heeger, 2012) as are models of how the recent history of offers alters choice and valuation in decision making (LoFaro et al., 2014, Khaw et al., 2017). Interestingly, these models and measurements of neurons show reference dependence that almost exactly 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

---

[1] 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)

around the expectation of future rewards. For this reason some neurobiological writers have treated the reference point as endogenous to this class of models while others have followed Kahneman and Tversky in treating the reference point as exogenous to the model - even if it is directly observable in the brain.

Perhaps the clearest example of an effort to endogenize the reference point comes from biological studies of the dopaminergic reward prediction error system which guides some classes of reward-related behavior. This particular neurobiological system computes a reference point (in a similar way), by taking time-weighted average of past rewards (Bayer and Glimcher, 2005). Because this is the most widely employed endogenization of the reference point in these classes of models, in this paper we employ this specific functional form for M_t, when we explore the feasibility of reference point endogenization.

$$M_t = \sum_{\tau=0}^{t-1} \gamma(1-\gamma)^{t-1-\tau} x_{i,\tau} \tag{3}$$

In this formula, $\gamma < 1$ is the forgetting rate that captures the observation that more recent rewards have a bigger impact on the reference point than older rewards. In Appendix B we briefly summarize how this reference point has been derived using standard dynamic programming approaches. M_t in our presentation (when endogenous) is modelled on these neurobiological objects and can thus be interpreted as a reference point computed as the experience-based expectation about average reward in the current environment, based on the rewards in the recent past. Notice critically that the subjective value of the same prize x_i may be different at different points in time if the reference point (M_t) changes, a key feature of choice called out earlier by Köszegi and Rabin (2006).

We note critically that this widely used reference point is relatively simple (too simple for many environments) because it derives from the neurobiologically studied reference points. One can say that in these models current and past experiences and nothing else form the reference point. Of course, we know that the reference point is much more complicated than this implies (e.g. Tymula and Plassmann (2016)). In this paper we do not aim to provide the ultimate definition of the reference point for choice. The empirical literature in economics suggests that not only received rewards but also information on rewards[2], social comparisons, emotions, other alternatives in the choice set[3] and other factors likely influence

---

[2] Although recent evidence on the endowment effect provides suggestive evidence that information alone cannot shift the reference point and generate endowment effect and instead physical experience of the good is necessary to affect valuation (Heffetz and List, 2014, Khalil and Wu, 2017)

[3] See Appendix C to learn how divisive normalisation model captures the impact of other options in the choice set on valuation. The main conclusion from the study of such "spatial" normalization is that the size of the choice set and value of its items affect individual's ability to pick the best alternative in very specific, and confirmed in the data, ways. However, changing other alternatives in the current choice set does not affect individual's risk attitude (in any other way than through making the individual a more noisy chooser). For tractability we do not include the other options in the reference point in out paper.

reference point, even though these factors play little role in the neurobiology of sensory systems. Perhaps individuals are even able to foresee how their decisions alter their future reference point and if yes, maybe they include this knowledge in their decision making that involves delayed outcomes. More work is needed to establish how these additional features combine with the reference point we employ here. But at this time no other endogenization has been well explored in the neurobiological choice literature. For this reason, we conclude that it may often be appropriate to render M_t exogenous to this model. This has the clear advantage of rendering the model more general and closer to many previous models, and it is for this reason that we are rather exploratory in our treatment of M_t.

Finally, we note that a keay feature of this class of models is the lack of a discontinuity at the reference point, which allows for easy estimation of the reference point and the factors that affect it (see Appendix A for an example) when M_t is treated as exogenous.

2.2 Predisposition (α)

Predisposition is a free parameter. From the empirical work in neuroscience, we know that it is a necessary component of the model that drastically improves empirical model fits. In visual cortex it has usually been estimated to be between 1.0 and 3.5 (Albrecht and Hamilton, 1982, Sclar et al., 1990, Busse et al., 2009, Carandini and Heeger, 2012). Neuroscientists agree that this exponent captures the nonlinear response of the neuronal activity to inputs (Heeger, 1992). More recently, Webb et al. (2016) related it to the resource constraints placed on neural systems and fit behavioral data to α. More theoretical and empirical work is necessary to understand the normative foundations of predisposition and whether we should think about it as individual-specific or environment-specific parameter.

2.3 Model illustration

Figure 1 plots examples of subjective value functions for a range of predisposition values and reference points. The model can produce a variety of familiar value functions: a value function concave for all $x_{i,t}$ that looks like the utility function introduced by Bernoulli, as well as the more modern S-shapes used in Prospect Theory (Kahneman and Tversky, 1979) and other reference-dependent models.

To give the reader a quick overview of the role of the predisposition and reference point in the subjective value function, in Figure 1 A we manipulate predisposition while keeping the reference point constant, and in Figure 1 B we manipulate the reference point while keeping the predisposition constant.

In Figure 1 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

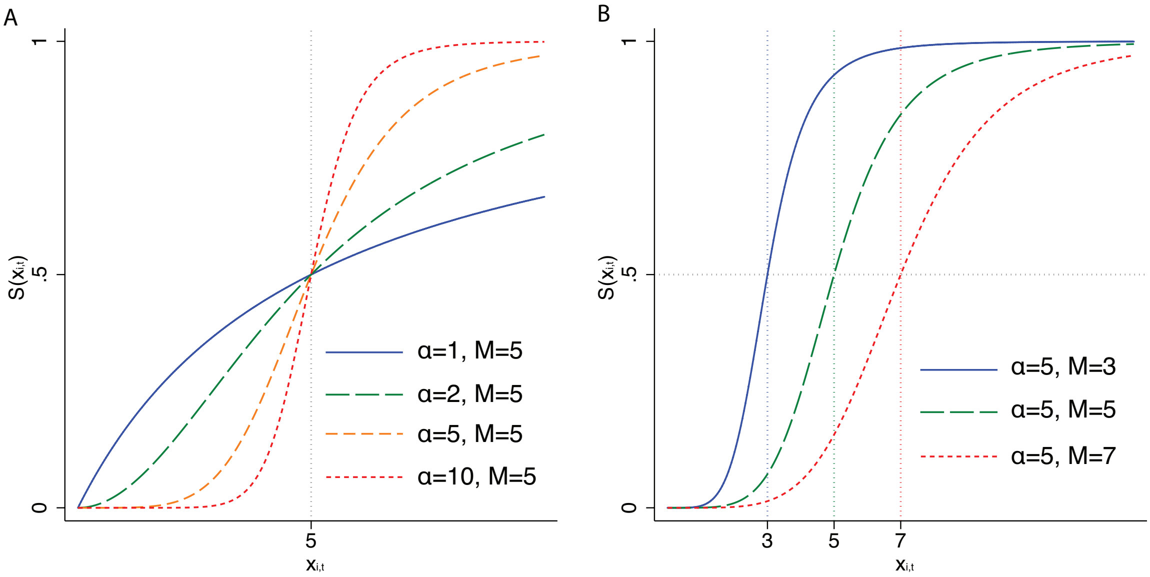

Figure 1: Examples of typical subjective value functions. In each graph, the vertical dashed lines are drawn at $x_{i,t} = M_t$. A: The effect of changes in the predisposition on the subjective value function. B: The effect of changes in the reference point (M_t) on the subjective value function.

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.

To illustrate the role of the reference point, Figure 1 B plots three different value functions with different reference points and the same predisposition. As the reference point 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, $S_t(M_t) = 0.5$. This can be easily checked by substituting M_t in place of $x_{i,t}$ in Equation 1. For every α, whenever $x_{i,t} = M_t$, then $S_t(x_{i,t}) = 0.5$. This property is consistent with behavioral and neural evidence that people evaluate the prizes, rewards, and other stmuli as departures from the reference point rather than in absolute terms. The fact that the subjective value function does not shift in a strictly parallel fashion preserves diminishing marginal value as we explain later in subsection 2.5.1. We take this as a significant feature which distinguishes this class of models from models where the reference point shifts the value function rigidly, sacrificing diminishing marginal value.

2.4 Risk Taking

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

$$RRA_t = -\frac{x_{i,t}S''}{S'} = \frac{(1+\alpha)x_{i,t}^{\alpha} - (\alpha-1)M_t^{\alpha}}{x_{i,t}^{\alpha} + M_t^{\alpha}} \tag{4}$$

The index can take both positive (risk aversion) and negative (risk seeking) values. In traditional reference-dependent models, individuals are assumed to be risk-averse (seeking) whenever the reward is larger (smaller) than the reference point. 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 reference point, but also by predisposition.

Theorem 1. *The individual with ESVT preferences is risk averse if*

$$RRA_t > 0 \Leftrightarrow x_{i,t} > \sqrt[\alpha]{\frac{\alpha-1}{\alpha+1}}M_t \tag{5}$$

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

Proof. Follows from the definition of Arrow-Pratt relative risk attitude index (Equation 4). ■

2.4.1 Adaptive Reference Point and (In)stability of Preferences

Since the reference point, M_t, varies in the model from period to period, the individual risk attitude for a given $x_{i,t}$, as determined by Equation 4 and Proposition 1, can fundamentally change over time. This observation reveals that these neurobiological models thus interact with 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 this model, the fact that the reference point 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. Perhaps surprisingly, the 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 overall preference structure.

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

Proposition 1. *For any given gamble, an individual is more risk tolerant the higher is her reference point.*

Proof. To show: $\frac{\partial RRA_t}{\partial M_t} < 0$
$\frac{\partial RRA_t}{\partial M_t} = -\frac{2\alpha^2 M_t^{\alpha-1} x_{i,t}^{\alpha}}{(x_{i,t}^{\alpha}+M_t^{\alpha})^2} < 0$ which completes the proof. ■

The role of reward history in explaining risk attitudes in the model goes beyond the pure wealth effects that economists have been modelling for centuries. It is not only the total income earned, but also the timing of 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 with an endogenous M_t, however, makes a different prediction. Since the more recent events have a larger impact on M_t, Oskar will have a higher reference point. We can now use Proposition 1 to predict that in period t Oskar will be more risk tolerant.

2.4.2 Predisposition and Risk Attitude

Proposition 1 implies that an individual with a very low-valued predisposition ($\alpha = 1$) is risk averse for all $x_{i,t}$, just as in the traditional utility model. At the other extreme, when the individual's predisposition has high values ($\alpha \to \infty$), the model predicts risk aversion for gains ($RRA_t > 0$ iff $x_{i,t} > M_t$) 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 the value function as its extreme cases without much complexity. One might suspect that most decision makers lie somewhere in between these two extremes: their subjective value function starts off as convex and as x increases switches to concave. The inflection in S occurs for some $x_{i,t} < M_t$ because $\sqrt[\alpha]{\frac{\alpha-1}{\alpha+1}} \leq 1$.

2.4.3 Reflection Effects

One of the three key elements of Prospect Theory — *the reflection effect* — 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). The neurobiological model, however, can do both: capture the generally observed reflection effects on the aggregate level while allowing for reflection in risk attitudes to occur or not occur on the level of individual.

We find that the model generally favors reflection effects, because individuals with ESVT preferences are less risk tolerant for higher $x_{i,t}$'s. (This can be easily verified by showing that for all α and for all M_t $\frac{\partial RRA}{\partial x_{i,t}} > 0$.) However, this is not sufficient to guarantee that reflection in risk attitude occurs for every individual. Intuitively, for some combinations of the predisposition parameter and the reference point, the same individual can be risk averse (or risk seeking) for both a gain and a loss amount relative to what is considered population-level reference point. If there are enough of such individuals in a population, then when we correlate individuals' risk attitudes in gains with their risk attitudes in losses, we will find a positive relationship which is in contradiction to Prospect Theory's assumptions.

Figure 2 provides a graphical demonstration of this feature of the model. To generate the figures in panels A and B, we simulated 1,000 individuals who varied in their predisposition and reference points. In these simulations, the mean predisposition was set equal to 5 and the mean reference point was set equal to 10. In panel A, individual predispositions and reference points were generated from uniform distributions ($\alpha \sim U(0, 10)$ and $M \sim U(0, 20)$). In panel B, they were drawn from normal distributions ($\alpha \sim N(5, 2)$ and $M \sim N(10, 2)$). Then, for each of our simulated individuals we calculated 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-gains is an individual's Arrow-Pratt index calculated at $x_{i,t} = 15$ and RRA-losses is the index calculated at $x_{i,t} = 5$. 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 in figure 2, horizontal (vertical) green lines indicate risk neutrality in gains (losses). The black lines and the gray quadrant represent the areas consistent with Prospect Theory (risk aversion in gains and risk seeking in losses). 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.

The first thing to notice is that in both panels all observations fall in the upper left parts of the graphs. This reflects a key property of the model that individuals are always more risk averse for higher x. This feature is consistent with a substantial experimental literature on the effect of stakes on risk attitudes in gains and losses summarized in Vieider et al. (2012).

The second thing to notice is that in panel B we reproduce perfect reflection effects (OLS regression coefficient = -0.765, $p < 0.001$), but in panel A, while we do see reflection effects on aggregate, we see the opposite of reflection effects on the individual level (OLS regression coefficient = 0.287, $p < 0.001$). The individuals contributing to the lack of reflection effect are those in the upper right and lower left quadrant. Individuals in the upper right quadrant have low predisposition and/or low reference points and are therefore risk averse for both losses and gains relative to the population's mean reference point. Individuals in the lower left quadrant are risk seeking for both 5 and 15. They must have a sufficiently high reference point and predisposition. Existence of both of these groups contributes to the overall positive correlation between risk attitude in gains and in losses.

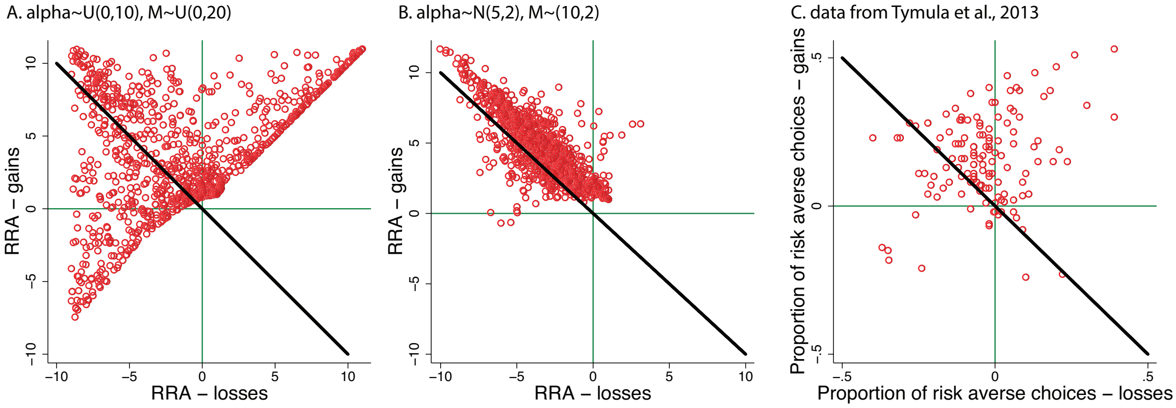

Figure 2: A-B plots of individual Arrow-Pratt relative risk aversion index in gains against the index in losses for 1,000 individuals simulated from different distributions of α and M. C: individual risk attitudes in gains against risk attitudes in losses, adapted from Tymula et al. (2013). Positive values in all graphs indicate risk aversion; negative risk seeking. Green lines indicate risk neutrality.

Finally, to demonstrate what kind of relationship one can expect between risk attitudes in gains and in losses, in Figure 2C 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 reference points. Consistent with the simulations of our model, in this diverse sample we found evidence for reflection effects on the aggregate level: on average, people were risk averse in gains and risk seeking in losses with the majority (79 out of 128) individuals in the gray quadrant. However, on the individual level we found evidence against the reflection effect as risk attitudes in gains are weakly positively correlated with risk attitudes in losses (correlation coefficient = 0.254, p = 0.003).

2.5 Loss Aversion

Loss aversion is a 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 value function that makes it steeper for losses than for gains. The most commonly used utility of a loss averse individual over reward $z \in \mathbb{R}$ is:

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

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, a point made previously by Blavatskyy and Pogrebna (2009), Ert and Erev (2013). In a similar way, the value function in the neurobiological model can be seen to produce loss aversion 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 (Abdellaoui et al., 2007):

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

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

$$\lambda_t = \frac{(M_t^\alpha - (M_t - k)^\alpha)(M_t^\alpha + (M_t + k)^\alpha)}{((M_t - k)^\alpha + M_t^\alpha)((M_t + k)^\alpha - M_t^\alpha)} \tag{7}$$

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.

Theorem 2. *An individual with ESVT preferences is always loss averse. The strength of loss aversion depends on stake size and predisposition. https://www.tenants.org.au/tag/granny-flats*

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 the neurobiological 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). Heterogeneity in loss aversion both on the individual and population levels is consistent with the variation in estimates found in the emprical literature.

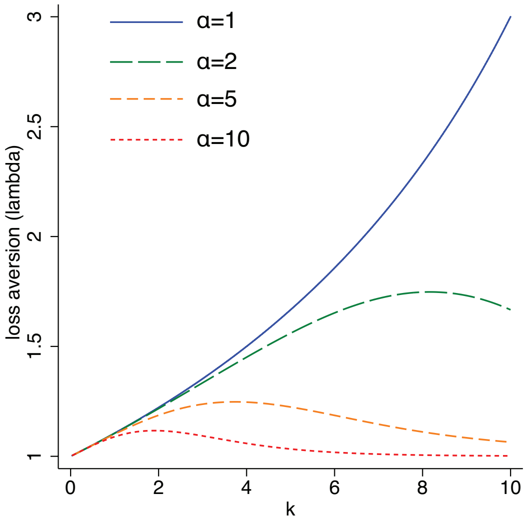

Figure 3: Loss aversion (λ) as a function of the gain (loss) size (k) at different values of predisposition (α).

Figure 3 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 higher predispositions lead to less loss aversion and that loss aversion first increases and then decreases in stake size. The model is thus in line with the mixed findings in the literature on the effect of stake size on loss aversion (Pope and Schweitzer, 2011, Blavatskyy and Pogrebna, 2007).

It is important to notice that in the 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 test 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 employ neurobiological foundations and normative features (Steverson et al., 2016) for the lack of a relationship between the inflection point, the kink and loss aversion in the utility function.

2.5.1 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. The loss aversion theory of the endowment effect (Kahneman et al., 1991) has been challenged by findings that the endowment effect does not show up in market-exprienced traders (List, 2004), in laboratory subjects who were trained to trade (Engelmann and Hollard, 2010), and when using certain experimental procedures (Plott and Zeiler (2007), but see Isoni et al. (2011), Plott and Zeiler (2005, 2011) for full discussion).

The model of Köszegi and Rabin (2006) provided a new theory of the endowment effect, however also crucially based on loss aversion. The difference with respect to a Prospect Theory explanation was that in Koszegi and Rabin's model the gap between willingness to pay (WTP) and willingness to accept (WTA), depended entirely on an individual's expectation about the likelihood-to-trade. If a trade is fully expected, even the most loss averse individual will not show an endowment effect in the Koszegi and Rabin model. The more unexpected the trade its, the larger should be the endowment effect. Individuals who are not loss averse, will not show an endowment effect in Köszegi and Rabin (2006) model under any likelihood-to-trade. Empirical tests of this alternative theory of the endowment effect have yielded mixed results. Ericson and Fuster (2011) endow individuals with a mug and find that manipulating the probability with which subjects are permitted to trade affects valuation. However, Heffetz and List (2014) do not find this manipulation to be effective in a significantly larger sample when there is no endowment. Khalil and Wu (2017) provide evidence that expectations alone cannot explain endowment effect and that the endowment is a necessary condition for the endowment effect to occur.

Bordalo et al. (2012a) suggested another explanation for the endowment effect that does not rely on loss aversion. Their model yields the novel predictions that the endowment effect occurs because different attributes of the goods are salient for goods that are owned and the same goods that are not owned, that endowment effect should not be observed when people are asked to trade the same goods, and that reverse endowment effect should occur for *bads*. These predictions received mixed support in the experimental literature (Brenner et al., 2007, Dertwinkel-Kalt and Köhler, 2016).

We observe here that the neurobiological model provides a novel explanation of the endowment effect that like Bordalo et al. (2012a) does not rely on an explicit loss aversion parameter but can still capture the regularities, referred to as "loss aversion", observed in the empirical literature.

Proposition 2. *Let* $x_{i,1}$ *be the individual's initial ownership level of the good under consideration. Let* $x_{j,2} > x_{i,1}$ *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* $M_2 < \frac{x_{i,1}x_{j,2}}{M_1}$*. The endowment effect does not occur if* $(M_1^\alpha - M_2^\alpha)(M_1^\alpha M_2^\alpha - x_{i,1}^\alpha x_{j,2}^\alpha) = 0$*.*

Proof. The endowment effect occurs whenever:

$$WTA = \frac{x_{j,2}^\alpha}{(M_2^\alpha + x_{j,2}^\alpha)} - \frac{x_{i,1}^\alpha}{(M_2^\alpha + x_{i,1}^\alpha)} > \frac{x_{j,2}^\alpha}{(M_1^\alpha + x_{j,2}^\alpha)} - \frac{x_{i,1}^\alpha}{(M_1^\alpha + x_{i,1}^\alpha)} = WTP$$

Which is equivalent to:

$$WTA = \frac{(x_{j,2}^\alpha - x_{i,1}^\alpha)M_2^\alpha}{(M_2^\alpha + x_{j,2}^\alpha)(M_2^\alpha + x_{i,1}^\alpha)} > \frac{(x_{j,2}^\alpha - x_{i,1}^\alpha)M_1^\alpha}{(M_1^\alpha + x_{j,2}^\alpha)(M_1^\alpha + x_{i,1}^\alpha)} = WTP$$

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

$$(M_1^\alpha - M_2^\alpha)(M_1^\alpha M_2^\alpha - x_{i,1}^\alpha x_{j,2}^\alpha) > 0$$

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

$$M_1^\alpha M_2^\alpha - x_{i,1}^\alpha x_{j,2}^\alpha < 0 \Leftrightarrow M_2 M_1 < x_{i,1}x_{j,2}.$$

The endowment effect does not occur whenever

$$WTA = WTP \Leftrightarrow (M_1^\alpha - M_2^\alpha)(M_1^\alpha M_2^\alpha - x_{i,1}^\alpha x_{j,2}^\alpha) = 0$$

■

Let us intuitively understand what this proposition implies. There are two situations when the endowment effect does not occur: First, when there is no change in the reference point after an endowment $M_1 = M_2$, for empirical documentation of such caes see the results in List (2004), Heffetz and List (2014), Khalil and Wu (2017). Second, when the reference point is equal to the status quo ($M_1 = x_{i,1}$ and $M_2 = x_{j,2}$). The second case is likely to occur when a sufficiently long time (or perhaps a sufficient number of offers) has passed between the endowment and an elicitation of willingness to accept; a period of time sufficient for the individual to fully adjust her reference point to the new ownership level (Figure 4 A).

To intuitively understand when the endowment effect occurs in the model, let us concentrate on two cases: First, let us assume that before the endowment occurs, the individual's reference point coincides with the ownership level $M_1 = x_{i,1}$ (as illustrated by blue curves in Figure 4). Then the endowment effect will be observed only if the reference point after the endowment does not fully update to the new ownership level immediately, $M_2 < x_{j,2}$ (as is illustrated by red curve in Figure 4 B). 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. As long as the reference point remains between the new and old ownership levels, the endowment effect occurs. At first it gets stronger as time since endowment passes (consistent with evidence in Strahilevitz and Loewenstein (1998)), and then it begins to weaken.

As a second case, let us assume that after the endowment, sufficient time (or offers) passes and the reference point updates fully to the new ownership level, such that $M_2 = x_{j,2}$ (as illustrated by the red curve in Figure 4 B). Then the endowment effect will occur only if the

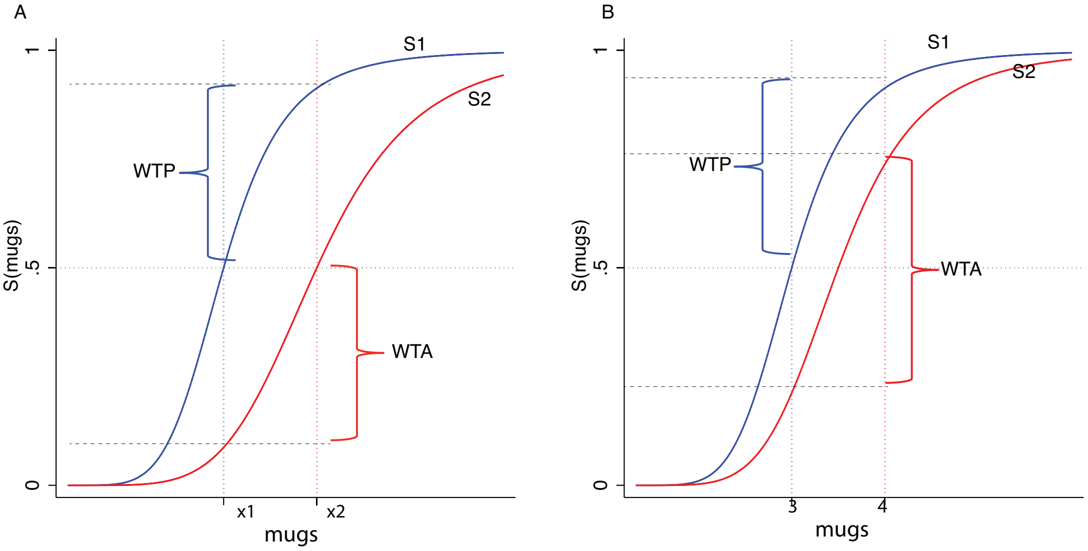

Figure 4: Endowment effect with full (in A) and incomplete (in B) adjustment of the reference point to current ownership. Blue (red) represents subjective value function before (after) endowment.

reference point before the endowment was below the ownership level $M_1 < x_{i,1}$. This could happen if the individual had only recently gained x up to the level of $x_{i,1}$ prior to the initial event, and her reference point was not fully adapted to that ownership level ($M_1 < x_{i,1}$) when she was asked how much she was willing to pay for another unit of x. In our Figure 4, the blue curve would have to shift to the left.

This framework can be used to derive more cases of when the endowment effect will or will not occur or even when it will be reversed. Importantly, the predictions of our model are consistent with the empirical findings that: 1) an “active” endowment is necessary for endowment effect to occur, 2) a reference-point based on prior ownership is a key determinant of the endowment effect, and 3) the duration of ownership since endowment modulates the endowment effect.

We note, that the existing evidence on the endowment effect provides us with important but *mixed* evidence about the accuracy of the endogenous form of the reference point and whether it should or should not be updated to include information about probabilistic future rewards (Ericson and Fuster, 2011, Heffetz and List, 2014).

2.6 Reference point and diminishing marginal value

More traditional reference-dependent models that are based on taking a difference between the reward and the reference point make, we believe, a surprising prediction of constant marginal utility with respect to the reference point. To illustrate the point, assume a simple Prospect Theory reference dependent model where the reference point is given by the status quo and the utility is given by $U = u(x - r)$, where x is the reward and r is the reference point. In such a simple model, the marginal utility from an additional unit of reward (e.g. a mug) is constant with respect to how many units of the reward individual already owns (see Figure 5A). In the endowment effect context, this implies that if we keep endowing an individual with additional mugs, one at a time, and ask for her willingness to pay for another mug after each endowment, she should be willing to pay the same amount of money for one more mug no matter whether it is her third, fourth or hundredth mug! We observe that the neurobiological model does not make this prediction (see Figure 5 B). An individual's willingness to pay for an extra mug always declines after an endowment in the model, which we prove in the remark below.

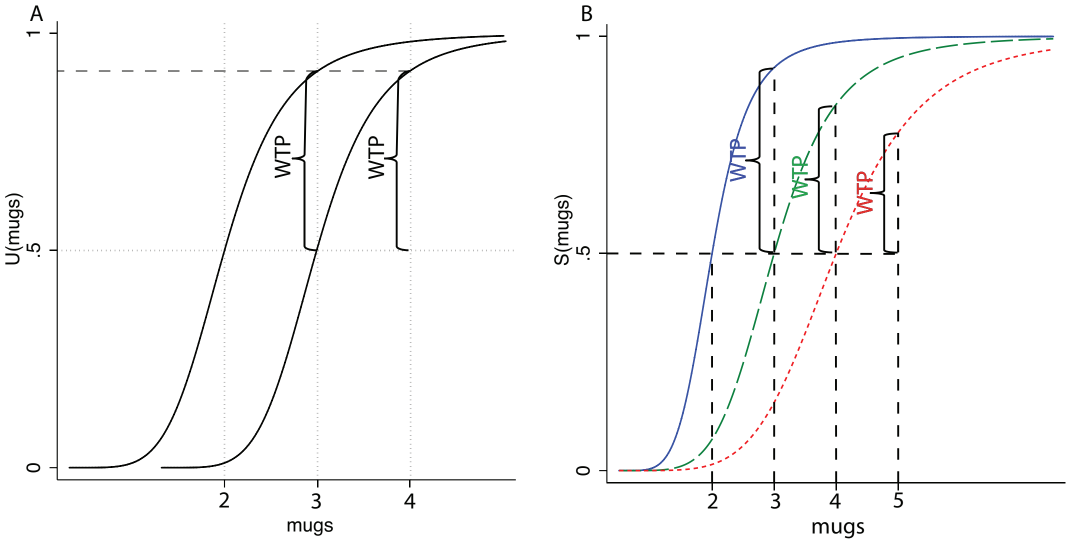

Figure 5: Marginal value in difference-based reference point model (A) and in divisive normalization model (B).

Remark 1. *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 $x_{1,t}$ to $x_{2,t}$ is given by $S_t(x_{2,t}) - S_t(x_{1,t}) = \frac{(x_{2,t}^{\alpha} - x_{1,t}^{\alpha})M_1^{\alpha}}{(M_1^{\alpha} + x_{2,t}^{\alpha})(M_1^{\alpha} + x_{1,t}^{\alpha})}$

To show: $S_t(x_{2,t} - x_{1,t}) > S_{t+1}(x_{3,t+1} - x_{2,t+1})$ (diminishing marginal utility condition), where $x_1 < x_2 < x_3$ and $x_2 = \frac{x_1 + x_3}{2}$

Given that $\forall x_{i,t} > M_t$, the marginal value of the additional unit of $x_{i,t}$ decreases, the willingness to pay for another unit of reward is the highest when the reference point coincides with the current level of ownership. Therefore, $S_{t+1}(x_{3,t+1} - x_{2,t+1})$ is maximized and the diminishing marginal utility condition is hardest to satisfy when $M_{t+1} = x_{2,t+1}$. If we show that this condition is satisfied for $M_{t+1} = x_{2,t+1}$, then it holds for all other M_{t+1}.

Given $M_t = x_{1,t}$ and $M_{t+1} = x_{2,t+1}$ we rewrite the diminishing marginal utility condition as

$S_t(x_{2,t}) - S_t(x_{1,t}) = \frac{(x_{2,t}^{\alpha} - x_{1,t}^{\alpha})x_{1,t}^{\alpha}}{(x_{1,t}^{\alpha} + x_{2,t}^{\alpha})2x_{1,t}^{\alpha}} > \frac{(x_{3,t+1}^{\alpha} - x_{2,t+1}^{\alpha})x_2^{\alpha}}{(x_{2,t+1}^{\alpha} + x_{3,t+1}^{\alpha})2x_{2,t+1}^{\alpha}} = S_{t+1}(x_{3,t+1}) - S_{t+1}(x_{2,t+1})$

Dropping the time indexes, 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.

■

2.7 Probability Weighting

To this point, we have devoted none of our paper to a discussion of the subjective probability function and its relation to the neurobiological model. Driven by experimental evidence, by Savage's insight (Savage, 1954), and by Allais paradox (Allais, 1953), Kahneman and Tversky (1979) proposed that people do not perceive probabilities objectively. Instead, they conjectured that people systematically 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). We note that existing models do not capture this heterogeneity at the individual level. Emotions, to take another example, have been found to affect an individual's subjective probabilities (Fehr-Duda et al., 2011), but a coherent economic theory that would explain individual differences in probability weighting functions does not exist.

The neurobiological model model, of course, leads to a complicated and non-stationary pattern of risk preferences. This arises because of the movement of the subjective value function induced by changes in the reference point. The model, however, as almost always employed in the literature does obey the Independence Axiom. In contrast, in Prospect Theory variations in risk attitudes reflect both the stationary value function and the curvature of the probability weighting function. Due to the stark difference between the methods by which these two models describe risk preferences, we felt it might be useful to directly compare the models with regard to the representation of probability. We therefore conducted a simple numerical exercise. First, we used the neurobiological model generatively to produce a series of choices in response to a set of lotteries. To do this, we simply implemented our value function, varying both the reference point and the predisposition. We then fit that choice data with traditional Prospect Theory.

We thus simulated nine individuals with different predispositions ($\alpha_{low} = 1$, $\alpha_{med} = 2$, and $\alpha_{high} = 5$) and different reference points ($M_{low} = 40$, $M_{med} = 60$, and $M_{high} = 80$) using our theory to determine their choices. Given these predispositions and reference points and assuming ESVT preferences and no distortion in probability, 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 made by our nine ESVT agents 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 6 presents the results of this analysis.

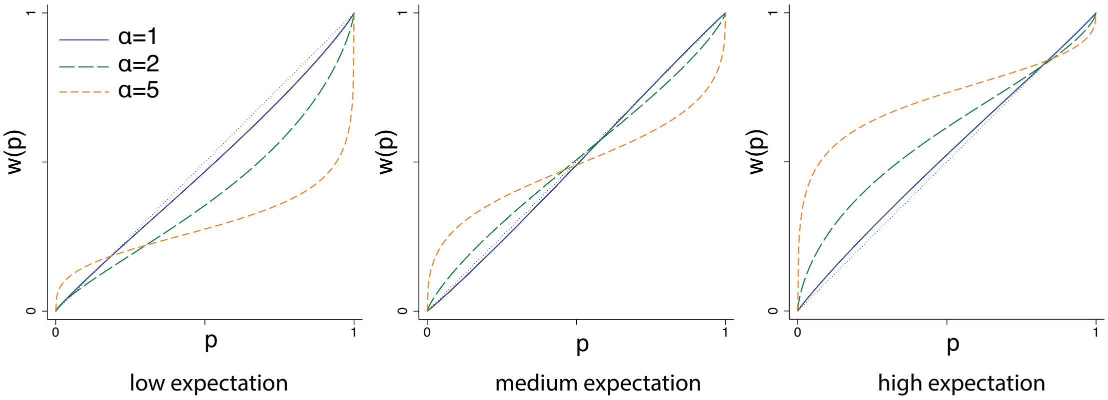

low expectation

medium expectation

high expectation

Figure 6: Probability weighting functions of ESVT choosers estimated with $w(p) = \frac{\delta p^{\gamma}}{\delta p^{\gamma}+(1-p)^{\gamma}}$ and CRRA utility functions

What is striking is that our fit of Prospect Theory to the ESVT-generated data replicates previously observed probability weighting function patterns, even though our simulated individuals do not generatively distort probabilities. In our example, individuals with the highest-valued predisposition show the most noticeable “probability distortions,” which depend critically on reference points. 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.* Intuitively, the reason for such distortions is that in ESVT, individuals with high reference point take more risks. Since the usual Prospect Theory fitting approach does not allow for the individuals to vary in their reference point, it captures the increased risk taking by allowing for more probability overweighting. The smallest probability distortions occur for individuals with low-valued predisposition parameters. This is expected because their value function closely resembles CRRA (with which it is fit) thus no significant distortions in probability are inferred by the Prospect Theoretic analysis.

It is not trivial to analytically establish a relationship between the primitives of ESVT (predisposition and reference point) and the parameters in Prospect Theory model because there is a non-unique mapping between parameters. To address the relationship between these sets of parameters in the face of this analytic limitations to some degree, in Figure 7 we present estimated probability weighting functions in a case where we assume a linear utility function (instead of CRRA) and the standard two parameter probability weighting function of Prospect Theory. Here, we observe an underestimation of probabilities which "compensates" for the lack of curvature in the utility function.

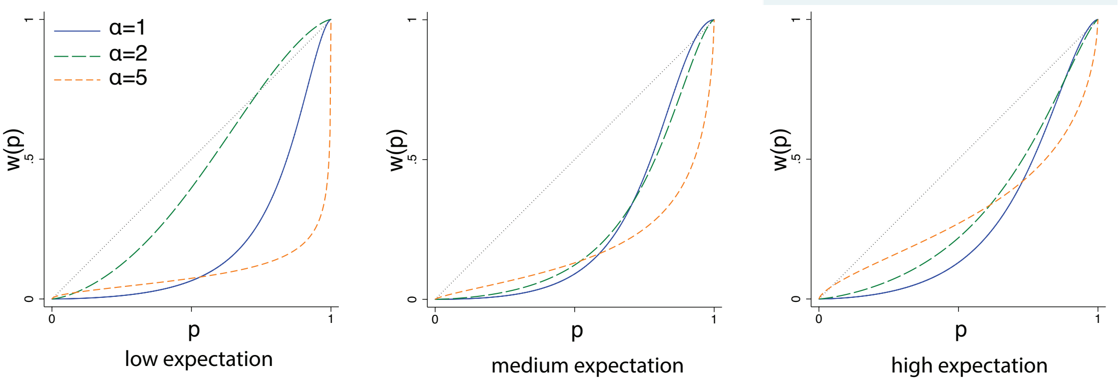

low expectation

medium expectation

high expectation

Figure 7: Probability weighting functions of ESVT choosers estimated with $w(p) = \frac{\delta p^{\gamma}}{\delta p^{\gamma}+(1-p)^{\gamma}}$ and linear utility function

These simulations suggest an interesting possibility, that the estimated probability weighting functions widely inferred from studies in the human population *may* be related to the assumed functional utility form. Observed heterogeneity in *estimated* probability weighting might in fact be due to differences in individuals' predisposition and reference point. *If choosers generated their choice behavior using the ESVT model but that choice behavior was fit using Prospect Theory, one might incorrectly conclude that probability distortions were more widespread than is actually the case.* Furthermore, the model raises the possibility that differences in the estimated probability weighting functions may be related to unaccounted individual differences in reference points.

3 Discussion

In this paper we have presented an analysis of the risk and preference structures produced by Expected Subjective Value Theory, assessing its strengths and weaknesses as a descriptive model of human choice behavior. Expected Subjective Value Theory rests on the optimal coding *Normalization Models* of neuroscience (Carandini and Heeger, 2012) which have been gaining popularity in Neuroeconomics (Louie et al., 2013) and which have recently been shown to have normative foundations in choice theory (Steverson et al., 2016). The analysis presented in Section 2 of this paper suggests that many of the key behavioral phenomena that Prospect Theory was designed to capture 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. Specifically, ESVT provides a new explanation of the endowment effect, explains the existence of the reflection effect in risk attitudes on the aggregate but not necessarily individual level, captures variation in loss aversion between and within individuals, and offers theoretical explanation for the heterogeneity in estimated probability weighting functions. It is worth noting that a simple model, created to describe the optimal representational frameworks in the brain, can capture such a broad variety of human choice behaviors for which it was not explicitly designed. In Appendix A we show that in an empirical dataset ESVT performs better than Prospect Theory (and Expected Utility Theory).

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 examine how and where one can compare the performance of Prospect Theory and ESVT in behavioral data. Section 2 serves as a guide to that process. We believe it suggests that ESVT compares surprisingly well with Prospect Theory in many regards, though we acknowledge that much work remains and that neurobiologists have not yet converged on a single functional form.

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 one possible use of theory in the domain of policy. 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. ESVT suggests the testable hypothesis that these differences are driven by lower reference points among women. This finding clearly suggests that a policy intervention that increases women's reference point might make women willing to take as many risks as men.

3.1 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* and probability weighting function (Kahneman and Tversky, 1979). This was accomplished by introducing two curvature terms (one for gains and one for losses), a loss aversion term, a reference point, and subjective distortions in probability perception (one or two additional parameters). 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, loss aversion, and subjective probabilities. These four ideas (utility curvature, a reference point, loss aversion, and probability weighting) form the core primitives of value and choice in their theory and are the four (or six) exogenous elements with which choosers are provisioned by the theory, although the reference point is sometimes viewed as endogenous.

In contrast, ESVT's value function *emerges* from the interaction between two core primitives of its theory: *predisposition* and *reference point.* 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; it would be the name of a measurement.

3.2 Related Models of Choice and Future Directions

The core idea underlying divisive normalization models 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), Bordalo et al. (2012b), 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. The model we analyze here is not the first one in economics to study how limited neural resources affect individual's attitudes towards risk.

Woodford (2012), Robson and Whitehead (2016), Rayo and Becker (2007) all come to the conclusion that the efficient use of finite processing resources would result in an S-shaped value function. In these models, known statistics about the reward structure in the environment determine the exact shape of the efficient subjective value function. Such adaptive coding is also present in the model of Padoa-Schioppa and Rustichini (2014). The normative divisive normalization models we examined can be thought of as a neural implementation of the ideas conveyed in these earlier theoretical models, related to the idea of limited neural resources are the models of attention.[4] For example, Bordalo et al. (2012b) show that assuming limited attention can account for many of the empirical phenomena for which Prospect Theory cannot account. We leave attention out of our analysis for now, but see the scope for incorporating ideas such as in Bordalo et al. (2012b) into our model in the future as some neurobiologists have begun to do for this class of model (e.g. Reynolds and Heeger (2009).

We wish to restate that we cannot assess the usefulness of a final endogenous definition of reference point in our current theory. The neurobiological researchers have simply not provided such a model for economic analysis. Though the endogenized reference point that we use is uncontroversially the dominant functional form in neuroeconomics and neuroscience, we have no doubt that the reference point in choice is a more complicated construct than in the theory presented here, which is why we remain agnostic about its endogenization. Some existing models, like Köszegi and Rabin (2006), assume that each possible outcome serves as a reference point to which other options are compared. In a similar way, divisive normalization models also allow for the value to be normalized by other options in the choice set, a possibility we discussed elsewhere. Such normalization leads to very specific violations of the independence of irrelevant alternatives (Louie et al., 2011, Louie and Glimcher, 2012, Louie et al., 2013, 2015) but does not affect risk attitudes, hence for tractability we leave them outside the scope of this paper. Also for tractability, we do not incorporate social reference points (Rabin, 1993, Fehr and Schmidt, 1999) in our paper.

Another popular element of economic models of choice that is left out of our current presentation is the probability weighting function. 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. In this paper, we have highlighted how unacounted for normalisation in value could explain the heterogeneity in the estimated probability weighting functions, a phenomenon that has escaped explanation in the economics literature so far. However, 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

---

[4] For a review of models of attention see Caplin (2016).

of these areas.

Finally, we note that it is quite unsatisfying to us that the central free parameter of the model, "predisposition", is not specified in some deep philosophical form. We acknowledge that neurobiologists have only made cursory attempts to explain what predisposition means. Heeger (1992) argues that it reflects the nonlinear response to inputs. Webb et al. (2016) relates it to the capacity constraints of the nervous system. We find these suggestions interesting but believe more theoretical and experimental research is needed to both define what predisposition is and render it observable.

References

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

Abdellaoui, M., H. Bleichrodt, and C. Paraschiv (2007, oct). Loss Aversion Under Prospect Theory: A Parameter-Free Measurement. *Management Science 53*(10), 1659–1674.

Abeler, J., A. Falk, L. Goette, and D. Huffman (2011, apr). Reference Points and Effort Provision. *American Economic Review 101*(2), 470–492.

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.

Albrecht, D. G. and D. B. Hamilton (1982, jul). Striate cortex of monkey and cat: contrast response function. *Journal of neurophysiology 48*(1), 217–37.

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.

Bartra, O., J. T. McGuire, and J. W. Kable (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage 76*, 412–427.

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.

Bayer, H. M. and P. W. Glimcher (2005, jul). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron 47*(1), 129–41.

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.

Bordalo, P., N. Gennaioli, and A. Shleifer (2012a, may). Salience in Experimental Tests of the Endowment Effect. *American Economic Review 102*(3), 47–52.

Bordalo, P., N. Gennaioli, and A. Shleifer (2012b, aug). Salience Theory of Choice Under Risk. *The Quarterly Journal of Economics 127*(3), 1243–1285.

Bordalo, P., N. Gennaioli, and A. Shleifer (2013, oct). Salience and Consumer Choice. *Journal of Political Economy 121*(5), 803–843.

Brenner, L., Y. Rottenstreich, S. Sood, and B. Bilgin (2007, oct). On the Psychology of Loss Aversion: Possession, Valence, and Reversals of the Endowment Effect: Table 1. *Journal of Consumer Research 34*(3), 369–376.

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

Busse, L., A. R. Wade, and M. Carandini (2009, dec). Representation of concurrent stimuli by population activity in visual cortex. *Neuron 64*(6), 931–42.

Caplin, A. (2016, oct). Measuring and Modeling Attention. *Annual Review of Economics 8*(1), 379–403.

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.

Dertwinkel-Kalt, M. and K. Köhler (2016, feb). Exchange asymmetries for bads? Experimental evidence. *European Economic Review 82*, 231–241.

Engelmann, D. and G. Hollard (2010). Reconsidering the effect of market experience on the endowment effect. *Econometrica 78*(6), 2005–2019.

Ericson, M. and A. Fuster (2011, oct). Expectations as Endowments: Evidence on Reference-Dependent Preferences from Exchange and Valuation Experiments. *The Quarterly Journal of Economics 126*(4), 1879–1907.

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, E. and K. M. Schmidt (1999, aug). A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics 114*(3), 817–868.

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.

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.

Gill, D. and V. Prowse (2012, feb). A Structural Analysis of Disappointment Aversion in a Real Effort Competition. *American Economic Review 102*(1), 469–503.

Glimcher, P. (2010). *Foundations of Neuroeconomic Analysis.* Oxford University Press.

Glimcher, P. (2015, jan). Understanding the Hows and Whys of Decision-Making: From Expected Utility to Divisive Normalization. *Cold Spring Harbor symposia on quantitative biology 79*, 169–76.

Glimcher, P. W. (2005, sep). Indeterminacy in Brain and Behavior. *Annual Review of Psychology 56*, 25–56.

Glimcher, P. W. (2011, sep). Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences of the United States of America 108 Suppl*(Supplement_3), 15647–54.

Gneezy, U. (2005). Experimental Business Research. Chapter Updating t. Boston, MA: Springer US.

Gneezy, U., L. Goette, C. Sprenger, and F. Zimmermann (2017, aug). The Limits of Expectations-Based Reference Dependence. *Journal of the European Economic Association 15*(4), 861–876.

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

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.

Hare, T. A., J. O'Doherty, C. F. Camerer, W. Schultz, and A. Rangel (2008, may). Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *The Journal of neuroscience : the official journal of the Society for Neuroscience 28*(22), 5623–30.

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., J. L. Knetsch, and R. H. Thaler (1991, feb). Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias. *Journal of Economic Perspectives 5*(1), 193–206.

Kahneman, D. and A. Tversky (1979). Prospect Theory. *Econometrica 47*(2), 311.

Kahneman, D. and A. Tversky (1984). Choices, values, and frames. *American Psychologist 39*(4), 341–350.

Khalil, E. and K. Wu (2017, may). Explicit vs implicit proprietorship: Can endowment effect theory explain exchange asymmetry? *Economics Letters 154*, 117–119.

Khaw, M. W., P. W. Glimcher, and K. Louie (2017, nov). Normalized value coding explains dynamic adaptation in the human valuation process. *Proceedings of the National Academy of Sciences 114*(48), 12696–12701.

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.

Kőszegi, B. and M. Rabin (2007, aug). Reference-Dependent Risk Attitudes. *American Economic Review 97*(4), 1047–1073.

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.

Levy, I., J. Snell, A. J. Nelson, A. Rustichini, and P. W. Glimcher (2010, feb). Neural representation of subjective value under risk and ambiguity. *Journal of Neurophysiology 103*(2), 1036–47.

List, J. A. (2003). Does Market Experience Eliminate Market Anomalies? *The Quarterly Journal of Economics 118*(1), 41–71.

List, J. A. (2004). Neoclassical Theory versus Prospect Theory: Evidence from the Marketplace.

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.

Padoa-Schioppa, C. and A. Rustichini (2014, may). Rational Attention and Adaptive Coding: A Puzzle and a Solution. *American Economic Review 104*(5), 507–513.

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 (2007, sep). Exchange Asymmetries Incorrectly Interpreted as Evidence of Endowment Effect Theory and Prospect Theory?

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.

Rabin, M. (1993). Incorporating Fairness into Game Theory and Economics.

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).

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.

Sclar, G., J. H. Maunsell, and P. Lennie (1990). Coding of image contrast in central visual pathways of the macaque monkey. *Vision research 30*(1), 1–10.

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. *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.

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 (2016, 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

As a demonstration, we fit the 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 a changing probability of winning this reward (13%, 25%, 38%, 50%, and 75%) (Levy et al., 2010). 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, we treat expectations as a free parameter rather than estimate it from the history of rewards because of the presence of ambigous rewards that we do not want to include inn the analysis. 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 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.

Thus in this gains-only dataset, ESVT performs well as an alternative to Prospect Theory. Of course much more detailed comparisons between Prospect Theory and ESVT will need to be made. Unfortunately, few datasets spanning losses and gains with so-called “mixed lotteries” and enough variation in probability exist that can provide the within-subject level

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	
α	1.153
	(0.019)
M	12.248
	(0.351)
noise	0.084
	(0.002)
AIC	55427
BIC	55456

B. Prospect Theory	
r	0.185
	(0.014)
δ	3.332
	(0.300)
γ	0.661
	(0.019)
noise	0.198
	(0.020)
AIC	56813
BIC	56849

C. Expected Utility	
r	0.456
	(0.004)
noise	0.739
	(0.014)
AIC	57632
BIC	576511

data required for a more rigorous test. Most datasets on probability weighting are collected using choice lists (many options presented at the same time in rows on one sheet of paper or computer screen) rather as one choice at a time. Unfortunately, with a choice list, we lose control over the order in which individual considers options and the length of time spent considering each one of them and hence cannot estimate the endogenous expectation of ESVT which relies on a knowledge of recent choice history. Second, we cannot use any data where current offers depend on previous choices, such as a trade-off method, as this introduces a bias to the endogenous expectation. Other existing datasets, to reduce the number of choices per individual were gathered across subjects not allowing us to estimate the models on the subject level. We are currently developing a new dataset for this purpose which will be the subject of a detailed comparison. Our goal in this simple appendix was first to demonstrate the basic empirical plausibility of ESVT.

Our next goal in this appendix is to 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. Of our 2,530 participants, 1,287 are male. The average age in the sample 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).[5] Using a standard CRRA utility function we find that men in our sample are more

---

[5] $r = r_{constant} + r_{male} male + r_{age} age$, $\alpha = \alpha_{constant} + \alpha_{male} male + \alpha_{age} age$ 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	
α	
male	0.056^{***}
	(0.007)
age	0.000
	(0.000)
constant	0.429^{***}
	(0.010)
σ	
constant	0.736^{***}
	(0.015)
No. of obs	62928

B. ESVT	
α	
male	0.050
	(0.041)
age	-0.005^{***}
	(–0.001)
constant	1.302^{***}
	(0.064)
M	
male	1.687^{**}
	(0.581)
age	0.146^{***}
	(0.025)
constant	6.478^{***}
	(0.852)
σ	
constant	0.082^{***}
	(0.002)
No. of obs	62928

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 observe an alternative possibility. We find that the gender variable does not significantly affect predisposition but that the difference in risk taking is explained, in ESVT, by differences in expectations. In our data, men have significantly higher expectations according to ESVT and therefore behave in a more risk tolerant manner in line with Proposition 1. 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 that suggests that such a conclusion might be premature. Consistent with lifetime wealth effects we find that expectations increase with age. 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 subjective value function underlying these choices changes.

Appendix B Computation of the Reference Point

Following Prospect Theory, in the main body of this report we place minimal restrictions on how the reference point, or expectation (M_t) is computed. We do, however, have strong neurobiological insights into how expectations are computed in the human brain and this data guides our selection of an initial functional form for the term M_t, when endogenization 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:

$$P_t = P_{t-1} + \gamma(x_{i,t} - P_{t-1}) \tag{8}$$

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 this iterative form imposes simpler requirements for data storage in the brain):

$$P_t = (1-\gamma)^t P_0 + \sum_{\tau=1}^{t} \gamma(1-\gamma)^{t-\tau} x_{i,\tau} \tag{9}$$

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:

$$M_t = P_{t-1} = \sum_{\tau=0}^{t-1} \gamma(1-\gamma)^{t-1-\tau} x_{i,\tau} \tag{10}$$

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 reason to believe that they might (Caplin and Dean, 2008, Rutledge et al., 2010). For the purposes of this paper that may not be critical. What is critical is that when ESVT is implemented as an alternative to Prospect Theory, this simple functional form for computing expectation with a single free parameter $\gamma \in (0, 1)$ can be employed to endogenize reference point.

Appendix C Normalization Model-Based Representations of Value

Appendix C.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 C1). 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 C1. 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.

The reason this is so important is that action potentials and neurons that generate them are incredibly costly in terms of calorie consumption. The cost of neural activity (measured in action potentials per neuron carrying them 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 (or precision of our representations) 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 neu-

Figure 1

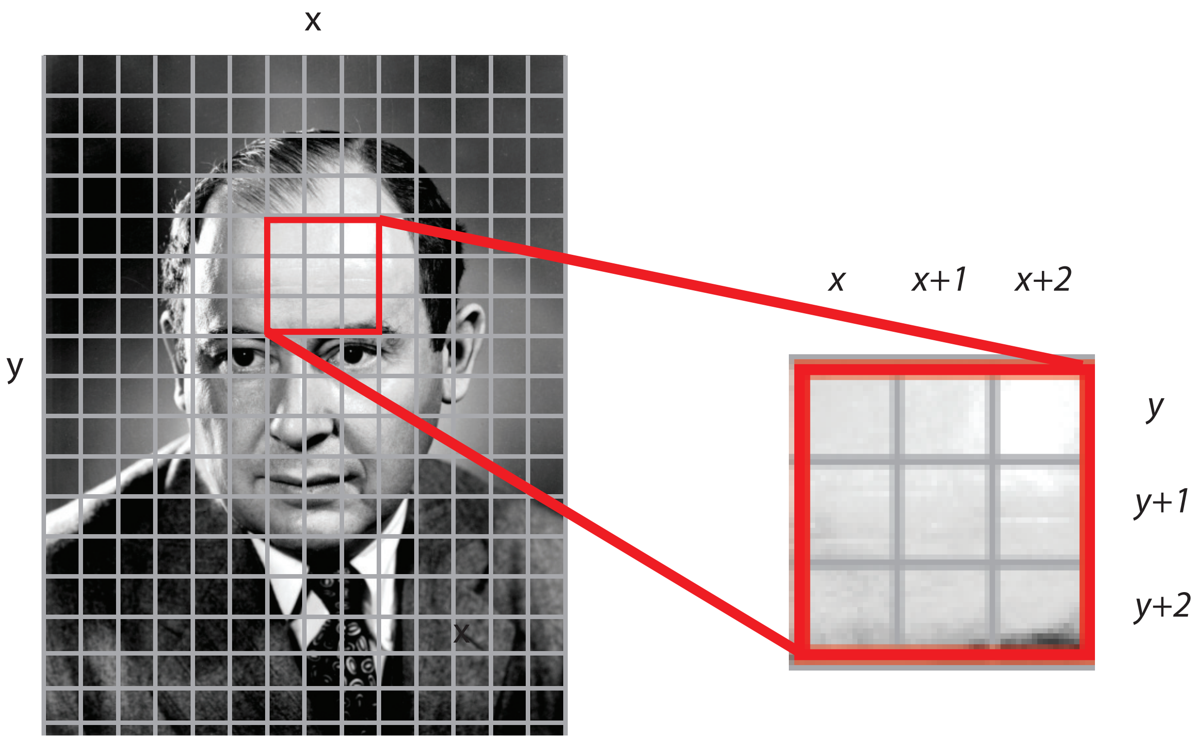

Figure C1: An array of local pixels from a photograph.

ronal 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.[6]

---

[6] 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).

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 (e.g. Steverson et al. (2016)).

Appendix C.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 bounded 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.[7] To a first approximation, this means that mean firing rate and variance are correlated at roughly 1.0. If we say that the mean observed value of F_1 is equal to 10 then the observed variance of F_1 is approximately 10 as well. In a more rigorous treatment of the firing rates in the foregoing section, the equations would have been presented as descriptions of means with specified variances —- a feature which has received tremendous attention in the neurobiological and neuroeconomic literatures and which is also related to the costs of producing each action potential, see Glimcher (2005).

For our purposes this is an important point because it makes the limited capacity of the system hugely problematic in finite time. If the entire set of neurons encoding an image has only an average aggregate action potential rate of 100 to spend and each element (neuron) has a variance proportional to its mean, then the information carried by the set can indeed be quite limited. For an economist it is probably useful to then think about these kinds of representations through the lens of random utility models where each neuron encodes a mean value plus a variance term. More on this approach to neural theories of representation can be found in Webb et al. (2016).

---

[7] 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.

Appendix C.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).

Appendix C.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.[8] *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

---

[8] 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 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.

Appendix C.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 a degenerate one-prize choice set —- a situation which has been 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:

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

where $D(t)$ is a weighting function such that the term on the left in the denominator $(\sum_0^{t-1} D(t)x_t)$ is a time-discounted average of the prizes previously encountered by the system. In the language of reference-dependent theories of choice, *it is the endogenous 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 earlier (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, a goal which has also been recently accomplished (Khaw et al., 2017). The dynamics of these 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.

Appendix C.6 Summary

In this brief literature review we have provided insight into the origins of the functional form we examine in the paper. 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.