Rationalizing Context-Dependent Preferences: Divisive Normalization and Neurobiological Constraints on Decision-Making

(Job Market Paper)

Ryan Webb*
Paul W. Glimcher
Kenway Louie
Center for Neuroeconomics
New York University

November 14, 2013

Abstract

Biology places constraints on the form of neural computation that ultimately characterizes choice behaviour. We identify a canonical neural computation, *divisive normalization*, which predicts that the composition and size of the choice set will adversely influence choice behaviour, including increased stochasticity, novel violations of the axiom of Independence of Irrelevant Alternatives, and a decreased likelihood of choosing a higher valued alternative. Evidence for such context-dependent choice behaviour is provided from two behavioural experiments, and these violations of IIA are more accurately captured by a choice model incorporating normalization than alternative econometric specifications. Finally, we address the normative implications of divisive normalization in choice behaviour. We demonstrate that normalization implements an efficient means for the brain to represent valuations given the constraints of a neural system, yielding the fewest choice errors possible. These results emphasize both the positive advances offered by developing choice models grounded in neuroscience, and the normative role neurobiological constraints can play in the study of decision-making.

---

[*] Center for Neuroeconomics, NYU, 4 Washington Pl., NY, NY, 10003, `webbr@nyu.edu`. This research was funded by the Army Research Office (W911ND-11-1-0482), and the experiment was approved by an IRB. The authors wish to thank A. Caplin, N.G. Martineau and T. LoFaro for helpful comments, as well as seminar participants at Harvard University, the University of Zurich, and Maastricht University. A recent version can be found at `http://www.cns.nyu.edu/%7Ewebbr/WebbJMP.pdf`

1 Introduction

The question of whether stochastic choice behaviour can be rationalized, and how such behaviour depends on the choice set, has long been of interest in economics. Block and Marschak (1960) first laid out a necessary *regularity* condition for the existence of a random utility representation, under which the probability of choosing an alternative must not increase as more alternatives are added to a choice set. Similarly, an early sufficient condition held that relative preferences (i.e. relative choice probabilities) between two alternatives should not depend on other alternatives in the choice set, formally stated as the axiom of Independence of Irrelevant Alternatives (IIA; Luce, 1959). However it was quickly recognized that IIA places strong conditions on choice behaviour (Debreu, 1960) and violations have been well-documented empirically, consisting of patterns of substitution between alternatives as the composition of the choice set is varied (Rieskamp, Busemeyer and Mellers, 2006). As a result, the structure of these substitution patterns, namely *how* relative preferences depend on the choice set, is of great interest in both the economic theory of the consumer (McFadden, 2005) and evaluating responses to economic policy in applications like transportation planning, energy usage, and consumer choice (McFadden, 1974; Hausman and Wise, 1978; Ben-Akiva and Lerman, 1985; Revelt and Train, 1998; Brownstone and Train, 1999; McFadden and Train, 2000; Train and Winston, 2007; Hess and Rose, 2009; see McFadden, 2001, for a review).

In this article, we examine novel forms of IIA violation that are necessarily predicted by constraints known to be imposed by neurobiology. To establish this point, we extend a Neural Random Utility Model (NRUM; Webb et al., 2013) to directly incorporates these neurobiological constraints. This model predicts patterns of substitution between alternatives that depend on the size and composition of the choice set, particularly a decrease in the likelihood of choosing a higher valued alternative. Evidence for such context-dependent choice behaviour is provided from two behavioural experiments, and we show that these violations of IIA are more accurately captured by an NRUM incorporating neurobiological constraints than alternative econometric specifications. Finally, while context-dependent choice behaviour might be considered irrational or anomalous – particularly since the probability of choosing a higher valued alternative is reduced – we demonstrate that this behaviour is in fact optimal given the constraints faced by neural systems.

The observation that biology places constraints on choice behaviour, and that these constraints should explain economic characteristics, is not novel in the economic literature (Robson, 2001; Robson and Samuelson, 2010; Netzer, 2009). This includes the observation that a constrained decision process implies behaviour we might strictly term irrational in the absence of the constraint (Simon, 1979). In particular, the role of constraints on information processing (Shannon, 1948) have been emphasized and shown to have important economic implications (Sims, 2003; Woodford, 2012). For instance, Rayo and Becker (2007) argue that the utility representation must be adaptive and depend on the distribution of the domain of utility, given constraints on the representation similar to those we identify here. Our contribution is to provide a neural mechanism for how such dependence arises from neurobiological constraints, explore its predictions for substitution patterns between choice alternatives, provide behavioural evidence for its existence, and demonstrate that it implements the normative solution to the neurobiological choice problem.

Just like constraints traditionally described in economics, neurobiological constraints arise from scarcity: resources in the brain must be allocated to various neural systems for various tasks. Moreover, there is some mechanism which selects more efficient allocations and rejects less efficient ones. In the case of the brain this mechanism is evolution, and we can identify a basic constraint that directly impacts how the decision-making process in the human brain has evolved. Since neurons require energy, and energy resources in the human body are finite, this puts a limit both on the total number of neurons in the brain on the number that can be allocated to any given task (since neurons are not perfectly substitutable). This basic capacity constraint on information processing in neural systems has two implications that will be particularly relevant for choice behaviour. First, the brain is forced (via the mechanism of evolution) to implement computational algorithms which compress the objective state of the world into the activity of a finite number of neurons for subjective perception, computation, and ultimately, choice. Second, neural activity is universally observed to be stochastic due to the small-scale, thermodynamic processes involved in the synaptic transmission of information between neurons. While this randomness can be partially mitigated through increasing the number (or size) of neurons that participate in any task, this is both costly and limited: neurons are not fully independent in their activity, therefore stochasticity remains. Ultimately, this neural stochasticity results in stochastic choice behaviour (Glimcher, 2005).

From our perspective, these two constraints imply limitations on the form of the decision-making process (i.e. the types of computations) the brain can implement. In recent years, neuroeconomists have been studying this process in detail, examining the structure and operation of the neural computations underlying choice (Glimcher, 2011; Fehr and Rangel, 2011). In this article, we focus on a pervasive computation in cortex, *divisive normalization*, in which the inputs to the computation are scaled (divided) by the magnitude and number of inputs. This computation therefore normalizes each input with respect to the entire set of inputs, yielding a relative relationship between inputs. It is this relative relationship that we will explore in the context of choice behaviour.

Originally observed in the cortical regions of visual perception (Heeger, 1992), divisive normalization has now been observed in multiple forms of sensory processing and across species ranging from invertebrates to primates (Carandini and Heeger, 2012). From a normative standpoint, it has also been shown to yield an efficient coding of perceptual information in a constrained neural system. Divisive normalization reduce the mutual information in a sensory stimulus, eliminating the redundant information found in natural stimuli (Schwartz and Simoncelli, 2001; Wainwright, Schwartz and Simoncelli, 2001; Sinz and Bethge, 2013). These results suggest normalization is a canonical neural computation, and perhaps unsurprisingly it has recently been observed in value signals that are believed to implement a decision (Louie, Grattan and Glimcher, 2011). However, the normative implications for behaviour have remained unclear, particularly because divisive normalization predicts a novel form of context-dependent choice behaviour that can be strictly termed inefficient and inconsistent: it necessarily leads to relative preferences that depend on the size and composition of the choice set (Louie, Khaw and Glimcher, 2013). In this article, we show that divisive normalization, when inserted into a neuroeconomic choice model, predicts such context-dependent choice behaviour by implementing a normative solution to the decision-making problem faced by constrained neural systems.

From an empirical standpoint, there is now growing evidence that the context in which a decision is made has implications for economic behaviour. Of particular relevance to our study, the size of the choice set has been shown to alter relative preference, the consistency of preferences, and the likelihood of making a purchase (Iyengar and Lepper, 2000; DeShazo and Fermo, 2002). The model we present here captures such effects by focusing on normalization at the level of choice alternatives. In doing so, it maintains an ordinal ranking over alternatives. However it is often observed that adding an alternative to a choice set which is either dominated, or similar, along the attribute dimension can alter relative choice towards one alternative — sometimes referred to as *asymmetric dominance*, or *attractor* effects, respectively (Huber, Payne and Puto, 1982; Tversky and Simonson, 1993; Rieskamp, Busemeyer and Mellers, 2006; Soltani, De Martino and Camerer, 2012).[1] To capture such phenomena, the extension of a model with normalization to the attribute dimension is a natural next step.

This article proceeds as follows. In section 2, we present a general theoretical framework of decision-making in the brain, the Neural Random Utility Model (Webb et al., 2013), that will allow us to investigate the role of neurobiological constraints. We then introduce the divisive normalization computation, review its positive and normative application in perceptual neuroscience, and review the neural evidence in areas of the brain known to implement decision-making. In section 3, we explore the positive predictions of divisive normalization, specifically patterns in choice behaviour that depend on the composition and size of the choice set. We then present behavioural evidence for normalization from two laboratory experiments and propose a method for incorporating normalization into an econometric specification.[2] Finally, in section 4 we address the normative issue by describing the neurobiological constraints on decision-making, and demonstrate that the divisive normalization computation implements choice behaviour that is optimal given these constraints. These results emphasize both the positive advances offered by developing choice models grounded in neuro-

---

[1] Possible theoretical explanations of these "decoy" effects, and of the general relationship between context and choice, have been raised (Thaler, 1985; Busemeyer and Townsend, 1993; Thaler, 1999; Bordalo, Gennaioli and Shleifer, 2012, 2013).

[2] Results from one of these experiments, which varies the composition of the choice set, have previously been reported in Louie, Khaw and Glimcher (2013). In this article, we bring our econometric specification to this dataset in order to compare it to established econometric methods, and investigate the prediction that some subjects should exhibit more normalization behaviour than others depending on the construction of the choice set.

science, and the normative role neurobiological constraints can play in the study of decision-making.

2 Normalization in a Neural Choice Model

Over the past fifteen years, the empirical literature has pointed towards a basic "two-stage" working hypothesis for decision-making in the primate brain (Glimcher, 2011; Fehr and Rangel, 2011). It is widely reported that brain regions in the frontal cortex and basal ganglia carry information about the relative valuations of choice alternatives.[3] It is also reported that stochastic neural systems located in regions of the frontal and parietal cortex implement a comparison to determine which of these valuations is largest (Padoa-Schioppa and Assad, 2006; Gold and Shadlen, 2007; Louie and Glimcher, 2010), with the result of this comparison passed to motor control regions for action.

These empirical results can be captured by an extension of the random utility framework to neural systems termed the Neural Random Utility Model (Webb et al., 2013). The NRUM is a reformulation of the traditional Random Utility Model (RUM; Becker, DeGroot and Marschak, 1963; McFadden, 1974) in which the unobservable latent variables are replaced with neural quantities.[4] As such, it attempts to capture the features of what is currently known about decision-making in the brain in a familiar econometric framework, and can be shown to be a reduced form of more structural, dynamic models found in neuroscience (Webb, 2013).

Because the NRUM is designed to make empirical claims given neural quantities, it provides a convenient empirical framework for exploring the implications of neural computation on choice behaviour and investigating the role of neurobiological constraints. In section 2.1, we begin by adapting the NRUM to a multi-alternative choice environment which explicitly allows for the normalization of value. In sections 2.2 and 2.3, we will describe in more detail the divisive normalization computation, its origins in the sensory neuroscience literature, and its extension to the domain of value and decision-making.

2.1 Neural Random Utility Model for Multi-Alternative Choice

We begin by defining the observable neural quantity we call *subjective value*. Subjective value is a measurement of neural activity in the frontal cortex of the human brain, and has been demonstrated to carry cardinal information about an individual's valuation of choice alternatives (Webb et al., 2013). Two recent meta-studies (Bartra, McGuire and Kable, 2013; Levy and Glimcher, 2011) now

---

[3] Elliott, Friston and Dolan (2000); Delgado et al. (2000); Knutson et al. (2001, 2003); Padoa-Schioppa and Assad (2006); Kable and Glimcher (2007); Plassmann, O'Doherty and Rangel (2007); Chib et al. (2009); Tusche, Bode and Haynes (2010); Levy et al. (2010, 2011); Levy and Glimcher (2012)

[4] We discuss the relationship between the NRUM and the RUM in more detail in section 3, and in Webb et al. (2013).

unambiguously indicate that this neural activity, particularly in the medial prefrontal cortex, is tightly correlated with every known economic method for estimating the values subjects place on choice objects. This includes behaviourally derived preferences over attributes (Tusche, Bode and Haynes, 2010; Levy et al., 2011), risk and ambiguity (Levy and Glimcher, 2012; Levy et al., 2010), and time (Kable and Glimcher, 2007, 2010). We denote the subjective value of each item in a choice set N by the vector $\mathbf{v} = [v_1, \ldots, v_N] \in \mathbb{R}_+^{\|N\|}$, with the valuation of each item i denoted v_i. Since subjective value is encoded by stochastic neural activity, it is a random variable with distribution $g(\mathbf{v})$.

Once subjective values $\mathbf{v}$ are instantiated in the activity of neurons in the frontal cortex, they must be compared for a choice to be executed. The electrophysiological evidence suggests that regions of the frontal and parietal cortex receive subjective valuations as input and determine which of them is largest by means of a competition between neural quantities (Sugrue, Corrado and Newsome, 2004; Louie and Glimcher, 2010).[5] Importantly, this competition is both stochastic and dynamic (Gold and Shadlen, 2007; Kiani, Hanks and Shadlen, 2008; Milosavljevic et al., 2010), and the structural models of this process found in neuroscience (e.g. the drift diffusion model) can be captured by an additive noise vector we denote $\eta \in \mathbb{R}^{\|N\|}$ (Webb, 2013). For now, we will assume a general distribution $f(\eta)$.

In this article, we hope to describe how the brain transforms subjective valuations as they move to frontal/parietal regions for comparison. To allow for this, we denote the transformation $\mathbf{z} = [z_1, ..., z_N] = Z(\mathbf{v}) : \mathbb{R}_+^N \to \mathbb{R}_+^N$ with the resulting distribution $h(\mathbf{z})$. The function $Z()$ allows for a relationship between an alternative and the composition of the choice set. Specifically, it allows the transformed valuation of an alternative, z_i, to depend on the entire vector of valuations for all alternatives $\mathbf{v}$. Shortly, we will explore behavioural predictions when $Z(\mathbf{v})$ takes takes a particular form, namely that of divisive normalization. All together, this yields the decision variable

$$\mathbf{u} = Z(\mathbf{v}) + \eta. \tag{1}$$

The decision-maker chooses option i such that

$$u_i > u_j \quad \forall j \neq i$$
$$z_i + \eta_i > z_j + \eta_j \quad \forall j \neq i$$

---

[5] For instance, in the Lateral Intra-Parietal area (LIP) – a neural circuit that guides the selection of eye movements (or saccades) in monkeys – groups of neurons each represent a saccade to a particular region of visual space. A decision between two saccadic options is instantiated as a competition between two groups of neurons representing the two possible eye movements (Roitman and Shadlen, 2002; Kiani, Hanks and Shadlen, 2008). In the LIP, value modulation biases the competition towards the better option and implements value-guided decision making. Like frontal cortical areas, activity recorded in these areas varies with both the magnitude and probability of reward associated with specific actions in lotteries (Platt and Glimcher, 1999), in strategic games (Dorris and Glimcher, 2004), and with discounting over time delays (Louie and Glimcher, 2010).

yielding a probability of choosing i

$$\begin{aligned} P_i &= P\left(z_i - z_j > \eta_j - \eta_i \quad \forall j \neq i\right) \\ &= \int\int \mathbb{1}\left[z_i - z_j > \eta_j - \eta_i, \quad \forall j \neq i\right] h(\mathbf{z}) f(\eta)\, dv\, d\eta \end{aligned} \tag{2}$$

assuming that $\mathbf{v}$ (therefore $\mathbf{z}$) and η are independent, with $\mathbb{1}[\cdot]$ indicating the statement in brackets holds.

To isolate the behavioural implications of normalization, we make two simplifying assumptions. First, we are only modeling normalization that takes place between cortical regions which encode subjective value and cortical regions which compare subjective values, as represented by the function $Z()$. We strongly suspect that normalization occurs in other stages of the decision making process, for instance in the construction of subjective value $\mathbf{v}$ from objective sensory attributes (Barron, Dolan and Behrens, 2013), but do not model it here.[6] Second, it will be useful to assume no stochasticity in subjective value $\mathbf{v}$, so the distribution $h(\mathbf{z})$ is degenerate and the probability of choosing i is given by

$$P_i = \int \mathbb{1}\left[z_i - z_j > \eta_j - \eta_i, \quad \forall j \neq i\right] f(\eta)\, d\eta. \tag{3}$$

In essence, we are assuming that we observe a particular realization of subjective value $\mathbf{v}$ and we are tracing its evolution through the normalization $Z()$ and the stochastic process captured by η.

2.2 Divisive Normalization in Sensory Systems

To understand how divisive normalization influences choice behaviour, it will be useful to understand its origins in the neuroscience literature. A fundamental goal in systems neuroscience is to understand how the brain processes information at the level of individual neurons. Much of this work has focused on the sensory domain in which an environmental stimulus is transduced through the sensory organs and passes through a series of processing layers (i.e. spatially and functionally distinct circuits of neurons), each of which has a particular role. In the visual system, for instance, light makes initial contact with photoreceptors in the retina and passes through subsequent stages of hierarchical neural processing in retinal ganglion cells, the thalamus, primary visual cortex, and so on. Generally speaking, sensory neurons in these regions are "tuned" for a particular stimulus, for instance the intensity of light in a particular region of visual space, or a particular direction of motion.

At each of these levels, initial theories of neural processing hypothesized that sensory neurons perform a linear computation: the activity of a neuron is simply a weighted sum of the inputs to that neuron. However, it was soon discovered

---

[6] In fact, the existence of some form of normalization in these regions is virtually a necessity considering the neural constraints we will highlight. This raises the possibility that attributes may be independently normalized, yielding behavioural effects like asymmetric dominance.

that sensory neurons display activity that cannot be solely explained by linear models. One prominent non-linearity is the saturation of neural activity at high input levels. A second is the phenomenon of *suppression*, where the neural activity elicited by a stimulus (for which the neuron is tuned) is reduced in the presence of a stimulus for which it is not tuned. The divisive normalization model was proposed to explain such nonlinear phenomena in primary visual cortex (Heeger, 1992). The critical feature of this model is a divisive rescaling of the activity of a neuron by the total activity of a larger pool of neurons, each receiving different inputs. The normalization model is defined, in its most general form, by the equation:

$$R_i(x_1, \ldots, x_N) = \kappa \frac{x_i^{\alpha}}{\sigma + (\sum_n \omega_n x_n^{\beta})^{\frac{1}{\gamma}}}, \tag{4}$$

where the response $R_i(x_1, \ldots, x_J)$ of a neuron i is a function of both the driving input x_i to that neuron and the summed inputs to a large number of neurons N termed the normalization pool. This general response function is governed by a number of parameters: κ denotes the maximum possible activity level, σ determines how responses saturate with increased driving input, ω_n weights each neuron in the normalization pool, and α, β and γ mediate the form of the normalization. This model can be generalized to neural populations, so that for each input i, the x_i represent the aggregate activity of the neurons tuned for i.

The divisive nature of the normalization implements several features consistent with the physiology of visual cortex. First, because the input to a given neuron is also included in the normalization pool, the model produces a saturating response to increases in driving input. This is consistent with the observation that neurons are biologically constrained to maximum firing rates. Second, division expresses the quantity encoded by a single neuron in relative terms, scaled to the pooled activity of other inputs. This relative representation of information is a general feature of sensory coding. In the retina, for example, light intensity is encoded relative to the average ambient illumination; this process of light adaptation is why black type on white newspaper appears the same in a darkened room and under bright sunlight.

Originally proposed to explain non-linearities in primary visual cortex, normalization has been identified in multiple stages of visual processing from the retina to downstream visual areas, and in species ranging from the fly to primates. In addition to vision, normalization characterizes neural responses in multiple modalities such as audition and olfaction, and extends to higher-order cognitive processes including attention and multi-sensory integration (Carandini and Heeger, 2012). This ubiquity suggests that normalization may be a general feature of information processing in cortical systems.[7]

---

[7] We should note that the normalization model describes a computational algorithm rather than a specific biophysical implementation. Researchers have identified a number of potential neural mechanisms for normalization and it is likely that the normalization computation is mediated by different processes in different species and systems. Moreover, since evolution selects on the fitness of the output of a particular neural system, the parametrization of the

2.2.1 Normative Implications in Sensory Systems

Given the ubiquity of divisive normalization in sensory systems, the question of its normative role in neural coding has been raised. In systems neuroscience, the benchmark normative framework for information processing is the *efficient coding* hypothesis (Barlow, 1961) based on Shannon's work on information transmission via finite capacity channels. Because neural systems face necessary biological constraints (e.g. maximum firing rates or numbers of neurons), sensory stimuli should be represented in a manner that minimizes the redundancy in encoded information; that is, neural responses should be statistically independent. Crucially, the statistics of the natural environment are decidedly not independent. For instance, regions of visual space with intense light occur in clusters, unlike the random "snow" displayed on an untuned analog television set. Under the efficient coding hypothesis, sensory systems should account for such dependencies and transform this information into a less redundant representation.

While linear weighting functions can remove some redundancies, the statistics of natural images are too complex for linear models to produce completely independent responses. However, the non-linear divisive normalization model markedly reduces higher-order correlations in responses to both natural images and sounds (Schwartz and Simoncelli, 2001; Lyu, 2011; Sinz and Bethge, 2013). Thus, normalization serves a specific normative role in implementing efficient information coding in sensory systems.

2.3 Divisive Normalization in Value Systems

Given the pervasive observation of divisive normalization in sensory cortex, and the general similarity of information processing in cortical systems, this has raised the question of whether a form of normalization is also taking place in decision-related cortical regions. Specifically, are the values of alternatives in a choice set represented in a normalized form relative to the valuation of available choice alternatives?

Recent neural evidence from primates suggests that this is the case (Rorie et al., 2010; Louie, Grattan and Glimcher, 2011; Pastor-Bernier and Cisek, 2011). For example, Louie, Grattan and Glimcher (2011) examined the activity of LIP neurons responsible for a particular region of visual space, termed the *receptive field* (RF). In this experiment, two targets were displayed on a screen, and the value of the RF target (in ml of water) was held constant while the value of the extra-RF target was varied. In the absence of normalization, varying the extra-RF target should elicit no change in activity from the recorded neuron. Nonetheless, the measured neural activity varied inversely with the value of the alternative target, consistent with model in which the neural activity coding the

---

normalization equation will depend on the role of the particular system and statistical features of its input. The widespread occurrence and varying implementations indicate that it is the normalization computation, rather than the specific mechanism and parameterization, which is critical to neural systems.

value of a specific action is suppressed by the value of alternative actions.

Further experiments show that this relative value representation is implemented via a divisive normalization computation. Returning to our model in section 2.1, we suggested that the cortical regions involved in decision-making transform subjective valuations according to the function $Z(\mathbf{v})$. Analogous to normalization in sensory processing, a general form of value normalization by the function $Z(\mathbf{v})$ can be written,

$$z_i = \kappa \frac{v_i^{\alpha}}{\sigma + \left(\sum_n \omega_n v_n^{\beta}\right)^{\frac{1}{\gamma}}}, \tag{5}$$

where the activity z_i is a nonlinear function of the value v_i and the values of the other possible actions. To test this hypothesis, Louie, Grattan and Glimcher (2011) implemented a three-target task and systematically varied the number of targets and their values, allowing a test of various relative value coding algorithms. Monkeys fixated on a central cue and were presented with either one, two, or three targets, each of which was associated with a different magnitude of water reward. Each randomized target array provided a unique combination of values associated with the targets in and out of the RF. The authors demonstrate that a restricted form of the above equation in which $\alpha = \beta = \gamma = 1$,

$$z_i = \kappa \frac{(\text{ml of water})_i}{\sigma + \sum_n (\text{ml of water})_n}, \tag{6}$$

captured the activity of LIP neurons better than either a simple linear form or other potential models (e.g. the difference in value).

Normalization in value regions has implications for both both positive and normative theories of choice behaviour because it transforms the neural activity that drives choice. Since the value of alternatives in a choice set are scaled by an identical term, ordinality is preserved, however the precise activity level corresponding to a given alternative depends on the other alternatives in the choice set. Given variability in value representation, this rescaling can have interesting implications for choice behaviour. The prevalence of divisive normalization in cortical regions also suggests a normative role related to efficient coding in the face of neurobiological constraints, but we will hold this question for section 4 and begin with the positive predictions.

3 Behavioural Implications of Normalization

The normalization of subjective valuations in cortex yields a number of behavioural predictions. In this section, we will incorporate normalization into the NRUM and explore its implications for context-dependent choice behaviour. In order to validate these predictions empirically within a standard behavioural framework, we then replace the neural variables in the model with behavioural correlates, namely the BDM bids of choice alternatives, and present evidence from two laboratory experiments: one which varies the composition of the choice set, and another which varies the size of the choice set. We demonstrate that the observed behaviour is accurately captured by the NRUM incorporating normalization and compare its performance to existing econometric techniques.

3.1 Substitution Patterns and Violations of IIA in Trinary Choice

The behavioural predictions from normalization stem from the interaction between scale and stochasticity. Consider the NRUM for a trinary choice set, depicted in Figure 1, and described formally by incorporating a restricted version of the divisive normalization equation (5),

$$z_i = \kappa \frac{v_i}{\sigma + \omega(\sum_n v_n^\beta)^{\frac{1}{\beta}}}, \tag{7}$$

directly in the NRUM from section 2.1. This yields a probability of choosing alternative i given by

$$P_i = \int \mathbb{1}\left[z_i - z_j > \eta_j - \eta_i, \quad \forall j \neq i\right] f(\eta)\, d\eta. \tag{3}$$

$$\int \mathbb{1}\left[\kappa \frac{v_i - v_j}{\sigma + \omega(\sum_n v_n^\beta)^{\frac{1}{\beta}}} > \eta_j - \eta_i, \quad \forall j \neq i\right] f(\eta) d\eta. \tag{8}$$

In a trinary choice set, normalization introduces a particular pattern of substitution between alternatives because it scales valuations by a function of all alternatives in the choice set, $(\sum_n v_n^\beta)^{\frac{1}{\beta}}$. As the value of the lowest alternative is increased, the valuations of the two highest items are scaled down; but the highest alternative is scaled more, in absolute terms, than the second-highest. Given a constant variance for the distribution $f(\eta)$, this increases the likelihood that the decision-maker will choose the second alternative relative to the first, yielding a relationship between relative choice probabilities and the choice set.

To understand the substitution patterns induced by normalization, it will be useful to return, for the moment, to the NRUM without normalization (i.e. $z_i = v_i$ in equation 3). The substitution patterns induced in this model depend on the distribution of the error term η_i. For instance, if we assume that η_i is distributed extreme value (type 1), then the NRUM takes the form of the familiar multinomial logit model (McFadden, 1978),

$$P_i = \frac{e^{v_i}}{\sum_n e^{v_n}}.$$

The multinomial logit model has the property that the ratio $\frac{P_i}{P_j} = \frac{e^{v_i}}{e^{v_j}}$ does not depend on the additional alternatives $k \neq i, j$ in the choice set, and this defines the axiom of Independence of Irrelevant Alternatives (IIA; Luce, 1959).[8]

---

[8] There is a weaker notion of IIA in which only the order of choice probabilities for op-

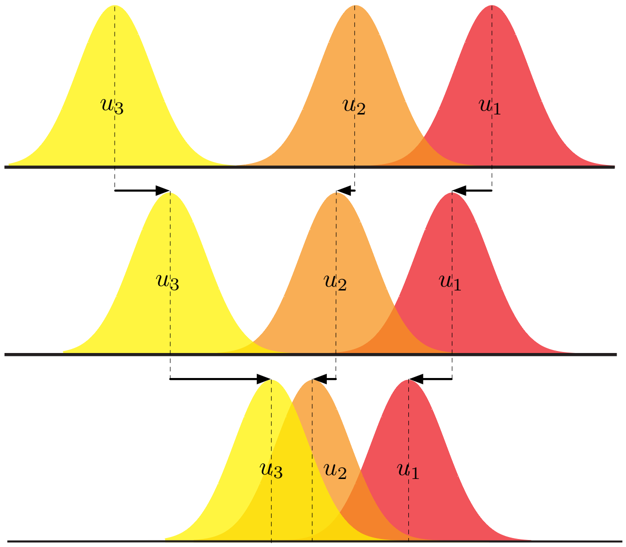

Figure 1: Divisive normalization in an NRUM for trinary choice. An increase in the value of alternative 3 affects the relative likelihood of choosing between alternatives 1 and 2.

Though it is now well documented that IIA is violated in a variety of choice environments (Rieskamp, Busemeyer and Mellers, 2006), it still provides a useful benchmark for exploring the implications of the choice set on substitution patterns. In particular, incorporating normalization in the NRUM predicts a relationship between relative choice probabilities and the choice set which violate IIA. This pattern of substitution is present even when the error distributions are extreme value, and can be seen clearly via simulation in Figure 2A: the ratio of $\frac{P_1}{P_2}$ declines as the subjective value of the third alternative increases.[9]

If we instead assume a normal distribution for the stochastic elements, the pattern of IIA violation predicted by normalization changes in an important way (Figure 2B). As the value of the third alternative increases, the ratio $\frac{P_1}{P_2}$ declines until the third alternative starts capturing significant choice probability from the first and second alternatives. Then the ratio $\frac{P_1}{P_2}$ increases and the first alternative is chosen relatively more often (Figure 2B).[10] Intuition for this result

---

tions i and j must remain constant in the presence of alternative options k. Violation of this condition has been observed in laboratory experiments (Rieskamp, Busemeyer and Mellers, 2006), notably the *asymmetric dominance* effect. Divisive normalization does not display this prediction since it preserves the ordinal ranking of the alternatives. To capture this property, an NRUM would have to include an explicit model of how attributes are (asymmetrically) weighted in the construction of subjective value (Simonson and Tversky, 1992; Bordalo, Gennaioli and Shleifer, 2013).

[9] Formally, $\frac{\partial \frac{P_i}{P_j}}{\partial v_k} = \frac{\partial e^{\frac{v_i - v_j}{\sum_n v_n}}}{\partial v_k} < 0, \forall k \neq i, j$.

[10] This is observed with both the linear form of divisive normalization ($\alpha = \beta = \gamma = 1$) and

can be found in Figure 1. As the distribution for the third alternative rises, it overlaps with the second alternative before the first. This pulls more probability from the second alternative being chosen relative to the first, and sends the probability of choosing the second alternative to zero faster. This results in a pattern of substitution between alternatives which has a characteristic "u-shape."

3.2 Alternative RUM specifications and their relationship with normalization

Many adaptations of the RUM framework have been proposed in order to relax the IIA axiom, all working through the specification of the error distribution, and it is interesting to consider whether these specifications can capture the substitution patterns predicted by normalization. For instance, if the stochastic elements of a RUM (without normalization) are distributed multivariate normal, yielding the multinomial probit model (Hausman and Wise, 1978), particular patterns of IIA violation are observed. If the distribution is independent, the ratio $\frac{P_1}{P_2}$ increases as the subjective value of the third alternative increases (Figure 2C). The assumption of independence between items in the choice set can be relaxed, allowing for a wider variety of IIA violations. For example, the simulation in Figure 2D allows correlation in the stochastic terms for the first and third alternative, yielding a ratio $\frac{P_1}{P_2}$ which decreases to zero as the value of the third alternative increases.

The fact that the multinomial probit model has a fully parameterized covariance matrix has led to the assertion that the model can "represent any substitution pattern" (Meyer, 1991). While it is true that this model relaxes substitution patterns, extensive simulations suggest that even a fully parameterized covariance matrix of the multivariate normal distribution can not capture the patterns of IIA violations predicted by cortical normalization.[11] This suggests that the multinomial probit model, without normalization, will not yield the pattern of relative choice probabilities in Figure 2B. In the trinary choice experiment which follows, we will indeed verify that the multinomial probit model performs worse than a model with normalization.

There is one additional observation to make about the relationship between the NRUM with normalization and current RUMs. The regularity condition, which states that the probability of choosing an alternative cannot increase as more alternatives are added to a choice set, is a necessary condition for a random utility representation. We find that some parameterizations of the NRUM with normalization (namely the variance of η) violate regularity, and we provide a proof in Appendix A for the case of the extreme value distribution. Some intuition for this result can be garnered from the trinary choice simulations in Figures 2A and 2B. Note that for very low valuations of the third alternative (where the probability of it being chosen is virtually zero), the probability of

---

less restrictive forms (e.g. $\alpha = 1$, $\beta = \gamma = 2$, see section 4.1.1)

[11] A proof is elusive in the absence of a closed form for the normal distribution.

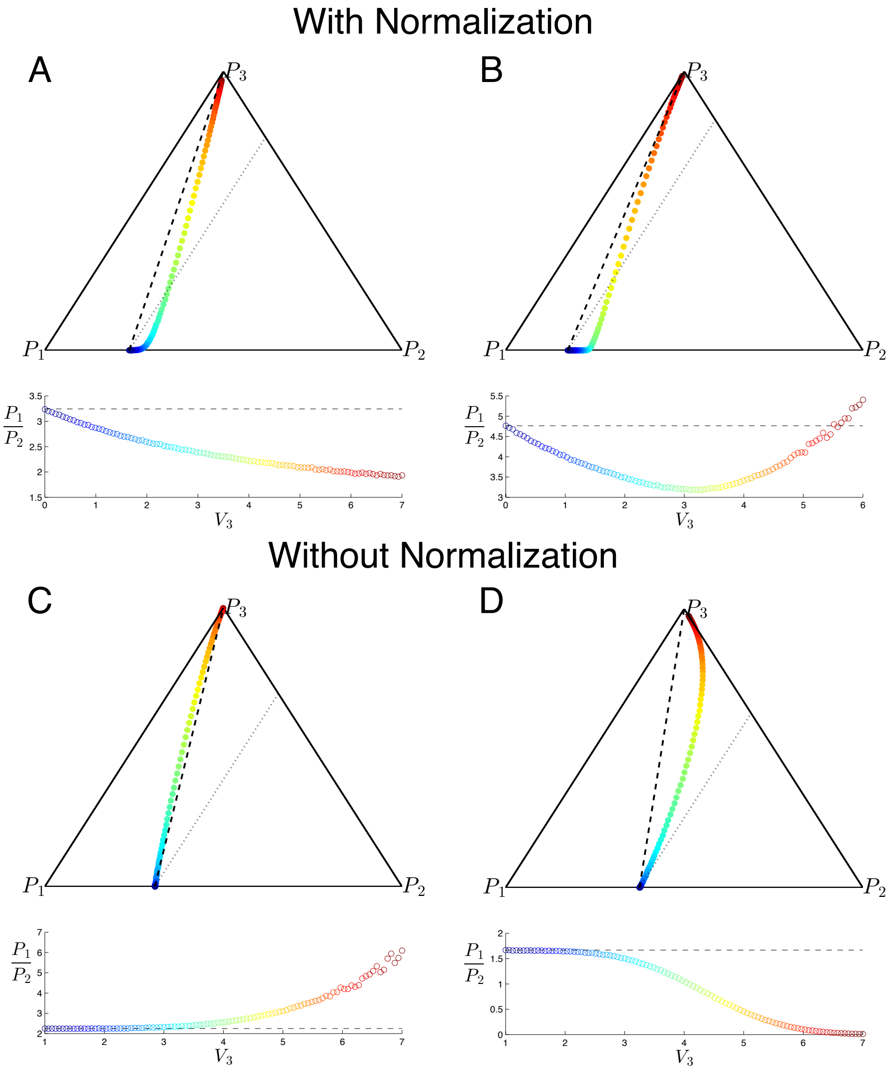

Figure 2: Simulated choice probabilities from an NRUM with (A,B) and without (C,D) normalization. The IIA property is denoted by the dashed line. The iso-probability line for P_2 is denoted by the dotted line. Panel A: Error term is distributed extreme value. Panels B,C: Error term is distributed independent multivariate normal. D: Error term is distributed multivariate normal with correlation.

choosing the second alternative increases in absolute terms. We will examine this prediction in more detail in section 3.4 when we present our experiment which increases the size of the choice set, but for now let us address the implications of this observation.

Since an NRUM with normalization violates regularity, it means that, technically, it is no longer a RUM. An even more general class of empirical model, referred to as mixed logit models (McFadden and Train, 2000), has been shown to approximate both any random utility model as well as some models which violate regularity. At this time we do not know if the mixed logit approach can approximate the NRUM with normalization. We do note that an independent multivariate normal distribution, together with the normalization equation, induces some distribution over the utility vector **u**. It is possible that some mixture of distributions can approximate it. However the mixed logit (or any mixture model) still requires a particular specification for the mixing distribution. As noted by Train (2009), "the researcher's task is simply to find variables and a mixing distribution that capture the [non-idiosyncratic] parts of utility, namely the parts that are correlated over alternative or heteroskedastic".

By incorporating divisive normalization in an NRUM, we accomplish this task by imposing a direct specification on the relationship between alternatives. We constrain the model by grounding it in the neural constraints underlying decision, and let these constraints induce patterns of substitution between alternatives. We find these patterns are modulated by a combination of at least two factors: the distribution of the stochastic terms (including its covariances), and the presence of normalization. The predictions of an NRUM with and without normalization therefore allow observed substitution patterns to distinguish between these factors in an empirical specification.

3.3 Trinary Choice Experiment

We have just seen that divisive normalization predicts particular violations of IIA when the value of a third "irrelevant" alternative is manipulated. In order to test this prediction in a familiar behavioural setting, Louie, Khaw and Glimcher (2013) developed a two-stage valuation and trinary choice task involving common snack-food items. The first stage elicited the valuation of each good (in isolation) using an incentive-compatible auction mechanism (Becker, DeGroot and Marschak, 1964). The second stage of the experiment varied the value composition of trinary choice sets, and had subjects choose a single alternative from these sets. In our analysis, we will use the behavioural measures of value recorded in the first stage in place of the neural measures proposed in the NRUM. Empirical evidence that subjective value activity correlates with BDM bids can be found in (Plassmann, O'Doherty and Rangel, 2007). This will allow us to fit the normalized model to the trinary choice dataset, compare it to the multinomial probit model, and examine subject-specific differences in normalization behaviour.

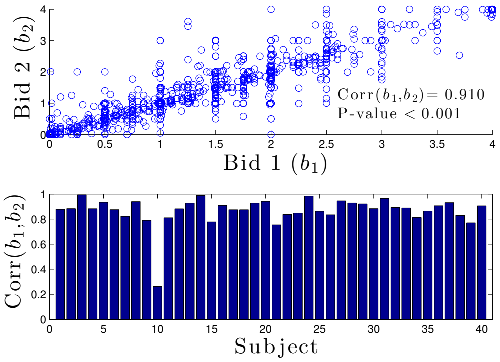

Figure 3: Top: The two bids for each snack food alternative from the first stage of the experiment, for all subjects. Bottom: The correlation coefficient for each subject.

3.3.1 The Trinary Choice Task

In the first stage of the experiment, 40 subjects performed 60 bid trials to establish subject-specific valuations for the array of experimental goods. In each bid trial, subjects viewed a picture of a single snack food item and reported their maximum willingness-to-pay for that item (\$0-\$4 in \$0.01 increments). To implement incentive compatibility, if a bid trial was chosen for realization at the end of the session, the outcome was determined via an auction mechanism (Becker, DeGroot and Marschak, 1964). Each of 30 items was shown twice in a randomized order, and items were ranked for each subject according to their mean bids. We denote the mean bid for each item b_i, and repeated bids for the same item were highly correlated ($\rho = 0.9105$, $p < 0.001$, Figure 3). The 10 highest-valued items (ranks 1-10) were termed *targets* and 10 items sampled evenly from the 20 lowest-valued items (ranks 12, 14, ..., 30) were termed *distracters*.

In the second stage of the experiment, subjects performed 250 choice trials. In each choice trial, subjects selected their preferred option from a trinary choice set. To investigate the effect of the composition of the choice set, each trinary choice grouped two target items and one distracter item. Choice sets were constructed to include varying target value differences and distracter values; each presented triplet was unique and not repeated. If a choice trial was chosen for realization, subjects received whichever good they selected in that trial.

Subjects were requested to fast for four hours before the experiment and required to stay in the lab one hour after completion of the session, during which time the only food they could consume was any item received from the experiment. Along with a possible snack item, each subject received a \$40 show-up fee and was additionally endowed with \$4 for use in the bid trials. At the conclusion of the experiment, a single trial from the session (bid or choice) was randomly selected for realization.

3.3.2 Trinary Choice – Pooled Analysis

We begin by reviewing a result reported in Louie, Khaw and Glimcher (2013), namely whether relative preference between the two high-valued targets varied as a function of the distracter value. In their analysis, distracter value was normalized to the mean value of the presented targets on each trial; this metric ranges from 0 (small value) to 1 (high value) and facilitated comparison across subjects with differing value ranges.

As discussed in the previous section, an NRUM with normalization predicts a particular bi-phasic relationship between relative target preference $\frac{P_1}{P_2}$ and distracter value (Figure 2B): relative choice probabilities decrease when b_3 is small, but then increase as b_3 approaches the value of b_1 and b_2. To quantify choice behaviour, the small number of trials in which the distractor was chosen were dropped and a logistic choice model fit to the data segregated by distracter value; this approach models relative preference as a binary choice between the two targets, restricted to regimes with similar distracter values. As shown in Figure 4, this analysis revealed a "u-shape" in relative choice probabilities: the logistic slope parameter initially decreases as the value of the distracter grows, decreasing the probability of selecting the largest alternative by $\sim 15\%$ (from 0.67 to 0.57) at the mean of $b_1 - b_2$. However the logistic parameter increases back to 0.68 at sufficiently large distracter values. As noted in section 3.1, this pattern likely can not be captured by the multinomial probit model with a fully parameterized covariance matrix. Instead, divisive normalization must be explicitly accounted for in an empirical specification.

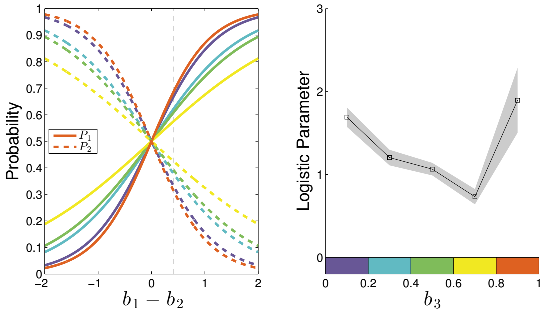

Figure 4: Logistic fit of binary choice between targets, for the range of distracter value segregated into five bins. The relative preference between two items in a choice set is given by the ratio of the probabilities for a given $b_1 - b_2$ (e.g. the mean of $b_1 - b_2$, dashed line). The bi-phasic effect on relative choice probability can be observed in the estimates of the logistic parameter for the five distracter values.

To fully address this claim, we estimate the model in equation 8 for our pooled sample in an effort to compare it to the multinomial Probit model. The pooled sample consisted of 250 trials from 40 subjects, for a total of 10000 observations. In our specification, the parameter ω measures the influence of the choice set on the the overall variance in the model (the denominator) and the relative choice probabilities. This yields the choice probability on each trial given by,

$$P_{i,t} = \int \mathbb{1}\left[\frac{\kappa(b_{i,t}-b_{j,t})}{\sigma+\omega(\sum_{n=1}^{3} b_{n,t}^{\beta})^{\frac{1}{\beta}}} > \eta_{j,t}-\eta_{i,t}, \quad \forall j \neq i\right] f(\eta_t)d\eta_t. \tag{9}$$

The parameters σ and ω were restricted to be positive, and following standard practice the distribution $f()$ and the parameter κ were appropriately normalized so that $f(\eta_j - \eta_i)$ was the standard normal distribution. The resulting maximum likelihood estimates for the entire pooled sample are reported in Table 1 for three nested specifications. We also report the cholesky factorization of the independent normal covariance matrix, differenced with respect to the third alternative, $\left(\begin{smallmatrix}1 & 0\\ l_{2,1} & l_{2,2}\end{smallmatrix}\right)$.

Table 1: Maximum likelihood estimates of the normalization model and the Multinomial Probit model for the pooled sample of the trinary choice experiment.

	Probit	Normalization		Multinomial Probit	
$\hat{\kappa}$	1	1	1	0.702	0.670
				(0.023)	(0.024)
$\hat{\sigma}$	0.996	0.114	0.012	1	1
	(0.004)	(0.002)	(0.001)		
$\hat{\omega}$	0	0.177	0.412	0	0
		(0.034)	(0.001)		
$\hat{\beta}$	1	1	511.5	1	1
			(0.001)		
$\hat{l}_{1,1}$	1	1	1	1	1
$\hat{l}_{2,2}$	0.866	0.866	0.866	0.526	0.792
				(0.016)	(0.033)
$\hat{l}_{2,1}$	0.5	0.5	0.5	0.5	0.431
					(0.022)
LL	-8153.55	-7944.39	-7881.21	-7995.36	-7969.81

In the unrestricted specification, the estimate of ω is significantly different from zero ($p < 0.001$) and the Log-Likelihood decreases significantly compared to the restricted specifications, particularly the model without normalization ($p < 0.001$). This suggests that, at the pooled level, we are observing behaviour in which the composition of the choice set affects the degree of variance in the model, and therefore the relative probabilities.[12]

To gauge whether an alternate econometric specification could capture the substitution patterns we observe in this dataset, we also estimated a multinomial probit model with a full covariance matrix, for a total of three parameters (κ, $l_{2,1}$, and $l_{2,2}$). The log-likelihood from this model was lower than even our restricted normalization model. Clearly, the behaviour we are observing is more accurately captured by an NRUM which includes normalization.

3.3.3 Trinary Choice – Subject Analysis

The normalization model can also be applied at the level of each subject to gauge heterogeneity in context-dependent choice behaviour. Our hypothesis – that divisive normalization is a canonical computation in value cortex – essentially requires that all human brains exhibit normalization. But divisive normalization does yield a prediction about *how much* context-dependent choice behaviour we should observe for a given range of choice sets. For instance, if trinary choice sets were composed of a Porsche, a Ferrari, and some lowly snackfood, we should not expect to see much substitution as the snack food item is varied. Therefore the heterogeneity in the preferences of subjects for the various snack food items in our sample can be used as an additional test of the model.

The resulting maximum likelihood estimates for each subject are in Figure 5.A, in which the estimate $\hat{\omega}$ is multiplied by the average value of the choice set to give a sense of scale relative to $\hat{\sigma}$. The statistical significance of the estimate $\hat{\omega}$ is calculated via a likelihood ratio test from the restricted model where $\omega = 0$.

For each subject, the estimated model either proportions variance to σ or to the normalization term governed by ω. This division of our subject pool can also be observed if we repeat the logistic analysis from Louie, Khaw and Glimcher (2013) on just those subjects for which $\hat{\omega}$ is significant, and those for which it is not, reported in Figure 5.B. The subset of subjects which yield a significant $\hat{\omega}$ exhibit a strong contextual effect on the choice probabilities, in contrast to the remaining subjects (who still exhibit the characteristic increase induced by the multivariate normal distribution).

So what to make of the observation that some subjects appear to be exhibiting contextual effects while others are not? Recall that the behavioural predictions from the divisive normalization equation arise because the difference in value between items $b_1 - b_2$ is scaled down by the valuation of the third item, b_3, in the denominator. If there is little exogenous variation of b_3 in our choice sets, or if the magnitude of $b_1 - b_2$ is small to begin with, then we should not observe much effect on the relative choice probabilities. For instance, in Figure 5.C we highlight the item set composition of three subjects, two of which (in red) the model identifies as not normalizing. One of these subjects has a small range of valuations of distractor items (b_3), while the other has a small range of

---

[12] Since our estimation specification replaces the variance parameter in a standard RUM with the parameters σ and ω in the denominator, a positive and significant relationship between P_i and $b_i - b_j$ occurs when the denominator approaches zero from above. This means that the larger our estimates of ω and σ, the more stochastic is choice behaviour.

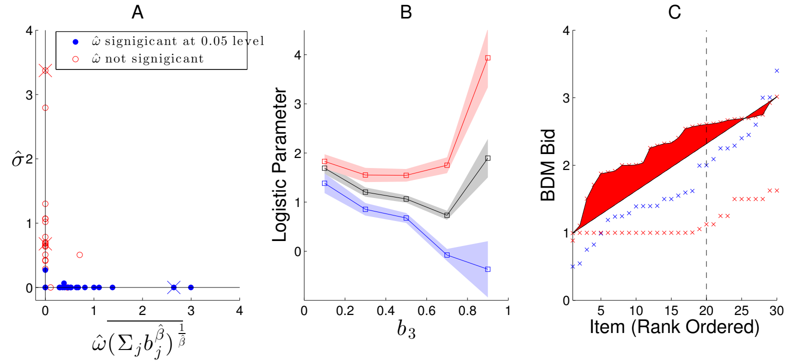

Figure 5: (A) The estimates of σ vs. ω for each subject. Heterogeneity in the sample is clearly present, with some subjects demonstrating more normalization in their choice probabilities (in blue) than others (in red). (B) The logistic fit re-estimated for each sub-sample of subjects. (C) Heterogeneity is predicted by a normalization model depending on a subject's range of valuations over the content of the item set (shown for the 3 subjects denoted with an '×').

valuations of target items (b_1, b_2), relative to the normalizing subject (in blue). Note that the bid spectrum of the normalizing subject has roughly equal range for both distractors and targets, therefore is roughly linear compared to the non-normalizing subjects.

To test this possible explanation for the discrepancy in our subject sample, we constructed a measurement of how much variation is present in valuations of the item set. For each item, we calculated the square of the difference between the rank-ordered valuations and a line which ranges from the minimum bid to the maximum bid. A subject for which the target and distractor set shared a similar range would exhibit a small measurement of variation and a large amount of normalization. A subject for which either the distractor items or target items yielded differing ranges of valuations would exhibit a large measurement and little normalization. We then averaged this metric over the two subsets of subjects defined by our regression results, and found it was significantly larger ($p < 0.05$) for those subjects which did not exhibit normalization. This suggests that the lack of normalization observed in a subset of our subjects resulted from our experiment presenting a set of choice alternatives which were not conducive to generating contextual behaviour in those subjects.

3.4 Set Size Experiment

The behavioural predictions of divisive normalization naturally extend to larger choice sets. Increasing the number of choice alternatives changes the total value of the choice set, and increases the scaling of the model (Figure 6). As in trinary choice, the relative probabilities shift and choice becomes more stochastic, in particular, the likelihood of choosing the highest ranked item again decreases.

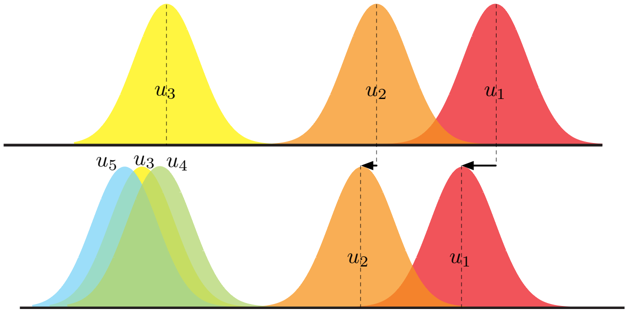

Figure 6: Increasing the size of the choice set shifts the relative probability of choosing the highest ranked item.

Of course, any stochastic choice model which obeys regularity will predict that choice becomes more stochastic as the number of choice alternatives increases. But divisive normalization predicts that we will observe *more* stochasticity than a model with a constant variance: as more alternatives are added, valuations are scaled downward and the variance in the model increases. In sections 3.4.1 and 3.4.2, we present a behavioural experiment designed to test these predictions.

As mentioned in section 3.2, normalization also predicts violations of regularity under some parameterizations of the model and for some sets of alternatives. The difference between the decision variables for the highest and second highest ranked alternative (u_1 and u_2) can decrease quite quickly when multiple low value alternatives are added to the choice set. With a small variance for the stochastic elements η, and low-enough value alternatives added to the set, this can generate an absolute increase in the probability of choosing the second alternative (Appendix A). In section 3.4.3 we will try and observe this prediction in the data.

3.4.1 The Set Size Task

To try and capture the behavioural effects of increasing set size, we developed a two-stage experimental setup based on the trinary choice experiment. In the initial valuation stage, 30 subjects performed 60 bid trials where they reported their maximum willingess-to-pay for 30 individual snack food items. For each subject, items were then ranked according to their mean bid value, and the 10 highest-valued and 10 lowest-valued items were denoted *targets* and *distracters*, respectively.

In the second stage of the experiment, subjects performed 270 choice trials. In each trial, subjects selected their preferred option from a choice set consisting of two target items and a variable number of randomly chosen distracter items for set sizes totalling $N_t \in \{2, 4, 6, 8, 10, 12\}$ items on choice trial t. Choice sets were constructed to include varying target value differences, with 45 different trials in each set size condition.

After completion of the session, subjects were required to stay in the lab, during which time the only food they could consume was any item received from the experiment. A single experimental trial was chosen for realization. If a bid trial was selected, the outcome was determined via BDM auction; if a choice trial was selected, subjects received whichever good was selected in that trial. Payment consisted of a \$40 showup fee, a \$4 endowment for use in bid trial realization, and a possible snack food item. In the analysis that follows, the bids for each subject were normalized by the subject's mean bid over all alternatives to facilitate pooling the data.

3.4.2 Set Size Results

As one might expect, increasing the size of the choice set decreases the probability that the highest ranked item is chosen. In our largest set size of 12 alternatives, the decrease in probability is as high as 20% (Figure 7). In order to determine if this decrease was larger than predicted by the standard RUM, we estimated our normalized NRUM from equation (8) with the only difference lying in the number of choice elements N_t now varying over trials. This yields the choice probabilities on trial t,

$$P_{i,t} = \int \mathbb{1}\left[\frac{\kappa(b_{i,t}-b_{j,t})}{\sigma+\omega(\sum_{n\in N_t} b_{n,t}^{\beta})^{\frac{1}{\beta}}} > \eta_{j,t}-\eta_{i,t}, \quad \forall j \neq i\right] f(\eta_t)d\eta_t. \tag{10}$$

As in the trinary case, the distribution $f()$ was assumed independent normal, κ was appropriately normalized to 1, and both the parameters σ and ω were restricted to be positive. The resulting maximum likelihood estimates for the entire pooled sample are reported in Table 2 for three nested specifications.[13]

Table 2: Maximum likelihood estimates for set size experiment.

	Probit	Normalization	
$\hat{\sigma}$	0.924	0.924	0.001
	(0.025)	(0.002)	(0.040)
$\hat{\omega}$	0	0.001	0.442
		(0.010)	(0.022)
$\hat{\beta}$	1	1	18.85
			(1.961)
LL	-9969.76	-9969.76	-9723.77

Again, we observe that in the unrestricted specification, the estimate of ω is significant ($p < 0.001$) and the Log-Likelihood decreases significantly compared

---

[13] Since our estimation specification replaces the variance parameter in a standard RUM with the parameters σ and ω in the denominator, a positive and significant relationship between P_i and $b_i - b_j$ occurs when the denominator approaches zero from above. This means that the larger our estimate of ω from 0, the more stochastic is choice behaviour.

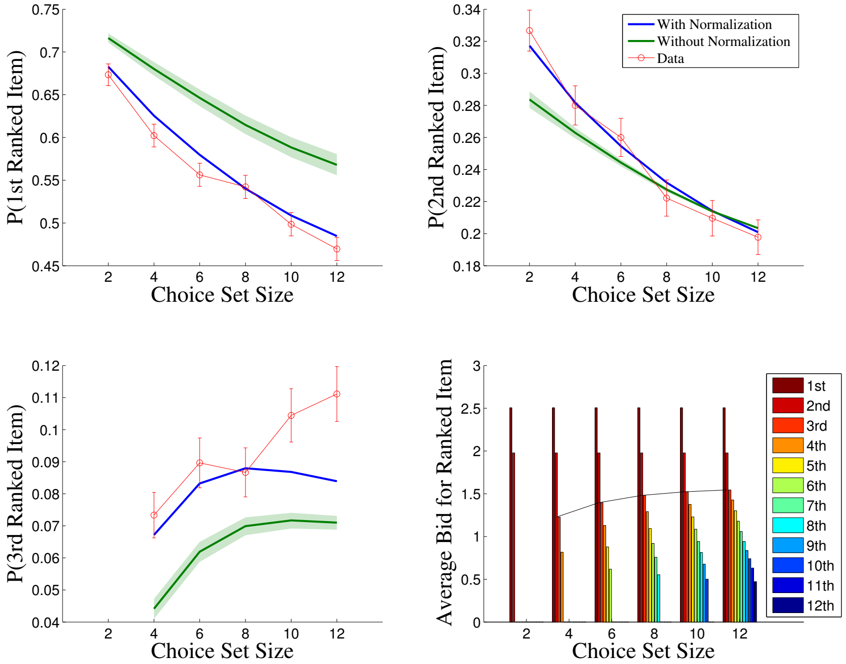

Figure 7: The probability of choosing first, second, and third ranked item, and the fitted probabilities derived from model estimates. The probabilities are evaluated at the average bids in each choice set.

to the unrestricted specification ($p < 0.001$). This suggests that, at the pooled level, we are observing choice behaviour in which the size of the choice set affects the degree of variance in the model, and therefore the relative probabilities.

In order to quantify the effect of set size, we have reported the fitted probabilities from our estimates in Figure 7 evaluated at the average bid values of the choice sets. The specification which allows normalization does a remarkable job of capturing the relative choice probabilities in the data, with the slight exception of under-predicting the probability of choosing the third ranked item in large set sizes.[14] This is in contrast to the RUM without normalization (i.e. Probit). In order to match the choice probabilities (for all items) with a constant variance over all sets, it over-predicts the probability of choosing the highest

---

[14] The increase in the predicted probability of choosing the 3rd ranked item arises from the construction of our choice set. Our ”distractor sets” were randomly sampled from a set of 10 items, which means that smaller choice sets were more likely to have a larger disparity between the third item and the remaining distractors. Notice in Figure 7 that the difference between average bid values for each distractor item is initially large, but shrinks as set size increases.

Table 3: Probability of choosing highest ranked item (evaluated at average bids in each choice set)

	$P_{N=2}$	$P_{N=12}$	$P_{N=12}$ - $P_{N=2}$
Data	0.67	0.47	0.20
Without normalization ($\hat{\sigma} = 0.924$)	0.72	0.57	0.15
With normalization ($\hat{\sigma} = 0.001$, $\hat{\omega} = 0.442$, $\hat{\beta} = 18.85$)	0.69	0.49	0.20

ranked item in the choice set and under-predicts the probability of choosing the lower ranked items. From these fitted probabilities, we can observe that for a set size of 12 items, the highest ranked item is chosen 10% less than predicted by a RUM with a constant variance.

3.4.3 Observing Regularity Violations

As noted earlier, the NRUM with normalization (for some parameter values) predicts violations of regularity. Specifically, the probability of choosing the second ranked alternative should increase when the number of alternatives added to a choice set are of low value and the variance in the model is small.

To investigate this prediction, we report the probability of choosing the second ranked alternative for four subsets of the data (in blue) in Figure 8. The subsets are generated by selecting the trials (from each subject) in which all of the distracters were in the 50th, 55th, 60th, and 70th of bids for that subject. Ideally, we would examine subsets lower than this range since this is where we would expect the largest number of violations, but unfortunately the design of our experiment severely underweights sampling from this range of alternatives. This significantly reduces the number of trials which fall into these subsets and limits the power of our dataset to identify the type of regularity violations predicted by normalization. As can be observed in Figure 8, there are some cases in which the average probability of choosing the second item increases as the set size grows, but these results are not statistically significant.

4 Neurobiological Constraints and Optimal Decision-Making

We have just observed that the ability of subjects to select their highest ranked alternative degrades as both the composition and the size of the choice set is altered. Additionally, we saw that a choice model which incorporates divisive normalization is able to capture these behavioural patterns which violate seemingly normative choice behaviour. Given that we observe these effects in behaviour, this raises the natural question of why a neural system – selected

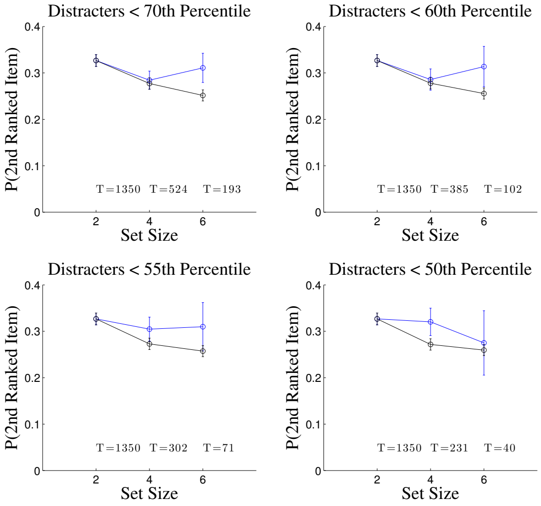

Figure 8: The probability of choosing the second ranked alternative for four subsets of the data (in blue) for choice sets of increasing size. Each subset is composed of all trials for each subject in which the distracters were less than the 50th, 55th, 60th, and 70th percentile of the distracter bids for that subject. The number of trials (over all subjects) in the subset are reported, and the probabilities calculated from the remaining trials are depicted (in black).

over tens of millions of years of evolution for the human brain alone – would exhibit adverse choice performance?

To provide some insight, we consider a neural system which is tasked with representing the subjective value of each element in the choice set for the purpose of ultimately determining which of them is largest. This refocuses our inquiry onto the question of how the brain should represent possible valuations given a set of neurobiological constraints on the precision and resolution with which it can encode quantities. Our discussion of optimality will essentially amount to an argument that divisive normalization leads to as few choice errors as possible given these neurobiological constraints.[15]

The argument assumes that evolutionary processes have allocated a neural resource of bounded size to decision-making, leaving open only the question of how decision-making is implemented given this bound. This assumption rests on the empirical observation that the resources allocated to decision-making

---

[15] Our focus on *choice errors* arises from the economic literature which has previously used errors as a metric of optimality (Robson, 2001; Rayo and Becker, 2007; Robson and Samuelson, 2010). While we do not argue explicitly in terms of the evolutionary criteria of “expected loss” (Netzer, 2009), which incorporates the magnitude of a choice error, we will see that the general divisive normalization form does enable a degree of monotonicity in the transformation of subjective values, such that “more valuable” choice sets lead to fewer errors than less valuable choice sets.

cannot be endogenously controlled by the decision-maker. Some background on the structure of the brain might be helpful in justifying this approach.

The total pool of neurons that make up the cortex of the brain can be divided into non-overlapping subsets of neurons, called Broadmann areas, and each of these areas is distinct in structure and organization. Empirical evidence suggests that any given Broadmann area serves a limited set of functional roles – which are the same across individuals – and a small subset of these areas have been identified as having a role in decision-making. It is essentially required by this logic that the amount of brain devoted to these decision-making areas must balance the advantages of devoting more neural resources to decision-making against the costs of that allocation to other functions, such as processing incoming sensory information or managing the heartbeat.

So why is a decision-maker limited in its ability to recruit more resources, even if it might be momentarily advantageous? The reason is that Broadmann areas are distinct in their connections with the rest of the brain, distinct in their cellular structure, and distinct in many of their biophysical properties. These specializations which distinguish the Broadmann areas and make them appropriate for particular functions develop over years, often over the first 20 years of life, and cannot be changed quickly. This imposes a significant constraint on the ability of the decision-maker to endogenously control the allocation of neural resources to a task. For these reasons, we take the allocation of neural resources to be effectively fixed at a particular moment in time, at least for the time horizon in which decisions are typically made by an individual.[16] Of course, a full normative treatment of the resource allocation problem would require a complete accounting of the costs of decision-related errors to the chooser and the metabolic costs of relaxing the constraints we identify — on an evolutionary timescale.

Now that we have set the stage, we will begin by considering the general normalization equation

$$z_i = \kappa \frac{v_i}{\sigma + \sum_n (v_n^\beta)^{\frac{1}{\beta}}} \tag{11}$$

under the assumption that $\sigma = 0$ to aid exposition. We will introduce two neurobiological constraints in section 4.1 and demonstrate that the normalization equation implements the normative solution to these constraints under particular parameterizations. We will then relax our assumption on σ in section 4.2 and explore its implications.

4.1 Neurobiological Constraints

Since neurons require energy, and energy resources in the human body are finite, there is a limit on the total number of neurons in the brain, the size of these

---

[16] There is a large literature in neuroscience and psychology which looks at the time spent on a decision and the prevalence of errors, known as the speed/accuracy tradeoff. This literature finds that one can improve accuracy with longer decision times, but that this increase in accuracy is bounded. The fact that increasing the time devoted to a decision cannot overcome this resource limitation problem has been both widely observed and biologically explained.

neurons, and on the number that can be allocated to any given task. This basic observation has two neurobiological implications that we can incorporate into our choice model through two constraints.

Our first constraint arises from the stochasticity of neural systems. As noted in section 2.1, the dynamic neural process which compares valuations is stochastic, and this system yields more errors when valuations are close together. This can be visualized for a binary choice set in Figure 9. When two alternatives are close together in value (i.e. close to the 45° line), this system is more likely to choose the lower valued item.

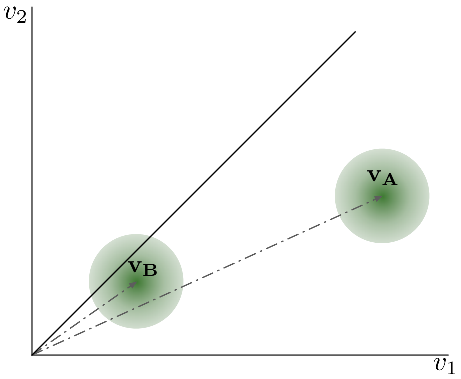

Figure 9: A binary choice for two different choice sets, A and B. Stochasticity in the comparison between v_1 and v_2 means the lower alternative (alternative 2) might be deemed highest, and this is more likely in choice set B.

Our second constraint amounts to a resource constraint on neural activity. As mentioned, we take the size of the brain area, and the number of neurons allocated to the representation of value, as fixed. We also impose a maximal firing rate for a neuron due to its refractory period. Together, this implies that the neural activity for the neurons allocated to representing z_i is bounded above at some level $\bar{z}$. Since increasing neural activity up to this boundary requires rivalrous resources, we initially impose the strong assumption that average firing rate within the brain area where valuations are represented is bounded, or equivalently, $\sum_i z_i \leq \bar{z}$. We explore this constraint in greater detail and consider weaker versions in section 4.1.1. Relaxing this assumption will not change the general result of our argument, only its magnitude.

Now we can observe the effect of normalization in our neural system. Our assumption that the average firing rate is bounded can be implemented through a particular form of the divisive normalization equation (Figure 10). If $\beta = 1$, then the normalization function $Z(\mathbf{v})$ scales the vector $\mathbf{v}$ to the vector $\mathbf{z}$ along a ray to the origin by means of the transformation:

$$z_i = \kappa \frac{v_i}{\sum_n v_n}. \tag{12}$$

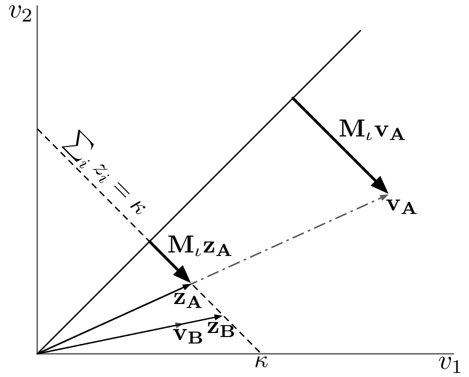

Figure 10: Divisive normalization scales $\mathbf{v}$ to the resource constraint. This scaling can lead to more choice errors (relative to $\mathbf{v}$) if the distance between $\mathbf{z}$ and the 45° line is smaller.

By re-interpreting the gain control parameter κ as the maximum neural activity $\bar{z}$, divisive normalization effectively proportions neural resources between the alternatives along the resource constraint

$$\sum_i z_i = \sum_i \frac{\kappa v_i}{\sum_n v_n} = \kappa \frac{\sum_i v_i}{\sum_n v_n} = \kappa. \tag{13}$$

The rescaling implemented by divisive normalization has important implications for choice behaviour. Let $\mathbf{M}_\iota \mathbf{v} = \mathbf{v} - \frac{\iota \cdot \mathbf{v}}{\iota \cdot \iota} \iota$ denote the image of the projection of the vector $\mathbf{v}$ to the unit vector ι (the 45° line). Therefore the length of the vector $\|\mathbf{M}_\iota \mathbf{v}\|$, relative to the variance of the noise η in the argmax operation, determines the probability of choosing a lower valued alternative in the absence of normalization.

What is the effect of normalization on this error probability? Since normalization scales all vectors $\mathbf{v}$ to the boundary of the firing rate budget, the length $\|\mathbf{M}_\iota \mathbf{z}\|$ is made as large as possible given the constraint on neural activity. If $\mathbf{v}$ is inside this constraint, normalization amplifies the relative magnitudes and $\|\mathbf{M}_\iota \mathbf{z}\| > \|\mathbf{M}_\iota \mathbf{v}\|$, therefore an error is made with lower probability. If $\mathbf{v}$ is above the neural resource constraint, normalization scales these quantities down to the boundary of the resource constraint. While this results in $\|\mathbf{M}_\iota \mathbf{z}\| < \|\mathbf{M}_\iota \mathbf{v}\|$, therefore a larger probability of error, this probability is made as small as possible since $\mathbf{z}$ does not lie in the interior of the constraint.

To summarize, divisive normalization rescales valuations to lie along a resource constraint which bounds the activity of neural systems, and this re-scaling can lead to seemingly sub-optimal choice behaviour (i.e. intransitivities due to stochastic choice). This phenomenon is magnified in cases where normalization scales quantities down, leading to more choice errors than would otherwise be observed. As more choice alternatives are added, or the value of some choice alternatives are modulated, this adverse effect on choice behaviour is amplified. However we should emphasize that these errors are optimal given the neural constraints facing the system (in particular a fixed resource κ), since divisive normalization makes $\|\mathbf{M}_\iota \mathbf{z}\|$ as large as possible given the constraints, limiting choice errors as much as possible.

4.1.1 Congested Resource Constraints in Neural Systems

In constructing our original resource constraint, we imposed the strong assumption that average neural activity is bounded within a brain area. Under this assumption, an increase in the neural activity of one alternative necessitates an equivalent decrease in the activity of the other alternative(s), yielding a linear constraint. While we should emphasize that the metabolic resources of neurons are rivalrous (i.e. oxygen consumed by one neuron cannot be consumed by another), we believe that the assumption of pure rivalry and a linear resource constraint is too strong for a neural system, primarily due to the spatial organization of a neural substrate. Moreover, our empirical analysis found estimates of β much larger then required by a linear constraint ($\beta = 1$). To account for these facts, we propose that a degree of non-rivalry in neural resources can be well-modelled by a relaxed constraint which incorporates congestion.

Consider the following thought experiment depicted in Figure 11. There are two neurons A and B, and each neuron requires one unit of a resource in order to produce one output (action potential) at a given point in time. Each neuron is located at a distinct location of neural substrate, and can only collect resources from within some radius of this location. The grid denotes the locations at which metabolic resources are available, and the small black squares denote resources that are available for collection by at least one neuron.

In this thought experiment, the spatial organization of the neurons determines the degree to which metabolic resources are rivalrous. If the neurons are organized so that the area in which they can collect resources do not overlap, then each neuron is free to collect the maximum amount and discharge at its maximum firing rate at a given point in time. In this case, the resources are non-rivalrous and the resource constraint only binds at the maximum output.

At the other extreme, if the two collection areas completely overlap then the neurons compete for the same, rivalrous, resources. Here, the resource constraint takes the familiar linear form in which an increase of activity for one neuron necessitates a decrease for the other. This is the restrictive constraint we originally assumed.

We believe the appropriate constraint lies between these two extremes. If the neurons are close enough so that their “collectable areas” only partially overlap, some of the neural resources are rivalrous, and some are not. This leads to a kink in the resource constraint. A simple extension to incorporate random motion of resource location, or a probability of collection that varies inversely with distance, will smooth the resource constraint to lie interior of the two extremes. This yields a constraint akin to a congested resource in the public goods literature.

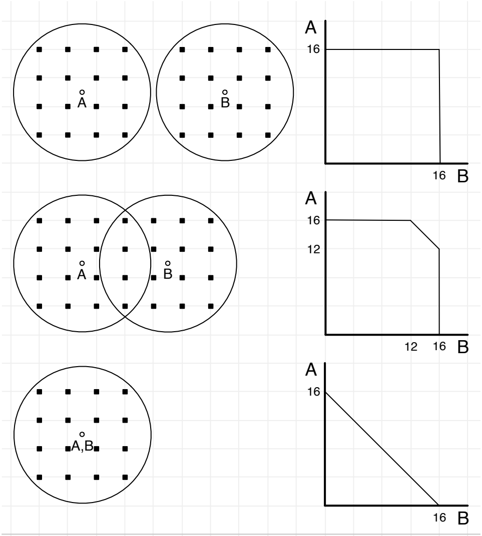

Figure 11: Congestion of Neural Resources

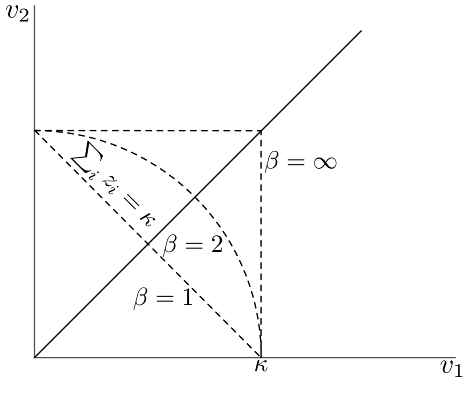

Figure 12: Three possible neural resource constraints that vary in their degree of rivalry. If neural resources are perfectly rivalrous, the constraint is linear. At the other extreme, if resources were non-rivalrous all neurons could be maximally active simultaneously. The intermediate case of a congested resource lies between the two. The paramterization of the normalization equation (β) which implements each constraint is depicted.

Just as in the linear case, the divisive normalization equation can map subjective values to a congested resource constraint if $\beta > 1$ (i.e. the denominator is a norm of order larger than one[17]). Normalization then maps $\mathbf{v}$ to a geometric figure known as a *super-ellipse* (Figure 12). For instance, if we maintain $\sigma = 0$ and set $\beta = 2$, the divisive normalization equation scales $\mathbf{v}$ to the circle with radius κ, as in Figure 12. Setting $\beta = 3$ maps to a hyperellipse, while $\beta = \infty$ maps to the square of size κ. Since this family of curves lies outside the original linear constraint, we can verify that $\|\mathbf{M}_\iota \mathbf{z}\|$ is now larger. This yields fewer choice errors than in the case of the overly-restrictive linear constraint, and still reduces choice errors as much as possible. Therefore the adverse normative implications are in a similar direction, but less pronounced, than the linear case.

The congested resource constraint has a number of nice features. First, it captures the idea that metabolic resources are rivalrous, but not perfectly so. Second, when the valuations are closer together, average neural activity is higher than when the difference between them is easier to distinguish. This captures the intuition that the total neural activity increases for decisions that are more difficult. Finally, it provides a normative account for our empirical estimates of relatively large values for β.

4.2 Relaxing σ

We now relax the assumption that $\sigma = 0$. This re-introduces monotonicity into the transformed values (Figure 13), but in a form which does not improve choice performance. For a given $\sigma > 0$, choice sets with lower total valuations are mapped to a line which intercepts lower than κ.

$$\sum_i z_i = \sum_i \frac{\kappa v_i}{\sigma + \sum_n v_n} = \kappa \frac{\sum_i v_i}{\sigma + \sum_n v_n} < \kappa. \tag{14}$$

With a positive σ, the length of the vector $\mathbf{z}$ now depends monotonically on the length of the vector $\mathbf{v}$. This yields the prediction that choice errors are reduced for more valuable choice sets than for less valuable chose sets. However, choice errors increase relative to the case in which $\sigma = 0$, therefore it seems that a neural system for decision-making in *static choice environments* should not have a positive σ. This is confirmed in our empirical results, in which our estimates of σ were near zero.[18]

5 Conclusion

In the study of the human brain and its functional relationship to behaviour, it has long been recognized that the biology of neural systems places constraints

---

[17] A norm of order p, or p-norm, is defined as $\|x\|_p \equiv \left(\sum_{n=1}^{N} |x_n|^p\right)^{1/p}$

[18] There is reason to believe this might not be the case in dynamic choice environments in which an average estimate of the choice set, or reference point, can prove useful for achieving optimal choice, suggesting that σ may in fact be a dynamic variable which varies with expectations.

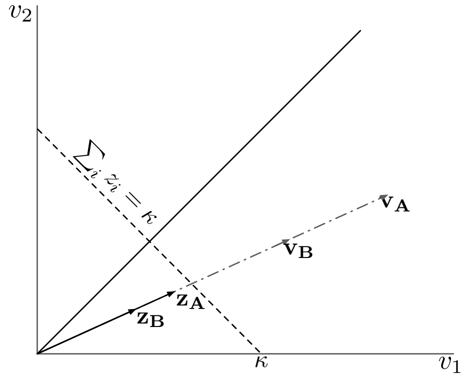

Figure 13: Divisive normaliztion when $\sigma > 0$.

on the form of neural computation. To abide by these resource constraints, information about the objective world is approximated, compressed, and transformed into quantities that can be represented in finite neural activity. It is only natural to examine what these constraints imply for choice, and whether theories of economic behaviour need incorporate these insights. In this article, we pursue this goal directly. We establish that a canonical neural computation, divisive normalization, predicts context-dependent choice behaviour. We verify these predictions empirically in two behavioural laboratory experiments, and demonstrate normalization should be considered optimal given the neurobiological constraints on the decision problem implemented by neural systems.

It is now established that measurements of neural quantities in the frontal cortex of the human brain (which we term subjective value) correlate with all known incentive-compatible methods for measuring value and predicting choice behaviour (Bartra, McGuire and Kable, 2013; Levy and Glimcher, 2011). This has led to the development of neuroeconomic models which take neural quantities as their primitives, and allow for neural constraints to be imposed on a model of economic behaviour (Webb et al., 2013; Fehr and Rangel, 2011). In this article, we identify a particular neural computation, divisive normalization, which is a canonical feature of cortex observed across all sensory modalities, a vast range of species, and recently in decision-making areas of the primate brain (Carandini and Heeger, 2012; Louie, Grattan and Glimcher, 2011). In sensory processing, this computation has been demonstrated to be normative (Schwartz and Simoncelli, 2001), however in the domain of decision-making the normative question has remained unclear because divisive normalization can lead to adverse choice performance.

The key insight of divisive normalization is that the subjective value of a particular choice alternative is suppressed, or normalized, by the total value of the choice set. This creates a relationship between the variables used to predict choice behaviour and the composition and size of the choice set, yielding context-dependent choice. The effects of this relationship include increased stochasticity, particular patterns of violation of the Independence of Irrelevant Alternatives axiom, and a decreased likelihood of choosing a higher valued alternative. We present evidence for such behaviour from two laboratory experiments, one which varies the composition of the choice set and one which varies its size, and demonstrate that a model which incorporates divisive normalization more accurately captures the patterns of substitution that we observe compared to existing econometric methods.

The fact that normalization predicts such adverse choice behaviour, and that we observe it in our two choice experiments, returns us to the normative question of why such behaviour is observed. Our answer is that the neural system which implements decision-making operates under a neurobiological constraint with two implications: neural systems are stochastic, and there is a bound (resource constraint) on neural activity. We demonstrate that a neuroeconomic choice model employing normalization scales valuations to the resource constraint, providing an efficient means for the brain to represent valuations because it yields the fewest possible choice errors.

All together, our results speak to the role of incorporating insights from biology into economic discourse. While developing an understanding of the brain is intriguing in its own right, the goal of neuroeconomic research should be predicting novel patterns in choice behaviour and producing a normative explanation for the behaviour we observe. This challenge to neuroeconomics is stated explicitly by Bernheim (2009):

> *"Provide an example of a novel economic model derived originally from neuroeconomic research that improves our measurement of the causal relationship between a standard exogenous environmental condition – one with which economists have been historically concerned – and a standard economic choice."*

We would argue that divisive normalization does just this. Our results emphasize both the positive advances offered by choice models grounded in neuroscience, and the normative role neurobiological constraints can play in the study of decision-making.

A Divisive Normalization and Regularity

We prove that the NRUM with divisive normalization will technically violate the regularity condition for some, but not all, scalings of the model (i.e. variance) and values. Our proof focuses on the extreme value distribution since it has closed-form probabilities. We have not yet characterized the entire range of distributions and range of valuations for which regularity fails, but extensive simulation suggests it is small.

We begin with a definition of Random Utility. Let $P_{i,N}$ denote the probability that a subject chooses alternative i from a finite set N.

Definition 1. *(Random Utilities) There exists a random vector* $U=(U_1,\ldots,U_n)$*, unique up to an increasing monotone transformation such that for any* $i \in M \subseteq N$*,*

$$Pr\{U_i \geq U_j, \forall j \in M\} = P_{i,M}.$$

A necessary condition for a random utility representation states that the probability of choosing an alternative decreases as more alternatives are added to the choice set (Block and Marschak, 1960), often referred to as a regularity condition.

Definition 2. *(Regularity) If* $L \subseteq M \subseteq N$*,*

$$P_{i,M} \leq P_{i,L}.$$

We now show that a model with normalization violates the regularity condition for some parameter values, therefore cannot be rationalized by a Random Utility model. We consider two choice sets of different sizes, the three alternative choice set {1,2,3}, with corresponding valuations $v_1 \geq v_2 \geq v_3 \geq 0$, and the two alternative subset {1,2}. Since we are assuming an extreme value distribution, the choice probabilities of the normalized model are given by the closed form equation,

$$P_{i,M} = \frac{e^{\kappa \frac{v_i}{\sum_{n \in M} v_n}}}{\sum_{m \in M} e^{\kappa \frac{v_m}{\sum_{n \in M} v_n}}}.$$

We will focus on the probability of choosing the second (ranked) item from the set. For our simple choice sets, these probabilities can be re-written

$$P_{2,\{1,2,3\}} = \frac{1}{e^{\frac{\kappa(v_1-v_2)}{v_1+v_2+v_3}} + e^{\frac{\kappa(v_3-v_2)}{v_1+v_2+v_3}} + 1},$$

and

$$P_{2,\{1,2\}} = \frac{1}{e^{\frac{\kappa(v_1-v_2)}{v_1+v_2}} + 1}.$$

We will show regularity is violated as the choice set is expanded from two to three alternatives.

Result 1. *An NRUM which includes divisive normalization is not Regular for some parameter values, therefore cannot be represent by Random Utilities. Specifically, there exists some* $\kappa > 0$ *such that* $P_{2,\{1,2,3\}} > P_{2,\{1,2\}}$*.*

Proof. We proceed by contradiction. Suppose $P_{2,\{1,2\}} > P_{2,\{1,2,3\}}$ for all κ.

Therefore,

$$e^{\frac{\kappa(v_1-v_2)}{v_1+v_2+v_3}} + e^{\frac{\kappa(v_3-v_2)}{v_1+v_2+v_3}} + 1 > e^{\frac{\kappa(v_1-v_2)}{v_1+v_2}} + 1$$

$$e^{\frac{\kappa(v_1-v_2)}{v_1+v_2+v_3} - \frac{\kappa(v_1-v_2)}{v_1+v_2}} + e^{\frac{\kappa(v_3-v_2)}{v_1+v_2+v_3} - \frac{\kappa(v_1-v_2)}{v_1+v_2}} > 1.$$

Now we must find values for κ, v_1, v_2, and v_3 for which this statement is false. Choose $\kappa = v_1 + v_2$, and $v_3 = v_2 = \frac{1}{2}v_1$. The above statement reduces to

$$\begin{aligned} e^{\frac{(v_1+v_2)(v_1-v_2)}{v_1+v_2+v_3}-(v_1-v_2)} + e^{\frac{(v_1+v_2)(v_3-v_2)}{v_1+v_2+v_3}-(v_1-v_2)} &> 1 \\ e^{\frac{(v_1+v_2)(v_1-v_2)}{v_1+2v_2}-(v_1-v_2)} + e^{-(v_1-v_2)} &> 1 \\ e^{-\frac{1}{8}v_1} + e^{-\frac{1}{2}v_1} &> 1 \end{aligned}$$

As v_1 grows, both elements on the left-hand side tend to zero, and we have a contradiction. □

B Estimation of Divisive Normalization

We report monte-carlo simulations of our estimates from equation 10, restated here.

$$P_{i,t} = \int \mathbb{1}\left[\frac{v_{i,t} - v_{j,t}}{\sigma + \omega(\sum_{n \in N_t} v_{n,t}^{\beta})^{\frac{1}{\beta}}} > \eta_{j,t} - \eta_{i,t}, \quad \forall j \neq i\right] f(\eta_t) d\eta_t. \tag{10}$$

This multi-dimensional integral is reduced to a single dimension through the assumption of an independent normal distribution for $f()$, and is approximated via Gaussian quadrature (Genz and Betz, 2009).

Mean parameter estimates from a simulated model with and without normalization are reported in Table 4. There are 2000 trials per choice set size, for a total of $T = 12000$. Note than under the null hypothesis ($\omega = 0$), the parameter β is not identified, and the log-likelihood does not decrease as the specification is relaxed. This is in contrast to the log-likelihood of the fully normalized specification. It decreases significantly relative to a model in which the variance is restricted to be constant and not normalized ($\omega = 0$).

Table 4: Mean parameter estimates from 100 choice simulations with and without normalization. Mean squared errors are reported.

	True	Simulation Without Normalization			
$\bar{\hat{\sigma}}$	1	1.001		1	0.497
		(0.001)			(0.258)
$\bar{\hat{\omega}}$	0	0		0.001	0.126
				(0.001)	(0.016)
$\bar{\hat{\beta}}$	1	1		286.74	295.72
				(81648.4)	(88834.05)
$\overline{LL}$		-9926.71		-9926.71	-9926.71

	True	Simulation With Normalization			
$\bar{\hat{\sigma}}$	1	18.09	1.026	1	1.745
		(1473.60)	(0.113)		(4.642)
$\bar{\hat{\omega}}$	1	0	1	1.002	0.892
				(0.099)	(0.118)
$\bar{\hat{\beta}}$	1	1	1	1.000	0.948
				(0.003)	(0.026)
$\overline{LL}$		-20731.94	-20629.54	-20629.00	-20628.68

References

Barlow, Horace B. 1961. “Possible principles underlying the transformation of sensory messages.” In *Sensory Communication.* , ed. W A Rosenblith. MIT Press.

Barron, Helen C, Raymond J Dolan, and Timothy E J Behrens. 2013. “Online evaluation of novel choices by simultaneous representation of multiple memories.” *Nature Neuroscience*, 16(10): 1492–1498.

Bartra, Oscar, Joseph T McGuire, and Joseph 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.

Becker, Gordon M, Morris H DeGroot, and Jacob Marschak. 1963. “Stochastic models of choice behavior.” *Behavioral Science*, 8(1): 41–55.

Becker, Gordon M, Morris H DeGroot, and Jacob Marschak. 1964. “Measuring utility by a singleresponse sequential method.” *Behavioral Science*, 9(3): 226–232.

Ben-Akiva, Moshe E, and Steven R Lerman. 1985. *Discrete Choice Analysis. Theory and Application to Travel Demand*, MIT Press.

Bernheim, B Douglas. 2009. “On the potential of neuroeconomics: a critical (but hopeful) appraisal.” *American Economic Journal: Microeconomics*, 1(2): 1–41.

Block, H D, and Jacob Marschak. 1960. “Random Orderings and Stochastic Theories of Responses.” *Cowles Foundation Discussion Papers.*

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. “Salience theory of choice under risk.” *The Quarterly Journal of Economics*, 127(3): 1243–1285.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. “Salience and Consumer Choice.” *Harvard Working Paper.*

Brownstone, David, and Kenneth Train. 1999. “Forecasting new product penetration with flexible substitution patterns.” *Journal of Econometrics*, 89: 109–129.

Busemeyer, Jerome R, and James T Townsend. 1993. “Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment.” *Psychological Review*, 100(3): 432–459.

Carandini, Matteo, and David J Heeger. 2012. “Normalization as a canonical neural computation.” *Nature Reviews Neuroscience*, 13(1): 51–62.

Chib, Vikram S, Antonio Rangel, Shinsuke Shimojo, and John P O'Doherty. 2009. "Evidence for a common representation of decision values for dissimilar goods in human ventromedial prefrontal cortex." *The Journal of Neuroscience*, 29(39): 12315–12320.

Debreu, Gerard. 1960. "Individual Choice Behavior: A Theoretical Analysis by R. Duncan Luce." *The American Economic Review*, 50(1): 186–188.

Delgado, M R, L E Nystrom, C Fissell, D C Noll, and J A Fiez. 2000. "Tracking the hemodynamic responses to reward and punishment in the striatum." *Journal of Neurophysiology*, 84(6): 3072–3077.

DeShazo, J R, and German Fermo. 2002. "Designing Choice Sets for Stated Preference Methods: The Effects of Complexity on Choice Consistency." *Journal of Environmental Economics and Management*, 44(1): 123–143.

Dorris, Michael C, and Paul W Glimcher. 2004. "Activity in posterior parietal cortex is correlated with the relative subjective desirability of action." *Neuron*, 44(2): 365–378.

Elliott, Rebecca, Karl J Friston, and Raymond J Dolan. 2000. "Dissociable neural responses in human reward systems." *The Journal of Neuroscience*, 20(16): 6159–6165.

Fehr, Ernst, and Antonio Rangel. 2011. "Neuroeconomic Foundations of Economic Choice—Recent Advances." *Journal of Economic Perspectives*, 25(4): 3–30.

Genz, Alan, and Frank Betz. 2009. *Computation of Multivariate Normal and* t *Probabilities.* Berlin:Springer.

Glimcher, Paul W. 2005. "Indeterminacy in Brain and Behavior." *Annual Review of Psychology*, 56(1): 25–56.

Glimcher, Paul W. 2011. *Foundations of Neuroeconomic Analysis.* Oxford University Press.

Gold, Joshua I, and Michael N Shadlen. 2007. "The neural basis of decision making." *Annual Review of Neuroscience*, 30: 535–574.

Hausman, Jerry A, and David A Wise. 1978. "A Conditional Probit Model for Qualitative Choice: Discrete Decisions Recognizing Interdependence and Heterogeneous Preferences." *Econometrica: Journal of the Econometric Society*, 46(2): 403–426.

Heeger, David J. 1992. "Normalization of cell responses in cat striate cortex." *Visual Neuroscience*, 9(2): 181–197.

Hess, Stephane, and John M Rose. 2009. "Allowing for intra-respondent variations in coefficients estimated on repeated choice data." *Transportation Research Part B: Methodological*, 43: 708–719.

Huber, Joel, John W Payne, and Christopher Puto. 1982. "Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis." *Journal of Consumer Research*, 9(1): 90–98.

Iyengar, Sheena S, and Mark R Lepper. 2000. "When choice is demotivating: can one desire too much of a good thing?" *Journal of Personality and Social Psychology*, 79(6): 995–1006.

Kable, Joseph W, and Paul W Glimcher. 2007. "The neural correlates of subjective value during intertemporal choice." *Nature Neuroscience*, 10(12): 1625–1633.

Kable, Joseph W, and Paul W Glimcher. 2010. "An 'As Soon As Possible' Effect in Human Intertemporal Decision Making: Behavioral Evidence and Neural Mechanisms." *Journal of Neurophysiology*, 103: 2513–2531.

Kiani, Roozbeh, Timothy D Hanks, and Michael N Shadlen. 2008. "Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment." *The Journal of Neuroscience*, 28(12): 3017–3029.

Knutson, Brian, Charles M Adams, Grace W Fong, and Daniel Hommer. 2001. "Anticipation of increasing monetary reward selectively recruits nucleus accumbens." *The Journal of Neuroscience*, 21(16): RC159.

Knutson, Brian, Grace W Fong, Shannon M Bennett, Charles M Adams, and Daniel Hommer. 2003. "A region of mesial prefrontal cortex tracks monetarily rewarding outcomes: characterization with rapid event-related fMRI." *NeuroImage*, 18(2): 263–272.

Levy, Dino, and Paul W Glimcher. 2011. "Comparing Apples and Oranges: Using Reward-Specific and Reward-General Subjective Value Representation in the Brain." *The Journal of Neuroscience*, 31(41): 14693–14707.

Levy, Dino J, and Paul W Glimcher. 2012. "The root of all value: a neural common currency for choice." *Current Opinion in Neurobiology*, 22(6): 1027–1038.

Levy, Ifat, J Snell, A J Nelson, Aldo Rustichini, and Paul W Glimcher. 2010. "Neural Representation of Subjective Value Under Risk and Ambiguity." *Journal of Neurophysiology*, 103(2): 1036–1047.

Levy, Ifat, Stephanie C Lazzaro, Robb B Rutledge, and Paul W Glimcher. 2011. "Choice from non-choice: predicting consumer preferences from blood oxygenation level-dependent signals obtained during passive viewing." *The Journal of Neuroscience*, 31(1): 118–125.

Louie, Kenway, and Paul W Glimcher. 2010. "Separating value from choice: delay discounting activity in the lateral intraparietal area." *The Journal of Neuroscience*, 30(16): 5498–5507.

Louie, Kenway, Lauren E Grattan, and Paul W Glimcher. 2011. "Reward value-based gain control: divisive normalization in parietal cortex." *The Journal of Neuroscience*, 31(29): 10627–10639.

Louie, Kenway, Mel W Khaw, and Paul W Glimcher. 2013. "Normalization is a general neural mechanism for context-dependent decision making." *Proceedings of the National Academy of Sciences*, 110(15): 6139–6144.

Luce, R. Duncan. 1959. *Indvidual Choice Behaviour.* John Wiley ands Sons, Inc.

Lyu, Siwei. 2011. "Dependency reduction with divisive normalization: justification and effectiveness." *Neural computation*, 23(11): 2942–2973.

McFadden, Daniel. 1974. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Economic Theory and Mathematical Economics.* , ed. Paul Zarembka. Academic Press, Inc.

McFadden, Daniel. 1978. "Modeling the choice of residential location." In *Spatial Interaction Theory and Planning Models.* , ed. A Karlqvist, L Lundqvist, F Snickars and J Weibull, 75–96. North-Holland, Amsterdam.

McFadden, Daniel. 2001. "Economic choices." *The American Economic Review*, 91(3): 351–378.

McFadden, Daniel, and K Train. 2000. "Mixed MNL models for discrete response." *Journal of Applied Econometrics*, 15: 447–470.

McFadden, Daniel L. 2005. "Revealed stochastic preference: a synthesis." *Economic Theory*, 26(2): 245–264.

Meyer, Robert J. 1991. "Probablistic Models of Consume Choice Behavior." In *Handbook of Consumer Behavior.* , ed. Thomas S Robertson. Prentice Hall, Inc.

Milosavljevic, M, J Malmaud, A Huth, Christof Koch, and Antonio Rangel. 2010. "The drift diffusion model can account for value-based choice response times under high and low time pressure." *Judgement and Decision Making*, 5(6): 437–449.

Netzer, Nick. 2009. "Evolution of Time Preferences and Attitudes toward Risk." *The American Economic Review*, 99(3): 937–955.

Padoa-Schioppa, Camillo, and J A Assad. 2006. "Neurons in the orbitofrontal cortex encode economic value." *Nature*, 441(7090): 223–226.

Pastor-Bernier, Alexandre, and Paul Cisek. 2011. "Neural correlates of biased competition in premotor cortex." *The Journal of Neuroscience*, 31(19): 7083–7088.

Plassmann, Hilke, John P O'Doherty, and Antonio Rangel. 2007. "Orbitofrontal cortex encodes willingness to pay in everyday economic transactions." *The Journal of Neuroscience*, 27(37): 9984–9988.

Platt, Michael L, and Paul W Glimcher. 1999. "Neural correlates of decision variables in parietal cortex." *Nature*, 400(6741): 233–238.

Rayo, Luis, and Gary S Becker. 2007. "Evolutionary Efficiency and Happiness." *Journal of Political Economy*, 115(2): 302–337.

Revelt, David, and Kenneth Train. 1998. "Mixed Logit with Repeated Choices: Households' Choices of Appliance Efficiency Level." *Review of Economics and Statistics*, 80(4): 647–657.

Rieskamp, Jorg, Jerome R Busemeyer, and B.A. Mellers. 2006. "Extending the bounds of rationality: evidence and theories of preferential choice." *Journal of Economic Literature*, 44(3): 631–661.

Robson, Arthur J. 2001. "The biological basis of economic behavior." *Journal of Economic Literature*, 39(1): 11–33.

Robson, Arthur J, and L Samuelson. 2010. "The Evolutionary Foundations of Preferences." In *Handbook of Social Economics.* , ed. J Benhabib, A Bisin and Matthew Jackson, 221–310. North Holland.

Roitman, Jaime D, and Michael N Shadlen. 2002. "Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task." *The Journal of Neuroscience.*

Rorie, Alan E, Juan Gao, James L McClelland, and William T Newsome. 2010. "Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey." *PloS one*, 5(2): e9308.

Schwartz, Odelia, and E P Simoncelli. 2001. "Natural signal statistics and sensory gain control." *Nature Neuroscience*, 4(8): 819–825.

Shannon, C E. 1948. "A Mathematical Theory of Communication." *Bell System Technical Journal*, 27(3): 379–423.

Simon, Herbert A. 1979. "Rational Decision Making in Business Organizations ." *The American Economic Review*, 69(4): 493–513.

Simonson, Itamar, and Amos Tversky. 1992. "Choice in Context: Tradeoff Contrast and Extremeness Aversion." *Journal of Marketing Research*, 29(3): 281–295.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics*, 50(3): 665–690.

Sinz, Fabian, and Matthias Bethge. 2013. "Temporal Adaptation Enhances Efficient Contrast Gain Control on Natural Images." *PLoS Computational Biology*, 9(1): e1002889.

Soltani, Alireza, Benedetto De Martino, and Colin F Camerer. 2012. "A range-normalization model of context-dependent choice: a new model and evidence." *PLoS Computational Biology*, 8(7): e1002607.

Sugrue, Leo P, Greg S Corrado, and William T Newsome. 2004. "Matching Behavior and the Representation of Value in the Parietal Cortex." *Science*, 304: 1782–1787.

Thaler, Richard. 1985. "Mental accounting and consumer choice." *Marketing science*, 4(3): 199–214.

Thaler, Richard H. 1999. "Mental Accounting Matters." *Journal of Behavioral Decision Making*, 12: 123–206.

Train, Kenneth E. 2009. *Discrete Choice Methods with Simulation.* . 2nd ed., Cambridge University Press.

Train, Kenneth E, and Clifford Winston. 2007. "Vehicle Choice Behavior and the Declining Market Share of U.S. Automakers." *International Economic Review*, 48(4): 1469–1496.

Tusche, Anita, Stefan Bode, and John-Dylan Haynes. 2010. "Neural Responses to Unattended Products Predict Later Consumer Choices." *The Journal of Neuroscience*, 30(23): 8024–8031.

Tversky, Amos, and Itamar Simonson. 1993. "Context-Dependent Preferences." *Management Science*, 39(10): 1179–1189.

Wainwright, Martin J, Odelia Schwartz, and E P Simoncelli. 2001. "Natural Image Statistics and Divisive Normalization: Modeling Nonlinearities and Adaptation in Cortical Neurons." In *Statistical Theories of the Brain.* MIT Press.

Webb, Ryan. 2013. "Dynamic constraints on the distribution of stochastic choice: Drift Diffusion implies Random Utility." *SSRN Working Paper.*

Webb, Ryan, Paul W Glimcher, Ifat Levy, Stephanie Lazzaro, and Robb B Rutledge. 2013. "Neural Random Utility and Measured Value." *SSRN Working Paper.*

Woodford, Michael. 2012. "Prospect Theory as Efficient Perceptual Distortion." *The American Economic Review*, 102(3): 41–46.