Neural Random Utility and Measured Value

Ryan Webb*
Paul W. Glimcher
Center for Neuroeconomics, NYU

Ifat Levy
Yale School of Medicine

Stephanie C. Lazzaro
Robb B. Rutledge
University College London

October 30, 2013

Abstract

We present a method for relating a neural measurement of value to choice behaviour. In a previous study, precisely targeted measurements of brain activity were made in the medial prefrontal cortex of subjects while they considered individual consumer goods. We present here two advances. First, we develop an empirical framework for relating this class of measured value data to choice prediction. Second, we apply a benchmarking tool to compare the predictive power of a measured value dataset with established techniques. We find that our measured neural activity cardinally encodes valuations and predicts choice behaviour, though a significant degree of measurement error affects prediction rates. Accounting for measurement error and combining neural data with standard observables improves predictive performance. We also note some potential normative implications of our measured value approach.

KEYWORDS: Neuroeconomics, Random Utility, Stochastic Choice.

JEL Classification: D87.

*Center for Neuroeconomics, NYU, 4 Washington Pl., NY, NY, 10003, webbr@nyu.edu.
R.W. and P.W.G. developed the model; I.L., S.C.L., R.B.R., and P.W.G. designed the experiment and constructed the neural dataset; R.W. performed analysis; R.W. and P.W.G. wrote the paper. This research was funded by NIA grant R01-AG033406 to I.L. and P.W.G, and the experiment was approved by an IRB. The authors wish to thank K. Menzel, A. Riedl, and M. Woodford for helpful comments.

To precise the ideas, let there be granted to the science of pleasure what is granted to the science of energy; to imagine an ideally perfect instrument, a psychophysical machine, continually registering the height of pleasure experienced by an individual, exactly according to the verdict of consciousness, or rather diverging therefrom according to a *law of errors*.

— Francis Y. Edgeworth, 1881

I. Introduction

Economics has traditionally relied on a mixture of theory and empirical measurement for the development of both positive and normative goals. Empirical techniques for measuring the value an individual places on goods and services have thus formed a natural subject for economists since the inception of the discipline. Since the mid-nineteenth century, for example, applied economics has relied on indirect empirical valuation methods derived from market and auction-pricing data (Fisher, 1892, 1927; Frisch, 1932; Hicks, 1942; Harberger, 1971).[1] Midway through the twentieth century, the *revealed preference* method (Samuelson, 1938; Houthakker, 1950), and its stochastic analog the *Random Utility Model* (RUM) (Marschak, 1960; Block and Marschak, 1960; Becker et al., 1963; McFadden, 1974, 2005), derived the conditions under which choices are consistent with utility maximization. This clarified what aspects of an economic model can and can not be identified by choice data alone, and laid a behavioural foundation for both positive and normative analysis based on indirect valuation methods. And, of course, more controversial techniques including stated preference and surveys have been explored and occasionally employed by economists (Boulier and Goldfarb, 1998; Carson and Hanemann, 2005).

Prior to the twentieth century, it was widely held that economics would be anchored to a hedonic object which guides choice behaviour. Some economists, such as Edgeworth, even foresaw that direct measurement of this hedonic object would be possible - a methodological approach we refer to as *measured value* methods. However, the discipline established its powerful indirect valuation methods while concurrently developing the "as-if" methodological viewpoint that an economic theory should be judged on its predictive power and not on the realism of its mechanism, *per se* (Friedman, 1953).[2] Some modern authors have interpreted this methodological stance as explicitly prohibiting a direct neural measurement of value in all economic discourse (Gul and Pesendorfer, 2008). We do not share this recent interpretation. An economic model indeed rests on its predictive performance, but in this paper we extend the positive

[1] For reviews of these, and more state-of-the-art methods, see Slesnick (1998); Colander (2007).

[2] The view that mechanism is not important to an economic model is a highly controversial one (Samuelson, 1963; Simon, 1979). Putting aside these arguments, here we proceed from Friedman's observation that a critical aspect of an economic model is its empirical performance.

aspect of economic methodology to a neurobiological dataset. We relate direct cardinal measurements of value-related signals in the human brain to choice prediction, and demonstrate how to judge the predictive power of a given neural dataset. We also note the potential normative implications of our measurement approach.

While it is the domain of future theoretical and empirical work to establish the usefulness of measured value methods in economic practice, this can be accomplished only if econometric techniques for relating these measurements to choice are first developed. We suspect, however, that these measured value methods may offer some unique advantages that will supplement existing methods. For example, such a method could supplement the stated-preference measurements economists are forced to rely on when incentive compatible mechanisms are unavailable, such as during contingent valuation for environmental goods (Carson and Hanemann, 2005). More broadly, the observation that choice behaviour is influenced by context and framing limits a simple indirect inference of value from choice (Caplin and Dean, 2008). A direct measurement of value offers a more structural account of how context affects valuation, as well as an ability to measure value in the absence of a choice set. Since a growing group of economists, psychologists and neurobiologists have already begun making (what are widely believed to be) direct measurements of the hedonic object of maximization behaviour – a literature which we briefly survey below – our goal is to lay out an econometric framework for such study.

This paper outlines a general formulation for relating neural value measurements to choice prediction which we call the Neural Random Utility Model (NRUM). The NRUM extends familiar aspects of the RUM framework to neural measurements, including maximization of stochastic decision variables, and we emphasize the relation between the neural and latent variable formulations.[3] A concrete example of subjects choosing over consumer items is developed in detail, demonstrating how these measurements can be made using existing brain scanning technology. We demonstrate that neural activity in a brain region called the medial Prefrontal Cortex (mPFC) cardinally encodes valuations and predicts choice behaviour. However, a significant degree of measurement error exists with current technology, adversely affecting both model inference and prediction rates. Conveniently, our application of the random utility model allows partial correction for the presence of error in neural measurements.

A means of benchmarking the predictive power of the measurements, with regard to choice, is also applied. In our dataset, neural activity by itself yields choice prediction results on par with a standard latent variable formulation. Combining neural data and standard observables improves predictive performance. To our knowledge, this is the first study establishing that a neural

[3]The neural basis of the random utility model has previously been conjectured (Glimcher, 2011; Fehr and Rangel, 2011). This is a natural hypothesis given that psychophysics provided the early inspiration for RUMs (McFadden, 2001), particularly the work of Fechner (1860), Thurstone (1927), and Luce (1959). In the economics literature, Hausman and Wise (1978) were the first to conjecture neural activity as the source of intra-individual stochasticity in a RUM.

measure of value can add predictive power to the toolset an economist would normally use in a similar choice problem.

I.A. Efforts to Measure the Neurobiological Object of Maximization Behavior

For an economist, the goal driving direct measurement of a "hedonic object" in the brain (known as "subjective value" in neuroscience) is to learn something about choice behaviour, or perhaps about welfare. But for neurobiologists working during the past decade that goal has largely been reversed; their ambition has been to use traditional methods for the measurement of value – or the representation of utility – to identify the anatomical and functional characteristics of subjective value-encoding signals in the human brain. In essence, they have worked to accomplish the engineering necessary for "Edgeworth's hedonometer".

To this end, nearly all economic methods for estimating value have been used in the neurobiological search for subjective value. Auction/willingness-to-pay-based methods (Plassmann et al., 2007; Chib et al., 2009), revealed preference-based methods (Platt and Glimcher, 1999; Dorris and Glimcher, 2004; Padoa-Schioppa and Assad, 2006; Glimcher et al., 2007; Levy et al., 2011; Smith et al., 2011), stated preference-based methods (e.g. Kringelbach et al., 2003; Hare et al., 2010) and even market price (Plassmann et al., 2008) have all been used as correlational probes in this search for the neural value signal. Strikingly, all of these methods have been found within the last decade to be correlated with brain activity measurements in two specific areas: the medial prefrontal cortex and the ventral striatum. Two recent meta-studies (Levy and Glimcher, 2012; Bartra et al., 2013) now unambiguously indicate that activity in these areas, particularly in the medial prefrontal cortex, is tightly correlated with every known economic method for estimating the values subjects place on choice objects - ranging from consumable goods, to money lotteries, to charitable donations, to durable goods, to social preferences, to political preferences. Activity in this brain area appears to complement traditional measurements of the values people place on the objects of choice.

In this paper we invert the neurobiological approach in an attempt to predict choice directly from neural measurements, taking it as a given that activity in the medial prefrontal cortex encodes the subjective values of choice objects under current consideration "or rather diverging therefrom according to a law of errors"[4]. We acknowledge that there are other brain areas that also carry information related to valuation and choice, and even acknowledge the possibility that some yet undiscovered area may carry a higher fidelity subjective value signal. Indeed, our goal is econometric: the tools we present are general enough that they can and should be used with regard to other brain areas and measurement methods whenever examining measured value data.

The model we propose to relate measured value data to choice behaviour, the NRUM, is one of many possible measured value methods. The first such method, *drift diffusion* (Fehr and Rangel, 2011), models the dynamic accu-

[4](Edgeworth, 1881)

mulation of a decision signal originally proposed in the psychology literature (Ratcliff, 1978). The choice probabilities depend on the slope of the accumulation, which in turn depends on the difference in value inputs. Neural evidence for the dynamics of this model has been uncovered both in psychophysical and economic choice tasks (Gold and Shadlen, 2007; Basten et al., 2010; Hare et al., 2011), as well as behavioural evidence for the role of decision dynamics and attention (Milosavljevic et al., 2010; Krajbich et al., 2010). A tight relationship exists between the two proposed methods: the NRUM is a reduced form of the drift diffusion model, with the particulars of the accumulation process impacting the distribution of random utility (Webb, 2013). While the more structural drift diffusion model will prove invaluable for exploring the decision process and restricting the NRUM, the reduced form brings a large econometric toolbox to bear for relating measured value to choice prediction.

In an example of this functionality, here we present an analysis of medial prefrontal cortex activity from a previously published laboratory experiment (Levy et al., 2011) that was divided into three stages. The first two stages of our experiment were performed inside a Magnetic Resonance Imaging (MRI) scanner, the third in a standard behavioural laboratory. In the first stage, subjects passively viewed the outcome of a series of small (consequential) lotteries over changes to their wealth. The purpose of this stage was to identify spatially discrete brain areas in each subject which encoded subjective values. In the second stage, subjects passively viewed 20 consumer items, one at a time, while intermittently performing an incentivized task. The purpose of this stage was to observe the activity in the areas identified in the first stage (to measure the subjective values) for the 20 consumer items. Immediately after the second stage, subjects were asked to perform a third stage outside of the scanner in which they made all possible binary choices over the set of items in an incentive compatible fashion. This procedure allowed an examination of the relationship between neural measurements of subjective value made in the scanner and the likelihood of choice outside the scanner. We now lay out the formal model for this relationship and describe our empirical strategy.

II. Neural Random Utility Model

Any measurement from which a value estimate will be derived suffers from some degree of error, and it is the job of the econometrician to develop tools for handling that error. In the physical sciences, errors were traditionally thought of only as randomly distributed perturbations in measurement: properties of the measurement tool rather than properties of the system actually being measured. While there are significant measurement errors associated with contemporary brain scanners (as we will see), we also encounter this latter form of variability, familiar to economists, in the neurobiological instantiantion of subjective value. First, it has been demonstrated empirically that the instantaneous perception of the attributes of an item is stochastic even when all properties of the item and state of the chooser are held constant (Fechner, 1860; Stevens, 1961). This stochasticity in subjective perception has been shown to be an obligate feature at all levels of sensory (attribute) processing (see Glimcher 2011, for an overview; Beck et al. 2012, for relation to optimality), and this necessarily leads to stochasticity in subjective value.[5] Second, the activity of all brain cells shows significant variation even under conditions in which measurement error can be shown to be near zero (Tolhurst et al., 1983; Churchland et al., 2010, 2011). For these reasons, the subjective value of an item is an inherently stochastic quantity, and we model it as a random variable in the RUM framework.

A unique issue that arises in developing the econometric specification for a neural measurement, however, is the existence of a second source of stochasticity which also must be accounted for. To understand this issue, one must recall that subjective value signals in the medial prefrontal cortex must affect choice through the intermediary of the rest of the brain - a series of physical systems which also induce stochasticity into choice. Such a view suggests that the least restrictive econometric specification of measured value should be composed of a stochastic valuation and a subsequent error term which is, in essence, strictly welfare decreasing. This then, is the most general form of the specification we develop below and then apply to a concrete example. We stress that it may, in some cases, be possible to restrict the econometric problem to a model with only one (or a linear combination) of these sources of error, but we present here the more general case as a starting point.

II.A. Relation Between Subjective Value and Choice

To capture these features of the neural choice process, we adapt the standard model for stochastic choice in economics, random utility maximization, to a neural form that explicitly treats subjective value as an observable stochastic object. We present the model for a binary choice situation in which we observe repeated choices from the same subject over a set of items and note that the extension of the model beyond binary choice is straightforward (Webb et al., 2013). The subjective value of an item $i \in I$ on binary choice trial $t \in T$ is defined to be an observable random variable $v_{i,t}$. The vector of subjective values on a choice trial is $\mathbf{v}_t$, composed of the $v_{i,t}$ for the items in the choice set. We assume the random vector $\mathbf{v}_t$ is independent over trials, but not necessarily over items.

Although we have yet to formally specify a distribution for $\mathbf{v}_t$, let us define $\nu_{i,t}$ as the difference between $v_{i,t}$ and its mean $E[v_{i,t}]$ for each item.

$$\nu_{i,t} \equiv v_{i,t} - E[v_{i,t}]. \tag{1}$$

The existence of a mean $E[v_{i,t}]$ requires further clarification. One possible interpretation of $E[v_{i,t}]$ is a 'core' value, instantiated noiselessly in neural firing rates somewhere else in the brain, which is represented with error in the neural

[5]To take one example, variance in the subjective value of a sweet tasting liquid can arise from variability in the behaviourally measured sensory experience of sweetness (which arises in turn from the stochasticity of neuronal signals) even when the objective sugar concentration is held constant.

substrate under observation.[6] This is not a view compatible with the biophysical properties of neural processes. Instead, $E[v_{i,t}]$ must be viewed as the limiting quantity of the sample mean of $v_{i,t}$ (i.e. the central tendency of $v_{i,t}$), and our definition of $\nu_{i,t}$ in an additive specification is for the purpose of exposition. We emphasize that $v_{i,t}$ is the only observable in equation (1) and we will provide a distributional assumption shortly.[7]

Once subjective values are instantiated in neural firing rates, they must be compared and a choice executed. This additional neural process, which we refer to as the choice mechanism, effectively transmits subjective values to the requisite circuitry for producing behaviour. A feature of the choice mechanism is that it is also stochastic, and the class of *bounded accumulation* or *drift diffusion* models from psychology and neuroscience is devoted to modelling it (Ratcliff, 1978; Gold and Shadlen, 2007). Therefore, we add to our NRUM an additive noise term $\eta_{i,t}$ which captures stochasticity in this argmax operation. This additive form has been demonstrated to be equivalent to the stochasticity found in bounded accumulation models (Webb, 2013), and adds an additional degree of stochasticity to choice behaviour. This yields the decision variable

$$u_{i,t} = v_{i,t} + \eta_{i,t}. \tag{2}$$

The subject chooses i vs. j on trial t if

$$\begin{aligned} u_{i,t} &> u_{j,t} \\ v_{i,t} + \eta_{i,t} &> v_{j,t} + \eta_{j,t}. \end{aligned}$$

If we denote this outcome $y_{ij,t} = 1$ (and 0 otherwise), this yields a probability of choosing i

$$\begin{aligned} P(y_{ij,t} = 1 \mid v_{i,t}, v_{j,t}) &= P\left(v_{i,t} - v_{j,t} > \eta_{j,t} - \eta_{i,t}\right) \\ &= P\left(\tilde{v}_{ij,t} > \tilde{\eta}_{ji,t}\right), \end{aligned} \tag{3}$$

where $\tilde{v}_{ij,t} \equiv v_{i,t} - v_{j,t}$ and the notation $\tilde{\cdot}_{ij}$ denotes the ijth item-pair difference throughout. If we were to measure subjective value in the neural substrate which computes the argmax operation, equation (3) would be the resulting conditional probability of choosing i given our measurement. Since the differences in measurements of subjective value determine these probabilities, the random utility framework is cardinal (Batley, 2008).

[6] We make here the assumption, widely held in neuroscience, that the distribution of $\nu_{i,t}$ reflects a fundamentally random process and not simply a high dimensional signal of zero stochasticity projected imperfectly into a low dimensional space. It should be noted that some scholars believe, that in some cases, 5% or more of the variance in $\nu_{i,t}$ may be formally non-stochastic under the appropriate analysis. For more on this issue see Rieke et al. (1997).

[7] Our definition of subjective value raises important issues about the stochastic specification of preferences that have remained unresolved (Loomes, 2005). Much of this debate involves the distribution of $v_{i,t}$ (or equivalently $\nu_{i,t}$), and whether its variance is constant and it is independent over i. We will assume the former, but not the latter. In addition, there is the question of whether the central tendency of subjective value is stable or if it can be manipulated through contextual effects; here we assume a stable mean over trials in this experiment. We touch further on these issues in section VI..

Imposing some distributional structure on our noise terms will bring us closer to a specification for use in our empirical setting. From this point forward, we assume that the difference in additive noise is independent over item-pair and trial, and distributed normally $\tilde{\eta}_t \sim \mathcal{N}(0, \sigma^2_{\tilde{\eta}} I)$.[8] This yields a probability of choosing i

$$P(y_{ij,t} = 1 \mid v_{i,t}, v_{j,t}) = \Phi\left(\frac{\tilde{v}_{ij,t}}{\sigma_{\tilde{\eta}}}\right), \tag{4}$$

where $\Phi()$ is the standard normal CDF.

Alternatively, this probability can be defined conditional on the mean of subjective value $E[v_{i,t}]$. To demonstrate this, let us now place the distributional assumption $\nu_t \sim \mathcal{N}(0, \Omega_\nu)$ with covariance matrix Ω_ν. Since our experiment uses a binary choice environment, the realizations of $\tilde{\nu}_{ij,t}$ for different item-pairs must occur on different trials t. Therefore the $\tilde{\nu}_{ij,t}$ are independent over ij due to independence over trials, even for different item-pairs that share an item.[9] Therefore $\tilde{\nu}_{ij,t} \sim \mathcal{N}(0, \sigma^2_{\tilde{\nu}})$ and this yields a probability of choosing i,

$$\begin{aligned} P(y_{ij,t} = 1 \mid E[v_{i,t}], E[v_{j,t}]) &= P\left(E[v_{i,t}] - E[v_{j,t}] > \nu_{j,t} - \nu_{i,t} + \eta_{j,t} - \eta_{i,t}\right) \\ &= P\left(E[\tilde{v}_{ij,t}] > \tilde{\nu}_{ji,t} + \tilde{\eta}_{ji,t}\right) \qquad (5) \\ &= \Phi\left(\frac{E[\tilde{v}_{ij,t}]}{\sigma_{\tilde{\nu}+\tilde{\eta}}}\right), \qquad (6) \end{aligned}$$

where $\sigma_{\tilde{\nu}+\tilde{\eta}}$ is the standard deviation of the sum of the two neural noise terms $\tilde{\nu}_t$ and $\tilde{\eta}_t$. Since $E[v_{i,t}]$ is not observable, equations (5) and (6) should be viewed as the limiting probabilities given a sample mean that approaches $E[v_{i,t}]$. The sample analog is

$$\begin{aligned} P(y_{ij,t} = 1 \mid \bar{v}_i, \bar{v}_j) &= P\left(\bar{\tilde{v}}_{ij} > \bar{\tilde{\nu}}_{ij} + \tilde{\nu}_{ji,t} + \tilde{\eta}_{ji,t}\right) \qquad (7) \\ &= \Phi\left(\frac{\bar{\tilde{v}}_{ij}}{\bar{\sigma}_{\tilde{\nu}+\tilde{\eta}}}\right), \qquad (8) \end{aligned}$$

where $\bar{\sigma}_{\tilde{\nu}+\tilde{\eta}} \to \sigma_{\tilde{\nu}+\tilde{\eta}}$ as $\bar{\tilde{\nu}}_{ij} \to 0$. This is the specification we will work from in our empirical setting.

We wish to highlight that the distinction between $\tilde{\nu}_t$ and $\tilde{\eta}_t$ may have critical normative implications. While $\tilde{\nu}_t$ reflects the variation in subjective values due to perception and their representation on a noisy neural substrate, $\tilde{\eta}_t$ reflects

[8] There is little known about the appropriate distribution of η_t at this level of aggregation, though some evidence suggests the variance may scale with the mean (Webb and Dorris, 2013; Webb, 2013). The assumption of independence over item-pair is only made for convenience, see footnote 9.

[9] The extension of the model beyond binary choice would have to account for a full covariance matrix for the vector composed of the $\tilde{\nu}_{ij,t}$ on each trial (similarly for the $\tilde{\eta}_{ij,t}$). In principle, a full covariance matrix should be identifiable for such a dataset (Hausman and Wise, 1978; Train, 2009) and the results that follow would have to be argued in terms of this full matrix. The assumption of normality for $\mathbf{v}_t$ is again made for convenience. To our knowledge no study has yet examined the distribution of the aggregate firing rates that make up subjective value.

error/stochasticity in the choice mechanism. If $\sigma_{\tilde{\eta}} = 0$, then all choice stochasticity is due to variation in subjective value and choice can be defined as optimal in the traditional economic sense because choosers then act to maximize their realized, albeit stochastic, subjective values. However, if $\sigma_{\tilde{\eta}} > 0$, then some choices can be classified as errors arising in the neural implementation of the argmax operation and the execution of the choice behaviour.[10] In most decision making contexts, people would therefore choose to minimize $\sigma_{\tilde{\eta}}$ if at all possible. Thus the relative sizes of $\sigma_{\tilde{\nu}}$ and $\sigma_{\tilde{\eta}}$ reflect the degree to which stochasticity in choice can be strictly viewed as welfare decreasing in any given measured value dataset.[11]

We should note that in all likelihood, ν and η are the product of realizations at multiple points in the human nervous system. The critical point that we seek to capture in this reduced form, however, is that two kinds of stochasticity can in principle arise in the neural substrate for choice: randomness that makes preferences stochastic, and randomness that adds to choice stochasticity. While we are unable to fully differentiate between these two sources of variance in this specific study because we do not make independent measurements at multiple stages along the pathways that represent subjective value, clearly this distinction is of significant importance and will doubtless be the object of future study. It is with this goal in mind that the model was formulated in precisely this manner. This then forms the core of our NRUM.

II.B. Comparison with Standard Latent Variable Modelling

The NRUM decomposes the uncertainty present in the standard RUM into biophysically distinct sources and yields the observable variable v on which to base choice prediction. This allows us to investigate, as a benchmark for our measurement, the potential benefit of using neural data to predict choices compared to a dataset of only standard economic observables. In particular, we focus on specification error in the standard approach due to the modeller's inability to observe all the attributes (of alternatives and decision makers) that make up subjective value (Manski, 1977). Formally, on a given trial the econometrician only observes a partition, $X_{i,t}$, of the full vector of attributes, $Z_{i,t}$, which make up subjective value for item i.

In the standard formulation of the RUM, the partitioning of the dataset matters since the econometrician does not observe the subjective value (or rather, the utility) of item i. Instead, the latent variable $u_{i,t}$ must be indirectly specified. The components of subjective value that are observed, $X_{i,t}$, are related to this latent variable as a linear combination, $X_{i,t}\beta$, while the components of $u_{i,t}$ that are unobserved are bundled in to an error term $\varepsilon_{i,t}$.

[10] We note that this specification thus allows one to handle a range of welfare-related stochastic specifications which would not be possible with a single random term.

[11] Evidence from perceptual neuroscience (in which there is an objectively "correct" answer) identifies that most of the variance in choice stochasticity can be attributed to brain areas encoding stimulus value, suggesting less than 15% of choice stochasticity can be attributed to downstream neural circuitry which implements the choice (Michelson et al., 2013; Drugowitsch et al., 2013).

Given our NRUM, we can decompose $\varepsilon_{i,t}$ into three sources. For the sake of this argument, we follow the standard approach and assume that subjective value is related to the arguments Z or X through the linear function $V(X_{i,t};\beta) = X_{i,t}\beta + \nu_{i,t}$.[12] The difference between the full specification $V(Z_{i,t};\beta)$ and the partitioned specification $V(X_{i,t};\beta)$, which we will refer to as specification error, is denoted $\omega_{i,t}$. Together with the stochasticity in subjective value and the choice mechanism, this yields a decision variable in which $\varepsilon_{i,t} \equiv \nu_{i,t} + \omega_{i,t} + \eta_{i,t}$ bundles together the three sources of uncertainty in our NRUM.

$$\begin{aligned}
v_{i,t} &= V(Z_{i,t}, \beta) \\
v_{i,t} &= V(X_{i,t}, \beta) + \omega_{i,t} \\
v_{i,t} + \eta_{i,t} &= X_{i,t}\beta + \nu_{i,t} + \omega_{i,t} + \eta_{i,t} \\
u_{i,t} &= X_{i,t}\beta + \nu_{i,t} + \omega_{i,t} + \eta_{i,t}.
\end{aligned}$$

As before, we can derive choice probabilities after imposing normality assumptions,

$$\begin{aligned}
P(y_{ij,t} = 1 \mid X_{ij,t}) &= P\left(X_{i,t}\beta + \tilde{\omega}_{ij,t} > \tilde{\nu}_{ji,t} + \tilde{\eta}_{ji,t}\right) && (9) \\
&= P\left(X_{ij,t}\beta > \tilde{\varepsilon}_{ij,t}\right) \\
&= \Phi\left(\frac{X_{ij,t}\beta}{\sigma_{\tilde{\varepsilon}}}\right), && (10)
\end{aligned}$$

where the variable $\tilde{\varepsilon}_{ij,t}$ aggregates all of the differenced error terms and $\sigma^2_{\tilde{\varepsilon}} = \sigma^2_{\tilde{\omega}} + \sigma^2_{\tilde{\nu}+\tilde{\eta}}$.

An obvious implication is that the latent variable model with non-zero specification error (10) will have the worst predictive power relative to the two neural specifications (4) and (6) since $\sigma^2_{\tilde{\eta}} \leq \sigma^2_{\tilde{\nu}+\tilde{\eta}} < \sigma^2_{\tilde{\varepsilon}}$. The latent variable formulation introduces error into the specification due to an inability of the modeller to fully explain subjective value with observables in the dataset. Observing a neural measure of subjective value removes this source of error, provided we can obtain a suitable neural measurement.

III. The Example Experiment

Our laboratory experiment was divided into three stages.[13] The first two stages were performed inside an MRI scanner. In the first stage, subjects passively viewed the outcome of a series of small lotteries over changes to their wealth.

[12] In practice, this function must be non-linear because v is bounded above and below. Additionally, there is evidence that $V()$ takes the entire vector X as its argument, yielding subjective values which depend on the composition of the choice set (Louie et al., 2011, 2013; Webb et al., 2013). Both of these issues result in misspecification error if unaccounted for. While the first issue can be easily dealt with in a standard RUM, the second requires careful attention (Webb et al., 2013). Regardless, both of these issues are not encountered if observing v directly, further bolstering our argument.

[13] The empirical portion of the experiment, but none of the analyses reported here unless noted, have been previously published in the neuroscientific literature in Levy et al. (2011). This paper also contains a complete description of the experimental methods.

The purpose of this stage was to identify the areas of the brain which encoded the subject's subjective values, $v_{i,t}$. In the second stage, subjects passively viewed 20 consumer items while intermittently performing an incentivized task so as to maintain subject engagement. The purpose of this stage was to repeatedly measure the subjective values of these items. Immediately after the second stage, subjects performed a third stage outside of the scanner in which they made all possible binary choices over this set of items in an incentive compatible fashion. Before leaving the subject also received a $25 show-up fee in cash.

III.A. Localization of Subjective Value in Medial Prefrontal Cortex

The first stage of the experiment was designed to identify an area in the brain of each subject which encodes subjective value. For brain measurements, we employed functional MRI (fMRI) using standard techniques (as in Caplin et al., 2010; Levy et al., 2011). These techniques indirectly measure brain activity over a 2 second interval in each of about 250,000 $3mm \times 3mm \times 3mm$ cubes (voxels) tiling the human brain. The product of this process is thus a time-series, in 2 second increments, of activation levels in each voxel.

The measure of activation is derived from the paramagnetic properties of the hemoglobin molecule and is known as the Blood-Oxygenation Level Dependent (BOLD) signal. This measurement has been demonstrated to be strictly monotonic in the average of the neural activity within the voxel, and most studies indicate that BOLD approximates a linear transformation of neural activity (Logothetis et al., 1999, 2001; Kahn et al., 2011).

A serious constraint arises from the sheer number of time-series fMRI generates and the statistical challenges imposed by determining which voxels/timeseries to study (Vul et al., 2009). We first restricted our analysis to a region of the brain known to encode subjective value-like signals, the medial prefrontal cortex .[14] We then performed an initial experiment aimed at independently 'localizing' subjective value encoding voxels within the mPFC, with the intention of conducting the analysis of our main experiment upon a time-series derived by averaging over these localized voxels.

In this initial stage of the experiment each subject was endowed with $40. On ensuing trials a lottery with equal probability of gaining or losing $2 was presented visually to the subject in the scanner. The outcome of the lottery was then revealed to the subject and the result was added to or deducted from the subject's wealth. In total, 128 trials of this kind were presented.[15] For each mPFC voxel, the difference in average activity between winning and losing was calculated. For each subject, voxels which showed a statistically significant difference were identified as our region of interest for encoding subjective

[14]We restrict our analysis to the mPFC since this location has been related to subjective value in previous studies by our group (Kable and Glimcher, 2007, 2010; Levy et al., 2010) and others (see Levy and Glimcher, 2012, for a review). It can also be demonstrated that other areas of the brain, such as the Striatum, are simultaneously, or perhaps in combination, encoding subjective value. We wish to emphasize that we are simply attempting to derive a measure of subjective value in the mPFC, not claim the singularity of this region.

[15]This task is a non-choice version of the task previously developed in Caplin et al. (2010).

valuation.

III.B. *Recording the Subjective Value of Items*

Immediately following the first stage, subjects completed a second stage in the scanner intended to measure the subjective values of 20 consumer items. Subjects completed six 7-minute brain scans over the course of 45 minutes, each consisting of 40 trials, for a total of 240 trials. In each of these trials, subjects passively viewed an image of one of 20 different items, including four DVD movies, two books, four art posters, three music CDs, two pieces of stationery, and five monetary lotteries represented by pie charts. Each lottery offered a 50% chance of receiving a designated amount of money ($10, $15, $20, $25, $30) and a 50% chance of receiving $0. All items were presented 12 times in a random order to each subject. Subjects were instructed that when they saw an item they should think about how much it was worth to them in a dollar amount.

To keep subjects alert, on 20 randomly selected trials (one for each of the 20 items), subjects were asked whether they preferred the item they had just seen or a randomly selected amount of money (ranging from $1 to $10). Subjects were told that one of these question trials would be randomly realized at the end and they would receive their selection on that trial - the item or the money. These 20 question trials were excluded from all behavioural and neural analysis. During the scanning stage, subjects did not know they would subsequently be offered an opportunity to choose between these same items after the scanning process was complete.

III.C. *Choice Task*

Following the second scanning stage, subjects were asked to perform a choice task outside of the scanner. Subjects were presented with a complete series of binary choices between the 20 items previously presented in the scanner. Each possible binary comparison (190 choices) was presented twice (switching the left-right location on each repetition), in random order, for a total of 380 choices. The result of one of these choices was randomly selected for realization.

The choices of subjects were largely consistent, with $96 \pm 2\%$ of triplets transitive and subjects switching their selection in only $9 \pm 1\%$ of choice repetitions. Choices were also highly idiosyncratic across subjects such that the individual preferences of a given subject could not be predicted from preferences exhibited by other subjects (mean correlation of ranking between pairs of subjects, excluding lotteries: $r = 0.1 \pm 0.3$).[16]

[16] We also verified that the random amounts of money used in the question trials in the scanner did not bias subjects' choices outside of the scanner.

IV. Measuring the Subjective Value of Consumer Items

We now wish to establish the NRUM as an econometric toolset for measuring subjective value in our experimental dataset. In section IV.A., we restate the main result from Levy et al. (2011) demonstrating that the ordering of our neural measure of subjective value correlates with choice likelihood. In section IV.B. we apply our model of neural random utility maximization to the combined dataset of choices and neural activity. The role measurement error plays in the relationship between our behavioural and neural observables is examined in section IV.C. and tested in IV.D.

In our analysis, we treat the item-pair and the two choices made in each pair as the dimensions of our dataset, and pool item-pairs over subjects. Essentially we are treating different subjects viewing the same item-pairs as equivalent to the same subject viewing different item-pairs. While this method allows each subject's preferences, therefore subjective valuations, to be idiosyncratic, it does contain the implicit assumption that the relationship between subjective valuation, the BOLD measure, and the choice likelihood is the same across subjects. We attempt to relax this assumption in section IV.E. at the expense of a reduced sample size.

IV.A. Ordinal Measure: Rank-Ordered Distance

The first step in demonstrating that measured neural activity encodes subjective value is determining whether the ordering of the neural activity associated with each of our 20 items can predict pairwise choices between those items. To demonstrate this, Levy et al. (2011) averaged neural activity from the 11 measurements of each item and then ordered the items by this mean activity. Subjects were then predicted to choose the item with the highest ranking in this order. Across all choices, they found that in $59 \pm 1\%$ of trials subjects chose according to this ordering.

With further analysis, however, a distinct pattern was observed in the ordinal neural ranking which suggests a random-utility-like representation. Figure I segregates prediction accuracy according to the rank-distance in neural activity between two items. Items with an ordinal distance of 19 represent the items with the highest and lowest neural activity, while items with an ordinal distance of 1 are items that are adjacent in the ranking. In choice pairs consisting of the highest ranked item versus the lowest ranked item, the predictions were accurate $83 \pm 8\%$ of the time. Pairs that are adjacent in the neural ranking are at chance. These previously published observations make two points. This is the first evidence that the magnitude of neural activity measured by fMRI, as evidenced by the ordinal distance between two goods, can predict later choice outcomes *even though these measurements of subjective value were made in a non-choice setting.* Additionally, the observation that ordinal distance matters to prediction accuracy suggests a degree of stochasticity in the neural representation, or its measurement, that invites analysis by a random utility model.

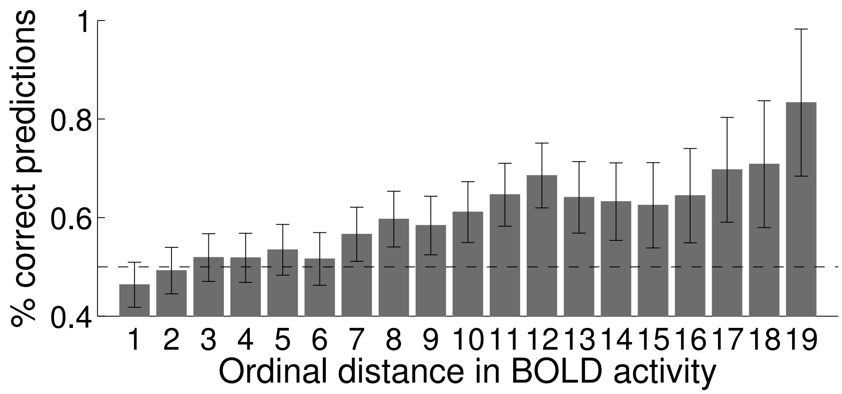

Figure I: Choice prediction results for items ordered by neural activity in mPFC, pooled over subjects. Items with higher neural activation (within subject) are predicted to be chosen. This is a reproduction of Figure 5 in Levy et al. (2011), with corrected 95% confidence intervals.

IV.B. Cardinal Measure: Testing a NRUM with Behavioural and Neural Measurements

The random utility model allows us to determine whether the difference in subjective values influences choice likelihood, therefore whether subjective value is a cardinal quantity. As we noted in section II.B., observing $v_{i,t}$ in synchrony with the choice of our subjects and using specification (4) would yield both the best predictive results and sharpest inference. However, the goal of this experiment was to determine whether subjective value measured in the *absence* of choice can be used to predict *later* choices, establishing the existence of subjective value in the absence of choice. By design we do not observe the realization of subjective value $v_{i,t}$ on the trial t in which the choice was made, therefore specification (4) is inappropriate for our analysis.

Instead, we measured $v_{i,m}$ in 11 scanning trials which preceded and were independent of the 2 choice trials of interest. We use the time index m to denote these measurement trials and assume a linear form for the relationship between our BOLD measurement $B_{i,m}$ and subjective value.

$$B_{i,m} = a + \gamma v_{i,m} + \mu_{i,m}.$$

The error term $\mu_{i,m} \sim N(0, \sigma_{\mu}^2)$ reflects the error present in measuring neural activity in an MRI scanner, therefore our neural measure of subjective value $B_{i,m}$ has two sources of variance: the fluctuation in subjective value on our measurement trials and measurement error. To arrive at a measure for predicting choice between items i and j on an independent trial t, we average over our

11 measurements and then take the difference.

$$\bar{B}_i = a + \gamma \bar{v}_i + \bar{\mu}_i \tag{11}$$

$$\tilde{\bar{B}}_{ij} = \gamma \tilde{\bar{v}}_{ij} + \tilde{\bar{\mu}}_{ij}. \tag{12}$$

Initially, we proceed under the assumption that there is no sampling and measurement error, $\tilde{\bar{B}}_{ij} = \gamma E[\tilde{v}_{ij,t}]$. While this assumption is clearly not valid, it does provide us some helpful intuition about the model which will prove useful when we relax it in section IV.C.. Specifically, assuming an error free measure of the mean of subjective value allows us to use specification (6). Substituting in yields a probability of choosing i,

$$P(y_{ij,t} = 1 \mid E[\tilde{v}_{ij,t}]) = \Phi\left(\frac{E[\tilde{v}_{ij,t}]}{\sigma_{\tilde{\nu}+\tilde{\eta}}}\right) \tag{6}$$

$$= \Phi\left(\frac{\gamma^{-1}}{\sigma_{\tilde{\nu}+\tilde{\eta}}}\tilde{\bar{B}}_{ij}\right). \tag{13}$$

Under this specification, the NRUM makes two predictions about the likelihood our subject will choose item i. First, as $\tilde{\bar{B}}_{ij}$ increases the subject should be more likely to choose item i. Second, recall that subjects made choices over each item-pair twice. If we segregate our item-pairs into those pairs in which the subject chose item i twice, once, or never at all as a function of $\tilde{\bar{B}}_{ij}$, the NRUM would predict $P(twice) > P(once) > P(never)$ for a positive difference in measured subjective value. We can visualize this prediction in Figure II, in which we simulated choices according to our model then fit the number of *twice*, *once*, and *never* observations using an ordered probit model.

Table I presents the estimates from bringing (13) to our dataset with the normalization $\sigma_{\tilde{\nu}+\tilde{\eta}} = 1$. We also included a specification with a constant term c predicted to be zero by the model. Since the estimate for γ^{-1} is positive, we can observe that the relationship between the difference in neural measurement ($\tilde{\bar{B}}_{ij}$) and the probability of choosing an item is indeed monotonic, as shown in Figure III. However the second prediction does not fare as well. The fit of the ordered probit model to the number of observed choices has a clear misordering since subjects are more likely to choose an item twice, than never, than once for positive $\tilde{\bar{B}}_{ij}$. Puzzlingly, we observe too few *once* choices when $\tilde{\bar{B}}_{ij}$ is small, too many when it is large, and far too many *never* choices when $\tilde{\bar{B}}_{ij}$ is large and positive (similarly for *twice* when it is large and negative).

IV.C. Accounting for Measurement Error

This apparent contradiction of our NRUM arises because we have assumed no error in our BOLD measurement and the construction of our neural measure $\tilde{\bar{B}}_{ij}$. We can identify at least three source of such error. First, since we are not measuring $v_{i,t}$ on the choice trial, the realizations of $v_{i,m}$ we do measure are not the ones causally related to choice on trial t. This component of our

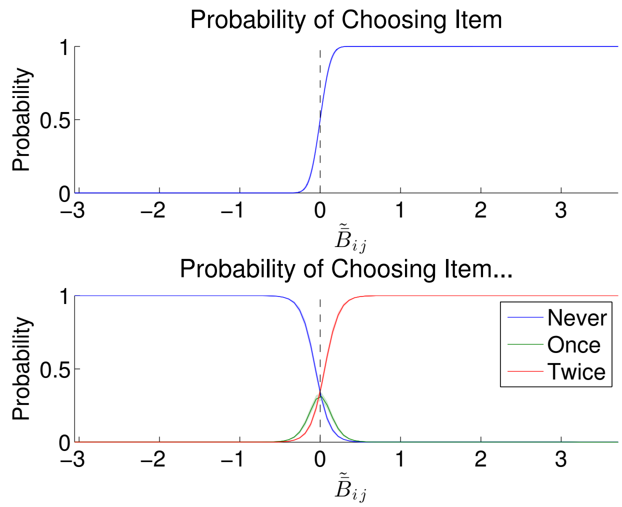

Figure II: Analysis of simulated choices with $\gamma^{-1} = 10$, $\sigma_{\tilde{\nu}+\tilde{\eta}} = 1$. Top Pane: The fit of the Probit model from (13). Bottom Pane: Fit of an ordered Probit model for the probability of observing the ith item in an ij pair chosen *twice*, *once*, and *never*.

Coefficient	No Constant	Constant
γ^{-1}	0.24	0.24
	(0.10)	(0.10)
c		-0.01
		(0.08)

Table I: Estimates of Probit Model (13) using difference in neural measurement between items. Clustered standard errors are in brackets.

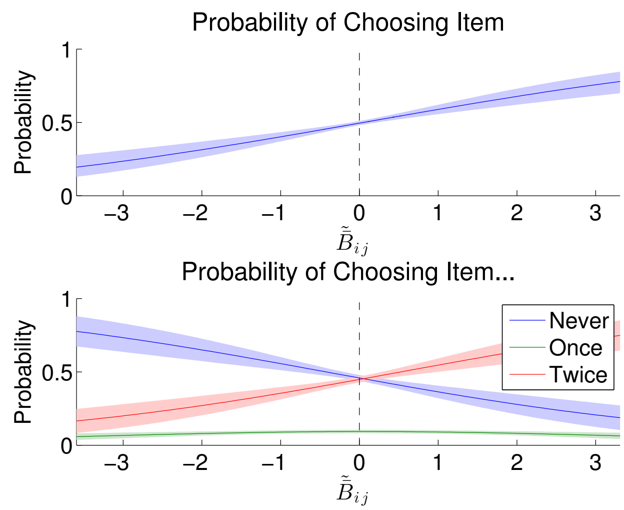

Figure III: Analysis of data assuming $\tilde{\bar{B}}_i = \gamma E[\tilde{v}_{ij,t}]$. Top Pane: The fit of the Probit model from (13). Bottom Pane: Fit of an ordered Probit model for the probability of observing the ith item in an ij pair chosen *twice*, *once*, and *never*.

measurement error is the sampling error present in $\bar{v}_i$ and is denoted $\bar{\nu}_i$. Second, we should also allow for error in our ex-ante procedure for identifying and constructing a single neural time-series from the 250,000 we measured. The degree to which the mean of our identified voxels capture the neural encoding of subjective value for consumer items depends on our ex-ante restriction to the mPFC and the accuracy with which our first procedure identifies these voxels. This source of variability is captured in $\mu_{i,m}$. A third source of noise doubtlessly results from the technical limitations imposed by measuring neural activation with an MRI scanner, which is also captured in $\mu_{i,m}$.

The effect of measurement error in non-linear models (such as Probit) is larger than in the linear model, but generally follows the same intuition: the data is over-dispersed along the dimension of the independent variable and the slope parameter is biased towards zero (Yatchew and Griliches, 1985). Formally, we can no longer work directly from specification (6) since $P(y_{ij,t} = i \mid \tilde{\bar{B}}_{ij})$ is no longer equivalent to $P(y_{ij,t} = i \mid E[\tilde{v}_{ij,t}])$. This means our estimate of γ^{-1} in section IV.B. is biased towards zero and the severity of this bias increases in the degree of measurement error. Since our hypothesis predicts a positive value for γ^{-1}, the inference performed on this biased estimate is still valid, though pursuing a less biased estimate will yield improved choice prediction.

Conveniently, bio-statistics provides some guidance on how to estimate non-linear longitudinal models with this form of measurement error (Carroll et al., 2006; Wang et al., 1998). Recalling equation (11), measurement error enters our specification as an item-specific i.i.d. error term.[17] If we proceed with a specification derived from substituting in our measured neural activation into the sample analog (7), the conditional probability of choosing i is

$$P(y_{ij,t} = i \mid \tilde{\tilde{B}}_{ij}) = P\left(\gamma^{-1}(\tilde{\tilde{B}}_{ij} - \tilde{\tilde{\mu}}_{ij}) > \tilde{\tilde{\nu}}_{ij} + \tilde{\nu}_{ji,t} + \tilde{\eta}_{ji,t}\right) \tag{14}$$

$$= P\left(\gamma^{-1}(\tilde{\tilde{B}}_{ij} - e_{ij}) > \tilde{\nu}_{ji,t} + \tilde{\eta}_{ji,t}\right), \tag{15}$$

where our sources of measurement error are grouped in the variable e_{ij}.

This probability has a form similar to a random-effects model, a standard method for dealing with subject-level idiosyncracy which we apply here at the level of the item-pair. The fact that subjects chose between each item-pair twice means that e_{ij} is constant over the two choice trials t. We can use this correlation pattern to achieve more efficient (and less biased) estimates of γ^{-1} and the variance of e_{ij} (relative to $\sigma^2_{\tilde{\nu}+\tilde{\eta}} = 1$) provided we specify and integrate out a distribution for e_{ij}. For instance, if we assume $e_{ij} \sim \mathcal{N}(0, \sigma_e^2)$, our specification takes the form of a random-effects Probit model with two important caveats.[18]

First, $\tilde{\tilde{B}}_{ij}$ and e_{ij} are not independent. This means that the random-effects Probit estimate of γ^{-1} will also be biased towards zero, though not as severely as a Probit with no random-effect, therefore we can only partially correct for the bias introduced by measurement error.[19] Simulation-based techniques for an unbiased estimate exist in the bio-statistics literature (Carroll et al., 2006, Chapter 5), but our simulations suggest that σ_μ would need to drop by roughly a factor of 4 for them to be applicable.

Second, the e_{ij} are not independent over choice pairs. Since the neural measurement takes place at the level of the individual item, when differencing the measurement for an item-pair there is correlation in the random effect e_{ij} between item-pairs that share an item. For instance, e_{12} and e_{13} are correlated because they share the measurement of item 1. This common issue with paired-data makes a random-effects estimate inefficient and will bias our standard errors towards zero if not controlled for. Multi-way clustering techniques have been developed to account for this pattern in the off-diagonal elements of the error covariance matrix (Cameron et al., 2011), but can not be used to improve the efficiency (or reduce the bias) of the estimate.

Here, we pursue a hybrid approach in which we estimate the random-effects

[17] This form of measurement error is referred to as "classical measurement error" since the error is additive and independent of the unobserved quantity (Carroll et al., 2006). It specifies that our neural measurement $\tilde{\tilde{B}}_{ij}$ has a larger variance than the unobserved quantity of interest, a natural assumption in the context of measuring neural activity with a noisy fMRI signal.

[18] A random-effect model is robust to the distributional assumption for the random-effect (here, measurement error) provided it is not highly asymmetric (Neuhaus et al., 2011).

[19] In addition, the estimate of the variance of the random-effect σ_e^2 will be biased positively (Wang et al., 1998). In our model we are assuming that the true random-effect variance is zero (there is no unobserved item-pair-specific term beyond measurement error); therefore. a positive variance estimate of σ_e^2 reflects this bias.

model clustered at the level of the item-pair (to capitalize on the common measurement error over choice trials within an item-pair, partially reducing the bias and achieving more efficient estimates), then correct our standard errors for inference using a multi-way clustering approach (to account for the non-independence of the differenced measurement errors).

The item-pair level likelihood is then given by

$$P\left(y_{ij,1}, y_{ij,2} | \tilde{\tilde{B}}_{ij}\right) = \int_{-\infty}^{\infty} \frac{\mathrm{e}^{-e_{ij}{}^2/2\sigma_w^2}}{\sqrt{2\pi}\sigma_w} \left[\prod_t F\left(y_{ij,t}, \tilde{\tilde{B}}_{ij}\right)\right] de_{ij}, \tag{16}$$

where $F(y,x) = \Phi\left(\frac{\gamma^{-1}(x-e_{ij})}{\sigma_{\tilde{\nu}+\tilde{\eta}}}\right)^y \left[1 - \Phi\left(\frac{\gamma^{-1}(x-e_{ij})}{\sigma_{\tilde{\nu}+\tilde{\eta}}}\right)\right]^{1-y}$.

Results from approximating this integral with Gaussian-Quadrature and estimating via Quasi-Maximum Likelihood are reported in Table II. We also included a specification with constant term c predicted to be zero by the model. After accounting for measurement error, the relationship between the difference in neural activity and the probability of choosing the item is now larger than in the specification which assumed no measurement error. We also note that over $\frac{\sigma_e^2}{\sigma_e^2+\sigma_{\tilde{\nu}+\tilde{\eta}}^2} = 95\%$ of the variance in the model can be attributed to the (positively biased) estimate of measurement error.

Figure IV shows the fitted probability of choosing item i as a function of the difference in neural activity assuming the random-effect $e_{ij} = 0$. This results in a significant improvement in the magnitude of the relationship between neural activity and choice probability when compared to our earlier analysis in section IV.B. which did not account for measurement error. To establish the cardinality of our neural measure, we also verified that the difference in neural activity yielded improved model fit compared to a simple ordinal ranking of the BOLD activity.

Coefficient	Probit		RE Probit	
	No Constant	Constant	No Constant	Constant
γ^{-1}	0.24	0.24	1.16	1.16
	(0.10)	(0.10)	(0.52)	(0.51)
c		-0.01		-0.06
		(0.08)		(0.37)
$\frac{\sigma_e^2}{\gamma^2}$			22.36	22.36
			(3.49)	(3.50)

Table II: Estimates of Random-Effects Probit Model (16). Clustered standard errors are in brackets.

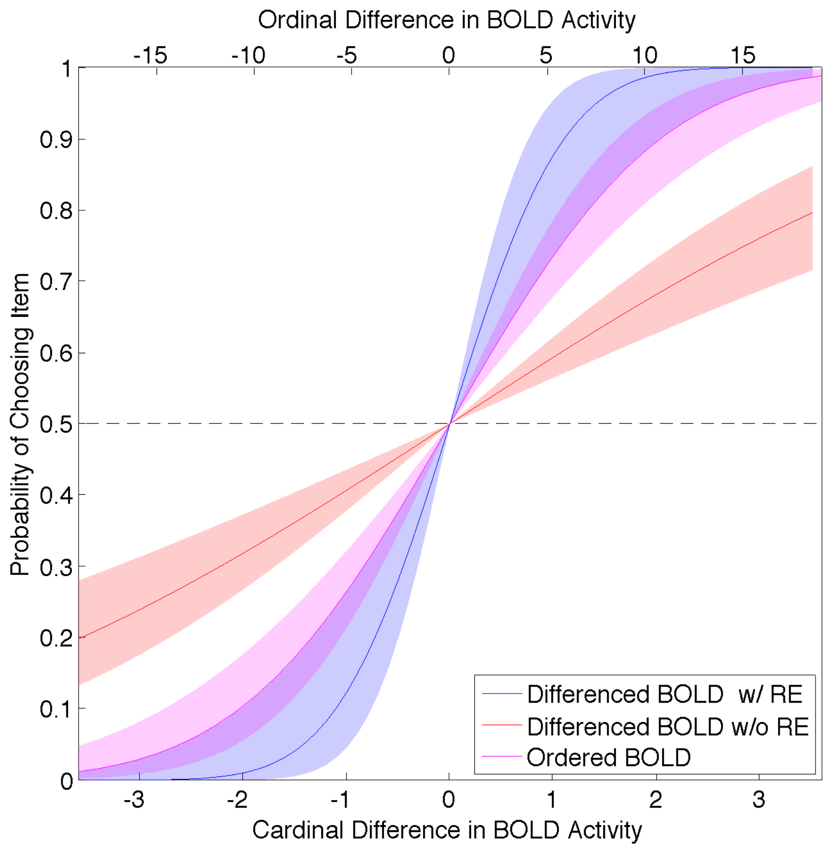

Figure IV: The probability of choosing an item depends on the difference in neural activity between items. The fitted probabilities are generated using the standard Probit estimate for γ^{-1}, the RE Probit estimate for γ^{-1} (assuming the random-effect is zero), and for a standard Probit estimate of choice on the ordinal difference in the BOLD ranking. 95% confidence interval for the fitted probabilities are depicted by the shaded areas.

IV.D. Implication of Measurement Error for Mis-ordered Choice Frequencies

To verify that measurement error is generating the results observed in section IV.B. and Figure III, we introduced measurement error into the simulated data reported in Figure II and repeated the analysis that did not take measurement error into account. The results are presented in Figure V. Introducing measurement error in to our simulation yields theoretical results which now match our empirical findings, predicting the observed relationship between $\tilde{\tilde{B}}_{ij}$ and choice probability. This occurs because measurement error has the effect of smearing out the observed *once* choices over the range of observed $\tilde{\tilde{B}}_{ij}$. A choice pair in which the distributions of subjective value are close together (small $E[\tilde{v}_{ij,t}]$), likely resulting in a *once* outcome, could yield a large $\tilde{\tilde{B}}_{ij}$ because of measurement error. Therefore the degree of measurement error has no effect on the number of *once* choices observed, only on where they appear on the $\tilde{\tilde{B}}_{ij}$ axis.

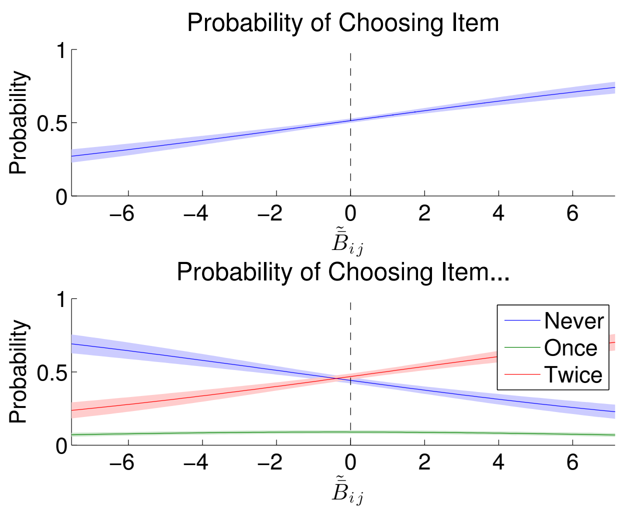

Figure V: Analysis of simulated dataset in which neural measurements are generated with error. The analysis are those run in section IV.B. which do not account for measurement error. Top Pane: The fit of the Probit model from (13). Bottom Pane: Fit of an ordered Probit model for the probability of observing the ith item in an ij pair chosen $Twice$, $Once$, and $Never$.

IV.E. Subject-Specific Analysis

Relaxing our assumption of pooling over subjects s would yield the relation

$$B_{s,i,m} = a + \gamma_s v_{i,m} + \mu_{s,i,m} \tag{17}$$

$$\tilde{\tilde{B}}_{s,ij} = \gamma_s \tilde{\tilde{v}}_{s,ij} + \tilde{\tilde{\mu}}_{s,ij}, \tag{18}$$

where γ_s is a subject specific relationship between our neural measurement and subjective value. We estimated γ_s^{-1} through a subject-$\tilde{\tilde{B}}_{ij}$ interaction term using specification (16) on the full sample, but this meant we only had 380 observations per subject to estimate this relationship.

Breaking up our sample into so few observations per subject reveals the inefficiency of the estimation method described above. Six of the subjects yield positive and significant estimates of γ_s^{-1}, while six are not significantly different from zero (Table III).[20] Commensurate with existing data and previous fMRI studies (Logothetis, 2003), some subjects yield a steeper mapping between the neural measurement we use here (BOLD) and neural activity than do others. The bulk of this difference is widely assumed to reflect a technical feature of the interaction between the scanner and the subject: the subject-specific coefficient describing the coupling of neural activity to the blood flow rate measured by fMRI. In any case, our subject-specific analysis makes clear that pooling data across subjects can impact the power of NRUMs as predictors of behaviour.

Coeff	Est.	Std. Err.	P-Val	Coeff	Est.	Std. Err.	P-Val
c_1	0.03	1.14	0.98	γ_1^{-1}	-1.17	1.07	0.27
c_2	-0.15	1.25	0.91	γ_2^{-1}	0.66	2.89	0.82
c_3	-0.07	1.27	0.95	γ_3^{-1}	-3.25	2.36	0.17
c_4	-0.34	1.17	0.77	γ_4^{-1}	10.14	2.90	0.00
c_5	0.08	1.22	0.95	γ_5^{-1}	1.39	0.57	0.02
c_6	-0.07	1.22	0.95	γ_6^{-1}	-3.23	2.50	0.20
c_7	-0.14	1.30	0.91	γ_7^{-1}	2.78	3.30	0.40
c_8	0.41	1.22	0.73	γ_8^{-1}	10.39	3.53	0.00
c_9	-0.18	1.18	0.88	γ_9^{-1}	4.98	2.38	0.04
c_{10}	0.69	1.24	0.58	γ_{10}^{-1}	5.01	1.39	0.00
c_{11}	0.07	1.23	0.95	γ_{11}^{-1}	2.61	3.18	0.41
c_{12}	-0.44	1.14	0.70	γ_{12}^{-1}	13.04	3.80	0.00
σ_e^2	20.49	3.46					

Table III: Subject specific RE Probit estimates of relationship between neural measure and choice likelihood.

V. Choice Prediction: Benchmarking Measured Preferences

The role of neural data in choice prediction has been the subject of much recent debate (Caplin and Schotter, eds, 2008). Bernheim (2009) noted that a positive neuroeconomic model must ultimately be equivalent to a positive behavioural

[20] Monte carlo simulations verified the loss in efficiency due to reducing observations. Simulated choice and neural data with $\gamma_s^{-1} = 10$ and measurement error from section IV.D. leads to ~5% of the γ_s^{-1} estimates less than, but not significantly different from, zero (in a total of 1000 simulations).

model which uses only choice data. Our view has been expressed in Caplin and Dean (2008), and Glimcher (2011), and Webb (2011). Briefly, we acknowledge that the two approaches are equivalent, but only when the behavioural model is correctly and completely specified. When the behavioural model is incorrect, or is missing explanatory variables, there is no reason the neuroeconomic model cannot outperform the behavioural model in its predictive power *with regard to choice*. As we demonstrated in section II.B., the fact that our neural measure of subjective value does not contain the specification error present in the latent variable approach implies that a neural measure of value can outperform the latent variable specification, all else equal.

However, the presence of measurement error, as noted in section IV.C., complicates matters because our estimated random-effects specification is no longer equivalent to the specification (6) which assumes we observe the mean of subjective value. Therefore the ordering of specifications (from lowest to highest variance) in which we observe subjective value on a choice trial (4), observe the mean (6), and observe only standard economic observables (10), may not hold in this dataset. If measurement error is high enough, and latent variable specification error low enough, the degree of bias in our estimate will lead (10) to out-predict our sample analog (8). Thus the technological frontier plays a critical role in the practicality of a neuroeconomic model, leaving as an empirical exercise the question of whether any given set of neural measurements outperform the latent variable approach with respect to choice prediction. To clarify this point: one goal of this paper is to state the conditions under which a given technological frontier permits a neuroeconomic model to outperform a purely behavioural model, a point not taken up by Bernheim (2009) in his assessment of positive neuroeconomics.

Table IV presents the choice prediction rates resulting from our neurally measured subjective value. We arrived at these rates by using the predicted probabilities for each item-pair (assuming the random-effect is zero) to simulate two trials of a binomial draw. If the predicted frequency with which an item was chosen from a pair (*never*, *once*, or *twice*) matched the data, the prediction was considered a success. In such a simulation, the prediction rates arrived at by chance depend on the distribution of *never*, *once*, or *twice* in the dataset. For our entire sample, the frequency of *never* is 46%, *once* is 9%, and *twice* is 45%. If each individual choice were predicted at chance, we would predict *never* on $\frac{1}{4}$ of trials, *once* on $\frac{1}{2}$, and *twice* on $\frac{1}{4}$, and we would be correct $\frac{1}{4} \times 46 + \frac{1}{2} \times 9 + \frac{1}{4} \times 45 \approx 27\%$ of the time.

The first two columns employ the biased random-effect estimates (without a constant term) from the population and subject, respectively. While prediction rates using the population level estimates are slightly above chance, when we allow the estimates to vary at the subject level we achieve a 43% prediction rate.[21] To account for the bias still present in our random effects estimates due to measurement error, we also examined prediction rates under the assumption

[21] These are comparable to a similar study conducted by Smith et al. (2011) which explores the scope of predicting choices within subjects, across subjects, and across items using fMRI data.

that $\gamma^{-1} = 10$ for all subjects (the value found in our calibrations in the previous sections). Imposing this restriction improves our prediction rate for the population to 46%, highlighting again the impact of measurement error on our estimation. Using half the sample as a training sample to estimate and predict the second half yielded no change in prediction rates.

	BOLD			Amazon*		Price		A+P*	P+B	A+P+B*
	RE		$\gamma^{-1} = 10$	RE		RE		RE	RE	RE
	Pop	Sub	Pop	Pop	Sub	Pop	Sub	Sub	Sub	Sub
chance	27	27	27	27	27	27	27	27	27	27
pop	31	43	46	47	46	53	52	52	57	60
sub_1	29	36	36	55	60	60	63	62	62	62
sub_2	30	28	47	38	26	54	55	27	55	47
sub_3	24	49	29	33	35	46	40	44	51	45
sub_4	32	53	53	46	45	62	66	56	71	65
sub_5	45	48	59	65	72	54	54	79	57	77
sub_6	26	40	35	65	70	59	61	71	64	75
sub_7	28	33	45	44	39	41	29	65	33	50
sub_8	30	49	49	47	45	50	47	45	56	70
sub_9	35	50	53	41	35	59	62	47	64	59
sub_{10}	33	47	51	48	48	45	42	48	52	54
sub_{11}	30	33	41	43	33	57	59	46	60	46
sub_{12}	32	51	49	42	37	48	47	38	56	62

Table IV: Choice prediction rates (%) resulting from 1000 simulated samples generated by our estimates. Prediction rates are calculated for both (Pop)ulation and (Sub)ject-based estimates, and prediction rates are shown for the (pop)ulation as a whole and for each (sub)ject. Prediction rates are also calculated using both (A)mazon and (P)rice observables, (P)rice and the (B)OLD measure, and all three predictors. *Amazon ratings were not available for the five lotteries, so choice pairs with the lotteries were excluded for these sets of predictions.

In an effort to compare with the standard latent variable approach we also attempted to find more traditional value measures on which to base predictions. Such an approach often seeks out attributes on which to construct the latent variable, however the consumer items in our experiment had no obvious attribute dimensions on which they all could be categorized (e.g. length, time, fidelity, square footage, etc...). Instead, we found two aggregate level valuation measures: the price of the item (a market-based method) and its 'Amazon star' rating (a stated-preference method).[22] Both of these measurements have the drawback of being population level variables which represent (to some degree)

[22]We used our purchase price of each item at the time of the experiment. The 'Amazon star' rating is the aggregation of user ratings that can be found on the item's description on amazon.com.

the aggregation of preference across all consumers, limiting their ability to predict individual choices. However, both of them were significant predictors. The Amazon rating varied positively with the choices of our subjects, suggesting some homogeneity in the preferences of New York University undergrads, while prices varied negatively with choice.[23]

For better or worse, our expensive neural measurements perform almost as well as the publicly available population level variables (Table IV; e.g. 43% vs. 46% vs. 52% prediction rates in the population for subject-level estimates; or 46% vs. 47% vs. 53% using the calibrated neural estimate at the population level). There is no escaping the fact that this measured value dataset developed with fMRI methods available in 2011 and analyzed with our general-purpose econometric tool, results in a neuroeconomic model which just matches the performance of a coarse behavioural model. This observation stresses the importance of the technological frontier to the practical use of positive neuroeconomics and measured value methods. Movements of the frontier in measurement technology, including advances in scanner resolution and signal analysis in the two years since the Levy et al. (2011) dataset was collected (Gross et al., 2012; Smith et al., 2012; De Martino et al., 2013), may play an important role in the forward development of measured value methods.

Of course, if the goal is to increase choice prediction rates in this setting, the best prediction rates (60%) arise from combining both standard and neural observables. This emphasizes that there is additional information in the neural measurements, likely due to the individual nature of the measured value data. To our knowledge, this is the first study establishing that a neural value measure can add predictive power in a choice prediction problem.

VI. Discussion of the Model and Measurement

VI.A. Implications of Measurement Error for Analysis and Experiment Design

Our analysis confirms a large degree of measurement error present in the six-year-old "3T" MRI technology used in this study. While scanner technology is currently being developed which may reduce sources of this error, such as the introduction of high resolution "7T" technology (De Martino et al., 2013), the effectiveness of these measurement advances with regard to subjective value is still an open question. As other less costly - but also less precise - neural measurement methods (e.g. Electroencephalography "EEG") are adopted, it should be clear from our analysis that measurement error is of concern when relating neural measurements to choice prediction. While the frontier in measurement technology is still in flux, we have introduced a set of econometric tools for relating all forms of neural measurements to choice prediction.

[23]This is somewhat surprising since one might expect subjects to be choosing high priced goods (which they receive at no monetary cost in the experiment), but likely reflects the popularity of the CDs in our choice set, a relatively inexpensive item.

Our econometric technique has implications for the design of future experiments which attempt to relate neural measurements to choice behaviour. While increasing the number of measurement repetitions per item will decrease the degree of measurement error directly, typical neural measurement experiments (such as brain scanning) are usually time constrained, limiting potential precision in this dimension. The random utility framework, together with a random effects specification to account for measurement error, offers another option. Since the trial dimension (per choice pair) is used to estimate the variance of measurement error, increasing this dimension (which usually takes place outside of the scanner and is less time constrained) will improve the efficiency of all estimates, improving choice prediction. We should also note that all of the econometric techniques used in this analysis are available in standard software packages.

VI.B. Context and Stability

In its general form, the NRUM places no restriction on the mapping from the sensory/economic environment to subjective value or about the stability of the resulting subjective value distribution. However our empirical specification does place such a restriction because we assume $E[v_{i,t}]$ does not vary over trials. The fact that we are able to predict choices while placing such a restriction is evidence that the mapping of the sensory environment to the distribution of subjective values, and the mapping of subjective value to choice, is in fact stable *in our choice experiment*.

There is mounting evidence, however, that these mappings can be subject to alteration through context effects mediated by choice set size, composition, and wealth. Such a contextual mapping is essentially required for the entire range of possible choice sets to fit within the finite activation rates of neurons (Louie et al., 2011). In the economics literature, the implications/constraints that perceptual and contextual manipulations place on choice data are beginning to be explored (Koszegi and Rabin, 2006; Caplin and Martin, 2012; Bordalo et al., 2013), notably rational inattention (Sims, 2003; Woodford, 2012). In the neuroscience literature, the physiological process by which subjective values are influenced by context has also begun to be described (Louie et al., 2011, 2013). In particular, this research identifies a neural computation in which subjective values are normalized by the size and composition of the choice set.

Since our NRUM decomposes the choice process into bio-physically defined systems, it provides a natural framework for understanding how the objective choice context is mapped to subjective value and then on to choice. For instance, in this version of the NRUM we have assumed that the decision variable for an item depends linearly on the subjective value of that item. This relationship could be generalized through a normalization function $Z(\mathbf{v}) : \Re_+^{\|I\|} \to \Re_+^{\|I\|}$ which maps the subjective value vector of all items, $\mathbf{v}$, to the decision variables, $\mathbf{u} = Z(\mathbf{v}) + \eta$. Depending on how one specifies the function $Z()$, this formulation could yield choice behaviour which depends on the size and composition of the choice set (Webb et al., 2013). This demonstrates how the framework can be extended and/or restricted to incorporate a deeper neurobiological understanding of how the choice environment influences behaviour.

VI.C. Distribution of Subjective Value and Normative Implications

In our version of an NRUM, we have attempted to formulate subjective value as generally as possible so that it might encompass the two predominant views about stochastic choice in the economic literature. One interpretation of the random vector of subjective values $\mathbf{v}_t$, in particular its mean, follows from the view that choices can be described by a single "core" preference relation (utility function) that is perceived or represented with independent error for each item (Hey and Orme, 1994). The random vector $\mathbf{v}_t$ would constitute the core valuation (its mean) plus this error. When these perturbed core valuations are compared, a choice in contradiction to the core valuation is possible, with the number of such "errors" governed by the magnitude of the difference between the core valuations. This cardinal model has been termed 'Fechnerian' in the economic literature due to its roots in psychophysics.

However the general formulation of a RUM places no such restriction on the distribution of utilities (Becker et al., 1963). A second class of models, *random preference* models, posits a set of preference relations (utility functions) for which each choice is represented by a single utility function drawn from this set (Loomes and Sugden, 1995, 1998). This approach allows for preferences to vary from trial to trial as new preference relations are drawn, but in a manner which is internally consistent with the axioms underlying the utility representation: each item in a choice is processed together and in synchrony according to a particular realized preference function. This has important implications for both model-testing and normative analysis. For instance, an expected utility maximizer with stochastic preferences of this type would never violate first-order stochastic dominance (Loomes, 2005).

The NRUM is general enough to allow for both of these views; the difference between them arising in the covariance matrix of $\mathbf{v}_t$. If $v_{i,m}$ and $v_{j,m}$ are independent, we have the Fechnerian model.[24] A particular non-zero covariance structure for $\mathbf{v}_t$ would yield a random preference model. In this study, we allow for random preferences since we put no restriction on the covariances of $\mathbf{v}_t$. The subjective values of items were measured independently, in isolation, and on different trials; therefore we can safely assume that $v_{i,m}$ and $v_{j,n}$ are independent over measurement trials m and n in our dataset.[25] We note, however, that a NRUM renders the covariance matrix of $\mathbf{v}_t$ empirically observable and would allow us to differentiate between these views, at least for subjective value, with an appropriate dataset.

[24] Though as mentioned in II.A., we disagree with interpreting the mean as 'core' preference. The view that a mean is encoded noiselessly in brain tissue is not a view compatible with the biophysical properties of neural processes. We view the mean simply as the limiting quantity of the sample mean of $v_{i,t}$.

[25] Making such an independence assumption in an alternative dataset in which the subjective values of both items were measured simultaneously (i.e. $m = n$) would preclude random preferences.

Even after allowing for a random preference specification for subjective value, however, our model still has some remnants of the psychophysical approach due to η_t, the stochastic element in the biophysical choice mechanism. In contrast with the assumption of independent ν_t, the noise vector η_t can result in choices which contradict the ranking of subjective values, $\mathbf{v}_t$. In such a hybrid model, the variance of η_t would determine how closely the reduced-form model resembles either class, a subject of current debate (Loomes, 2005; Loomes et al., 2012).

Regardless of the source, we observe choice behaviour that has features of the psychophysical specification: a larger difference in subjective value makes an item more likely to be chosen. Our own conviction, which stems from an amalgamation of the economic and neurobiological literature, is that a model which incorporates both classes of stochasticity will most closely approximate the structure of human choice behaviour. We note that anchoring our model to this conviction effectively posits a distinction between the fraction of choice stochasticity that can be attributed to stochasticity in preference and the fraction that can be attributed to errors induced by the choice mechanism. This distinction has clear welfare implications that would necessarily be of interest as more is learned about these sources of stochasticity in choice behaviour (Bernheim, 2009).

VII. Conclusion

We have presented a method for measuring the value an individual places on consumer goods without direct recourse to choice behaviour. The class of such methods, which we refer to as measured value methods, are established through the relation of direct value measurements to choice prediction. We have proposed an econometric framework, the Neural Random Utility Model, for doing just this. The NRUM is an extension of the standard econometric framework used in applied economics to neural measurements of value. This framework attempts to relate neural measurements to choice behaviour in as general a form as possible, and offers the advantage of an established econometric toolkit for analyzing this relationship. A concrete example of subjects choosing over consumer items was developed in detail, demonstrating how neural measurements can be made using existing brain scanning technology and how they are related to choice behaviour. We found that the magnitude of the difference in neural activity makes the higher item more likely to be chosen, implying our measurement is cardinal, and neural activity measured in isolation can predict choices.

Our modelling framework proposes that, in principle, choice predictions based on a neural measure will have lower variance than those based on a latent variable formulation. To check this, a tool for benchmarking the predictive power of the measurements, with regard to choice, was also developed. Our current analysis suggests that the NRUM, and the measurements possible at today's technological frontier, are competitive with standard latent variable specifications, however error in our neural measurement techniques limits the effectiveness of choice prediction. Econometric techniques available to the RUM framework mitigate some of the impact of measurement error, improving choice prediction results. Combining neural measurements and standard observables further improves choice prediction. To our knowledge, this is the first study establishing that a neural value measure can add predictive power to the toolset an economist would normally use in a similar choice problem. With that said, we acknowledge that this improvement comes at a high implementation cost for brain-scanning technology which currently limits the prevalence and usefulness of neural measurements.

There are many techniques and technologies available for measuring value in the brain, but the problem of how these measurements are related to choice prediction in an empirical framework is shared by all of them. We have proposed a method which can be used for all types of neural measurements, but also retains a relationship with more structural models of decision-making, including the predominant model of the dynamic processes underlying choice (e.g. drift diffusion; Fehr and Rangel, 2011). The NRUM is a reduced form of these dynamic models (Webb, 2013), and a deeper understanding of this process will further restrict the NRUM and offer advances in modelling choice behaviour.

Our approach to measured value thus offers four valuable features to the economic literature: It establishes the positive performance of one measured value method and defines clearly the technological frontier that will be required for future methods. It establishes that these measurements carry cardinal information about the relative values of alternatives. It offers a general framework for combined economic-neurobiological modelling from which both richer, more restrictive specifications can be developed. And finally, it lays out the basic welfare structure inherent in a neurobiological decision model.

VIII. *

References

Bartra, Oscar, Joseph T McGuire, and Joseph W Kable, "The valuation system: a coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value," *NeuroImage*, August 2013, *76*, 412–427.

Basten, U, G Biele, H R Heekeren, and C J Fiebach, "How the brain integrates costs and benefits during decision making," *Proceedings of the National Academy of Sciences*, 2010, *107* (50), 21767–21772.

Batley, R, "On Ordinal Utility, Cardinal Utility and Random Utility," *Theory and Decision*, 2008, *64* (1), 37–63.

Beck, Jeffrey M, Wei Ji Ma, Xaq Pitkow, Peter E Latham, and Alexandre Pouget, "Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability," *Neuron*, April 2012, *74* (1), 30–39.

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

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

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

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

Boulier, Bryan L and Robert S Goldfarb, "On the use and nonuse of surveys in economics," *Journal of Economic Methodology*, June 1998, *5* (1), 1–21.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Robust Inference With Multiway Clustering," *Journal of Business and Economic Statistics*, April 2011, *29* (2), 238–249.

Caplin, Andrew and Andrew Schotter, eds, *The Foundations of Positive and Normative Economics: A Handbook*, Oxford University Press, 2008.

_ **and Daniel Martin**, "Framing Effects and Optimization," *NYU Working Paper*, 2012, pp. 1–32.

_ **and Mark Dean**, "Economic Insights from "Neuroeconomic" Data," *The American Economic Review*, 2008, *98* (2), 169–174.

_ **,** _ **, Paul W Glimcher, and Robb B Rutledge**, "Measuring beliefs and rewards: a neuroeconomic approach," *Quarterly Journal of Economics*, 2010, *125* (3), 923–960.

Carroll, Raymond J, David Ruppert, Leonard A. Stefanski, and Ciprian M. Crainiceanu, *Measurement Error in Nonlinear Models: A Modern Perspective*, 2nd ed., CRC Press, June 2006.

Carson, Richard T and W Michael Hanemann, "Contingent Valuation," in K G Maler and J R Vincent, eds., *Handbook of Environmental Economics*, North Holland, 2005.

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

Churchland, Anne K, R Kiani, R Chaudhuri, Xiao-Jing Wang, Alexandre Pouget, and Michael N Shadlen, "Variance as a Signature of Neural Computations during Decision Making," *Neuron*, February 2011, *69* (4), 818–831.

Churchland, Mark M, Byron M Yu, John P Cunningham, Leo P Sugrue, Marlene R Cohen, Greg S Corrado, William T Newsome, Andrew M Clark, Paymon Hosseini, Benjamin B Scott, David C Bradley, Matthew A Smith, Adam Kohn, J Anthony Movshon, Katherine M Armstrong, Tirin Moore, Steve W Chang, Lawrence H Snyder, Stephen G Lisberger, Nicholas J Priebe, Ian M Finn, David Ferster, Stephen I Ryu, Gopal Santhanam, Maneesh Sahani, and Krishna V Shenoy, "Stimulus onset quenches neural variability: a widespread cortical phenomenon," *Nature Neuroscience*, March 2010, *13* (3), 369–378.

Colander, D, "Retrospectives: Edgeworth's hedonimeter and the quest to measure utility," *The Journal of Economic Perspectives*, 2007, *21* (2), 215–225.

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

Drugowitsch, Jan, Valentin Wyart, and Etienne Koechlin, "Inference rather than selection noise explains behavioral variability in perceptual decision-making," in "Computational and Systems Neuroscience (Cosyne) Meetings 2013" Salt Lake City 2013.

Edgeworth, Francis Y., *Mathematical Psychics; An Essay on the Application of Mathematics to the Moral Sciences* Reprint of Economic Classics (1961), A.M. Kelley, New York, 1881.

Fechner, G, *Elements of Psychophysics*, Vol. 1, Holt, Rhinehart and Winston, New York, 1860.

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

Fisher, Irving, *Mathematical Investigations in the Theory of Value and Prices*, Yale University Press, 1892.

_ , *A Statistical Method for Measuring "Marginal Utility" and Testing the Justice of a Progressive Income Tax* Economic essays contributed in honor of John Bates Clark., MacMillan, Inc., 1927.

Friedman, M, *Essays in Positive Economics*, The University of Chicago Press, 1953.

Frisch, Ragnar, *New methods of measuring marginal utility*, Mohr, 1932.

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

— , **Joseph W Kable, and Kenway Louie**, "Neuroeconomic Studies of Impulsivity: Now or Just as Soon as Possible?," *The American Economic Review*, January 2007, *97* (2), 142–147.

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

Gross, J, E Woelbert, J Zimmermann, S Okamoto-Barth, A Riedl, and R Goebel, "Value signals in the prefrontal cortex predict individual preferences within and across reward categories," in "Society for Neuroeconomics Annual Meetings" Key Biscayne, FL 2012.

Gul, Faruk and Wolfgang Pesendorfer, "The Case for Mindless Economics," in Andrew Caplin and Andrew Schotter, eds., *The Foundations of Positive and Normative Economics: A Handbook*, Oxford University Press, 2008.

Harberger, A C, "Three Basic Postulates for Applied Welfare Economics: An Interpretive Essay," *Journal of Economic Literature*, 1971, *9* (3), 785–797.

Hare, Todd A, Colin F Camerer, Daniel T Knoepfle, and Antonio Rangel, "Value computations in ventral medial prefrontal cortex during charitable decision making incorporate input from regions involved in social cognition," *The Journal of Neuroscience*, January 2010, *30* (2), 583–590.

— , **Wolfram Schultz, Colin F Camerer, John P O'Doherty, and Antonio Rangel**, "Transformation of stimulus value signals into motor commands during simple choice," *Proceedings of the National Academy of Sciences*, November 2011, *108* (44), 18120–18125.

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

Hey, John D. and Chris Orme, "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica: Journal of the Econometric Society*, November 1994, *62* (6), 1291–1326.

Hicks, J R, "Consumers' Surplus and Index-Numbers," *The Review of Economic Studies*, 1942, *9* (2), 126–137.

Houthakker, H, "Revealed preference and the utility function," *Economica*, January 1950, *17* (66), 159–174.

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

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

Kahn, Itamar, Mitul Desai, Ulf Knoblich, Jacob Bernstein, Michael Henninger, Ann M Graybiel, Edward S Boyden, Randy L Buckner, and Christopher I Moore, "Characterization of the functional MRI response temporal linearity via optical control of neocortical pyramidal neurons," *The Journal of Neuroscience*, October 2011, *31* (42), 15086–15091.

Koszegi, B and Matthew Rabin, "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics*, November 2006, *121* (4), 1133–1165.

Krajbich, Ian, Carrie Armel, and Antonio Rangel, "Visual fixations and the computation and comparison of value in simple choice," *Nature Neuroscience*, October 2010, *13* (10), 1292–1298.

Kringelbach, M L, J O'Doherty, E T Rolls, and C Andrews, "Activation of the human orbitofrontal cortex to a liquid food stimulus is correlated with its subjective pleasantness," *Cerebral Cortex*, October 2003, *13* (10), 1064–1071.

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

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

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

Logothetis, N K, H Guggenberger, S Peled, and J Pauls, "Functional imaging of the monkey brain," *Nature Neuroscience*, June 1999, *2* (6), 555–562.

— , J Pauls, M Augath, T Trinath, and A Oeltermann, "Neurophysiological investigation of the basis of the fMRI signal," *Nature*, July 2001, *412* (6843), 150–157.

Logothetis, Nikos K, "The underpinnings of the BOLD functional magnetic Resonance Imaging Signal," *The Journal of Neuroscience*, May 2003, *23* (10), 3963–3971.

Loomes, Graham, "Modelling the stochastic component of behaviour in experiments: Some issues for the interpretation of data," *Experimental Economics*, 2005, *8* (4), 301–323.

— **and Robert Sugden**, "Incorporating a stochastic element into decision theories," *European Economic Review*, January 1995, pp. 641–648.

— **and** — , "Testing Different Stochastic Specifications of Risky Choice," *Economica*, January 1998, pp. 1–22.

— **, Jose Luis Pinto-Prades, Jose Maria Abellan-Perpinan, and Eva Rodriguez-Miguez**, "Modelling Noise and Imprecision in Individual Decisions," *University of Warwick Working Paper*, 2012.

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

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

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

Manski, Charles F, "The structure of random utility models," *Theory and Decision*, July 1977, *8* (3), 229–254.

Marschak, Jacob, "Binary choice constraints on random utility indications," in Kenneth Arrow, ed., *Stanford Symposium on Mathematical Methods in the Social Sciences*, Stanford University Press, 1960.

Martino, Federico De, Jan Zimmermann, Lars Muckli, Kamil Ugurbil, Essa Yacoub, and Rainer Goebel, "Cortical depth dependent functional responses in humans at 7T: improved specificity with 3D GRASE," *PloS one*, 2013, *8* (3), e60514.

McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," in Paul Zarembka, ed., *Frontiers in Econometrics*, Academic Press, Inc., 1974.

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

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

Michelson, Charles, Jonathan Pillow, and Eyal Seidemann, "Perceptual decisions are limited primarily by variability in early sensory cortex," in "Computational and Systems Neuroscience (Cosyne) Meetings 2013" Salt Lake City, UT 2013.

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

Neuhaus, John M, Charles E McCulloch, and Ross Boylan, "A note on type II error under random effects misspecification in generalized linear mixed models," *Biometrics*, June 2011, *67* (2), 654–6; disucssion 656–60.

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

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

_ , _ , B Shiv, and Antonio Rangel, "Marketing actions can modulate neural representations of experienced pleasantness," *Proceedings of the National Academy of Sciences*, January 2008.

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

Ratcliff, Roger, "A theory of memory retrieval," *Psychological Review*, 1978, *85* (2), 59–108.

Rieke, Fred, David Warlan, Rob de Ruyter van Steveninck, and William Bialek, *Spikes*, MIT Press, 1997.

Samuelson, P, "A note on the pure theory of consumer's behaviour," *Economica*, January 1938.

_ , "Discussion," *The American Economic Review*, 1963, *53* (2), 227–236.

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

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

Slesnick, D T, "Empirical Approaches to the Measurement of Welfare," *Journal of Economic Literature*, 1998, *36* (4), 2108–2165.

Smith, Alec, B Douglas Bernheim, Colin F Camerer, and Antonio Rangel, "Neural Activity Reveals Preferences Without Choices," *Cal Tech Working Paper*, 2011.

Smith, D V, John Clithero, R McKell Carter, and Scott A Huettel, "Reward Anticipation Engenders Functional Parcellation of the Striatum," in "Society for Neuroeconomics Annual Meetings" Key Biscayne, FL 2012.

Stevens, S S, "To Honor Fechner and Repeal His Law: A power function, not a log function, describes the operating characteristic of a sensory system," *Science*, January 1961, *133* (3446), 80–86.

Thurstone, L.L., "A law of comparative judgment," *Psychological Review*, 1927, *34* (4), 273–286.

Tolhurst, D J, J Anthony Movshon, and A F Dean, "The statistical reliability of signals in single neurons in cat and monkey visual cortex," *Vision research*, 1983, *23* (8), 775–785.

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

Vul, Edward, C Harris, Piotr Winkielman, and H Pashler, "Puzzlingly High Correlations in fMRI Studies of Emotion, Personality, and Social Cognitionn," *Perspectives on Psychological Science*, May 2009, *4* (3), 274–290.

Wang, N, X Lin, RG Gutierrez, and Raymond J Carroll, "Bias Analysis and SIMEX Approach in Generalized Linear Mixed Measurement Error Models," *Journal of the American Statistical Association*, 1998, *93* (441), 249–261.

Webb, Ryan, "Does Neuroeconomics Have a Role in Explaining Choice?," *SSRN Working Paper*, 2011.

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

_ **and Michael C Dorris**, "A Neural Model of Stochastic Choice in a Mixed Strategy Game," *SSRN Working Paper*, 2013.

_ **, Paul W Glimcher, and Kenway Louie**, "Neural Foundations for Context Dependent Choice," in "47th Annual Conference of the Canadian Economic Association" June 2013.

Woodford, M, "Prospect Theory as Efficient Perceptual Distortion," *The American Economic Review*, 2012, pp. 41–46.

Yatchew, A. and Z. Griliches, "Specification error in probit models," *The Review of Economics and Statistics*, 1985, pp. 134–139.