Decision by sampling implements efficient coding of psychoeconomic functions

Rahul Bhui[1,2] and Samuel J. Gershman[1]

[1]Department of Psychology and Center for Brain Science

[2]Department of Economics

Harvard University

Address for correspondence:

Samuel Gershman

Department of Psychology

Harvard University

52 Oxford St., room 295.05

Cambridge, MA 02138

Phone: 773-607-9817

E-mail: gershman@fas.harvard.edu

Abstract

The theory of decision by sampling (DbS) proposes that an attribute's subjective value is its rank within a sample of attribute values retrieved from memory. This can account for instances of context dependence beyond the reach of classic theories which assume stable preferences. In this paper, we provide a normative justification for DbS that is based on the principle of efficient coding. The efficient representation of information in a noiseless communication channel is characterized by a uniform response distribution, which the rank transformation implements. However, cognitive limitations imply that decision samples are finite, introducing noise. Efficient coding in a noisy channel requires smoothing of the signal, a principle that leads to a new generalization of DbS. This generalization is closely connected to range-frequency theory, and helps descriptively account for a wider set of behavioral observations, such as how context sensitivity varies with the number of available response categories.

Keywords: decision by sampling, efficient coding, kernel smoothing, range-frequency theory

Descriptive accounts of decision making, such as expected utility theory and prospect theory (Kahneman & Tversky, 1979), are typically based on a stable set of "psychoeconomic" functions specifying the mental representations of gains, losses, probabilities, and delays. However, the psychological reality of such functions has been challenged by evidence that decisions are highly context-sensitive: the mental representation of an attribute changes depending on the choice set and other attribute values retrieved from memory (Vlaev, Chater, Stewart, & Brown, 2011). As an alternative, some accounts have proposed that decisions are based on more elementary cognitive operations, namely memory retrieval and comparison (Johnson, Häubl, & Keinan, 2007; Marchiori, Di Guida, & Erev, 2015; Stewart, 2009; Stewart, Chater, & Brown, 2006). One particularly influential account—*decision by sampling* (DbS; Stewart, 2009; Stewart et al., 2006)—attempts to reconcile these viewpoints by proposing that psychoeconomic functions can be derived from principles of memory retrieval and comparison. According to DbS, the shapes of these functions are malleable, reflecting both local context and long-term statistical regularities that constitute the database from which memories are sampled. Despite its simplicity, DbS has accounted for a wide range of empirical phenomena (Stewart, Chater, Stott, & Reimers, 2003; Stewart, Reimers, & Harris, 2014; Ungemach, Stewart, & Reimers, 2011; Walasek & Stewart, 2015).

The basic idea of DbS is that attributes are sampled from memory and ordinally compared to the attributes of the current prospect. By tallying these ordinal comparisons, a decision maker computes the rank of the prospect's attribute relative to the distribution of attributes in memory. Preference is then determined by comparing the ranks across prospects.[1] While DbS is a psychological process model, this paper shows that the same set of ideas can be arrived at through a normative analysis. In particular, we derive DbS from the principle of *efficient coding*, which has a long history in the study of perceptual systems (Atick &

[1]Stewart (2009) assumed that probability and reward ranks are combined additively, in contrast to expected utility and prospect theory, where the attributes are combined multiplicatively. This assumption is not central to the present paper, which focuses on the psychoeconomic functions themselves rather than the decision process.

Redlich, 1992; Attneave, 1954; Barlow, 1961; Laughlin, 1981), and has more recently been applied to value representation in the brain (Louie, Grattan, & Glimcher, 2011; Louie, Khaw, & Glimcher, 2013; Louie, LoFaro, Webb, & Glimcher, 2014; Rangel & Clithero, 2012).

According to the efficient coding principle, the brain is designed to communicate information with as few spikes as possible, since spikes are metabolically expensive. As will be described in more detail below, this is accomplished by choosing a neural code that maximizes the mutual information between a neuron's inputs and outputs. When neurons are conceived as noiseless communication channels, maximizing mutual information is equivalent to minimizing redundancy, which can be achieved by recoding inputs according to their rank—precisely the operation implemented by DbS in the limit of an infinite number of samples. However, the channel becomes noisy when only a finite number of samples are drawn from memory, in which case some redundancy in the code is required to suppress noise. An approximation of the information-maximizing strategy is to smooth the samples prior to the rank transformation. This leads to modifications of DbS that were previously proposed to account for range effects on choice (Brown & Matthews, 2011; Parducci, 1995; Ronayne & Brown, 2016; Stewart et al., 2006).

The central contribution of our work is to clarify the computational design principles of DbS and related models, uniting them with an important strand of theoretical neuroscience. This paves the way for new behavioral predictions, insights into how DbS might be implemented in the brain, and a deeper understanding of the connections between information theory and decision making.

Decision by sampling

In this section, we present DbS formally and then review its applications to empirical phenomena.

Let $x \in [a, b]$ denote an attribute value in a psychoeconomic space (e.g., gains, losses, probabilities, delays). This attribute occurs in the environment with probability distribution $f(x)$. DbS samples a set of comparison values $\mathbf{x}_{1:N} = \{x_1, \ldots, x_N\}$ from $f(x)$ and then computes the rank of x relative to the sample:

$$\hat{F}(x; \mathbf{x}_{1:N}) = \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}[x \geq x_i], \tag{1}$$

where $\mathbb{I}[\cdot] = 1$ when its argument is true, 0 otherwise. In the infinite sample limit, the rank function converges with probability 1 to the cumulative distribution function (CDF), $F(x)$:

$$\lim_{N \to \infty} \hat{F}(x; \mathbf{x}_{1:N}) = \int_{-\infty}^{x} f(x') dx' = F(x). \tag{2}$$

The rank function (and the CDF) is monotonic in x, but importantly it exhibits steeper changes in high probability regions of the attribute space. Stewart et al. (2006) used this property to explain several well-known properties of psychoeconomic functions, using various proxy estimates of $f(x)$ for different attributes:

- Small gains and losses are more probable than large gains and losses, thus the rank functions for gains and losses are concave (diminishing marginal utility). Importantly, small losses are relatively more likely than small gains, implying that the rank function for losses is steeper, as proposed in prospect theory (Kahneman & Tversky, 1979).
- The distribution of temporal delays approximately follows a power law, giving rise to a power-law rank function. This subsumes hyperbolic discounting as a special case, but in fact the best-fitting rank function is sub-hyperbolic, consistent with several experimental studies (Myerson & Green, 1995; Simpson & Vuchinich, 2000).
- Very small and very large probabilities are more commonly encountered than mid-range probabilities, giving rise to an inverse S-shaped rank function (i.e., overweighting of low probabilities and underweighting of high probabilities), in agreement with the

probability weighting function derived from choice experiments (Gonzalez & Wu, 1999) and postulated by prospect theory.

One of the principal advantages of DbS over traditional approaches is that it can explain departures from these properties as the result of transient contextual information that distorts the long-term statistical regularities. Stewart et al. (2014) showed that the shapes of utility, discount, and probability weighting functions could be altered by exposing subjects to distributions of attribute values that varied in their skew. For example, a concave utility function could be converted to a convex function simply by populating the set of large gains more densely than the set of small gains, and the rate of discounting could be slowed by sampling delays from a uniform distribution (rather than an ecologically valid distribution with positive skew). Field studies by Ungemach et al. (2011) have recapitulated these observations, finding that choices between two lotteries were affected by incidental exposure to intermediate attribute values (supermarket prices), and choices between two delayed outcomes were affected by exposure to events occurring at intermediate delays. Large-scale studies of satisfaction as a function of income make the same point: relative income rank strongly determines satisfaction (Boyce, Brown, & Moore, 2010; Brown, Gardner, Oswald, & Qian, 2008).

An efficient coding perspective

What is the computational logic of the rank transformation? To shed some light on this question, let us view psychoeconomic functions as *communication channels*, taking as input an attribute value x and emitting as output a signal y drawn from the probability distribution $f(y|x)$. In designing such a channel, a basic problem is that the amount of information that can be reliably transmitted over a channel with fixed transmission rate (the channel capacity) is finite (Shannon & Weaver, 1949). A neuron consumes several orders of magnitude more energy during spiking compared to rest, such that the brain's energy budget can only afford to have around 1% of neurons active at any time (Lennie, 2003). Thus, the energy budget imposes a stringent constraint on the channel capacity of neurons, placing demands on neural codes to communicate information with as few spikes as possible. Consistent with this proposition, studies of many different neural systems suggest that economizing on spikes is a fundamental design principle (Laughlin, 2001).

There are two strategies to reduce the cost of information transmission. One is to reduce the signal-to-noise ratio (i.e., transmit lower precision messages); we will not consider this strategy further here, under the assumption that organisms need to maintain a certain level of precision for survival. The second strategy is to eliminate redundancy by recoding inputs. Intuitively, if an input can be predicted before the output has been observed, then the output is not conveying any information about the input—it is redundant with the receiver's prior knowledge. In other words, unpredictable outputs are more informative than predictable outputs. The most unpredictable output distribution is uniform; hence, the goal of redundancy reduction is to find a code that gets the output distribution close to uniform.

To make this idea more formal, let us define the mutual information between input x and output y (Cover & Thomas, 2006):

$$I(x; y) = H(y) - H(y|x), \tag{3}$$

where $H(y)$ is the (differential) entropy of the output, and $H(y|x)$ is the conditional entropy of the output given the input. Noise in the channel is captured by $H(y|x)$ which reflects the residual uncertainty in the response knowing the stimulus; it expresses how many bits are lost on average when transmitting x over the channel. Mutual information can equivalently be written as $I(x; y) = H(x) - H(x|y)$, meaning it tells us how much knowing the response reduces uncertainty about the stimulus. Both formulations describe the information that the response y carries about the stimulus x. The principle of efficient coding implies that stimuli should be encoded to maximize this quantity—that is, the mapping from x to y should maximize $I(x; y)$.

In the noiseless regime, $H(y|x)$ is 0, so maximizing mutual information is equivalent to maximizing output entropy (i.e., unpredictability). This is achieved by encoding x using the CDF, $y = F(x)$, also known as the *probability integral transform*, which guarantees that y is uniformly distributed (Laughlin, 1981). Since DbS approximates the probability integral transform, it can be understood as implementing efficient coding of psychoeconomic functions. In other words, DbS removes redundancies from the representations of gains, losses, probabilities, and delays, so that they can be represented with fewer bits (and thus presumably a lower metabolic cost). When the decision sample is large, the empirical rank $\hat{F}(x)$ will serve as a good approximation of the true rank $F(x)$.

We can extend this perspective by considering the noisy regime, in which case the probability integral transform is no longer optimal. When the channel is noisy, the resulting code will be corrupted, thereby increasing $H(y|x)$. In particular, when the decision sample is finite—which is necessarily the case under the inherent computational constraints that organisms face—the empirical rank estimate will suffer. Then the optimal code will have some redundancy in order to prevent information loss.

One can heuristically satisfy the conflicting demands of redundancy reduction and information transmission by first smoothing the inputs prior to computing the probability integral transform (Atick & Redlich, 1992):

$$\hat{F}_h(x; \mathbf{x}_{1:N}) = \frac{1}{N} \sum_{i=1}^{N} K\left(\frac{x - x_i}{h}\right), \tag{4}$$

where $K(z) = \int_{-\infty}^{z} k(z')dz'$ is an integrated kernel function and h is a bandwidth parameter. Smoothing alleviates sampling error by expanding the set of data contributing to the estimate. Because these data are more distant from the target, bias is introduced, but as has been extensively studied in the statistical literature on nonparametrics, this is a worthwhile cost to incur when it is outweighed by the corresponding reduction in estimator variance. This suggests that the principle of smoothing may guide the development and assessment of psychoeconomic models. In the rest of this paper, we show that this leads to new insight into generalizations of DbS along with Parducci's (1965; 1995) range-frequency theory.

Range sensitivity as kernel smoothing

One limitation of DbS, noted by Stewart et al. (2006), is that it does not capture the effect of attribute range on decisions and other economic judgments. For example, increasing the range of gains causes people to become more risk-seeking (Lim, 1995), and to be less satisfied with a given gain (Parducci, 1968). Range effects have also been documented in wage satisfaction judgments (Brown et al., 2008; Tripp & Brown, 2016), lottery pricing (Blavatskyy & Köhler, 2009), and job application decisions (Highhouse, Luong, & Sarkar-Barney, 1999; Rynes, Schwab, & Heneman, 1983).

Parducci's (1995) range-frequency theory (RFT) is the most well-known account of these effects, grounded in the psychophysics of perception. The theory states that the judged value (J) of the nth item is a convex combination of its rank ($F = (n-1)/(N-1)$, what Parducci refers to as its "frequency") and its position relative to the range of attribute values ($R = (x_n - x_\ell)/(x_h - x_\ell)$, where x_h and x_ℓ are the highest and lowest values):

$$J = wR + (1-w)F, \tag{5}$$

where w is a weighting parameter that determines the compromise between range and frequency.

We show that the range component of RFT can be derived as a kernel-smoothed estimate of the CDF using a uniform kernel. Both components of RFT are special cases based on the kernel's bandwidth, and the weighting parameter heuristically tunes the degree of smoothing. The frequency component emerges when the bandwidth is 0, and the range component emerges when the bandwidth is proportional to the sample range. We accordingly propose that while the frequency component of range-frequency theory can be viewed as implementing redundancy minimization for a noiseless channel, the range component can be viewed as assisting with efficient coding in the face of noise.

Just as kernel density estimates can be written as the average of kernel densities $k(z)$ centered around each data point, the distribution estimate can be similarly written as the average of the corresponding distributions $K(z) = \int_{-\infty}^{z} k(z')dz'$. For a uniform kernel $k(z) = 1/2$ for $|z| \leq 1$ and 0 otherwise, the distribution is $K(z) = (z+1)/2$ for $|z| \leq 1$, 0 for $z < -1$, and 1 for $z > 1$. Thus the CDF can be estimated smoothly by

$$
\begin{aligned}
\hat{F}_R(x) &= \frac{1}{N}\sum_i K\left(\frac{x - x_i}{h}\right) \\
&= \frac{1}{N}\sum_i K\left(\frac{x - x_i}{(x_h - x_\ell)/2}\right) \quad \text{setting bandwidth equal to half the range} \\
&\approx \frac{1}{N}\sum_i \left(\frac{x - x_i}{x_h - x_\ell} + \frac{1}{2}\right) \quad \text{with uniform kernel, approximate due to truncation of } K(\cdot) \\
&= \frac{x - \sum_i x_i/N}{x_h - x_\ell} + \frac{1}{2} \\
&\approx \frac{x - \left(x_\ell + \dfrac{x_h - x_\ell}{2}\right)}{x_h - x_\ell} + \frac{1}{2} \quad \text{approximating the mean by the mid-range} \\
&= \frac{x - x_\ell}{x_h - x_\ell}
\end{aligned}
$$

so $R \approx \hat{F}_R(x)$, meaning the range component of RFT inherits the properties of a kernel-smoothed estimate of the CDF.

Observe also that as bandwidth goes to 0, the estimate simply counts the number of data points less than the target value. If $x_i < x$ then $\lim_{h\to 0} K\left(\frac{x - x_i}{h}\right) = 1$, while if $x_i > x$ then

$\lim_{h\to 0} K\left(\frac{x-x_i}{h}\right) = 0$. (In the knife-edge case that $x = x_i$, $K\left(\frac{x-x_i}{h}\right) = 1/2$.) This means the estimate becomes the unsmoothed empirical rank, so both the range and frequency components of RFT can be characterized as special cases of the kernel-smoothed estimate based on bandwidth. Thus the RFT prediction can be written equivalently as a combination of smoothed and unsmoothed distribution estimates:

$$wR + (1-w)F \approx w\hat{F}_R(x) + (1-w)\hat{F}_0(x).$$

This derivation implies that the weighting parameter regulates the level of smoothing in place of the bandwidth parameter. If psychoeconomic functions are indeed attuned to the statistical properties of encountered attributes, the argument predicts that the range component should come to dominate as noise increases.

Why assume that the kernel is uniform and that bandwidth is proportional to range? These turn out to have optimal statistical properties for estimating CDFs. The uniform kernel minimizes mean integrated squared error when estimating distribution functions (Jones, 1990). Although other kernels are nearly as efficient, clearly the uniform is among those that work quite well. Optimal bandwidths (for both distribution and density estimation) are typically related to standard deviation. This occurs because when standard deviation rises, data becomes sparse and variance in the estimate increases. The benefit of variance reduction rises relative to the cost of increased bias from a higher bandwidth. Distribution-optimal plug-in bandwidths generally take the form $c\hat{\sigma}N^{-1/3}$ (Azzalini, 1981; Hansen, 2004) for various constants c. Standard deviation is approximately proportional to range, as illustrated by the "range rule" which states that the standard deviation is roughly the range divided by 4.

The plug-in bandwidth form also directly reveals that the sample size should influence bandwidth. Smoothing is only useful when samples are small, otherwise the bias it causes outweighs the variance reduction and increases error on net. This argument suggests that the weighting parameter which stands in for the bandwidth should vary also with the relevant sample size (whether of the observed context or the subsample drawn from memory). Although the mathematical expressions we present are not strictly optimal in an information-theoretic sense, they capture some important regularities, namely that smoothing should increase when the variability of the attribute values is large and should decrease with the number of values retrieved from memory. Moreover, the optimal level of smoothing may change according to other, less obvious factors such as the coarseness of response categories, which we illustrate in the following section.

Predicting range sensitivity

In the above characterization of RFT as a smoothed distribution estimator, the range weighting parameter w controls the degree of smoothing. This parameter is typically fixed at 0.5 when applying RFT to the data. However, the smoothing interpretation generates new predictions about what should cause it to vary. Here is one example of how this can help account for observed data.

When stimuli or responses are limited to a finite number of categories, people tend to partition the space so as to evenly use all bins (in other words, their limens are the distribution quantiles). This is just as redundancy minimization demands. Uniform use of categories has been found in judgments of loudness (Stevens, 1958), size (Parducci & Perrett, 1971), and could also explain the "numbers-of-levels" effect in marketing, whereby an increase in the number of attribute levels leads to an increase in the relative importance of that attribute (Verlegh, Schifferstein, & Wittink, 2002). However, Parducci and Wedell (1986) further find that the number of available response categories influences the effects of skewness and the apparent weighting of range and frequency components. As the number of categories increases, the range component becomes more dominant and skewness has a diminished effect on judgment.

Optimal smoothing provides a possible explanation for this result. Smoothing is most beneficial when attempting to distribute stimuli among many categories (Simonoff & Tutz, 2000). When there are few categories, the benefit is low because coarseness mitigates the effects of noise. To see how and when smoothing helps, consider that if the sample size is much smaller than the number of categories, the frequency-based limens do not change over large parts of the stimulus space.

For instance, suppose your sample consists of only a single draw from the context. The rank of all stimulus values below (or above) the draw will then be identical: values below will have rank 0, and values above will have rank 1. If you are trying to divide the distribution uniformly among two response categories, this is not especially harmful; you are effectively performing a median split, and the lone draw is a passable estimate of the median.

However, if you have 100 categories, the flatness of rank means there is little information with which to partition the rest of the stimulus space. The sample tells you about the median, but not about any other quantile. The noise from sampling variability is damaging even if the quantiles are interpolated. Thus when the response space is rich relative to the number of samples, a pure frequency estimate has little ability to discriminate between most categories. By contrast, the range estimate changes robustly and is sensitive across the full distribution (supposing as Parducci and Wedell do that the endpoints are privileged in memory). The tradeoff is that it may be too smooth, but when the sample size is small, this is acceptable.

To demonstrate these effects, we conduct simulations showing when uniform categorization performance is improved or impaired by the range component. These simulations assess how the error between the RFT quantiles and the true quantiles varies with the range weight w and the number of response categories. Following Parducci (1965), RFT limens are the weighted average of range and frequency limens. Range limens here are the quantiles of the standard uniform distribution, and frequency limens are the (linearly interpolated) quantiles of the sample. Which quantiles are chosen depends on the number of response

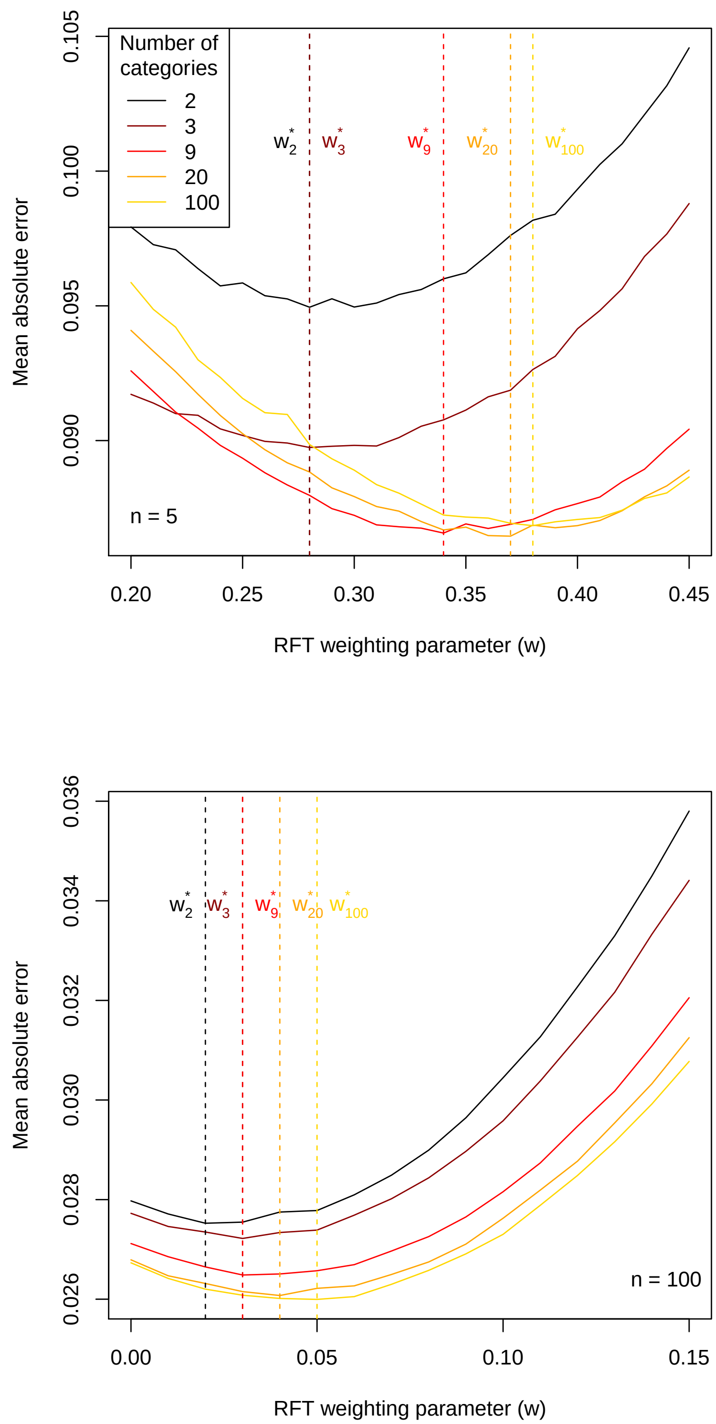

Figure 1: Error-minimizing range-frequency weights for different numbers of response categories. (top) Size of sample is small. Optimal values are $w_2^* = 0.28$, $w_3^* = 0.28$, $w_9^* = 0.34$, $w_{20}^* = 0.37$, and $w_{100}^* = 0.38$. (bottom) Size of sample is large. Optimal values are $w_2^* = 0.02$, $w_3^* = 0.03$, $w_9^* = 0.03$, $w_{20}^* = 0.04$, and $w_{100}^* = 0.05$.

categories. With three categories, for example, an even partition is formed by the 33rd and 67th percentiles. We pick as the context a skewed Beta distribution with parameters $\alpha = 1$ and $\beta = 2$ (or vice versa) to roughly resemble Parducci and Wedell's (1986) stimuli. The true quantiles are of this distribution.

Each line in Figure 1 depicts the mean absolute difference between the RFT quantiles and the true quantiles for values of w. The lowest point on each line represents the optimal (error-minimizing) value of w. The lines correspond to various numbers of response categories used by Parducci and Wedell (1986), between two and 100. We assume the sample is small, consisting of 5 draws from the context, to reflect natural constraints on working memory. As the figure reveals, the optimal value of w is increasing in the number of categories. That is, the range component helps accurately estimate the desired category boundaries especially when response categories are numerous.

Note also that when the sample is large, there is little need for smoothing. The utility of smoothing is a consequence of variability from finite samples. Hence with a large sample, the optimal values of w shrink substantially and the range component becomes negligible however many categories are available.

Smoothing as reduced discriminability

As suggested by Stewart et al. (2006), and later elaborated by Brown and Matthews (2011), range-like effects can be captured by DbS if one assumes that experienced attribute values are not perfectly discriminable in memory. In particular, many models of memory assume that the discriminability of items is inversely proportional to their density in attribute space (e.g., Brown, Neath, & Chater, 2007; Murdock, 1960). Thus, an item is less likely to be retrieved if other similar items enter into competition for retrieval. This mechanically flattens out the effective retrieval distribution, damping its original skewness. Estimated ranks are then based on a more uniform distribution, mimicking an increase in the relative importance of the RFT range component. Parducci and Wedell (1986) find they can also explain their category results described in the previous section using a model that acts like that of Brown and Matthews. Reducing stimulus discriminability and increasing range weight thus have similar explanatory capabilities.

This link might explain why range weighting appears to change depending on the salience of the contextual distribution. For example, when the distribution must be drawn from memory due to sequential rather than simultaneous presentation, the range component seems to become more important (Choplin & Wedell, 2014; Niedrich, Sharma, & Wedell, 2001; Qian & Brown, 2005). Recalling samples from memory injects noise and reduces discriminability. Similarly, customers who have been exposed to a trend in prices exhibit less range weighting (Niedrich, Weathers, Hill, & Bell, 2009), perhaps because the clearer structure of the contextual distribution enables more precise retrieval.

The efficient coding framework offers another perspective on this connection: imperfect discriminability may reflect part of a mechanism for introducing redundancy in order to reduce coding errors. Kernel smoothing from a sampling perspective can manifest as reduced discriminability.

Suppose that when you draw an item from memory, you feel some uncertainty about its true location. This entails that items won't be completely distinguishable, and those nearer each other will be harder to distinguish. These are the assumptions imposed by reduced discriminability models. The coarse binary comparisons of DbS are then replaced with graded assessments of order to allow some tolerance. Rather than simply determining whether the target is greater than each sample value, the differences between the target and the samples are judged as significant to varying degrees based on the level of uncertainty.

This smoothed comparison is exactly what a kernel encodes, illustrated in Figure 2. Each sample value x_i contributes $K(\frac{x-x_i}{h})$ to the rank estimate of x. The level of uncertainty is

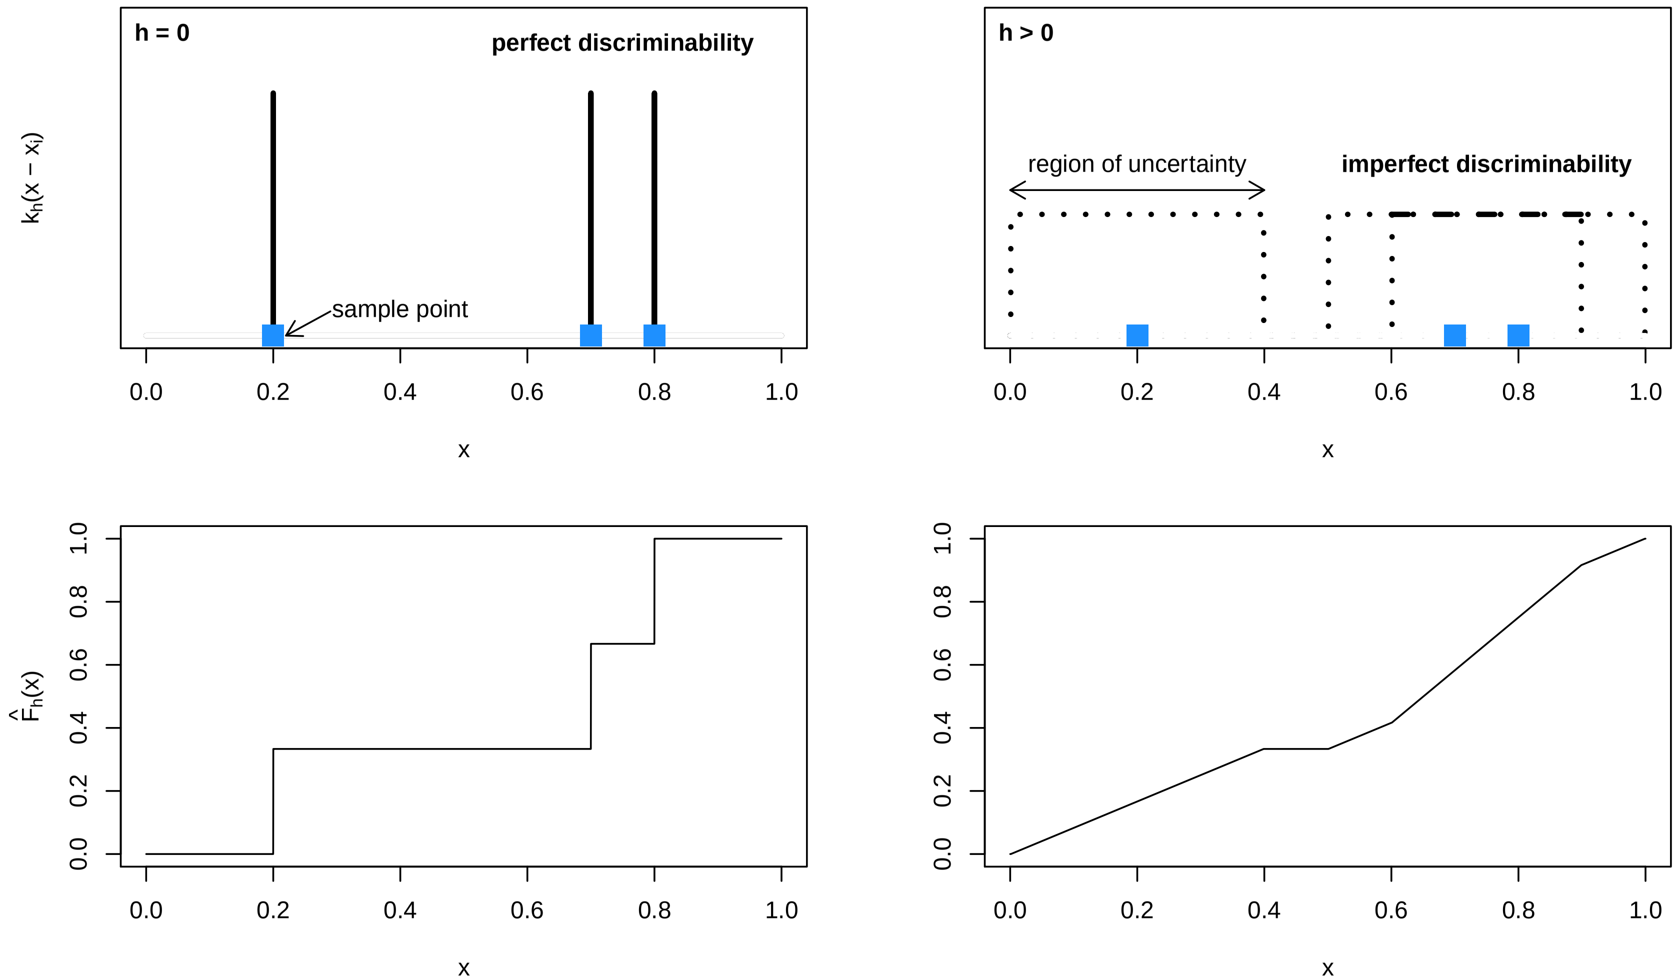

Figure 2: Kernel smoothing manifests as imperfect discriminability. (left) When items are completely distinguished from each other, bandwidth is zero, and no distribution smoothing occurs. (right) When some tolerance is allowed for uncertainty in location, bandwidth is positive, and the distribution is smoothed.

represented by the bandwidth h, and accordingly controls the degree of smoothing. If this is small, the contribution of x_i boils down to $\lim_{h\to 0} K(\frac{x-x_i}{h}) = \mathbb{I}[x \geq x_i]$, a pure binary comparison. As uncertainty grows large, this becomes $\lim_{h\to\infty} K(\frac{x-x_i}{h}) = K(0)\ \forall\ x, x_i$. No comparisons can be made at all, so the effective distribution becomes more uniform, just as in the reduced discriminability models. The optimal amount of smoothing to account for sampling variability is somewhere in the middle, informed by range as discussed earlier. In sampling terms, range provides information about how much comparison tolerance should be allotted.

We note further that the region of uncertainty can itself be instantiated by resampling each retrieved item, suggesting how range sensitivity could be implemented via purely binary comparisons. If uncertainty is high, resamples will be spread out, some of which will be greater than items otherwise higher on the scale. This may be more cost-effective than drawing fresh samples if resamples are cheaper to obtain, which is plausible when they can be anchored to their originally retrieved estimates. As the cognitive simplicity of DbS is a key part of its motivation, placing extensions on equal footing contributes to their justification. In addition, the range component of RFT requires recall of only the distribution endpoints, and thus may be cognitively undemanding in that sense. Kernel-smoothed rank need not be more taxing to compute than empirical rank.

Thus the assumption of imperfect discriminability can be directly tied to the implementation of kernel smoothing. From the efficient coding perspective, the relationship between discriminability and range sensitivity is not so surprising. They are alternative facets of the same distribution smoothing phenomenon.

Smoothing and context effects

The explanatory power of DbS stems from its ability to predict how judgment is influenced by the contextual distribution. Context distorts value judgments in ways that produce violations of classical expected utility theory such as those described earlier. Incorporating smoothing enables DbS to capture an even wider class of context effects. We have discussed some of these in a single-attribute setting under the banner of range sensitivity.

In a multi-attribute setting, Ronayne and Brown (2016) show how DbS with a form of local sampling can predict the attraction, compromise, and similarity effects. The compromise effect is most closely linked to locality so we focus on this. It refers to a scenario in which the addition of an extreme option makes the newly intermediate option more likely to be chosen. Figure 3 illustrates a kernel smoothing version of Ronayne and Brown's setup. Options A and B are focal price-quality pairs, and R_A and R_B denote the regions they uniquely dominate. C_A denotes the third, extreme option that makes A the compromise and hence preferred. The model supposes that a finite sample is drawn from memory, and the subjective value of each option is the number of samples in the region that it dominates. Hence their unique dominance regions play a key role in the relative values of options.

Crucially, Ronayne and Brown assume that regions close to the presented options in attribute space are preferentially sampled; the support of the effective retrieval distribution is represented by the shaded region. As a result, C_A contributes more probability mass to the solo-dominance region of its neighbor A than it does to that of B. Thus the presence of C_A benefits A relative to B, producing the compromise effect.

This locality assumption can be interpreted as a form of kernel smoothing. If the available options constitute draws from the context, the location they represent for purposes of rank comparison is inexact. Along similar lines to the previous section, the resulting region of uncertainty instantiates a kernel which trails off when it gets farther from the option. If this

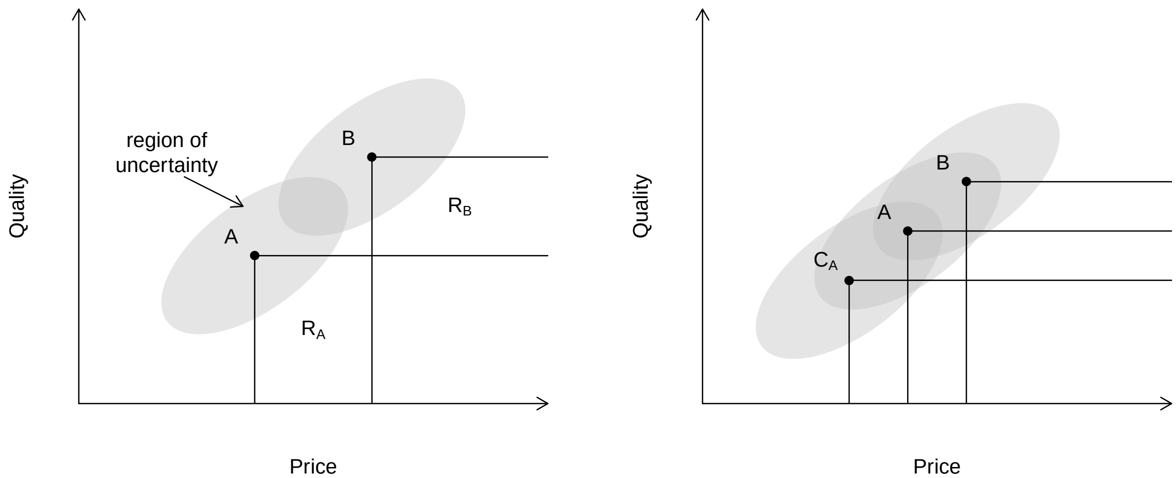

Figure 3: Compromise effects can be generated by multi-attribute DbS with locality. (left) A and B are the focal options, and R_A and R_B represent their unique dominance regions. Locality of sampling can be interpreted as a kernel around each option. (right) The third option C_A contributes more mass to R_A than R_B, benefiting A over B.

is the case, optimal smoothing predicts that the size of the kernel should grow according to the scale of the option distribution, naturally providing a reason for compromise effects at all scales. The scope of the compromise effect should also be related to factors like the number of options available and the psychological noisiness of context retrieval. This could help explain why the compromise effect decreases when more options are presented (Gourville & Soman, 2007). Because a larger sample is available, smoothing is less necessary and bandwidth decreases.

Neural evidence and implementation

There is much evidence of neuronal gain control that produces range sensitivity (e.g. Padoa-Schioppa, 2009; Tremblay & Schultz, 1999), though little research has attempted to pinpoint value encoding by option rank. Mullett and Tunney (2013) provide the most direct neural evidence to date. In their experiment, participants were faced with blocks containing either rewards of £0.10, £0.20, and £0.30 or rewards of £5.00, £7.00, and £10.00. Using fMRI, they found activity in ventromedial prefrontal cortex and anterior cingulate cortex to be linear in the global option rank rather than absolute value, meaning the difference between £5.00 and £0.30 was similar to the difference between £0.30 and £0.20. Interestingly, these regions encoded rank based on the set of stimuli presented across the whole experiment, while activity in the caudate and thalamus scaled according to the experimental block, exhibiting a more local context dependency.

Why have there not been more direct signs of rank encoding? We note two possible reasons why rank-based value representations may be difficult to discern. First, the point of smoothing is to gracefully approximate rank across the entire spectrum of available values. We have shown that this can lead to the range component of RFT, which exhibits linearity across the full distribution support and is not as choppy as the frequency component. Thus in some cases smoothing may conceal the effects of contextual skewness. Second, rank-based coding may take unusual forms not previously considered by those studying decision making. In practice, forms of rate and population coding have dominated the neuroeconomic literature. Alternatively, efficient coding could be naturally implemented in the temporal domain.

On a neuronal level, encoding based on the rank of spike timing—known as rank order coding (Thorpe & Gautrais, 1998)—has several benefits. Such schemes convey information more efficiently than standard rate codes while using simpler and more robust decoding mechanisms than precise timing codes (VanRullen & Thorpe, 2001). The most important information can be transmitted first and quickly decoded, enabling the kind of rapid responses to stimuli that may be essential for survival. Rank order coding is capable of transmitting information on the very quick timescales of the sensory domain (VanRullen, Guyonneau, & Thorpe, 2005), and there is evidence of its role in retinal processing (Portelli et al., 2016). We show how smoothing can be naturally integrated into rank order coding. Existing neuronal implementations permit neurally plausible extensions that would generate smoothing.

Consider the simple example circuit shown in Figure 4 modeled after Thorpe, Delorme, and Van Rullen (2001). Suppose neurons A, B, and C represent stimuli with attribute intensities $A > B > C$. They are to be compared using their rank-influenced cumulative outputs a, b, and c. Normally A will fire first because the intensity of its stimulation is highest, B will fire later, and C will fire later yet. So a should be the greatest, and this can be accomplished straightforwardly via the inhibitory interneurons I that attenuate the effects of later firing inputs. Because A fires first, the inhibition will grow to depress the effects of B, and subsequently C, reflecting their lower rank. Stronger inhibitory power sharply raises the relative value of the highest ranks, producing a value of a much greater than b and c.

In this mechanism, neural activity is implicitly assumed to decay instantly. However, a gradual decay may be more realistic. If active traces remain after a neuron's initial spike, subsequent neurons can fire before earlier traces completely decay. The inhibitory effects of higher ranked neurons are accordingly lessened since part of their activity occurs after lower ranked neurons spike. This means the active trace functions like an asymmetric kernel, with the decay distribution affecting the kernel shape and the decay rate controlling the bandwidth. A slow decay reflects a large bandwidth that smooths the encoded attribute ranks. Steep changes due to differences in rank are consequently diminished.

In summary, we have shown how efficient coding could be implemented naturally in the temporal domain, with smoothing implemented through gradual decay of activity. However, this hypothesis remains to be verified experimentally.

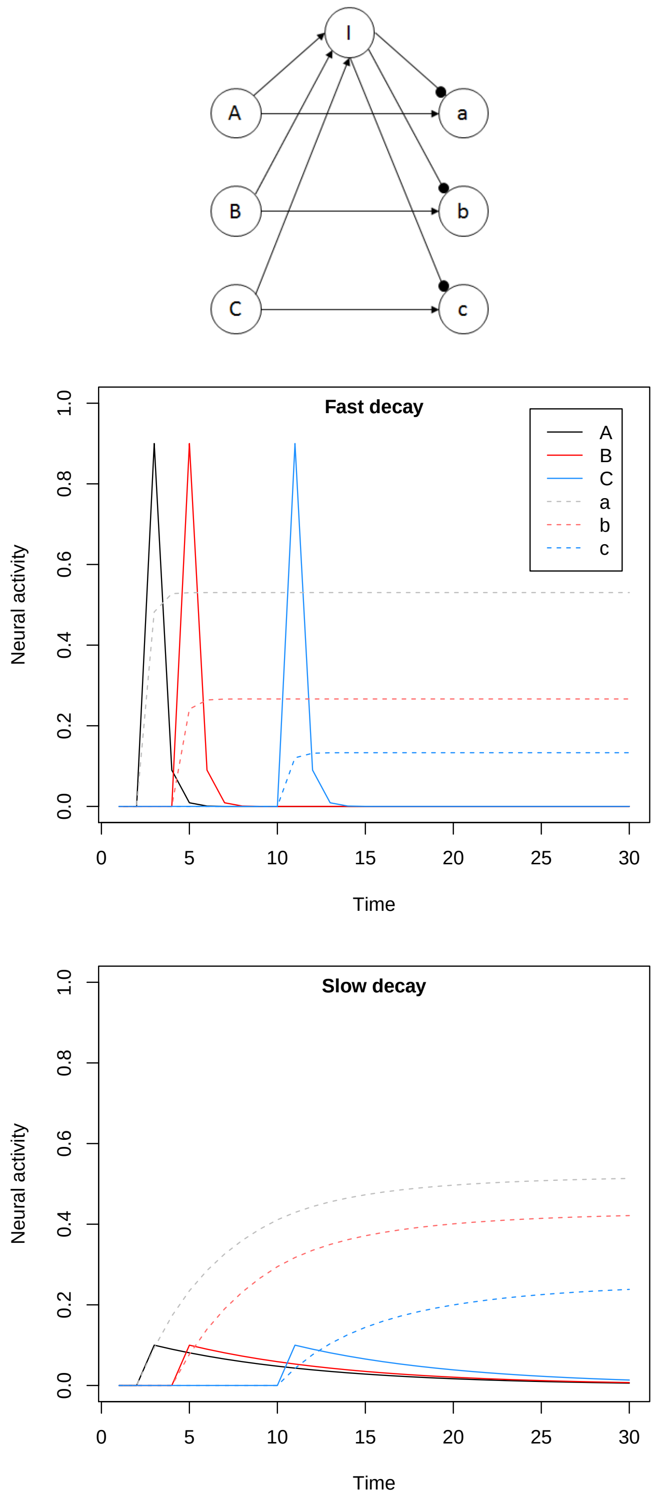

Figure 4: Slow post-spike activity decay implements smoothed rank order coding. (top) Diagram of simple network implementing rank order coding. Simulation assumes divisive inhibition which attenuates effect of inputs by a factor of $0.5^{I(t)}$. (middle) When decay is fast, output values are sharply sensitive to rank, as $a >> b, c$. (bottom) When decay is slow, output values are less sharply sensitive to rank, as the decay operates similar to an asymmetric kernel.

Neuroeconomic models and efficient coding

By grounding DbS in the efficient coding framework, we provide an adaptive justification for its existence. This also adds DbS to the class of several increasingly popular theories of decision making—and more—that are intimately linked to principles of optimality. Efficient coding seems to provide a unifying principle that spans a wide range of domains, from low-level perception to high-level judgment. Indeed, early research in behavioral economics was guided by the notion that cognitive processes underlying judgment resembled those underlying perception (Kahneman, 2002). This supposition has worked well historically and continues to do so. Research in cognitive science and computational neuroscience has yielded a wealth of precise, quantitative, and tractable characterizations of perceptual processing. These insights are being ported into the study of decision making and have already demonstrated great predictive power. We mention two such instances here, divisive normalization and sequential sampling, both of which were motivated in part by their optimal statistical properties.

Divisive normalization is a mechanism whereby neurons inhibit each other, and in so doing, exhibit responses that are normalized with respect to their pooled inputs. This produces adaptive gain control that calibrates the sensitivity of neuronal responses according to the contextual distribution. Across various species, divisive normalization has been observed in neural pathways for vision (Busse, Wade, & Carandini, 2009; Carandini, Heeger, & Movshon, 1997; Heeger, 1992), audition (Rabinowitz, Willmore, Schnupp, & King, 2011; Schwartz & Simoncelli, 2001b), olfaction (Luo, Axel, & Abbott, 2010; Olsen, Bhandawat, & Wilson, 2010), and even multisensory integration (Ohshiro, Angelaki, & DeAngelis, 2011). It is considered pervasive enough to be among the set of "canonical computations" (Carandini & Heeger, 2012). When applied to decision making, it predicts context dependence that is able to account for classic deviations from expected utility theory (Louie et al., 2011, 2013, 2014; Rangel & Clithero, 2012).

Divisive normalization was proposed to be a method of redundancy reduction in sensory processing (Schwartz & Simoncelli, 2001a; Wainwright, Schwartz, & Simoncelli, 2002), an efficient coding interpretation which has played an important role in its popularity. It has been shown theoretically and empirically to reduce higher-order redundancies. For example, divisive normalization gaussianizes the multivariate heavy-tailed distributions that characterize natural image statistics. This transformation renders them more independent when combined with filtering and helps to reduce their higher-order statistical dependencies (Lyu, 2010, 2011; Malo & Laparra, 2010). This highlights further that context sensitivity is a fundamental aspect of redundancy reduction, and should be widespread.

Sequential sampling models describe the process of decision making as the accumulation of evidence over time in favor of each option (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Busemeyer & Townsend, 1993; Laming, 1968; Ratcliff, 1978; Stone, 1960). An option is chosen when its accumulated evidence reaches some threshold. These models, typically applied to perceptual decision making, make precise predictions about the joint distribution of choices and response times. Sequential sampling has been shown to account quite well for behavioral data (e.g. Ratcliff, 2002; Ratcliff & Rouder, 1998; Smith, Ratcliff, & Wolfgang, 2004), and moreover appears to match patterns of neural activity (Gold & Shadlen, 2002; Hanes & Schall, 1996; Ratcliff, Cherian, & Segraves, 2003; Shadlen & Newsome, 2001; Smith & Ratcliff, 2004). The approach has has been successfully applied to economic decisions supposing the accumulation process is driven by the value difference between options (Krajbich, Armel, & Rangel, 2010; Krajbich, Hare, Bartling, Morishima, & Fehr, 2015; Krajbich, Lu, Camerer, & Rangel, 2012; Krajbich & Rangel, 2011; Milosavljevic, Malmaud, Huth, Koch, & Rangel, 2010).

From the efficient coding standpoint, sequential sampling models elaborate on the alternative strategy touched upon earlier for reducing the cost of information transmission by changing the signal-to-noise ratio. They were motivated by efficient, Bayes-optimal statistical algorithms for estimating the state of the world from a sequence of incoming data (specifically the sequential probability ratio test; Arrow, Blackwell, and Girshick 1949; Wald 1947; Wald and Wolfowitz 1948). More recently, Woodford (2014, 2016) developed a model of optimal evidence accumulation that formally incorporates information theory. He describes the internal processing of signals as a sensor with finite channel capacity. The agent is assumed to choose a stopping rule that maximizes expected value subject to a bound on mutual information. In other words, given that increasing signal precision is costly and the channel is constrained, the agent must decide how much precision suits their needs. Woodford's model fits behavioral data possibly better than traditional models, and opens up further vistas for the analysis of stochastic choice and response times.

There are good reasons why efficient coding is a powerful unifying principle, and why perception research has much to offer judgment research. Natural selection imposes pressure on organisms to efficiently represent the information they need to survive and reproduce. The brain spends a tremendous amount of energy, accounting for 20% of resting oxygen consumption in adult humans, most of which is directly required for signaling (Laughlin, 2001). Given the enormous metabolic costs of neural activity, wasteful encoding of information would produce a steep drop in fitness, and should therefore be sharply curtailed by selection pressures. This argument applies generally across organisms and types of processing. The domain of perception provides a low-level testbed in which computational descriptions of problems to be solved, and hence the nature of optimality within them, can be more transparently specified. Thus the development and assessment of theories, including those based on efficient coding, can progress at a faster rate. This creates an arbitrage opportunity for the study of decision making.

We note that Woodford's model incorporates information theory in a manner which is directly sensitive to economic preferences. This is atypical for work on efficient coding but seems essential in economic domains. In other domains too, a greater focus on costs and benefits might enable us to quantify efficiency in more sophisticated ways and derive new results (see for example Levy and Baxter 1996). Park and Pillow (2017) generalize classical efficient coding in a Bayesian framework which enlarges the space of permissible loss functions, and can give rise to a dramatically different set of efficient codes. This hints that theoretical practice could profit from alternative modeling styles that move beyond the standard approach of simply maximizing mutual information. Research on the efficient coding hypothesis stands to gain from a richer integration of information theory with the prospective costs and benefits faced by the agent. This may be particularly important for value-based decision making, where variation in incentives is widespread and plays a crucial explanatory role. We leave this as a direction for future research.

Conclusion

Abundant evidence demonstrates context sensitivity in judgment and decision making that deviates from classic theoretical models. Decision by sampling was intended to account for such data, proposing that the value of an attribute is encoded as its rank in a contextual distribution drawn from memory. This can be computed by tallying ordinal comparisons between the target attribute level and samples from the context, and entails that psychoeconomic functions are intrinsically malleable. Here, we ground DbS in an influential strand of theoretical neuroscience which posits that activity patterns in the brain efficiently represent information. This hypothesis of efficient coding predicts that neural processing, and therefore the psychophysical functions it generates, should be adapted to natural stimuli faced by the organism.

We identified DbS (and equivalently, the frequency component of range-frequency theory) as an implementation of information-theoretic redundancy minimization with a noiseless communication channel. Redundancy minimization requires values to be encoded so that they are uniformly distributed across the bounded response range, which is achieved by the rank transformation of DbS. However, noise is introduced when only a finite sample can be drawn from memory, as is necessitated by inherent computational constraints. If the adaptive purpose of DbS is premised on estimating rank, this argument suggests that it can be extended by incorporating kernel smoothing which would alleviate the error caused by finite samples.

We drew out the implications of this kernel smoothing, showing that under certain assumptions that reflect optimal smoothing, the range component of RFT can be derived as a kernel-smoothed estimate of rank. This derivation revealed that RFT implements efficient coding in a previously unrecognized fashion. It also suggests that principles of optimal smoothing enable us to predict variation in the RFT weighting parameter, and we demonstrated that this can help account for past data on how judgment is affected by the number of available response categories. Psychologically, kernel smoothing can manifest as reduced discriminability of retrieved items, which sheds light on how previous extensions of DbS and RFT that assume imperfect discriminability capture range sensitivity. Similarly, extensions of DbS based on locality of sampling that capture context effects observed in decision making can be interpreted as kernel smoothing, providing further rationale for such assumptions.

These insights into DbS open up many directions for future research. First, our analysis indicates that the optimal amount of smoothing depends on sample size; thus, can we predict smoothing from individual differences in memory capacity? Second, can we manipulate the degree of smoothing by changing parameters of the task? For example, can we increase smoothing by placing individuals under cognitive load? Third, can we find direct evidence for adaptive smoothing in the brain? Answering this question will require measurement techniques with high temporal resolution (such as single unit recordings) in order to test the predictions of the temporal coding scheme described in this paper. Finally, can we link this coding scheme to decision making behavior, such as context effects? We believe that efficient coding provides a powerful framework for addressing these questions.

References

Arrow, K. J., Blackwell, D., & Girshick, M. A. (1949). Bayes and minimax solutions of sequential decision problems. *Econometrica*, 213–244.

Atick, J. J., & Redlich, A. N. (1992). What does the retina know about natural scenes? *Neural Computation*, *4*, 196–210.

Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, *61*, 183–193.

Azzalini, A. (1981). A note on the estimation of a distribution function and quantiles by a kernel method. *Biometrika*, *68*(1), 326–328.

Barlow, H. B. (1961). Possible principles underlying the transformations of sensory messages. In W. Rosenblith (Ed.), *Sensory communication.* MIT press.

Blavatskyy, P. R., & Köhler, W. R. (2009). Range effects and lottery pricing. *Experimental Economics*, *12*, 332–349.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological review*, *113*(4), 700.

Boyce, C. J., Brown, G. D., & Moore, S. C. (2010). Money and happiness: Rank of income, not income, affects life satisfaction. *Psychological Science*, *21*, 471–475.

Brown, G. D., Gardner, J., Oswald, A. J., & Qian, J. (2008). Does wage rank affect employees well-being? *Industrial Relations: A Journal of Economy and Society*, *47*, 355–389.

Brown, G. D., & Matthews, W. J. (2011). Decision by sampling and memory distinctiveness: Range effects from rank-based models of judgment and choice. *Frontiers in Psychology*, *2*.

Brown, G. D., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psycho-*

logical Review, *114*, 539–576.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review*, *100*(3), 432.

Busse, L., Wade, A. R., & Carandini, M. (2009). Representation of concurrent stimuli by population activity in visual cortex. *Neuron*, *64*(6), 931–942.

Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, *13*(1), 51–62.

Carandini, M., Heeger, D. J., & Movshon, J. A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *Journal of Neuroscience*, *17*(21), 8621–8644.

Choplin, J. M., & Wedell, D. H. (2014). How many calories were in those hamburgers again? distribution density biases recall of attribute values. *Judgment and Decision Making*, *9*(3), 243.

Cover, T. M., & Thomas, J. A. (2006). *Elements of information theory.* John Wiley & Sons.

Gold, J. I., & Shadlen, M. N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, *36*(2), 299–308.

Gonzalez, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology*, *38*, 129–166.

Gourville, J. T., & Soman, D. (2007). Extremeness seeking: When and why consumers prefer the extremes.

Hanes, D. P., & Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science*, *274*(5286), 427–430.

Hansen, B. E. (2004). Bandwidth selection for nonparametric distribution estimation. *manuscript, University of Wisconsin*.

Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual neuroscience*, *9*(2), 181–197.

Highhouse, S., Luong, A., & Sarkar-Barney, S. (1999). Research design, measurement, and effects of attribute range on job choice: More than meets the eye. *Organizational Research Methods*, *2*, 37–48.

Johnson, E. J., Häubl, G., & Keinan, A. (2007). Aspects of endowment: a query theory of value construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 461–474.

Jones, M. (1990). The performance of kernel density functions in kernel distribution function estimation. *Statistics & Probability Letters*, *9*(2), 129–132.

Kahneman, D. (2002). Maps of bounded rationality: A perspective on intuitive judgment and choice. *Nobel prize lecture*, *8*, 351–401.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 263–291.

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature neuroscience*, *13*(10), 1292–1298.

Krajbich, I., Hare, T., Bartling, B., Morishima, Y., & Fehr, E. (2015). A common mechanism underlying food choice and social decisions. *PLoS Computational Biology*, *11*(10), e1004371.

Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in Psychology*, *3*.

Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, *108*(33), 13852–13857.

Laming, D. R. J. (1968). *Information theory of choice-reaction times.* Academic Press.

Laughlin, S. B. (1981). A simple coding procedure enhances a neuron's information capacity. *Zeitschrift für Naturforschung c*, *36*, 910–912.

Laughlin, S. B. (2001). Energy as a constraint on the coding and processing of sensory information. *Current Opinion in Neurobiology*, *11*, 475–480.

Lennie, P. (2003). The cost of cortical computation. *Current biology*, *13*, 493–497.

Levy, W. B., & Baxter, R. A. (1996). Energy efficient neural codes. *Neural computation*, *8*(3), 531–543.

Lim, R. G. (1995). A range-frequency explanation of shifting reference points in risky decision making. *Organizational Behavior and Human Decision Processes*, *63*, 6–20.

Louie, K., Grattan, L. E., & Glimcher, P. W. (2011). Reward value-based gain control: divisive normalization in parietal cortex. *Journal of Neuroscience*, *31*, 10627–10639.

Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences*, *110*, 6139–6144.

Louie, K., LoFaro, T., Webb, R., & Glimcher, P. W. (2014). Dynamic divisive normalization predicts time-varying value coding in decision-related circuits. *Journal of Neuroscience*, *34*, 16046–16057.

Luo, S. X., Axel, R., & Abbott, L. (2010). Generating sparse and selective third-order responses in the olfactory system of the fly. *Proceedings of the National Academy of Sciences*, *107*(23), 10713–10718.

Lyu, S. (2010). Divisive normalization: Justification and effectiveness as efficient coding transform. In *Advances in neural information processing systems* (pp. 1522–1530).

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

Malo, J., & Laparra, V. (2010). Psychophysically tuned divisive normalization approximately factorizes the pdf of natural images. *Neural computation*, *22*(12), 3179–3206.

Marchiori, D., Di Guida, S., & Erev, I. (2015). Noisy retrieval models of over-and undersensitivity to rare events. *Decision*, *2*, 82–106.

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgment and Decision Making*, *5*(6), 437–449.

Mullett, T. L., & Tunney, R. J. (2013). Value representations by rank order in a distributed network of varying context dependency. *Brain and cognition*, *82*(1), 76–83.

Murdock, B. B. (1960). The distinctiveness of stimuli. *Psychological Review*, *67*, 16–31.

Myerson, J., & Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the Experimental Analysis of Behavior*, *64*, 263–276.

Niedrich, R. W., Sharma, S., & Wedell, D. H. (2001). Reference price and price perceptions: A comparison of alternative models. *Journal of Consumer Research*, *28*(3), 339–354.

Niedrich, R. W., Weathers, D., Hill, R. C., & Bell, D. R. (2009). Specifying price judgments with range–frequency theory in models of brand choice. *Journal of marketing research*, *46*(5), 693–702.

Ohshiro, T., Angelaki, D. E., & DeAngelis, G. C. (2011). A normalization model of multisensory integration. *Nature neuroscience*, *14*(6), 775–782.

Olsen, S. R., Bhandawat, V., & Wilson, R. I. (2010). Divisive normalization in olfactory population codes. *Neuron*, *66*(2), 287–299.

Padoa-Schioppa, C. (2009). Range-adapting representation of economic value in the orbitofrontal cortex. *Journal of Neuroscience*, *29*(44), 14004–14014.

Parducci, A. (1965). Category judgment: a range-frequency model. *Psychological review*, *72*(6), 407.

Parducci, A. (1968). The relativism of absolute judgments. *Scientific American*, *219*, 84–90.

Parducci, A. (1995). *Happiness, pleasure, and judgment: The contextual theory and its applications.* Lawrence Erlbaum Associates, Inc.

Parducci, A., & Perrett, L. F. (1971). Category rating scales: Effects of relative spacing and frequency of stimulus values. *Journal of Experimental Psychology*, *89*(2), 427.

Parducci, A., & Wedell, D. H. (1986). The category effect with rating scales: number of categories, number of stimuli, and method of presentation. *Journal of Experimental Psychology: Human Perception and Performance*, *12*, 496–516.

Park, I. M., & Pillow, J. W. (2017). Bayesian efficient coding. *bioRxiv*, 178418.

Portelli, G., Barrett, J. M., Hilgen, G., Masquelier, T., Maccione, A., Di Marco, S., … Sernagor, E. (2016). Rank order coding: a retinal information decoding strategy revealed by large-scale multielectrode array retinal recordings. *eneuro*, *3*(3), ENEURO–0134.

Qian, J., & Brown, G. D. (2005). Similarity-based sampling: Testing a model of price psychophysics. In *Proceedings of the 27th annual conference of the cognitive science society* (pp. 1785–1790).

Rabinowitz, N. C., Willmore, B. D., Schnupp, J. W., & King, A. J. (2011). Contrast gain control in auditory cortex. *Neuron*, *70*(6), 1178–1191.

Rangel, A., & Clithero, J. A. (2012). Value normalization in decision making: theory and evidence. *Current Opinion in Neurobiology*, *22*, 970–981.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, *85*(2), 59.

Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic bulletin & review*, *9*(2), 278–291.

Ratcliff, R., Cherian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *Journal of neurophysiology*, *90*(3), 1392–1407.

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347–356.

Ronayne, D., & Brown, G. D. (2016). Multi-attribute decision by sampling: an account of the attraction, compromise and similarity effects.

Rynes, S. L., Schwab, D. P., & Heneman, H. G. (1983). The role of pay and market pay variability in job application decisions. *Organizational Behavior and Human Performance*, *31*, 353–364.

Schwartz, O., & Simoncelli, E. P. (2001a). Natural signal statistics and sensory gain control. *Nature neuroscience*, *4*(8), 819.

Schwartz, O., & Simoncelli, E. P. (2001b). Natural sound statistics and divisive normalization in the auditory system. In *Advances in neural information processing systems* (pp. 166–172).

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *Journal of neurophysiology*, *86*(4), 1916–1936.

Shannon, C. E., & Weaver, W. (1949). *The mathematical theory of communication.* University of Illinois Press.

Simonoff, J. S., & Tutz, G. (2000). Smoothing methods for discrete data. *Smoothing and regression: Approaches, computation, and application*, 193–228.

Simpson, C. A., & Vuchinich, R. E. (2000). Reliability of a measure of temporal discounting. *The Psychological Record*, *50*, 3–16.

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in neurosciences*, *27*(3), 161–168.

Smith, P. L., Ratcliff, R., & Wolfgang, B. J. (2004). Attention orienting and the time course of perceptual decisions: Response time distributions with masked and unmasked displays. *Vision research*, *44*(12), 1297–1320.

Stevens, J. C. (1958). Stimulus spacing and the judgment of loudness. *Journal of Experimental Psychology*, *56*(3), 246.

Stewart, N. (2009). Decision by sampling: The role of the decision environment in risky choice. *The Quarterly Journal of Experimental Psychology*, *62*, 1041–1062.

Stewart, N., Chater, N., & Brown, G. D. (2006). Decision by sampling. *Cognitive Psychology*, *53*, 1–26.

Stewart, N., Chater, N., Stott, H. P., & Reimers, S. (2003). Prospect relativity: how choice options influence decision under risk. *Journal of Experimental Psychology: General*, *132*, 23–446.

Stewart, N., Reimers, S., & Harris, A. J. (2014). On the origin of utility, weighting, and discounting functions: How they get their shapes and how to change their shapes.

Management Science, *61*, 687–705.

Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, *25*(3), 251–260.

Thorpe, S., Delorme, A., & Van Rullen, R. (2001). Spike-based strategies for rapid processing. *Neural networks*, *14*(6), 715–725.

Thorpe, S., & Gautrais, J. (1998). Rank order coding. In *Computational neuroscience* (pp. 113–118). Springer.

Tremblay, L., & Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. *Nature*, *398*(6729), 704.

Tripp, J., & Brown, G. D. (2016). Being paid relatively well most of the time: Negatively skewed payments are more satisfying. *Memory & Cognition*, *44*, 966–973.

Ungemach, C., Stewart, N., & Reimers, S. (2011). How incidental values from the environment affect decisions about money, risk, and delay. *Psychological Science*, *22*, 253–260.

VanRullen, R., Guyonneau, R., & Thorpe, S. J. (2005). Spike times make sense. *Trends in neurosciences*, *28*(1), 1–4.

VanRullen, R., & Thorpe, S. J. (2001). The time course of visual processing: from early perception to decision-making. *Journal of cognitive neuroscience*, *13*(4), 454–461.

Verlegh, P. W., Schifferstein, H. N., & Wittink, D. R. (2002). Range and number-of-levels effects in derived and stated measures of attribute importance. *Marketing Letters*, *13*(1), 41–52.

Vlaev, I., Chater, N., Stewart, N., & Brown, G. D. (2011). Does the brain calculate value? *Trends in Cognitive Sciences*, *15*, 546–554.

Wainwright, M., Schwartz, O., & Simoncelli, E. (2002). Natural image statistics and divisive normalization: Modeling nonlinearity and adaptation in cortical neurons. In *Probabilistic models of the brain: Perception and neural function.* MIT Press.

Walasek, L., & Stewart, N. (2015). How to make loss aversion disappear and reverse: Tests of the decision by sampling origin of loss aversion. *Journal of Experimental Psychology:*

General, *144*, 7–11.

Wald, A. (1947). *Sequential analysis.* New York: Wiley.

Wald, A., & Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, *19*(3), 326–339.

Woodford, M. (2014). Stochastic choice: An optimizing neuroeconomic model. *The American Economic Review*, *104*(5), 495–500.

Woodford, M. (2016). *Optimal evidence accumulation and stochastic choice.* unpublished manuscript, Columbia University.