Reference effects on decision-making elicited by previous rewards

Francesco Rigoli[1]

[1] *City, University of London, Northampton Square, London, EC1V 0HB, UK*

Correspondence:

Dr. Francesco Rigoli

City, University of London

Northampton Square,

EC1V 0HB, London (UK)

francesco.rigoli@city.ac.uk

ABSTRACT

Substantial evidence has highlighted reference effects occurring during decision-making, whereby subjective value is not calculated in absolute terms but relative to the distribution of rewards characterizing a context. Among these, within-choice effects are exerted by options simultaneously available during choice. These should be distinguished from between-choice effects, which depend on the distribution of options presented in the past. Influential theories on between-choice effects include Decision-by-Sampling, Expectation-as-Reference and Divisive Normalization. Surprisingly, previous literature has focused on each theory individually disregarding the others. Thus, similarities and differences among theories remain to be systematically examined. Here we fill this gap by offering an overview of the state-of-the-art of research about between-choice reference effects. Our comparison of alternative theories shows that, at present, none of them is able to account for the full range of empirical data. To address this, we propose a model inspired by previous perspectives and based on a logistic framework, hence called *logistic model of subjective value*. Predictions of the model are analysed in detail about reference effects and risky decision-making. We conclude that our proposal offers a compelling framework for interpreting the multifaceted manifestations of between-choice reference effects.

Keywords: decision-making; risk; reference effect; context effect; reward prediction error

1. INTRODUCTION

Consider an individual who, right before closing a deal for buying a flat, discovers that the price of the flat is £10 more than expected. Compare this with a case of a person who, before buying a coffee, is asked to pay £10 more than predicted. Apparently, an objectively equivalent extra-cost is experienced by both individuals. However, many would agree that the second person will be more upset. Examples like this describe well how our emotions, feelings, and behaviours are not just determined by the objective circumstances we are experiencing, but also by the frame we use to interpret these circumstances. Helson was one of the first scholars proposing that our judgements are rarely constructed in absolute terms, but they are usually built relative to a frame of reference (Helson, 1948; see also the work of New Look psychology for a similar perspective; e.g., Bruner & Goodman, 1947). For example, they arise from a comparison with the average stimulus encountered in the past. Similar ideas have flourished in research on perception, emotion, social relations, and affect (e.g., Cash et al., 1983; Clifford, 2002; Crosby, 1976; Diener et al., 2009; Sherman et al., 1978; Thompson, 1981; Webster, 2015). Different terminologies are sometimes used across domains, but the processes involved are analogous as all describe judgments that are formed based on the position of a stimulus relative to some other stimuli in a context. Hence, the general term *reference effect* appropriately characterizes the common nature of the different processes described in the literature.

One domain where reference effects have a dramatic impact is when individuals are inferring the subjective value (sometimes referred to also as incentive value or utility) of stimuli, and have to rely on this to make choice. The consideration that reference effects play a critical role in decision-making is already implicit in classical economic theories such as Expected Utility Theory (EUT; Camerer et al., 2011; Von Neumann & Morgenstern, 1944). However, an explicit and systematic treatment of reference effects was first proposed by a psychological account, namely Prospect Theory (PT; Camerer et al., 2011; Kahneman & Tversky, 1979), where the concept of reference point (intended as the status quo) plays a prominent role. Building on this and analogous ideas, theoretical and empirical work has probed reference effects further, for example examining the role of expectations (Hunter & Gershman, 2018; Kőszegi & Rabin, 2006; Rigoli et al., 2016a; 2017) memory retrieval (Brown & Matthews, 2011; Stewart, 2009; Stewart et al., 2006; 2015), and efficient coding (Louie et al., 2013; 2014; 2015; Rangel & Clithero, 2012). Moreover, analogous effects have been observed during multiattribute choice, generating conspicuous theoretical debate (Huber at al., 1982; Noguchi & Stewart, 2018; Rigoli et al., 2017; Roe et al., 2001; Ronayne & Brown, 2017; Simonson & Tversky, 1992; Soltani et al., 2012; Trueblood et al., 2014; Tsetsos et al., 2010; Tversky, 1972).

A useful taxonomy of reference effects in decision-making can be built by considering two dichotomies (Louie et al., 2015; Simonson & Tversky, 1992; Rigoli et al., 2017). One views between-choice effects (sometimes referred to as temporal effects (Louie et al., 2015) or background effects (Simonson & Tversky, 1992)), elicited by stimuli or options presented in previous trials, as opposed to within-choice effects (sometimes referred to as spatial effects (Louie et al., 2015)), which depend on stimuli and options currently available. The second dichotomy opposes multiattribute decision-making, occurring when a trade-off between multiple attributes is required, versus non-multiattribute contexts, where such trade-off is not at play. The combinations derived from the dichotomies generate four possible categories. The category that has attracted more attention comprises effects in within-choice and multiattribute contexts. On this, sophisticated theories exist that have been discussed in depth elsewhere (Noguchi & Stewart, 2018; Roe et al., 2001; Ronayne & Brown, 2017; Simonson & Tversky, 1992; Soltani et al., 2012; Trueblood et al., 2014; Tsetsos et al., 2010). The category comprising within-choice and non-multiattribute effects and the category comprising between-choice and multiattribute effects remain poorly investigated (Louie et al., 2013; Rustichini et al., 2017; Simonson & Tversky, 1992; Vlaev et al., 2009). The last category comprises between-choice and non-multiattribute effects. Especially in recent years, these have received substantial attention (e.g., Louie et al. 2014; Stewart et al., 2015; Rigoli et al., 2016c; 2016d), yet a systematization of the different theories and empirical evidence is lacking.

The goal of the present paper is to discuss the state-of-the-art of research about reference effects in between-choice and non-multiattribute contexts (other forms of reference effects are not examined here), and to contribute to extend this research further. First, we describe between-choice reference effects and how they have been documented empirically. Also, we discuss how decision-making models interpret between-choice reference effects, with a specific focus on classical theories such as EUT and PT, and on recent proposals such as Expectation-as-Reference models (EaR) (Hunter & Gershman, 2018; Kőszegi & Rabin, 2006; Rigoli et al., 2016a; 2017), Decision-by-Sampling theory (DbS) (Brown & Matthews, 2011; Stewart, 2009; Stewart et al., 2006; 2015; Walasek & Stewart, 2015; Vlaev, 2018; Vlaev et al., 2011), and Divisive Normalization theory (DNT) (Louie et al., 2013; 2014; 2015; Rangel & Clithero, 2012). It is surprising that, so far, previous literature has largely considered each of these models in isolation, and that a systematic analysis of their similarities and differences is absent. To address this gap, we compare predictions of contemporary perspectives with respect to their fit with empirical data. This examination reveals that, so far, none of the available models fully fit with available evidence. To address this, we propose a model that integrates previous accounts and aims at providing a comprehensive explanation of all available empirical data. The theoretical and empirical implications of the model are discussed, and we conclude that it provides an improved description of available evidence.

2. RESEARCH ON BETWEEN-CHOICEN REFERENCE EFFECTS

2.1 Empirical observations

Reference effects that depend on stimuli encountered in the past (i.e., between-choice effects) have been first documented in perceptive processes (e.g., Clifford, 2002; Jerger, 1957; Khon, 2007;

Schweinberger et al., 2008; Thompson, 1981; Webster, 2015; Watkinson et al., 2013). However, a large body of evidence supports their critical role also in emotional, affective, and social domains (e.g., Cash et al., 1983; Crosby, 1976; Diener et al., 2009; Sherman et al., 1978). For example, whether a facial stimulus is perceived as showing anger depends on whether preceding stimuli display angry faces or not (Webster et al., 2004). Specifically, the likelihood of perceiving anger increases if previous stimuli do not show angry faces. Analogous phenomena have been documented in social domains. For example, research has found that the self-reported well-being depends on a comparison between the current and the past level of wealth (Boyce et al., 2010; Brown et al., 2008; Diener et al., 1999; 2009; Rutledge et al., 2014) (or the wealth of other people in the reference group such as neighbours, colleagues, and individuals sharing the same ethnic origin (e.g., Crosby et al., 1976)).

Given the ubiquity of reference effects in affective processes, it is not surprising that they have emerged also during value attribution and choice. For example, a field study (Simonsohn & Loewenstein, 2006) has reported that movers arriving from more expensive cities rent pricier apartments than those arriving from cheaper cities (presumably, here previous cities play the role of context). Research has attempted to shed light on the nature of this sort of effects. To this aim, recent laboratory experiments have systematically manipulated the features of the distribution of reward presented in previous trials (a contextual reward distribution), and have assessed the impact of this manipulation on value attribution and choice (Hunter & Gershman, 2018; Rigoli et al., 2016a; 2016b; 2016c; 2018; Stewart et al., 2015; Walasek & Stewart, 2015). The simplest manipulation has targeted the average of the distribution. Evidence has shown that the same reward is attributed a higher incentive value when the contextual average is lower (e.g., Rigoli et al., 2016b; 2016c) (an influence sometimes referred to as contrast effect (Simonsohn & Loewenstein, 2006)). For example, a recent study (Rigoli et al, 2016c) has compared choice behaviour for a low-average context including £1, £3, and £5 as amounts against a high-average context including £3, £5, and £7 as amounts (fig 1a); note that £3 and £5 are common to both contexts. Data indicated that participants attributed higher subjective value to common amounts during the low-value context. More recently, researchers have observed that, in addition to the average, the variability of the distribution also affects value attribution and choice (Rigoli et al., 2016a). A recent study (Rigoli et al., 2016a) has compared choice behaviour for a low-variability context including £3 and £4 as amounts against a high-variability context including £2, £3, £4, and £5 as amounts (fig 1b); note that £3 and £4 are common to both contexts. When comparing the low-variability versus high-variability context, data indicate that participants attributed higher subjective value to £4, but lower subjective value to £3. This finding suggests that, with lower variability, distances across rewards amounts are magnified in such a way that, subjectively, two different reward amounts will appear as farther apart from each other.

In short, research has revealed at least two forms of reference effects on value attribution and choice, one dependent on the average and the other on the variability of the contextual reward distribution. A theory should be able to explain these two forms of effects. Below, we review classical and contemporary perspectives on reference effects in decision making and assess their ability to explain the two empirical effects introduced here.

FIG 1 around here

2.2 Theories

EUT is the most influential economic theory of decision-making, based on the simple idea that subjective utility (which is the variable driving choice) is a concave function of wealth (Camerer et al., 2011; Von Neumann & Morgenstern, 1944). An implication of this model is that, to produce the same utility experience, the increase of wealth required depends on the initial level of wealth. In other words, if the initial wealth is small, the same wealth increase will obtain a higher utility compared to when the initial wealth is large. This property of the model can be viewed as an implicit attempt to conceptualize reference effects, being the notion of initial wealth close to the notion of contextual average reward as formulated by research on reference effects. However, EUT presents at least two shortcomings. First, while analogies may be proposed between initial wealth and contextual average reward, EUT does not provide any insight for explaining an influence of contextual variability. Second, reflecting the actual goods and assets owned by an individual, the concept of wealth refers to an objective quantity and not to subjective beliefs. On the contrary, empirical data have emphasized the psychological nature of reference effects (Simonsohn & Loewenstein, 2006). For example, in the experiment on movers mentioned above (showing that movers arriving from more expensive cities rent pricier apartments than those arriving from cheaper cities), the reference effect observed was not mediated by the individuals' objective wealth (Simonsohn & Loewenstein, 2006). In other words, it was the psychological context, consisting in the representation of the apartments' prices, that counted, and not whether an individual was objectively richer or poorer.

An explicit emphasis on the psychological context characterizes PT, which is the most influential theory of decision-making proposed by psychologists (Camerer et al., 2011; Kahneman & Tversky, 1979). This framework is an eminent example of a model created primarily for explaining reference effects. Inspired by Helson's adaptation level model (Helson, 1948), PT proposes that the subjective value of a stimulus is not based on its absolute properties, but it is constructed by comparing the stimulus against a reference point. This allows the model to distinguish between gains and losses, experienced when the stimulus is better and worse than the reference point, respectively. The idea of reference point is analogous to the notion of contextual average reward, providing an elegant interpretation of reference effects elicited when the contextual average reward varies. However, two important shortcomings can also be identified for PT. First, as for EUT, PT is not sufficient to account for reference effects dependent on contextual variability. A second shortcoming is the notion that the reference point corresponds to the status quo. In other words, the reference point in

PT is conceived as corresponding to the current condition as assessed by an individual. This assumption has been criticized by EaR models for reasons described below.

EaR models rely on a perspective analogous to PT (Hunter & Gershman, 2018; Kőszegi & Rabin, 2006; Rigoli et al., 2016a; 2017). However, they do not interpret the frame of reference as the status quo, but as reflecting expectations about upcoming stimuli (Kőszegi & Rabin, 2006). This distinction is important, as sometimes individuals might believe that the status quo will remain invariant, but other times they might expect the future to be different. According to EaR, it is the beliefs about the future, and not the beliefs about the status quo, that play the role of frame of reference to which rewards are compared once they are experienced. In line with PT, a first version of EaR models conceives the frame of reference as a single point (Kőszegi & Rabin, 2006). According to this model, the subjective value of a stimulus corresponds to a reward prediction error (RPE), namely to the difference between the reward obtained and the reference point or expected reward. This enables the model to explain reference effects due to manipulations of the contextual average reward. However, this EaR model is unable to explain reference effects emergent when the contextual variability is manipulated. To address this shortcoming, a more recent EaR model has adopted a Bayesian account in which the reference frame corresponds to a Gaussian distribution, and therefore it is defined by the two parameters of average and variance (Rigoli et al., 2016a; 2017). In this way, the subjective value of a stimulus does not correspond simply to the difference between the reward and the reference point (as in PT and in previous EaR accounts), but it corresponds to this difference *weighted* by the variance, or to a weighted RPE. In its simplest form (assuming that the prior variance is equal to one; see Rigoli 2016a; 2017), the subjective value V_R associated with a reward amount R is computed similarly to a z-score and corresponds to:

$$V_R = \frac{R - \mu}{\sigma^2} \quad (1)$$

Where the parameter μ corresponds to the expected reward (usually equal to the average of the contextual reward distribution), and the parameter σ^2 corresponds to the uncertainty (usually calculated as the variance of the distribution). In line with empirical evidence, in the model a smaller variance corresponds to a larger weight of the RPE, implying that perceived differences among reward amounts will be magnified. This aspect allows the model to explain, in addition to effects dependent on the contextual average, also effects elicited by the contextual variability.

While EaR models adopt a perspective analogous to PT, a radically alternative approach characterizes DbS (Brown & Matthews, 2011; Stewart, 2009; Stewart et al., 2006; 2015; Walasek & Stewart, 2015; Vlaev, 2018; Vlaev et al., 2011). This theory can be described at two different levels. At a more abstract level, DbS proposes that relative ranking is the principle driving computation of subjective value. In other words, the subjective value V_R of a stimulus *R* depends on its ranking within a contextual stimulus set (calculated by counting how many stimuli are considered as worse than the target stimulus) and on the number of stimuli in the set *n*:

$$V_R = \frac{rank(R) - 1}{n - 1} \tag{2}$$

For example, in a set of eleven objects, a stimulus better than three alternatives (hence having ranking equal to four) will be associated with $V_R = 0.3$. This elegant idea is sufficient to explain both reference effects dependent on the contextual average and those dependent on the contextual variability. The higher level of analysis proposed by DbS is consistent with a lower level where the theory describes the fine-grained cognitive processes engaged. Specifically, when a set of options is offered, each option elicits retrieval from memory (in the form of random sampling) of stimuli encountered in the past, especially those associated with the current context. A set of binary comparisons follows between the option and the samples, and the number of comparisons in which the option is favoured over each sample is recorded. This number corresponds to the subjective value of the option and is computed for all options available, hence determining their relative preference. Since samples are drawn from memory, they depend on past experience and therefore reflect the distribution of options and outcomes characterizing the environment or context, providing the conditions for reference effects to emerge.

To understand DbS, it is helpful to consider a more general framework called Range-Frequency Theory (RFT) (Parducci, 1965; 1995). RFT is a highly influential model applied to a variety of domains, from perception to affect, regarding how we build our judgements about magnitudes. The underlying idea is that judgements derive from integrating a ranking influence similar to DbS (see equation 2) combined with an influence of the range of the contextual distribution. Applied to the calculation of subjective value, RFT can be formulated as (Brown & Matthews, 2011):

$$V_R = w\left(\frac{rank(R) - 1}{n - 1}\right) + (1 - w)\left(\frac{R - min}{max - min}\right) \tag{3}$$

Where *max* reflects the largest reward amount within a set of *n* amounts characterizing a given context, *min* reflects the smallest amount, and where *w* (bounded between zero and one) represents the relative weight of the ranking component over the range-related component. The weight *w* is multiplied by the relative ranking of *R* (see equation 2), and the weight *1-w* is multiplied by the range-related component.

A third contemporary perspective on reference effects is offered by DNT (Louie et al., 2013; 2014; 2015; Rangel & Clithero, 2012). Divisive normalisation was initially proposed in the sensory domain to explain phenomena such as neural adaptation within the retina to stimuli of varying intensity (Carandini & Heeger, 2012). This framework has been generalized to describe higher-order cognitive processes such as selective attention and perceptual decision-making (Carandini & Heeger, 2012; Cheadle et al., 2014). Its overarching principle is the notion of efficient coding, in other words the maximization of mutual information between the neural signalling and the statistics of the environment. Assuming noise in neural responding, it is postulated that an optimal signal-to-noise ratio is achieved if the sensitivity of neurons is tuned to the distribution of stimuli within a context. The specific way DNT realizes efficient coding is by adjusting the neuronal gain based on the average of the contextual distribution. Recently, this framework has been extended to interpret reference effects in value-based decision-making (Louie et al., 2013; 2014; 2015; Rangel & Clithero, 2012). With regard to between-choice reference effects (Louie et al., 2014; 2015), it has been proposed that the subjective value V_t of a reward amount R_t presented at trial *t* is obtained by:

$$V_t = V_{MAX} \frac{R_t + \beta}{\omega + (\Sigma_{j=1}^{t-1} \gamma^j R_j)/(t-1)} \tag{4}$$

Where V_{MAX}, β and ω are constant. The critical aspect of this proposal is that the denominator includes the mean of the rewards encountered in the past, calculated by weighting each reward by its temporal delay according to an exponential discounting governed by the parameter $0 < \gamma < 1$. This enables the model to explain reference effects elicited by manipulations of the contextual average reward. Note that equation 3 does not implement any influence of the variance of the distribution, and thus it does not explain reference effects exerted by the contextual variance. However, DNT offers a promising framework, given its interest in the notion of efficient coding and its explicit connections with neurophysiological processes.

In short, we have identified three main contemporary models that aim at explaining reference effects occurring in between-choice and non-multiattribute contexts. Below we consider the models' predictions in more detail, and assess their fit with empirical data.

2.3 Effects of skewness

For research on reference effects occurring in between-choice and non-multiattribute contexts, an important step is now to compare contemporary theories in terms of their specific predictions and fit with empirical data. Here, we consider such comparison. To this aim, we consider all between-choice non-multiattribute effects reported by empirical research we are aware of, and we ask whether each of these effects is compatible with predictions arising from the different models. This approach is common in the literature investigating between-choice non-multiattribute effects (e.g., Stewart et al., 2015). Note this is different from an approach where model free parameters are fitted to data and models are compared based on model-fitting indexes. The latter approach is not applicable here, because some of the models (such as DbS; see equation 2) have no free parameters (see also Palminteri et al., 2017).

Above, our focus has been on the influences of the contextual average and variability. We have seen that all models considered here explain the effect exerted by contextual average, but only the Bayesian version of EaR, DbS and RFT fit with an effect of contextual variability. Another important characteristic of a contextual reward distribution which has been recently examined is the skewness. That the latter exerts an influence is a specific prediction of DbS and RFT which is not shared by EaR and DNT. To our knowledge, to date one study alone has examined reference effects dependent on skewness (Stewart et al., 2015; experiment 1A, 1B and 1C). In one experiment of this study (1A), participants were offered monetary amounts ranging from £10 to £500, but one group experienced more often smaller amounts than a second group (amounts included £10, £20, £50, £100, £200, £500 for the first group and £10, £310, £410, £460, £490, £500 for the second group), resulting in a positively skewed distribution for the first group and a negatively skewed distribution for the second group (Stewart et al., 2015). Empirically, the subjective value attributed to these amounts was described by a concave function for the group associated with negative skew, and by a convex function for the group associated with positive skew (fig. 1c; similar observations emerged from study 1B and 1C of Stewart et al., 2015). These findings fit nicely with predictions of DbS and RFT, but they are hard to explain by EaR and DNT.

These findings raise an important question of whether effects of contextual average can always be explained by relative ranking as proposed by DbS. This question can be addressed by examining the same experiment (Stewart et al., 2015; Experiment 1A). Relying on the notion that ranking is the unique factor determining context effects, DbS predicts that in the experiment of Stewart et al. (2015) £10 and £500 (which are the two amounts common for both contexts) will be considered as equally valuable across contexts. This is because in both contexts these are the worst and the best amounts encountered, respectively, and the number of amounts is equal across contexts. The same prediction arises from RFT because the two contexts have equivalent amount range and because £10 and £500, in addition to having equivalent relative ranking, also occupy the same position within the amount range. On the contrary, EaR and DNT predict a higher subjective value will be attributed to £10 and £500 in the positive-skew compared to negative-skew context. This is because, in that experiment, the positive-skew context is also characterized by lower contextual average. Although this question was not addressed in the paper of Stewart et al. (2015), the associated empirical data are freely available as supplementary material of that paper, allowing us to address this question here. In experiment 1A of Stewart et al. (2015), on each trial participants (positive-skew group: n = 22; negative-skew group: n = 19; groups are obtained after discarding four participants who failed 10% of catch trials or more; see Stewart et al., 2015) were presented with a choice between two option A and B, being option A associated with either amount A_A with p_A or with amount zero otherwise, and option B associated with either amount A_B with p_B or with amount zero otherwise (see Stewart et al., 2015, for details). Following an approach similar to Stewart et al. (2015), we examined choice data (excluding catch trials from analysis as in Stewart et al., 2015) to assess the subjective value attributed by participants to different amounts. Specifically, for each participant, each specific amount *h* was associated with a free parameter ω_h estimated from choice data (hence, there were six free parameters overall for each participant). Considering a decision between £10 with p_A versus £500 with p_B as an example, choice was described by the following logistic regression:

$$log\left[\frac{P(choice\ A)}{1-P(choice\ A)}\right] = p_A\omega_{10} - p_B\omega_{500} \tag{5}$$

Note that the model fitted to data is unconstrained and does not correspond to any specific model such as DbS. Hence the model is suitable to assess predictions of reference effects models such as DbS in a qualitative way. Given our interest in comparing the subjective value for £10 and £500 across contexts, a two-way mixed ANOVA having context as between-subject factor and amount as within-subject factor was run (fig. 1c). This revealed a main effect of context ($F(1,35) = 6.56$, $p = 0.015$, $\eta_P^2 = 0.158$), indicating that subjective value for £10 and £500 was higher in the positive-skew compared to the negative-skew context. A main effect of amount also emerged ($F(1,35) = 82.66$, $p < 0.001$, $\eta_P^2 = 0.703$) (indicating that subjective value was higher for £500 compared to £10), with no interaction ($F(1,35) = 0.281$, $p = 0.599$, $\eta_P^2 = 0.008$). A larger subjective value in the positive-skew compared to the negative-skew context does not fit with DbS and RFT predictions. This is because the ranking of these two amounts and their position within the amount range do not change across contexts. On the contrary, EaR and DNT fit with this observation because the positive-skew context is also characterized by lower contextual average.

Note that the study of Stewart et al. (2015) includes other experiments (1B and 1C in Stewart et al., 2015) were the skewness of the distribution was manipulated. Also in these experiments the authors report a concave value function for the group associated with negative skew, and by a convex function for the group associated with positive skew. However, contrary to experiment 1A, these experiments are not suited for assessing any effect of average which is independent of relative ranking and of the relative position within the range. This is because in those experiments contexts had different numbers of amounts (e.g., in experiment 1B one context had six different amounts and the other context had seven different amounts) and hence the relative ranking and the relative position within the range of the best and worst amounts is not equal across contexts as in experiment 1A.

In summary, current EaR accounts fail to explain effects due to the skewness of the contextual distribution, and DNT has problems in accounting for effects elicited by either skewness or variability. Although DbS is able to explain both effects due to variability and skewness, the model sometimes appear to fail to account for the influence exerted by the average. The whole picture suggests that none of the theories considered so far is able to fully account for available empirical evidence. This motivated us to search for a model able to explain effects of average (beyond simple ranking), variability and skewness. Below we present a model which satisfies these requirements.

4. A LOGISTIC MODEL OF SUBJECTIVE VALUE

Our proposal aims primarily at explaining the full range of empirical data about reference effects in between-choice and non-multiattribute conditions. However, it is important to evaluate the proposal also outside this domain, especially in the realm of decision-making under risk, something which is discussed below. Our proposal builds on previous frameworks including PT, EaR, DbS and DNT, and aims at integrating them. Indeed, with just minimal adjustments, our model can be meaningfully interpreted within the frame of each previous theory. Analogies with the latter are highlighted below.

The basic idea of the model is that the calculation of the subjective value V_R associated with reward amount R depends on the following logistic function (hence the name Logistic Model of Subjective Value (LMSV)):

$$V_R = \frac{1}{1 + e^{-\lambda \frac{R-\mu}{\sigma}}} \qquad (6)$$

The model includes three parameters. As in equation 1, μ and σ are the expected reward and the uncertainty, respectively. If, in a given context, an agent is exposed to a sequence of reward amounts, the parameter μ can be usually interpreted as the average amount and the parameter σ as the SD. The parameter λ (being $\lambda > 0$) is a constant (i.e., it does not change with context) which can be used to capture spontaneous fluctuations, or individual differences, in choice stochasticity.

The proposal of a logistic function is primarily motivated by its simplicity, which (similar to a previous EaR model described in equation 1) allows one to interpret subjective value straightforwardly as a z-scored RPE constrained to be between zero and one. A similar idea has been proposed for studying of perception (De Gardelle & Summerfield, 2011; Vandormael et al., 2017), and also to characterize decision-making (Woodford (2012) has proposed an account with many analogies to the one developed here). In addition, it turns out that relying on a logistic function offers several advantages, which are highlighted below when comparing LMSV with previous theories.

The similarities between LMSV and PT are evident. In both, a prominent role is played by the reference point. In PT the reference point reflects the perception of the status quo, while, following EaR accounts, in LMSV it corresponds to the reward expected within a given context. However, in analogy with DbS, LMSV treats gains and losses as two separate attributes, and in this way it is able to account for reference effects affecting decision-making under risk (see below). In addition, LMSV does not view subjective value as an unbounded quantity, but as varying between zero and one. As in the EaR model described in equation 1, an important feature implemented in LMSV, but absent in PT, is a modulation of the steepness of the function, which is realized by weighting the RPE with the SD. We will see that this endows LMSV with the capacity of accounting for effects dependent on the variability of the contextual distribution. In short, LMSV can be viewed as an extension of PT built to explain a wider variety of reference effects found empirically (see below).

Clear similarities also exist between LMSV and EaR models, especially the version described in equation 1. Essentially, both adopt calculations similar to z-scoring, where a RPE score is divided by a measure of variability. A key difference is the use of a sigmodal function in LMSV, implying that the subjective value is bounded between zero and one. A critical consequence of this is that the function mapping the reward amount to subjective value is not linear as in equation 1. We will see that such non-linearity enables LMSV to explain effects of skewness, and not only effects of average and variability as previous EaR accounts do. Altogether, LMSV extends previous EaR models by relying on a logistic function to explain also effects of skewness.

We can also interpret LMSV within a DbS framework, modified with some specific assumptions. In a way analogous to DbS, LMSV can be conceived as relying on the notion that subjective value results from binary comparisons between a current reward and fictive reward samples. However, LMSV requires assuming that these samples are drawn from a logistic distribution described by the probability density function:

$$f(R, \mu, \sigma) = \frac{e^{-\lambda \frac{R-\mu}{\sigma}}}{\sigma \left(1 + e^{-\lambda \frac{R-\mu}{\sigma}}\right)^2} \tag{7}$$

This derives because equation 7 corresponds to the cumulative distribution function (CDF) of equation 6 (the same parameters are present in both equation 6 and 7) (Bhui & Gershman, 2018). Hence, processes including sampling and binary comparison, initially proposed in the context of DbS, are compatible as well with the computations underlying LMSV.

Finally, LMSV has important analogies with DNT. The latter was originally proposed as a description of how the brain realizes efficient coding, consisting in maximizing the mutual information between neuronal spiking rates and the distribution of relevant environmental stimuli (Carandini & Heeger, 2012). A similar rationale can be proposed for LMSV, provided we assume that the brain embodies believes that the probability of reward amounts follows the logistic distribution described by equation 7. This because equation 6, which calculates the subjective value according to LMSV, corresponds to the CDF associated with equation 7, and in general CDFs have been proved to realize efficient coding (Bhui & Gershman, 2018; Laughlin, 1981). In other words, spiking rates that reflect the subjective value as calculated by LMSV are consistent with efficient coding, under the assumption of a logistic distribution. We emphasize that, to establish a link between LMSV and efficient coding, brain regions implicated in the computation of subjective value must represent logistic distributions (note that this implication is not necessary for other regions, for example those implicated in abstract reasoning). This prediction remains to be empirically tested, as it is not yet fully understood in detail how subjective value is represented in the brain. In short, although relying on different algorithms and therefore implying distinct predictions, DNT and LMSV both represent possible realizations of efficient coding in the brain.

In sum, we have introduced a simple model of subjective value aimed at explaining a large set of reference effects. This model integrates features of previous accounts including PT, EaR, DbS and DNT. Below we explore how the model can account for empirical evidence on reference effects, and we also evaluate its implications for risky decision-making. In Supplementary Material, we explore three other interesting aspects of the new model. First, we derive some novel predictions of the model regarding reference effects. Second, we examine the normative principles potentially supported by the model, consisting in realizing optimal choice adaptation. Third, we extend the model to interpret between-choice-reference effects occurring in multiattribute contexts, namely when several attributes need to be traded-off against each other.

4.1 Reference effects

It is instructive to simulate LMSV to assess the ensuing general predictions. The model includes two parameters μ and σ, which reflect beliefs about the average and variability of the context, respectively. We assessed the impact of varying these parameters. Fig. 2a plots the subjective value as a function of reward amount, comparing contexts characterized by different average, captured by the parameter μ (the parameter σ was kept constant). The figure shows that the same reward amount becomes more valuable when the average contextual reward is smaller. This can potentially explain empirical observations about the effect of contextual average (Rigoli et al., 2016b, 2016c). Fig 2b compares the value function for contexts characterized by equal average (hence the parameter μ was kept constant) but different variability, captured by the parameter σ. The figure shows that the steepness of the function decreases with variability. This can potentially explain empirical data indicating that, with higher compared to lower variability, distances across different reward amounts decrease (Rigoli et al., 2016a).

A critical question is whether LMSV predicts any effect of contextual skewness (Stewart et al., 2015). We can address this question by simulating experiment 1A of Stewart et al. (2015) adopting LMSV. Remember this experiment compared a positive-skew contextual distribution, including the amounts £10, £20, £50, £100, £200, and £500, against a negative-skew distribution including £10, £310, £410, £460, £490, and £500. To model this scenario, we estimated the value function according to LMSV, assuming that the expected reward μ was equal to the contextual average (i.e., $\mu = 147$ and $\mu = 363$, respectively) and the uncertainty was equal to the contextual SD (SD is equal across contexts; hence $\sigma = 187$ for both). Focusing specifically on the £10-£500 range common across contexts, LMSV predicts a concave function for the positive-skew context, and a convex function for the negative-skew context (fig. 2c). Note also that, according to LMSV, the subjective value attributed to £10 and £500 (which are the two amounts common across contexts) is predicted to be higher in the positive-skew compared to the negative-skew context (fig. 2c). This is consistent with the novel analysis of empirical data reported above. These simulations show that LMSV predicts contextual effects of skewness. Ultimately, these effects are predicted as a consequence of changes in the contextual average. In experiments investigating the influence of skewness (Stewart e al., 2015), when the contextual skewness changes, also the contextual average changes (positive and negative skew implicate lower and higher average, respectively). This has important implications when LMSV is adopted. The sigmoid function postulated by LMSV is non-linear and presents one region which is convex (when the subjective value is close to zero) and one region which is concave (when the subjective value is close to one). When the average of the contextual distribution varies, the convex and concave regions will move. Therefore, a specific range of reward amounts may initially map to the concave region of the function. But, when the contextual average changes (specifically, when it increases), the same amount range may now map to the concave region of the function. In other words, in some conditions LMSV predicts that when the contextual distribution has lower average (and positive skew) one amount range may appear as being associated with a concave function; when the contextual distribution has higher average (and negative skew) the same amount range may appear as being associated with a convex function.

In addition to considering simulations, it is important to examine in detail how well LMSV fits with empirical data on between-choice and non-multiattribute reference effects. To this aim, we consider all such effects we are aware of as they have been described in the literature. These include an effect of the average of the contextual reward distribution (Rigoli et al., 2016b, 2016c), an effect of variability (Rigoli et al., 2016a), and one of skewness (Stewart et al., 2015). Fig. 3 shows empirical evidence on between-choice and non-multiattribute reference effects in conjunction with predictions derived from LMSV regarding this evidence. Here, LMSV is implemented by assuming that, for any context condition, μ and σ correspond to the actual average and SD, respectively, characterising that context in the experiment. A comprehensive list of between-choice and non-multiattribute reference effects is also provided in tab. 1. Also, for the other theories of choice discussed above, tab. 1 indicates whether their predictions fit with each specific reference effect. From fig. 3 and tab. 1, we can see that LMSV is able to replicate all empirical observations considered. Specifically, it is consistent with studies on the effect of contextual average (fig. 3a; Rigoli et al., 2016b, 2016c), contextual variability (fig. 3b; Rigoli et al., 2016a), and contextual skewness (fig. 3c; Stewart et al., 2015). When considering experiment 1A of Stewart et al. (2015), LMSV predicts higher subjective value for £10 and £500 in the positive-skew compared to negative-skew context (fig. 3c).

In summary, we argue that our discussion and simulations indicate that LMSV offers a comprehensive explanation of known reference effects in between-choice and non-multiattribute scenarios (see SI for novel predictions arising from LMSV regarding reference effects). In fact, this model predicts effects of skewness which are not contemplated by EaR and DNT, and it explains effects of average in scenarios where these effects emerge empirically but are not predicted by DbS nor by RFT. To date, in between-choice and non-multiattribute domains, we are not aware of any form of reference effect found empirically which does not fit with LMSV. In Supplementary Material, we have identified novel and specific predictions of LMSV about reference effects which remain to be examined by future research.

TAB 1, FIG 2, and FIG 3 around here

4.2 Decision-making under risk

For a theory of subjective value, a fundamental question regards its ability to reproduce empirical data about risk-sensitive decision-making. Here, we ask this question for LMSV, with the aim of offering a broad overview of this new theory. We do not discuss EaR and DNT in this context, though we consider DbS because the latter has offered a new compelling perspective on risky-decision making which also inspires LMSV (Stewart, 2009; Stewart et al., 2006; 2015; Walasek & Stewart, 2015; Vlaev, 2018; Vlaev et al., 2011).

Contemporary literature often relies on PT for interpreting decision-making under risk (Camerer et al., 2011; Kahneman, D., & Tversky, 1979). The specific value function proposed by this theory is concave for gains and convex for losses, hence predicting risk aversion for gains and risk seeking for losses. In addition, the value function is steeper for losses, generating loss aversion, expressed for example in a preference for a null outcome over a 50/50 gamble returning either a monetary amount gained or the same amount lost. A concave function for gains, a convex function for losses, and loss aversion have been reported by several empirical investigations (Camerer et al., 2011; Kahneman, D., & Tversky, 1979).

Although PT has been remarkably successful in explaining empirical data on decision-making under risk (Camerer et al., 2011), recent studies, originally aimed at testing predictions of DbS, have emphasized a new compelling perspective, where reference effects play a central role (Stewart, 2009; Stewart et al., 2006; 2015; Walasek & Stewart, 2015; Vlaev, 2018; Vlaev et al., 2011). This perspective relies on treating gains and losses as two separate attributes, meaning that decisions requiring a trade-off between gains and losses are considered in the realm of multiattribute choice. It follows that, for options characterized by both a potential gain and loss, the former is evaluated with respect to the contextual distribution of gains and irrespective of losses, and the latter with respect to the distribution of losses irrespective of gains. Eventually, the total option value is obtained by summing the subjective value of the gain and the value of the loss.

An analogous approach can be proposed for LMSV. We used simulations to assess the implications of this approach when applied to LMSV. First, we tested predictions for gain and loss contexts separately. While risk aversion and risk seeking have been usually observed empirically for gains and losses, respectively (Camerer et al., 2011; Kahneman, D., & Tversky, 1979), recent evidence raises the possibility that these findings are the product of reference effects (Stewart et al., 2015). This evidence corresponds to the study examined above when discussing reference effects dependent on skewness (Stewart et al., 2015), opposing a positive-skew contextual distribution including amounts £10, £20, £50, £100, £200, and £500 to a negative-skew distribution including £10, £310, £410, £460, £490, and £500. As seen above, a concave function describes subjective value in the positive-skew condition (fig. 1c). This implies risk aversion, as for example a sure gain will be favoured over a 50/50 gamble returning either a double gain or no gain. On the contrary, a convex function characterises subjective value in the positive-skew condition (fig. 1c). This implies risk seeking, as for example a sure gain will be avoided in favour of a 50/50 gamble returning either a double gain or no gain. In other words, the study of Stewart et al. (2015) implicates that risk sensitivity in the gain domain depends on the context. We have already seen that this scenario is captured by LMSV (fig 2c, fig 3c), which predicts a concave and convex function in the negative-skew and positive-skew condition, respectively. Hence this captures context-dependent changes in risk sensitivity as emerged empirically (Stewart et al., 2015).

Whether context affects risk sensitivity in the loss domain remains to be explored empirically. In this domain, LMSV offers a similar interpretation which we describe using simulations (fig. 5a). We simulate a scenario considering a negative-skew contextual distribution including losses -£10, -£20, -£50, -£100, -£200, and -£500, and a positive-skew distribution including -£10, -£310, -£410, -£460, -£490, and -£500. A convex function is obtained by LMSV in the negative-skew condition which implies risk seeking, as for example a sure loss is avoided in favour of a 50/50 gamble returning either a double loss or no loss (fig. 4a). A concave function is obtained in the positive-skew condition which implies risk aversion, as for example a sure loss is favoured over a 50/50 gamble returning either a double loss or no loss (fig. 4a). In short, inspired by a perspective pioneered by DbS, LMSV views risk sensitivity for both gains and losses as the product of reference effects. Therefore, the commonly observed risk aversion for gains and risk seeking for losses are interpreted not as intrinsic phenomena, but as emergent from the fact that, in most ecological contexts, smaller gains and losses are more frequent than larger gains and losses.

Next, we evaluated LMSV in contexts where information about gains and losses need to be integrated. A recent study has shown that reference effects are critical also in this domain (Walasek & Stewart, 2015). This study involved repeated choices between a null outcome and a 50/50 gamble returning either a gain of amount *x* or a loss of equivalent amount -*x*. Crucially, the range of the contextual distribution of gains and losses varied across conditions. Four conditions were compared: one involving a [£0 £20] range for both gains and losses, one having a [£0 £40] range for both gains and losses, a third having a [£0 £40] range for gains and a [£0 £20] range for losses, and finally one characterized by a [£0 £20] range for gains and a [£0 £40] range for losses. Evidence indicated that, when an equal range characterizes gains and losses, the gamble and the sure option were equally preferred, implying an equal weight for gains and losses (fig. 4b). Loss aversion, expressed as an avoidance of the gamble, emerged when the loss range was smaller than the gain range (fig. 4b). The opposite of loss aversion, expressed is favouring the gamble over the sure option, was evident when the loss range was larger than the gain range (fig. 4b).

We simulated this scenario using LMSV, estimating (for each option) the subjective value of the gain and of the loss using their respective contextual average and SD. The total value of the safe option was estimated as the subjective value associated with no gain plus the value associated with no loss. This was subtracted from the total subjective value of the gamble, calculated as the subjective value of the gain x plus the value of the loss $-x$ divided by two (to account for the 50/50 chance). Simulations replicate the empirical findings described above (fig. 4b). Interestingly, simulations also highlight a novel prediction that, when losses and gains have different range, there will be a specific monetary amount x (located between the average of losses and the average of gains) for which the value difference between the gamble and the sure option will be maximal (fig. 4c). In short, for LMSV, differences in the relative weight attributed to losses compared to gains are not intrinsic, but they are the product of reference effects. An implication is that the usual observation of loss aversion is interpreted as the consequence of a smaller variability of losses compared to gains characterizing many ecological contexts (Stewart, 2009; Stewart et al., 2006).

Altogether, the approach to risk sensitivity adopted by LMSV follows the approach proposed originally by DbS (which is extensively discussed elsewhere; Stewart, 2009; Stewart et al., 2006; 2015; Walasek & Stewart, 2015; Vlaev, 2018; Vlaev et al., 2011). The key idea consists in treating losses and gains as distinct attributes, and in viewing risk sensitivity not as intrinsic but as the consequence of reference effects. Although DbS and LMSV share this common perspective, nonetheless it is possible to identify specific predictions distinguishing the two models. For instance, consider the task described above involving choices between a null outcome and a 50/50 gamble returning either a gain *x* or loss -*x*. Keeping the context for gains constant, compare a low-variability context where the distribution of losses includes -£5, -£9, -£11, and -£15 against a high-variability context where the distribution of losses is -£5, -£6, -£14, and -£15. Given the influence of the SD postulated by LMSV, this model predicts that, in the low- compared to high-variability context, the -£5 loss will be perceived as relatively better, and the -£15 as relatively worse. No such prediction derives from DbS. Analogous scenarios remain to be examined empirically.

FIG 4 around here

----

5. GENERAL DISCUSSION

In this paper, we offer a general overview of research about reference effects occurring during decision-making in between-choice and non-multiattribute contexts. Although these effects were implicitly acknowledged already by EUT (Von Neumann & Morgenstern, 1944), their explicit and systematic treatment can be found for the first time in PT (Kahneman & Tversky, 1979). More recently, theoretical and empirical research on these effects has flourished. Although contemporary models already provide a compelling description of the underlying processes, we note that none accounts fully for empirical data. Motivated by shortcomings of previous accounts, we propose LMSV as a model able to explain a broader set of empirical data.

LMSV can be usefully interpreted within the perspective offered by each previous theory. Its core principles are inspired by PT and EaR (Hunter & Gershman, 2018; Kahneman & Tversky, 1979; Kőszegi & Rabin, 2006; Rigoli et al., 2016a; 2017), and consist in conceiving subjective value essentially as a “squashed” z-score. Different from PT and EaR, reliance on a logistic function enables LMSV to implement efficient coding in a similar way as DNT (Louie et al., 2013; 2014; 2015; Rangel & Clithero, 2012), and potentially to fit with memory sampling processes analogous to those proposed by DbS (Brown & Matthews, 2011; Stewart, 2009; Stewart et al., 2006; 2015; Walasek & Stewart, 2015; Vlaev, 2018; Vlaev et al., 2011). Thanks to these substantial connections with previous theories, LMSV represents a unifying framework integrating fundamental insights proposed originally by other accounts.

LMSV can be interpreted as belonging to a broad family of models which rely on a sigmoid function to explain context effects (Cheadle et al., 2014; Juechems et al., 2017; Li et al., 2018). Similar to LMSV, Cheadle at al. (2014) have proposed that perceptive judgements are based on transforming a stimulus through a sigmoid function postulated to change based on past experience. However, in Cheadle et al. (2014) past experience does not change the steepness of the sigmoid function. This is a critical difference from LMSV, where the steepness of the function varies reflecting the variability of the context. This characteristic of LMSV, absent in Cheadle et al. (2014), is critical here because it allows LMSV to explain context effects based on contextual variability. The model of Cheadle et al. (2014) has also been extended to explain phenomena in value-based decision-making such as within-choice context effects (Li et al., 2018) and how an overall propensity to gamble changes based on past reward experience (Juechems et al., 2017). Similar to LMSV, these proposals also rely on a sigmoid function to explain context effects. However, like Cheadle et al. (2014) but contrary to LMSV, they do not allow the steepness of the function to vary (hence they would fail to explain between-choice context effects dependent on variability).

It is important to emphasize that the level of analysis of LMSV is relatively abstract. This implies that LMSV is agnostic on the fine-grained psychological processes underlying computation of subjective value, though it provides constraints for theories speaking to such level of analysis. A possibility on the nature of fine-grained processes is inspired by DbS, and consists in memory sampling and binary comparisons. As discussed above, LMSV is consistent with this type of process, but does not necessarily imply it. A compelling alternative is that the brain directly entertains beliefs about the summary statistics of a contextual reward distribution (e.g., Friston, 2005), in the form of expected reward and uncertainty, and uses these during calculation of subjective value. Hybrid processes are also in line with LMSV, in which summary statistics are integrated with individual memory samples. With this regard, we note that value-based decision-making is a phylogenetically ancient function already present in simple organisms such as insects (Strausfel & Hirth, 2013). These animals' behaviour is driven by basic sensory-motor processes, and lacks higher-order memory capabilities. Therefore, at least for some animals, it is possible that subjective value computation, and associated reference effects observed empirically (Marsh & Kacelnik, 2002), relies primarily on basic summary statistics rather than complex memory sampling. This raises the possibility of a role for summary statistics, alone or together with memory sampling, also in humans.

From a comparison between LMSV and DbS, another important question arises: how does the brain represent distributions of variables? According to DbS, the brain represents the precise distribution of values characterising a context; this distribution can have any shape. On the contrary, according to LMSV, the brain assumes that the distribution of values characterising a context is logistic. This implies that in DbS the brain represents the real-world distribution faithfully, while in LMSV the brain transforms any distribution to a logistic distribution. This raises two questions. First, how can we test empirically whether the distribution is reflected faithfully (as in DbS) or whether it is approximated with a symmetrical bell-shaped function (such as a logistic function; as in LSMV)? Second, why would it be reasonable for the brain to approximate any distribution with a symmetrical bell-shaped function such as a logistic function? The answer to the first question, we argue, requires to tease apart different predictions of models such as DbS and LMSV and to assess their fit with empirical data (something partially done here in the context of between-choice non-multiattribute reference effects). For the second question, we propose the following answer. First, in the real-world variables may be sometimes negatively skewed, sometimes positively skewed, and other times characterised by zero skewness; implying that on average the level of skewness may be around zero. Therefore, given a limited computational capacity, the brain may simply assume a level of skewness equal to the average, namely zero (as in a logistic distribution). Second, neurophysiological evidence supports influential theories such as predictive coding where the brain represents variables adopting a symmetrical (in that case Gaussian) distribution (e.g., Friston, 2005). For all these considerations, it may be reasonable for the brain to approximate any distribution in the real world by assuming a priori that any distribution is a symmetrical bell-shaped function (such as in LMSV). Although we emphasise the plausibility of this reasoning, we also stress that a promising avenue for future research is to explore different functions and to assess their merits and shortcomings with respect to empirical data. With this regard, there is evidence that, in the real world, both gains and losses are distributed according to a power law function (Stewart et al., 2006), hinting to the possibility that the latter may also be a good candidate for explaining reference effects.

Most models (including DbS, DNT, RFT, and LMSV) assume that subjective value is a bounded quantity. This can potentially produce ceiling effects that are problematic. Consider an example where, during previous trials, an agent has been offered options comprising £1, £2, £3 as amounts. Next, at the current trial, an agent is offered an option returning £10000 with 1/100 chance versus an option returning £20000 with 1/100 chance. Because £10000 and £20000 both rank better, and are way higher, than all previous amounts, according to most theories they will both receive the highest subjective value (for DbS, LMSV or RFT, a subjective value equal to one). This predicts indifference for the two options, a prediction which is potentially problematic because real agents are likely to always choose the £20000 option. One potential way to address this is to assume that amounts present in the current trial also exert reference effects (these are forms of within-choice effect). If this is assumed, then £20000 would be evaluated with respect to £1, £2, £3, and £10000; while £10000 would be evaluated with respect to £1, £2, £3, and £20000. A preference for the £20000 option would now be predicted.

An important question is how reference effects arise out of neural mechanisms. Research has identified a specific brain network engaged during computation of subjective value and choice. Within this network, a central role is played by the dopaminergic midbrain (e.g., Glimcher, 2011; Schultz et al., 1997). A large body of evidence has revealed that neural activity in the dopaminergic midbrain is proportional to the subjective value (Schultz et al., 1997). In addition, studies manipulating the contextual reward distribution have found that this neural response is inversely correlated with the expected reward (Schultz et al., 1997). These findings have been interpreted as evidence of an expression of a RPE signal in this region (Glimcher, 2011; Schultz et al., 1997). Moreover, research has manipulated the range of the reward distribution, observing that two different reward magnitudes elicit a more distinct neural response when the range is smaller (Diederen et al., 2016; Tobler et al., 2005). Altogether, these data fit with the notion that a RPE weighted by uncertainty is signalled in the dopaminergic midbrain. Notably, this is analogous to the calculations proposed by LMSV, establishing a connection between this theory and neural processes. A critical empirical question is whether a link exists between neural and behavioural reference effects. Addressing this question, recent studies have shown that reference effects emerge together in dopaminergic midbrain and in choice behaviour (Rigoli et al., 2016b; 2016c; 2018). These findings provide compelling evidence in line with the possibility that behavioural reference effects are supported by neural adaptation processes occurring in the dopaminergic midbrain. Interestingly, a recent study has highlighted also a role for the hippocampus, as engagement of this region favours the emergence of reference effects (Rigoli et al., 2016c). The hippocampus is critical for processing contextual information in several domains (Holland & Bouton, 1999), and this research has extended this role to decision-making. Interestingly, this region is implicated also in memory processing

(Squire, 1992), raising the possibility that it may guide memory sampling underlying reference effects in value computation.

An important distinction in decision-making under risk is between decision from description and decision from experience (e.g., Ludvig & Spetch, 2011). In the former, participants are explicitly informed about the values at stake, while in the latter they have to learn these values from experience. Most empirical studies of between-choice non-multiattribute reference effects have focused on decision from description (Rigoli et al., 2016a; 2016b; 2016c; 2018; Stewart et al., 2015; Walasek & Stewart, 2015). However, we note also studies about decision from experience which have manipulated the contextual average (Hunter & Gershman, 2018; Rigoli et al., 2018). Results of these studies are analogous to those emerged during decision from description. However a systematic exploration of reference effects during decision from experience requires further research.

We have emphasized that the specific scope of this manuscript is research on value-based decision-making focusing on between-choice and non-multiattribute contexts. However, it is useful to briefly discuss the more general implications of our arguments, as reference effects are important also outside decision-making, and, within the latter, also outside the between-choice and non-multiattribute case. With respect to phenomena going beyond decision-making, evidence from some psychophysics contexts support models analogous to DNT, where the average stimulation, independent of the variability, regulates the discriminability of two stimuli (Carandini & Heeger, 2012). For example, if auditory stimuli are on average louder, sounds are usually harder to discriminate, even if their variability is constant (Jerger, 1957). This suggests that models similar to LMSV may be inappropriate to describe some perceptual phenomena, as in LMSV discriminability depends on the variability, and not on the average, of the distribution. However, recent studies about perceptive decision-making have shown that models analogous to LMSV sometimes offer a compelling explanation of reference effects (Cheadle et al., 2014; De Gardelle & Summerfield, 2011;

Vandormael et al., 2017). Further research is needed to establish, outside value-based decision-making, in which conditions processes analogous to those proposed by LMSV are at play.

Research has also examined reference effects affecting subjective judgement. For example, a study has explored how knowledge of other people's income affects satisfaction about the own income (Brown et al., 2008). This study found that the relative ranking of the own income (a concept analogous to DbS) affects satisfaction ratings. However, ranking effects provided an insufficient explanation, because the actual distance from other incomes was also critical. One way to interpret these and similar findings is to rely on RFT (Parducci, 1965; 1995), which combines ranking and range processes. Alternatively, an adaptation of DbS grounded on sophisticated memory retrieval can also explain these and similar observations (Brown & Matthews, 2011). By relying on a parametric framework which, as we have seen here, is capable of explaining some forms of skewness effects, LMSV may emerge as capable to capture both ranking and range influences observed in judgement (and possibly make also specific predictions as in the case of decision-making). This remains an open question for future research.

Within decision-making, the most studied reference effects are those characterizing within-choice and multiattribute contexts, and sophisticated theories have been proposed in this domain (Huber at al., 1982; Noguchi & Stewart, 2018; Rigoli et al., 2017; Roe et al., 2001; Ronayne & Brown, 2017; Simonson & Tversky, 1992; Soltani et al., 2012; Trueblood et al., 2014; Tsetsos et al., 2010; Tversky, 1972). Similar to PT (Kahneman & Tversky, 1979) and early EaR models (Kőszegi & Rabin, 2006), many theories assume that, for each relevant attribute, a stimulus is evaluated relative to a reference point corresponding to the average value of the options available (Roe et al., 2001; Soltani et al., 2012; Trueblood et al., 2014; Tsetsos et al., 2010; Tversky, 1972). Other theories are based on an extension of DbS (Noguchi & Stewart, 2018; Ronayne & Brown, 2017). Finally, one theory proposes an extension of the EaR model described in equation 1, emphasizing the influence of variability, in addition to the average, of the options available (Rigoli et al., 2017). Further theoretical analyses are needed to assess whether the concepts proposed here in relation with LMSV may contribute to improve our understanding of within-choice and multiattribute contexts. Finally, two additional categories of effects exist, between-choice and multiattribute effects on the one hand, within-choice and non-multiattribute effects on the other. While the latter remain to be explored under the framework offered by LMSV, in Supplementary Material we extend LMSV to the former, implying that this theory can be regarded as a general model of between-choice reference effects in decision-making.

In short, we have summarised and critically evaluated research on reference effects during decision-making, focusing on between-choice contexts. Although we emphasize several strengths of previous perspectives on this topic, we have noted shortcomings of each of these perspectives. Integrating important insights originally proposed by previous theories, we offer an alternative model which explains a broader set of data. A question for future research is how general the framework is, both within and outside decision-making.

6. ACKNOWLEDGEMENT

This study has been funded by a British Academy Small Research Grant (SG170621).

7. REFERENCES

1. Boyce, C. J., Brown, G. D., & Moore, S. C. (2010). Money and happiness: Rank of income, not income, affects life satisfaction. *Psychological Science*, *21*(4), 471-475.
2. 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*(3), 355-389.
3. 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*, 299.
4. Bruner, J. S., & Goodman, C. C. (1947). Value and need as organizing factors in perception. *The journal of abnormal and social psychology*, *42*(1), 33.
5. Bhui, R., & Gershman, S. J. (in press). Decision by sampling implements efficient coding of psychoeconomic functions. *Psychological Review*.
6. Camerer, C. F., Loewenstein, G., & Rabin, M. (Eds.). (2011). *Advances in behavioral economics*. Princeton university press.
7. Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, *13*(1), 51.

8. Cash, T. F., Cash, D. W., & Butters, J. W. (1983). "Mirror, Mirror, on the Wall...?" Contrast Effects and Self-Evaluations of Physical Attractiveness. *Personality and Social Psychology Bulletin*, *9*(3), 351-358.
9. Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., De Gardelle, V., Castañón, S. H., & Summerfield, C. (2014). Adaptive gain control during human perceptual choice. *Neuron*, *81*(6), 1429-1441.
10. Clifford, C. W. (2002). Perceptual adaptation: motion parallels orientation. *Trends in cognitive sciences*, *6*(3), 136-143.
11. Crosby, F. (1976). A model of egoistical relative deprivation. *Psychological review*, *83*(2), 85.
12. Cohen, J. D., McClure, S. M., & Angela, J. Y. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *362*(1481), 933-942.
13. Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Routledge.
14. Daw, N. D., O'doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*(7095), 876.
15. De Gardelle, V., & Summerfield, C. (2011). Robust averaging during perceptual judgment. *Proceedings of the National Academy of Sciences*, *108*(32), 13341-13346.
16. Diederen, K. M., Spencer, T., Vestergaard, M. D., Fletcher, P. C., & Schultz, W. (2016). Adaptive prediction error coding in the human midbrain and striatum facilitates behavioral adaptation and learning efficiency. *Neuron*, *90*(5), 1127-1138.
17. Diener, E., Suh, E. M., Lucas, R. E., & Smith, H. L. (1999). Subjective well-being: Three decades of progress. *Psychological bulletin*, *125*(2), 276.
18. Diener, E., Lucas, R. E., & Scollon, C. N. (2009). Beyond the hedonic treadmill: Revising the adaptation theory of well-being. In *The science of well-being* (pp. 103-118). Springer, Dordrecht.
19. Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *360*(1456), 815-836.
20. Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive neuroscience*, *6*(4), 187-214.
21. Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. *Cognition*, *173*, 34-42.
22. Glimcher, P. W. (2011). Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, *108*(Supplement 3), 15647-15654.
23. Helson, H. (1948). Adaptation-level as a basis for a quantitative theory of frames of reference. *Psychological review*, *55*(6), 297.
24. Herrnstein, R. J. (2000). *The matching law: Papers in psychology and economics*. Harvard University Press.
25. Holland, P. C., & Bouton, M. E. (1999). Hippocampus and context in classical conditioning. *Current opinion in neurobiology*, *9*(2), 195-202.
26. Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of consumer research*, *9*(1), 90-98.
27. Hunter, L. E., & Gershman, S. J. (2018). Reference-dependent preferences arise from structure learning. *bioRxiv*, 252692.
28. Jerger, J. F. (1957). Auditory adaptation. *The Journal of the Acoustical Society of America*, *29*(3), 357-363.
29. Juechems, K., Balaguer, J., Ruz, M., & Summerfield, C. (2017). Ventromedial prefrontal cortex encodes a latent estimate of cumulative reward. *Neuron*, *93*(3), 705-714.
30. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 278.
31. Kohn, A. (2007). Visual adaptation: physiology, mechanisms, and functional benefits. *Journal of neurophysiology*, *97*(5), 3155-3164.
32. Kőszegi, B., & Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, *121*(4), 1133-1165.
33. Marsh, B., & Kacelnik, A. (2002). Framing effects and risky decisions in starlings. *Proceedings of the National Academy of Sciences*, *99*(5), 3352-3355.
34. Laughlin, S. (1981). A simple coding procedure enhances a neuron's information capacity. *Zeitschrift für Naturforschung c*, *36*(9-10), 910-912.

35. Li, V., Michael, E., Balaguer, J., Castañón, S. H., & Summerfield, C. (2018). Gain control explains the effect of distraction in human perceptual, cognitive, and economic decision making. *Proceedings of the National Academy of Sciences*, *115*(38), E8825-E8834.
36. 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*, 201217854.
37. 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*(48), 16046-16057.
38. Louie, K., Glimcher, P. W., & Webb, R. (2015). Adaptive neural coding: from biological to behavioral decision-making. *Current opinion in behavioral sciences*, *5*, 91-99.
39. Luce, R. D. (1959). On the possible psychophysical laws. *Psychological review*, *66*(2), 81.
40. Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: is the description-experience gap in risky choice limited to rare events?. *PloS one*, *6*(6), e20262.
41. Noguchi, T., & Stewart, N. (2018). Multialternative decision by sampling: A model of decision making constrained by process data. *Psychological Review, 125(4), 512-544*.
42. Parducci, A. (1965). Category judgment: a range-frequency model. *Psychological review*, *72*(6), 407.
43. Parducci, A. (1995). *Happiness, pleasure, and judgment: The contextual theory and its applications*. Lawrence Erlbaum Associates, Inc.
44. Rangel, A., & Clithero, J. A. (2012). Value normalization in decision making: theory and evidence. *Current opinion in neurobiology*, *22*(6), 970-981.
45. Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological bulletin*, *114*(3), 510.
46. Rigoli, F., Friston, K. J., Martinelli, C., Selaković, M., Shergill, S. S., & Dolan, R. J. (2016a). A Bayesian model of context-sensitive value attribution. *ELife*, *5*, e16127.
47. Rigoli, F., Rutledge, R. B., Dayan, P., & Dolan, R. J. (2016b). The influence of contextual reward statistics on risk preference. *NeuroImage*, *128*, 74-84.
48. Rigoli, F., Friston, K. J., & Dolan, R. J. (2016c). Neural processes mediating contextual influences on human choice behaviour. *Nature communications*, *7*, 12416.
49. Rigoli, F., Rutledge, R. B., Chew, B., Ousdal, O. T., Dayan, P., & Dolan, R. J. (2016d). Dopamine increases a value-independent gambling propensity. *Neuropsychopharmacology*, *41*(11), 2658.
50. Rigoli, F., Mathys, C., Friston, K. J., & Dolan, R. J. (2017). A unifying Bayesian account of contextual effects in value-based choice. *PLoS computational biology*, *13*(10), e1005769.
51. Rigoli, F., Chew, B., Dayan, P., & Dolan, R. J. (2018). Learning contextual reward expectations for value adaptation. *Journal of cognitive neuroscience*, *30*(1), 50-69.
52. Ronayne, D., & Brown, G. D. (2017). Multi-attribute decision by sampling: An account of the attraction, compromise and similarity effects. *Journal of Mathematical Psychology*, *81*, 11-27.
53. Rustichini, A., Conen, K. E., Cai, X., & Padoa-Schioppa, C. (2017). Optimal coding and neuronal adaptation in economic decisions. *Nature Communications*, *8*(1), 1208.
54. Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *Proceedings of the National Academy of Sciences*, *111*(33), 12252-12257.
55. Schweinberger, S. R., Casper, C., Hauthal, N., Kaufmann, J. M., Kawahara, H., Kloth, N., ... & Zäske, R. (2008). Auditory adaptation in voice perception. *Current Biology*, *18*(9), 684-688.
56. Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593-1599.
57. Sherman, S. J., Ahlm, K., Berman, L., & Lynn, S. (1978). Contrast effects and their relationship to subsequent behavior. *Journal of Experimental Social Psychology*, *14*(4), 340-350.
58. Simonsohn, U., & Loewenstein, G. (2006). Mistake# 37: The effect of previously encountered prices on current housing demand. *The Economic Journal*, *116*(508), 175-199.
59. Simonson, I., & Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *JMR, Journal of Marketing Research*, *29*(3), 281.
60. Sokol-Hessner, P., Hsu, M., Curley, N. G., Delgado, M. R., Camerer, C. F., & Phelps, E. A. (2009). Thinking like a trader selectively reduces individuals' loss aversion. *Proceedings of the National Academy of Sciences*, *106*(13), 5035-5040.
61. Soltani, A., De Martino, B., & Camerer, C. (2012). A range-normalization model of context-dependent choice: a new model and evidence. *PLoS computational biology*, *8*(7), e1002607.
62. Squire, L. R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychological review*, *99*(2), 195.

63. Stewart, N. (2009). Decision by sampling: The role of the decision environment in risky choice. *The Quarterly Journal of Experimental Psychology*, *62*(6), 1041-1062.
64. Stewart, N., Chater, N., & Brown, G. D. (2006). Decision by sampling. *Cognitive psychology*, *53*(1), 1-26.
65. Stewart, N., Reimers, S., & Harris, A. J. (2015). On the origin of utility, weighting, and discounting functions: How they get their shapes and how to change their shapes. *Management Science*, *61*(3), 687-705.
66. Strausfeld, N. J., & Hirth, F. (2013). Deep homology of arthropod central complex and vertebrate basal ganglia. *Science*, *340*(6129), 157-161.
67. Thompson, P. (1981). Velocity after-effects: the effects of adaptation to moving stimuli on the perception of subsequently seen moving stimuli. *Vision research*, *21*(3), 337-345.
68. Tobler, P. N., Fiorillo, C. D., & Schultz, W. (2005). Adaptive coding of reward value by dopamine neurons. *Science*, *307*(5715), 1642-1645.
69. Tsetsos, K., Usher, M., & Chater, N. (2010). Preference reversal in multiattribute choice. *Psychological review*, *117*(4), 1275.
70. Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological review*, *79*(4), 281.
71. Trueblood, J. S., Brown, S. D., & Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological review*, *121*(2), 179.
72. 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*(1), 7.
73. Watkinson, P., Wood, A. M., Lloyd, D. M., & Brown, G. D. (2013). Pain ratings reflect cognitive context: A range frequency model of pain perception. *PAIN®*, *154*(5), 743-749.
74. Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, *428*(6982), 557.
75. Woodford, M. (2012). Prospect theory as efficient perceptual distortion. *American Economic Review*, *102*(3), 41-46.
76. Vandormael, H., Castañón, S. H., Balaguer, J., Li, V., & Summerfield, C. (2017). Robust sampling of decision information during perceptual choice. *Proceedings of the National Academy of Sciences*, 201613950.
77. Vlaev, I. (2018). Local Choices: Rationality and the Contextuality of Decision-Making. *Brain sciences*, *8*(1), 8.
78. Vlaev, I., Seymour, B., Dolan, R. J., & Chater, N. (2009). The price of pain and the value of suffering. *Psychological science*, *20*(3), 309-317.
79. Vlaev, I., Chater, N., Stewart, N., & Brown, G. D. (2011). Does the brain calculate value?. *Trends in cognitive sciences*, *15*(11), 546-554.
80. Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton university press.
81. Webster, M. A. (2015). Visual adaptation. *Annual review of vision science*, *1*, 547-567.
82. Palminteri, S., Wyart, V., & Koechlin, E. (2017). The importance of falsification in computational cognitive modeling. *Trends in cognitive sciences*, *21*(6), 425-433.

Source	Context manipulation	Effect	EaR	Bayesian EaR	DbS	DNT	RFT	LMSV
Rigoli et al (2016c) (see also Rigoli et al 2016b)	A: £1 £3 £5 B: £3 £5 £7	£3A > £3B £5A > £5B	↑	↑	↑	↑	↑	↑
Rigoli et al (2016a)	A: £2 £3 £4 £5 B: £3 £4	£3A > £3B £4A < £4B	↓	↑	↑	↓	↑	↑
Stewart et al Exp A (see also B and C) (2015)	A: £10 £20 £50 £100 £200 £500 B: £10 £310 £410 £460 £490 £500	A: concave B: convex	↓	↓	↑	↓	↑	↑
Stewart et al Exp A(2015)	A: £10 £20 £50 £100 £200 £500 B: £10 £310 £410 £460 £490 £500	£10A > £10B £500A > £500B	↑	↑	↓	↑	↓	↑

Tab. 1. Summary of reference effects that have been observed empirically in between-choice and non-multiattribute scenarios. Column one indicates where the empirical effects have been described in detail. Column two describes the manipulation adopted, reporting the monetary amount characterizing context A and context B. Column three describes the effect found empirically. Columns four-to-nine indicate, for each effect, whether it is consistent (↑) or not (↓) with predictions of each theory. Note that EaR refers to the model proposed by Kőszegi & Rabin (2006), Bayesian EaR refers to the model described in equation 1.

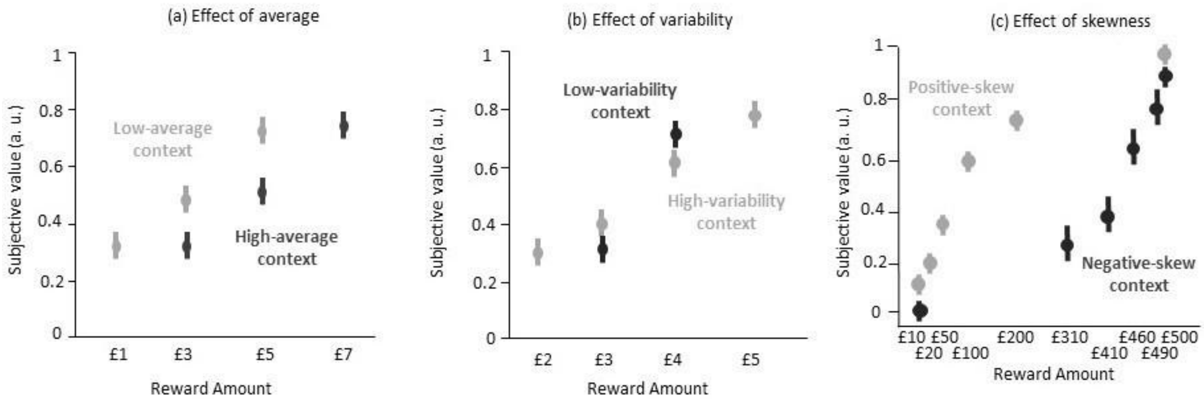

Fig 1. Empirical data on reference effects occurring in between-choice and non-multiattribute contexts. The vertical axis reports the subjective value estimated from choice behaviour (error bars represent standard errors). **A**: Data from Rigoli et al., (2016c) where a low-average context was characterised by £1, £3 and £5 and a high-average context was characterised by £3, £5 and £7. The figure is adapted from fig 1d of Rigoli et al. (2016c), pooling all participants together. **B**: Data from Rigoli et al., (2016a) where a high-variability context was characterised by £2, £3, £4 and £5 and a low-variability context was characterised by £3 and £4. The figure is adapted from fig 4a-b of Rigoli et al. (2016a), pooling all participants together. **C**: Data from Stewart et al. (2015; experiment 1A) where a positive-skew context was characterised by £10, £20, £50, £100, £200, and £500 and a negative-skew context was characterised by £10, £310, £410, £460, £490, and £500. The figure is adapted from fig. 4A of Stewart et al. (2015), rescaled based on the analysis about £10 ad £500 run here (see main text).

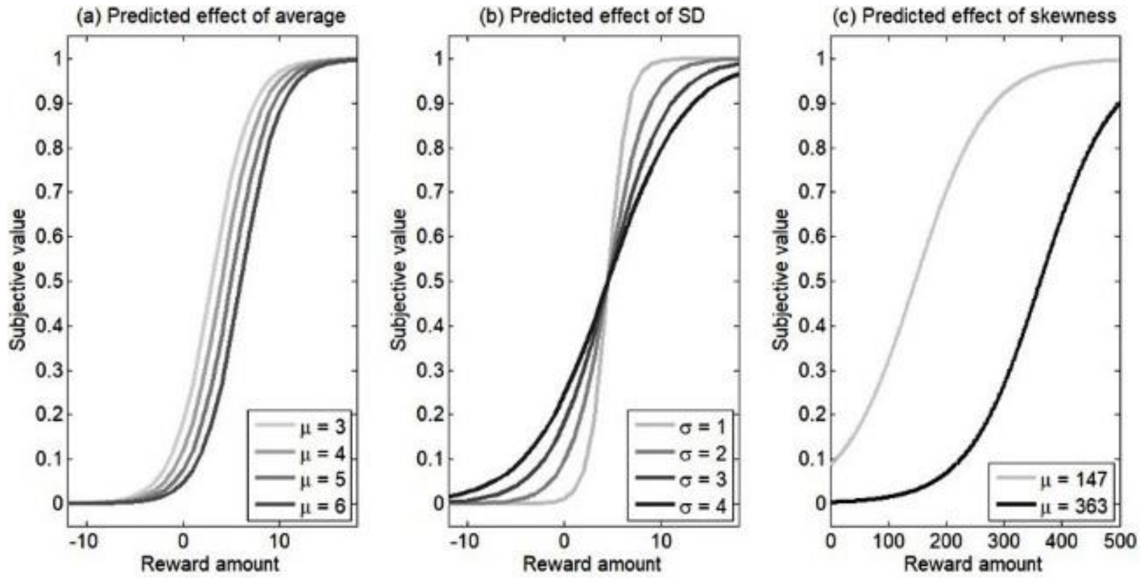

Fig. 2. Predictions of LMSV about reference effects in between-choice and non-multiattribute contexts. Different lines represent data simulated with different parameters. **A**: Predicted effect of average (σ = 2; λ = 1). **B:** Predicted effect of variability (μ = 4.5; λ = 1). **C**: Predicted effect of skewness (σ = 187; λ = 3).

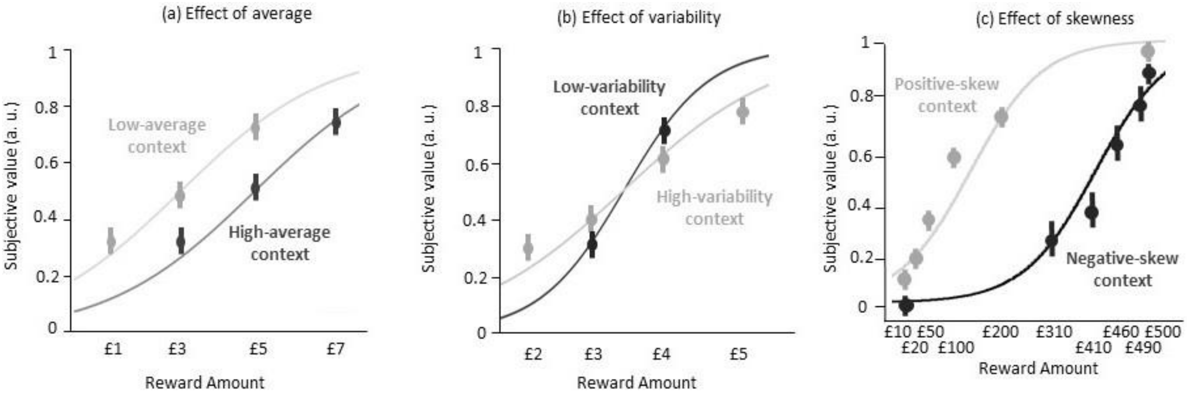

Fig 3. The same empirical data shown in fig. 1 are displayed together with predictions arising from LMSV. For each specific context condition, LMSV was fitted using parameters μ and σ which correspond to the average and SD, respectively, characterising that context condition.

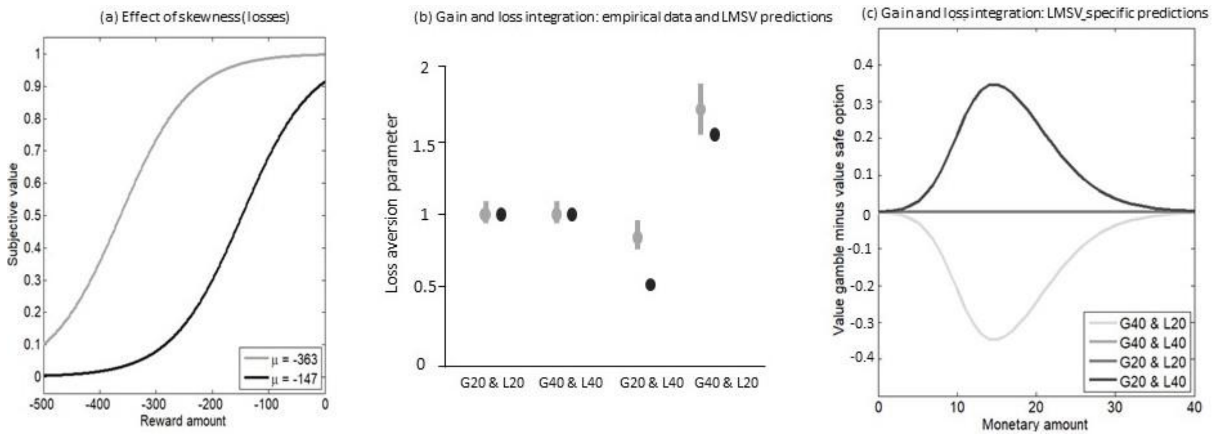

Fig. 4. Predictions of LMSV relative to reference effects occurring during decision-making under risk. Different lines represent data simulated with different parameters. **A**: Predicted reference effects regarding the value function of losses (σ = 187; λ = 3). **B**: Gray dots describe empirical data (with 95% confidence interval) about reference effects occurring during choices between a null outcome and a 50/50 gamble between a monetary amount gained and the same amount lost (Walasek & Stewart (2015)), adapted from tab. 1 relative to experiment 1a in that paper). The range of the contextual distribution is varied for gains and losses. For instance, G40 & L20 indicates that the range for gains is 40 and the range for losses is 20. The loss aversion parameter describes the ratio among the weight given to losses over gains in a logistic regression mode of choice (See Walasek & Stewart 2015 for details). Dark dots represent predictions of LMSV in this task. When the range is 40, μ = ±20 and σ = 11.98; when the range is 20, μ = ±10 and σ = 6.2. We adopted a parameter λ = 3 for all conditions. **C**: Predictions of LMSV in the same scenario regarding the value difference between options as function of monetary amount.