

REFERENCE POINTS AS INFORMATION*

Mark Dean Benjamin Enke Thomas Graeber Pietro Ortoleva

March 28, 2026

Abstract

We develop a model of reference dependence based on the idea that reference points provide *comparative information* to people who are uncertain how to value objects or actions. When people are uncertain how much they value an option, they may still be able to tell whether it is better or worse than a reference point on each dimension, thus providing information. In a Bayesian framework, value uncertainty and comparative information generate regularities such as thinking in gain-loss categories; income targeting and bunching; a preference for advantages over tradeoffs; avoiding losses; coherent arbitrariness; behavioral attenuation; contrast, assimilation, decoy, and range effects; joint-vs.-separate evaluation differences; simultaneous overreaction to small and underreaction to large changes; and—when allowing for caution regarding value uncertainty—the endowment effect. The model predicts that reference effects are most pronounced when value uncertainty is high and the reference point well-known. We then develop a theory of endogenous reference points based on which option is most informative for choice by reducing value uncertainty. This yields several novel predictions, including when the reference point is determined by the status quo, expected outcomes, similarity-based memory, or other options.

*We thank Vincent Marohl for outstanding research assistance. Dean: Department of Economics, Columbia University, mark.dean@columbia.edu. Enke: Department of Economics, Harvard University, and NBER, enke@fas.harvard.edu. Graeber: Department of Economics, University of Zurich, thomas.graeber@econ.uzh.ch. Ortoleva: Department of Economics and SPIA, Princeton University, pietro.ortoleva@princeton.edu.

1 Introduction

A central insight of behavioral economics is that valuation and choice are reference dependent. Standard models capture this by assuming that reference points directly alter the utility of outcomes—for example, that people *dislike* outcomes below the reference point, as in classic Prospect Theory (Kahneman and Tversky, 1979). In this class of models, people know how much utility they get from each option, and the reference point affects this utility.

This paper proposes an alternative view: reference dependence arises because reference points provide *comparative information* to decision makers who are unsure how to value objects or actions. We show that the simple ideas of value uncertainty and learning from reference points generate reference dependence (thinking in gain-loss categories), income targeting and bunching, and several other empirical regularities often treated as unrelated to reference dependence. This framework also yields a new theory of endogenous reference point formation, with several new predictions.

We start from the premise that people are often uncertain about the utility value of options. For instance, they may not know how to translate attributes into utility, they may be uncertain about the characteristics of a good, or the good may be too complex to evaluate. However, even when absolute utility assessments are difficult, it may still be feasible to determine whether an outcome is better than a reference point. For instance, I may be unable to assess the absolute utility of making \$90 today, but I know that it is less than the utility of the \$100 I typically make. The idea behind this paper is that such comparisons provide information that—in the presence of value uncertainty—generates reference dependence.¹

Consider the Tversky and Kahneman (1991) improvements vs. tradeoffs effect. Suppose I must choose between two new jobs relative to my current one (the reference). Job A pays more for fewer hours (two improvements). Job B pays even more, but for more hours (tradeoff). Because I know the utilities of my current wage and hours, comparing them to Job A allows me to rule out low-utility realizations, increasing my posterior evaluation of Job A in both dimensions. For Job B, instead, the tradeoff means that my utility beliefs get truncated from above in one dimension and from below in the other, producing conflicting informational “boosts.” That positive and negative dimension-by-dimension comparisons boost utility evaluations up and down is the hallmark characteristic of reference dependence, yet here it arises purely through an informational channel—the decision maker is uncertain about utility but can easily compare each option with the known reference point. This increases the probability of choosing Job A over B.

Deriving reference dependence from information theory offers three main advantages. First,

¹In their classic paper on the endowment effect, Kahneman et al. (1990) note that “The valuation ambiguity produced by . . . lack of commensurability is necessary . . . for both loss aversion and a buying-selling discrepancy.”

it highlights a tight relationship between reference dependence and value uncertainty, with associated new comparative statics predictions. Second, the model links reference dependence to a wide range of other behavioral phenomena—from contrast effects to coherent arbitrariness—within the same framework. Third, treating reference points as information allows us to endogenize them: in our model, the endogenous reference point is optimal in the sense of being the most informative comparator for subsequent decisions by reducing value uncertainty. This generates sharp predictions about when the reference point is given by expected outcomes, the status quo, similarity-based memory, or other options. Our informational view of reference dependence is not in conflict with reference-dependent preferences; both may exist, and while some predictions are identical, others are not, potentially encouraging future empirical tests.

Why comparative information? As we summarize in Appendix A, a core insight in cognitive psychology is that performance on absolute judgment tasks (evaluating magnitudes) is substantially poorer than on comparative tasks (judging which of two stimuli is brighter, longer, heavier, etc.), suggesting that cognition is optimized for detecting differences rather than encoding absolute levels (Laming, 2009, 1984; Braida and Durlach, 1972; Garner, 1953).² Moreover, this body of work shows that while comparisons are easier than absolute evaluations, ordinal comparisons are especially precise (Luce and Green, 1974; Parducci, 1965; Laming, 1997). This evidence has led to the view that even ostensibly absolute cardinal judgments are often constructed from underlying (partially ordinal) comparisons (Braida and Durlach, 1972; Stewart et al., 2005, 2006). Our model formalizes these key ideas: comparisons provide information when absolute assessments are imperfect, and ordinal comparisons are particularly precise.

Setup. The model rests on three core premises. First, the decision maker (DM) is uncertain about how to value option attributes. We are agnostic about whether this uncertainty reflects preference uncertainty, aggregation uncertainty, or something else. The DM holds a prior over the utility contribution of each attribute, which can be interpreted as the expected utility associated with that attribute before observing the option at hand.

Second, the DM has access to a reference point, the utility of which is fully or partially known. For example, the DM may have experienced an outcome, and thus (possibly imperfectly) remembers its utility. The DM receives both absolute and comparative information, dimension-by-dimension. First, he receives an absolute signal about the utility of the attribute. Second, a comparative ordinal signal about whether the outcome is better than the reference point in that dimension. This comparative, ordinal information is the novelty of our setup and it implies that

²While they model reference points as affecting utility, Kahneman and Tversky (1979) note that “the perceptual apparatus is attuned to the evaluation of changes or differences rather than to the evaluation of absolute magnitudes.”

our DM effectively partitions outcomes into gain-loss categories. We allow this ordinal signal to be noisy such that the DM gets the correct signal with a higher probability when the utility difference between the option and the reference point is large. This implies that, on average, our DM responds not only to ordinal comparisons but also to the underlying cardinal differences.

Third, our DM is Bayesian about his utility uncertainty. Our model thus builds on the “Bayesian brain” literature in cognitive science (Anderson, 2013; Oaksford and Chater, 2007; Gershman, 2021; Griffiths et al., 2024). However, what matters for our model is not so much that the DM is literally Bayesian, but rather that they combine imperfect utility information with comparative signals, and recognize the informational value of reference points.

Posteriors after comparative information. Focusing on Normally distributed priors and signals, we show that posterior beliefs admit a simple structure: the sum of an *absolute* utility evaluation of the outcome and a *relative*, reference-dependent component that depends on the distance between the outcome and the reference point. This reference-dependent term captures the positive or negative informational boost from the comparison.

This structure parallels classic models of reference dependence, in particular those combining absolute and relative utility (e.g. Sugden, 2003; Kőszegi and Rabin, 2006). However, there are several key distinctions. First, in our framework, reference dependence emerges purely through informational mechanisms—in the limit, a DM with no value uncertainty exhibits no reference dependence. This predicts that reference effects will be more pronounced when a decision is difficult or unfamiliar. Second, reference dependence typically weakens when there is limited information about the utility value of the referent. This formalizes the widespread intuition that reference points must “sink in” to affect behavior. Third, classical models in the prospect theory tradition predict a contrast effect with respect to the reference point—higher reference points imply lower utility evaluations. In our model, the same is true (for informational reasons). However, our setup also predicts an assimilation effect with respect to *evaluations* of the reference point; intuitively, when the utility of the reference point is believed to be high, a positive ordinal signal means that the utility of the option is probably also high.

Contrast, attenuation and coherent arbitrariness. Relative to the benchmark of no value uncertainty, the DM’s perceived utility is excessively sensitive around the reference point—where gain-loss categorization shifts—but is otherwise insufficiently sensitive to variation in outcomes. Therefore, our DM overreacts to small changes around the reference point but underreacts to large ones. Intuitively, uncertainty induces belief compression toward the prior, leading to attenuation, except around the reference point, where ordinal information produces a contrast effect. The model, therefore, naturally accounts for the sometimes-puzzling empirical evidence on the coexistence of insensitivity and overreaction (e.g., Augenblick et al., 2025; Graeber et al., 2025).

Furthermore, the model captures the well-known evidence on coherent arbitrariness. In our framework, initial valuations in the absence of a reference point are volatile, yet ordinal signals render subsequent valuations internally coherent. For instance, a DM may face high uncertainty regarding his willingness to pay (WTP) for 10 packets of pasta; yet, having estimated this value, he infers that his WTP for 11 packets must be higher. This yields stable relative behavior despite large underlying value uncertainty.

Avoiding losses, income targeting and bunching. While our model features no asymmetry between gains and losses, it nonetheless predicts a class of behaviors traditionally attributed to loss aversion, purely because in our model utility evaluations vary steeply around the reference point. A first example is the phenomenon of “avoiding losses.” A robust empirical regularity is that people appear to penalize options that involve a loss rather than only gains. For instance losing \$5 is treated as considerable worse than winning \$5, compared to how much worse winning \$1 is treated than winning \$11. If the reference point is \$0, our model trivially predicts this effect because—while the DM may not know the precise utility value of each payout—he does know that losing \$5 is less than his reference, producing a negative informational “boost.”

A second example is income targeting in labor supply decisions (e.g., Camerer et al., 1997; Abeler et al., 2011) or bunching at reference points (e.g., Pope and Schweitzer, 2011; Allen et al., 2017). Intuitively, the ordinal comparative signal relative to the reference point implies that the DM’s perceived marginal utility is higher around the reference point than elsewhere. As a result, over a range of wages, our DM will work “just enough” to exceed the reference income level and avoid the negative informational “boost” from getting less than what he knows. This mechanism produces bunching around the reference point and a locally downward-sloping labor supply curve—the key stylized facts documented in the literature. This mechanism of a higher perceived marginal utility right below compared to right above the reference point is not too dissimilar from how Prospect Theory generates targeting and bunching, except that here it purely reflects beliefs rather than true utility.

Theory of endogenous reference points. Because reference points in our model only reduce value uncertainty, we propose that the endogenous reference point is the alternative that is, on average, most useful for evaluating the current choice. This theory of endogenous reference points is an as-if account: it does not require individuals to consciously choose reference points. Rather, just as we typically bring to mind memories that are broadly relevant to the situation at hand, we may also bring to mind useful reference points. For example, when assessing the value of a high-quality wine, we naturally compare it to other high-quality wines we have tried, rather than to the worst ones.

Formally, this reduces to a problem of optimal information choice. In valuation problems with quadratic losses, the informational value of a reference point admits a simple closed-form characterization: when all potential reference points are equally well-known, the most informative reference point is the expected outcome. When the options near the expected outcome are not well known, there is a bias-variance trade-off: a well-known option, like the status quo, can be more informative if it is not too different from what typically happens. In addition, options that provide clearer ordinal signals are more informative. For example, potential reference points presented in the same format as the option will have an informational advantage.

We also link our view of reference points as information to work on similarity-based memory. We consider environments in which the outcome distribution is itself unknown, such as when the DM does not know the average utility of the objects in a store. In this setting, we prove that with Normally distributed signals, the most informative reference point in the presence of a contextual cue is given by a formula that *exactly* matches the canonical similarity-based memory functional form used in psychology and economics (Kahana, 2012; Bordalo et al., 2020). This shows that the assumptions of value uncertainty and informative comparisons provide a possible microfoundation for the influential work on memory-based reference points.

Discrete choice: Overweighting advantages and menu effects. The idea that DMs learn from comparisons applies not only to reference points but also to other choice options. We model choice assuming that the DM receives ordinal signals about the relative ranking of options within each dimension. We show that this naturally leads the DM to favor options with many advantages. Although this resembles familiar heuristics, it is Bayesian in our setup.

The model also predicts systematic context effects, such as the asymmetric dominance decoy effect, or the Kahneman-Tversky tradeoffs-vs.-improvements effect. Intuitively, introducing a dominated reference point (or decoy) produces additional positive comparisons for the target option, boosting up its evaluation. This mechanism parallels existing informational models of menu effects (Natenzon, 2019). However, unlike those models, ours predicts such effects even when a good does not globally dominate the decoy, as long as there are sufficiently many dimensions of asymmetric dominance.

Caution. Finally, we extend our model to incorporate caution—uncertainty aversion with respect to one’s own value uncertainty (Cerreia-Vioglio et al., 2015, 2024), akin to a form of ambiguity aversion (Gilboa and Marinacci, 2011). Under this additional assumption, our model captures the endowment effect and other forms of status quo bias.

Reference dependence as information and / or preferences? Our view is that when DMs are unsure about valuations, reference points contain information. We show that incorporating

this informational channel produces several widely documented behavioral patterns traditionally associated with reference dependence. Of course, this does not rule out the existence of genuine reference-dependent preferences or loss aversion. We deliberately exclude such preferences to isolate the explanatory power of information. We view empirically distinguishing the relative importance of these channels as an important topic for future research.

Related literature. Our paper contributes to the recent literature on the cognitive foundations of behavioral economics. We connect the literature on cognitive noise and uncertainty to work on reference dependence. Relative to existing models of value uncertainty or cognitive noise (e.g., Gabaix, 2014, 2019; Woodford, 2020; Khaw et al., 2021; Frydman and Jin, 2022; Vieider, 2024), we model reference points as information. The closest antecedent is Villas-Boas (2024), who also models a DM who receives an ordinal signal relative to a reference point, but does not examine the applications or endogenous reference points we consider. Our idea of learning from reference points is also inspired by Natenzon (2019) and subsequent work by Shubatt and Yang (2024), who model DMs that learn from other choice options through comparisons. Augenblick et al. (2025) study a setting in which people understand the direction in which they need to update beliefs, but not the magnitude. In Woodford (2012), efficient coding yields a steeper slope of perceived utility around the prior, yet his model does not study reference points. Wu (2024) presents a non-Bayesian model of contrast and assimilation. In a similar vein, our paper relates to the extensive literature on incomplete preferences or preference imprecision.³

Second, our paper builds on the enormous literature on reference dependence (Barberis, 2013), including work that has endogenized the reference point as expectations (Gul, 1991; Loomes and Sugden, 1986; Kőszegi and Rabin, 2006) or similarity-based memory (Bordalo et al., 2020). Unlike these approaches, we study reference effects from the perspective of value uncertainty, which also yields new predictions about what the reference point is.

Finally, there is some empirical evidence consistent with our model's implication that the magnitude of reference effects increases with value uncertainty. Graeber & Enke (2026) show that contrast and assimilation with respect to comparison points weaken when decision uncertainty is experimentally reduced. Graeber et al. (2025) show that investor responses to earnings surprises exhibit stronger reference effects when assets are more difficult to value.

The paper proceeds as follows. Section 2 illustrates the core intuition of our model. Sections 3 and 4 present the general model and apply it to different regularities. Section 5 endogenizes the reference point. Section 6 studies discrete choice, Section 7 caution and Section 8 concludes and discusses limitations of our model. The appendix includes additional results and all proofs.

³Among recent contributions, see Masatlioglu and Ok (2005, 2014); Sagi (2006); Butler and Loomes (2007, 2011); Ok et al. (2011); Chakraborty (2021); Nishimura and Ok (2021); Cerreia-Vioglio et al. (2024).

2 An Illustrative Example

To build intuition, we sketch a simple special case of our model here. Suppose a DM evaluates the utility of a one-dimensional option, $U(x) = u_x$. For instance, he may evaluate his willingness-to-pay for a car, where the relevant attribute is MPG. The DM is uncertain about the value of u_x . He is Bayesian and has a Normally distributed prior, $u_x \sim \mathcal{N}(\tilde{u}_x, \sigma_p^2)$. This is represented by the dashed line in the top left panel of Figure 1.

The DM receives two types of information about u_x . First, through deliberation, he gets a cardinal signal about the level of u_x , $s = u_x + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. Here, σ_ϵ^2 reflects the DM's uncertainty in assessing the utility. Second, the DM has a reference point r with utility u_r . The DM receives an ordinal signal o that tells him whether the option he evaluates is better or worse than the reference point, dimension-by-dimension: $u_x - u_r \gtrless \nu$. For now, assume that the utility of the reference point is known and that $\nu = 0$, meaning that the ordinal comparison is noiseless. As explained below, when ν is not constant, the ordinal signal is noisy in such a way that the probability of getting the correct comparison increases in the cardinal utility difference. Thus, in our full model, cardinal—rather than only ordinal—differences matter for average behavior

From the cardinal signal s , the DM obtains a posterior that follows the standard Normal-Normal shrinkage formula, with mean $\mathbb{E}[u_x|s] := \tilde{u}_x^s = \lambda s + (1-\lambda)\tilde{u}_x$. This posterior is illustrated in grey in the left panel of Figure 1. What happens with an ordinal signal? The belief gets truncated—giving us the density depicted in red in the top left panel of Figure 1. Naturally, this “boosts” the posterior mean, which is given by

$$\mathbb{E}[u_x|s, o] = \underbrace{\tilde{u}_x^s}_{\text{Ref.-ind. component}} + \underbrace{\sigma_x}_{\text{Weight of ref.-dep. comp.}} \underbrace{\psi(\tilde{u}_x^s - u_r, o)}_{\text{Gain-loss component}} \quad (1)$$

While the exact formula for the gain-loss component $\psi(\cdot)$ involves the inverse Mills ratio, it admits a very good piece-wise linear approximation:

$$\psi(\tilde{u}_x^s - u_r, o) \approx \begin{cases} \max \left\{ \sqrt{\frac{2}{\pi}} - \frac{2}{\pi\sigma_x} (\tilde{u}_x^s - u_r) ; 0 \right\} & \text{if } u_x > u_r \\ -\max \left\{ \sqrt{\frac{2}{\pi}} + \frac{2}{\pi\sigma_x} (\tilde{u}_x^s - u_r) ; 0 \right\} & \text{if } u_x < u_r. \end{cases} \quad (2)$$

The top right panel of Figure 1 plots the average posterior mean as the true value of u_x varies. Two key forces are evident in this figure: *attenuation* away from the reference point, and a *comparative boost* (contrast effect) around the reference point. These two effects reflect that the posterior mean in Eq. (1) is the sum of two components. First, an *absolute* (reference-

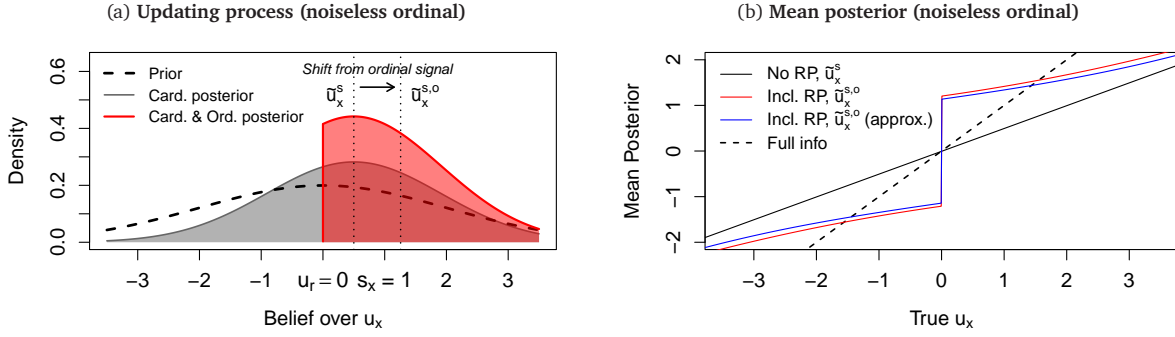


Figure 1: Illustration of the updating mechanism. The left panel shows a ‘snapshot’ of the prior and posterior distributions after receiving either a cardinal or cardinal and ordinal ($u_x > u_r$) signal. In this snapshot, the prior (dashed black) is updated by a cardinal signal $s_x = 1$ to form the posterior \tilde{u}_x^s (grey distribution, mean at left dotted line). The ordinal signal ($u_x > u_r$) truncates the posterior distribution, resulting in the red distribution. The arrow at the top illustrates the resulting ‘boost’ to the mean, which is shifted to $\tilde{u}_x^{s,o}$. The right panel plots the mean posterior as a function of u_x , where we fix $u_r = \tilde{u}_x = 0$. The figure compares the mean posterior with *only* the cardinal signal (black line, \tilde{u}_x^s), the posterior including the ordinal signal (red line, $\tilde{u}_x^{s,o}$), a piece-wise linear approximation (blue line), and the full-information benchmark.

independent) term that reflects the Bayesian shrinkage combining the cardinal signal with the prior. This generates the standard Bayesian attenuation of perceived utility to variation in the option’s objective value. Second, a *relative* or *comparative*, reference-dependent component. Intuitively, a signal that the utility of the option is above the utility of the reference point boosts the mean belief because it rules out low utility values. The weight of the gain-loss component endogenously reflects the DM’s value uncertainty—in our model, there is no reference dependence when the DM perfectly knows u_x .

The model sketched here assumes that the utility of the reference point is perfectly known and that the ordinal signal is always correct. When we relax these assumptions below, the mean posterior becomes smoother around the reference point, but the patterns of contrast and attenuation remain.

Which reference point? We propose that the endogenous reference point is the option that is *most informative* for a choice by reducing value uncertainty. This is an as-if formulation—clearly, individuals do not consciously choose their reference points. But just as people are known to retrieve memories that are roughly useful for the situation at hand, they may also retrieve reference points that are, on average, most useful for assessing the options they face.

As a simple illustrative example, suppose the DM knows that they will want to assess the utility u_x of an alternative and would like to minimize the Mean Square Prediction Error. The value of u_x is drawn from a Normal distribution of known mean \tilde{u}_x and variance σ_x^2 . Two reference

points are available. First, the status quo r_{SQ} , which has a known utility value but is biased—its value differs from the expected outcome \tilde{u}_x by some amount b . Second, an alternative reference point r_A , which is given by the expected outcome \tilde{u}_x , but its utility is not perfectly known, with variance σ_A^2 . Assume for now that the DM expects no further information and no ordinal noise (both assumptions are relaxed below). Which reference point provides more information? The status quo is the most informative reference point if and only if its bias satisfies

$$b \leq b^* \approx \frac{3}{2} \sigma_x \sqrt{\log \left(1 + \frac{5}{4} \cdot \frac{\sigma_A^2}{\sigma_x^2} \right)}.$$

The status quo is the more informative reference point whenever its utility is not too far from the expected outcome, and when the uncertainty about the expected outcome is sufficiently large. Conversely, when the expected outcome has no uncertainty attached to it, it is always the most informative reference point. The intuition is simple. Reference points that are too low or too high will almost always yield the same ordinal comparison, which would be uninformative. On the other hand, uncertainty about the reference point’s utility diminishes its informational value. These two forces create a bias-variance tradeoff, captured by the formula above.

Moreover, we show below that once the decision maker does not know the distribution of outcomes but has access to a contextual cue, a well-known formula from research on similarity-based memory turns out to characterize the informationally optimal reference point.

3 Formal Framework and Reference Points as Information

3.1 Framework: Objects of Uncertain Values

A decision maker (DM) faces one or more alternatives defined by n dimensions. Each option x returns a total utility of

$$U(x) = \sum_{i=1}^n u_{x,i},$$

where $u_{x,i}$ denotes the utility of x from dimension i , for $i = 1, \dots, n$. For example, the item x in question may be a car, and the various $u_{x,i}$ may be the utility from speed, seat comfort, *etc.*

The DM knows this structure, but is uncertain about each $u_{x,i}$.⁴ This uncertainty can have different sources. First, the DM may be uncertain about the actual characteristics of the car, such as its true speed. Second, the DM may be uncertain about the utility from a given characteristic:

⁴For evidence on such uncertainty see, among many others, Butler and Loomes (2007); Cubitt et al. (2015); Agranov and Ortoleva (2017); Enke and Graeber (2023); Enke et al. (2024); Halevy et al. (2023).

they may know that the car’s maximum speed is 280 km/h, but may be unsure of the utility this provides. Third, the DM may struggle to aggregate across dimensions. We remain agnostic about the source of the uncertainty.

The individual is Bayesian and holds priors about each unknown variable. Throughout the paper, we denote by \tilde{z} the mean of the belief about the random variable z , and \tilde{z}^A the posterior mean after information A . For tractability, we assume that the prior about each $u_{x,i}$ is distributed according to $\mathcal{N}(\tilde{u}_{x,i}, \sigma_{p,i}^2)$ and that values are independent across dimensions of the same good ($u_{x,i}$ and $u_{x,j}$ are independent for any i, j with $i \neq j$) and across goods ($u_{x,i}$ and $u_{y,i}$ are independent for all i and all x, y with $x \neq y$).

In addition to the knowledge implicit in the prior, the DM also receives *cardinal* information about the utility in each dimension, which we interpret as the outcome of a deliberation process. For each $i = 1, \dots, N$, the DM receives a signal $s_{x,i} = u_{x,i} + \varepsilon_{x,i}$, where $\varepsilon_{x,i} \sim \mathcal{N}(0, \sigma_{\varepsilon,i}^2)$.⁵ When no confusion arises, we will drop the subscript x . Thus far, this is a standard choice model with value uncertainty and cardinal signals.

Before proceeding, two remarks are in order. First, we focus on Normal priors and signals for simplicity. As will be clear in the discussion that follows, while the exact functional forms of our results rely on Normality—mostly, because truncated Normal distributions are well-understood and well-behaved—the core intuition generalizes. Second, we assume additively separable utility across dimensions, again for simplicity; our results naturally extend, suitably adapted.

3.2 Reference Points as Information

We study situations in which the DM evaluates option x while having (at least mental) access to another object, r , with utility $u_{r,i}$ in the same n dimensions. Here, we treat r as the reference point; later, we show how the same framework can be used to model situations in which r is something else, such as another option in the choice set. We maintain reference points as exogenous for now and endogenize them in Section 5.

We consider both situations in which the utility of r in each dimension is fully known, and when it is not. We assume that the DM believes $u_{r,i} \sim \mathcal{N}(\tilde{u}_{r,i}, \sigma_{r,i}^2)$ for all i , where $\sigma_{r,i}^2 = 0$ captures the case of no uncertainty. Uncertainty about the utility from the reference point may arise from limited experience with the reference point or from memory imperfections. Prior to receiving any further information, r and x are treated as independent.

Our main innovation is to assume that the DM receives *ordinal* signals, in each dimension, about whether it is the reference point r or the item at hand x that has higher utility. We allow

⁵Each $\varepsilon_{x,i}$ is assumed to be independent of all other variables in the model.

these signals to be noisy, and assume that the DM obtains signal $o_{x,i}$ with values

$$o_{x,i} = \begin{cases} + & \text{if } u_{x,i} - u_{r,i} \geq v_{x,i} \\ - & \text{if } u_{x,i} - u_{r,i} < v_{x,i} \end{cases} \quad (3)$$

where $v_{x,i} \sim \mathcal{N}(0, \sigma_{o,i}^2)$ for $i = 1, \dots, n$.⁶ (We will drop the subscript x when no confusion arises.) The parameter $\sigma_{o,i}^2$ captures the noise in the ordinal signal. Importantly, as we discuss in detail below, the noisiness of the ordinal signal operates in such a way that the probability of receiving the correct ordinal signal increases in the distance between the true utility of the option and the reference point. Not only does this correspond to the Probit functional form widespread in economics, but it also aligns with standard principles in psychology and economics regarding just noticeable differences.

As $\sigma_{o,i}^2 \rightarrow 0$, the ordinal signal becomes perfect, regardless of the distance between option and referent. In general, we should expect the difficulty of ordinal comparisons to vary across goods or dimensions: there may be little ordinal noise in comparing car MPGs, but more in qualitative dimensions, such as seat comfort.

To sum up, our model encompasses three types of uncertainty. Suppose you're estimating the value of a new TV, which has only one relevant dimension (resolution). You have a prior about the utility of owning a TV and get a cardinal signal by looking at the new TV, both of which have uncertainty ($\sigma_{p,i}^2$ and $\sigma_{\epsilon,i}^2$). You additionally have access to the utility from a comparison TV—the one you currently own—and that you think you like pretty well, but you're not entirely sure how well ($\sigma_{r,i}^2$). Moreover, you think that the target TV has higher resolution than your current TV, but again, you're not fully certain ($\sigma_{o,i}^2$).

Cardinal comparative information. In our model, the comparative information is ordinal: the DM can tell whether $u_{x,i} \geq u_{r,i}$ but not by how much. While this resonates with a large literature in psychology emphasizing that ordinal reasoning is especially simple and precise (see Appendix A), there are clearly also contexts in which people have a sense for *how much* better or worse an option is compared to a reference quantity. We now discuss to what extent our model already captures this intuition, and how it can easily be extended to capture it in full.

Our DM *does* engage in cardinal reasoning, in two ways. First, the absolute signal about $u_{x,i}$ is cardinal in nature. Second, and more importantly for the present discussion, the comparative ordinal signal is specified such that the DM receives the correct ordinal signal with a probability

⁶One may also want to consider the event in which $u_{x,i} - u_{r,i} = v_i$ as separate and associated to a signal “=” Since this is a probability-zero event (given our absolutely continuous distributions), to avoid extra notation and since nothing changes, for simplicity, we leave this possibility aside and associate this event with the signal “+.”

that increases in the utility difference between x and r , see eq. (3). As a result, average posterior beliefs directionally vary with the cardinal utility difference ($u_{x,i} - u_{r,i}$) in the same way as in a model with a cardinal comparative signal as specified below. This aspect of our formulation is deliberate. While we could instead have specified the probability of receiving the correct ordinal signal as independent of the cardinal utility difference, we chose to allow for cardinal differences to affect the comparative information (without giving up the importance of ordinal information) in a simple and tractable way.

However, should readers wish to incorporate a direct cardinal comparative signal into our setup, this is straightforward as well. Concretely, suppose that in addition to the ordinal signal, the DM also receives a cardinal comparative signal, $d_i = u_{x,i} - u_{r,i} + \zeta_i$, where $\zeta_i \sim \mathcal{N}(0, \sigma_{d,i}^2)$. First, it turns out that our current model is informationally *equivalent* to one in which the cardinal comparative signal d_i is also received, simply by adapting the parameters, under mild assumptions on the relevant covariance matrix. Second, it is also straightforward to extend our results to incorporate an unrestricted version of this signal in a simple and tractable way. Both results are included in Appendix B.2.

3.3 Perceived Utility with Reference Points

Given the additively separable structure of $U(x)$ and independence across dimensions, we can focus on each dimension separately, as the multi-dimensional case is simply the sum across dimensions; therefore, we suppress the subscript indicating the dimension (we write u_x for $u_{x,1}$).

A standard result in Bayesian updating gives us that the posterior after the cardinal signal is $N(\tilde{u}_x^s, \sigma_x^2)$, with $\tilde{u}_x^s = \lambda s + (1 - \lambda)\tilde{u}_x$, where $\lambda := \frac{\sigma_p^2}{\sigma_p^2 + \sigma_\epsilon^2}$, and $\sigma_x^2 := \frac{\sigma_\epsilon^2 \sigma_p^2}{\sigma_\epsilon^2 + \sigma_p^2}$. How does the ordinal signal affect it? The following proposition characterizes its effects on posterior means.⁷

Proposition 1. *The following is true:*

$$\mathbb{E}[u_x | o, s] = \tilde{u}_x^s + \frac{\sigma_x^2}{\sqrt{\sigma_x^2 + \sigma_r^2 + \sigma_o^2}} \psi \left(\frac{\tilde{u}_x^s - \tilde{u}_r}{\sqrt{\sigma_x^2 + \sigma_r^2 + \sigma_o^2}}, o \right) \quad (4)$$

$$\text{where } \psi(t, o) := \begin{cases} \frac{\phi(t)}{\Phi(t)} & \text{if } o = + \\ -\frac{\phi(-t)}{\Phi(-t)} & \text{if } o = -. \end{cases} \quad (5)$$

⁷Proposition 12 in Appendix B.3 characterizes the effects on the variance. The posterior variance is, on average, increasing in the distance from the reference point. Intuitively, being close to the reference point means being close to a point whose utility one knows, and the further one moves away, the larger the posterior variance becomes.

To illustrate, consider the case of a “+” ordinal signal. Then,

$$\mathbb{E}[u_x|+, s] = \underbrace{\tilde{u}_x^s}_{\text{Ref.-indep. component}} + \underbrace{\frac{\sigma_x^2}{\sqrt{\sigma_x^2 + \sigma_r^2 + \sigma_o^2}}}_{\text{Weight on ref.-dep.}} \underbrace{\frac{\phi\left(\frac{\tilde{u}_x^s - \tilde{u}_r}{\sqrt{\sigma_x^2 + \sigma_r^2 + \sigma_o^2}}\right)}{\Phi\left(\frac{\tilde{u}_x^s - \tilde{u}_r}{\sqrt{\sigma_x^2 + \sigma_r^2 + \sigma_o^2}}\right)}}_{\text{Ref.-dependent component}} \quad (6)$$

The posterior mean after the cardinal and ordinal signals is the sum of two components. First, an *absolute* (reference-independent) term that reflects the DM’s mean belief about the utility of the option combining the cardinal signal with the prior through a Bayesian shrinkage, giving rise to attenuation.⁸ Second, a *relative* or *comparative*, reference-dependent component: a “boost” in beliefs obtained from the “+” ordinal signal, where the weight of the boost depends on the DM’s uncertainty. This formula is the generalization of the one used in our example in the previous section, Eq. (1), allowing for uncertainty about u_r and noise in the ordinal signal, and giving the correct expression of the boost rather than an approximation.

Through the comparative boost, the mean belief is repelled (or contrasted away) from the reference point. Formally, this boost is given by the so-called inverse Mill’s ratio that has attracted much attention in the econometrics literature on correcting for sample selection bias (Heckman, 1979). Note how the boosts have the expected signs: it is immediate to show that $\psi(t, +) > 0$ and $\psi(t, -) < 0$ for all t .

Proposition 1 describes beliefs conditional on signal realizations. While this makes the mechanics clear, signal realizations may not be observed by the researcher. Perhaps a more empirically relevant metric is average beliefs across possible signal realizations: Figure 2 illustrates the average mean posterior and average reference-dependent component for different levels of σ_x . The figure fixes the reference point and thus implicitly shows comparative statics with respect to u_x . (Proposition 13 in Appendix B.3 provides a full characterization of average beliefs.)

Mean beliefs are given by the sum of the reference-dependent boost and the reference-independent beliefs based on the cardinal signal. Does this combination lead to over- or under-estimation? The answer is both, depending on how close u_x is to u_r . Suppose $u_x > u_r$ (the patterns are symmetric for the opposite case). As long as the ordinal noise is not too big and priors are aligned, the average posterior belief will be strictly above the true u_x when u_x is close to u_r , as the boost dominates. As the difference between them increases, the mean belief will eventually go below the true value—the boost becomes weaker and attenuation dominates. (This

⁸We understand attenuation as $\partial \mathbb{E}_{s,o} [\mathbb{E}[u_x|s, o]] / \partial u_x < 1$, which holds in this model as long as the ordinal signal doesn’t change.

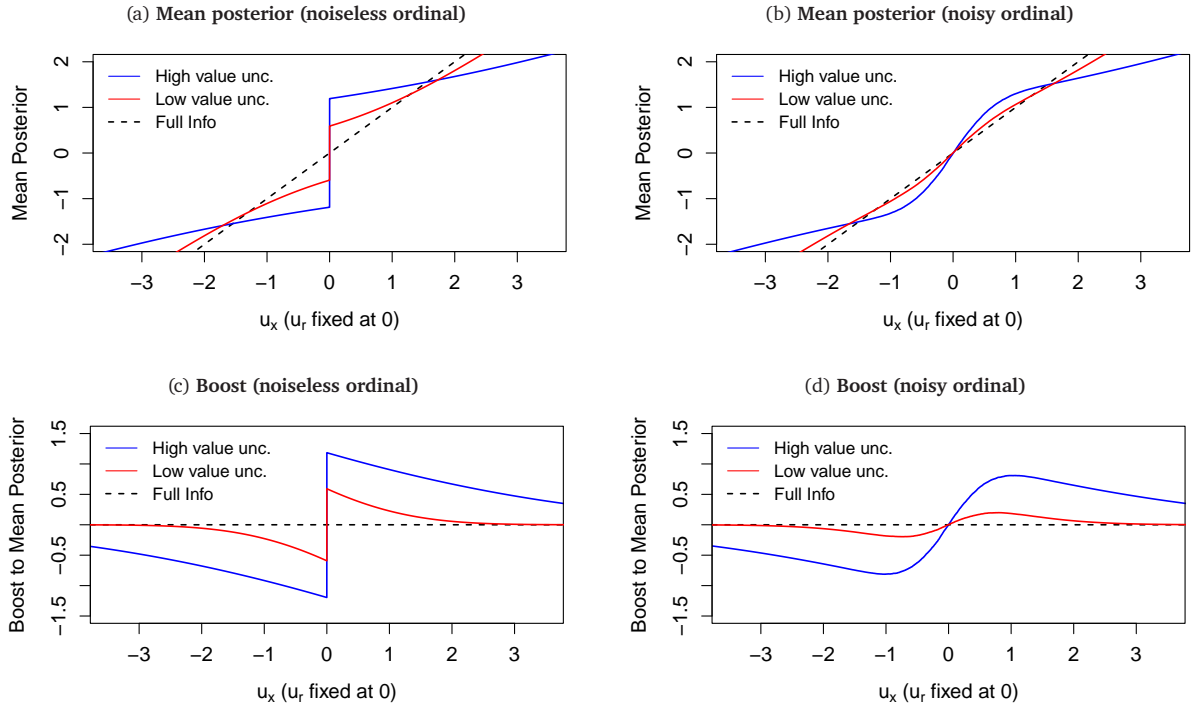


Figure 2: Illustration of the mean posterior (top panels) and its reference-dependent boost component (bottom panels) as a function of value uncertainty σ_ϵ , separately for noiseless and noisy ordinal signals. We set $\tilde{u}_x = u_r = \tilde{u}_r = 0$ and then vary u_x (and thus s_x). All panels compare a 'low value uncertainty' DM ($\sigma_\epsilon = 0.75$, red line) with a 'high value uncertainty' DM ($\sigma_\epsilon = 2$, blue line), where we hold constant $\sigma_p = 2$. The left panels (a, c) illustrate the case with a *noisy reference, noiseless ordinal* signal ($\sigma_r = 0.25, \sigma_o = 0$). The right panels (b, d) illustrate the case with a *noisy reference, noisy ordinal* signal ($\sigma_r = 0.25, \sigma_o = 0.5$). The top row (a, b) plots the mean posterior, $\tilde{u}_x^{s,o}$, while the bottom row (c, d) plots the reference-dependent 'boost', $\psi(\cdot)$. Black dashed lines indicate the full information benchmark. Simulation conducted over 10,000 draws.

intuition is formalized in Proposition 2 below.)

3.4 Comparative Statics

We now study how the mean posterior (which we also refer to as 'evaluation') varies as a function of (i) the utility of the option, u_x , (ii) the utility of the reference point, u_r , and (iii) the different types of uncertainty in the model, $\sigma_x^2, \sigma_r^2, \sigma_o^2$. We first discuss these results at an intuitive level and then provide formal statements.

Comparative statics with respect to the value of the option. An increase in the value of the option impacts beliefs through two channels: by changing the distribution of the cardinal signal received about u_x , and by changing the probability of receiving a positive ordinal signal.

Conditional on receiving a positive ordinal signal, as we increase u_x , \tilde{u}_x^s is (on average) higher, but the boost itself is (on average) *lower*. This is because $\psi(t, +)$ is a decreasing function of t ,

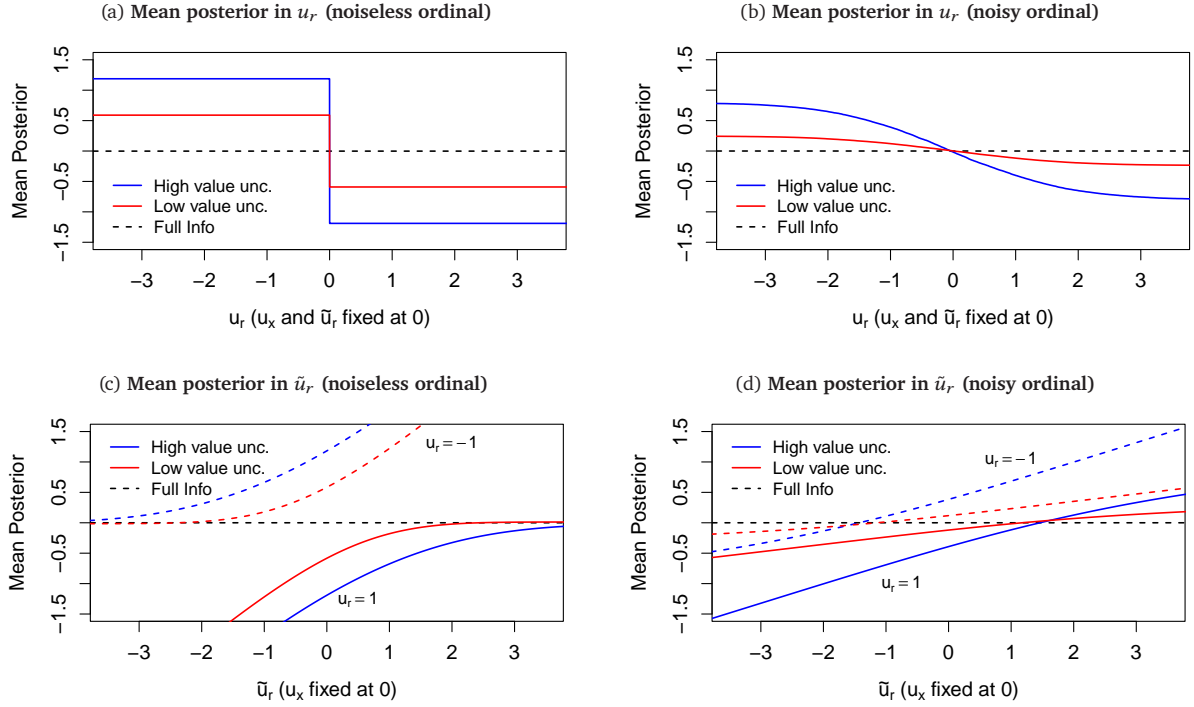


Figure 3: Illustration of comparative statics with respect to reference quantities. The top panels show the mean posterior as a function of the reference point u_r , holding constant $u_x = \tilde{u}_r = 0$. The bottom panels show the mean posterior as a function of the reference evaluation \tilde{u}_r , holding $u_x = 0$ and fixing $u_r \in \{-1, 1\}$ (solid lines: $u_r = 1$; dashed lines: $u_r = -1$). The left panels (a,c) illustrate the case with a *noisy reference, noiseless ordinal* signal ($\sigma_r = 0.25$, $\sigma_o = 0$). The right panels (b,d) illustrate the case with a *noisy reference, noisy ordinal* signal ($\sigma_r = 0.25$, $\sigma_o = 1.5$). All panels compare $\tilde{u}_x^{S,0}$ for a 'low value uncertainty' DM ($\sigma_\epsilon = 0.75$, red lines) with a 'high value uncertainty' DM ($\sigma_\epsilon = 2$, blue lines), where we hold constant $\sigma_p = 2$ and prior $\tilde{u}_x = 0$ and only vary σ_ϵ^2 . Black dashed lines indicate the full information benchmark. Simulation conducted over 10,000 draws.

and is also intuitive: the impact of learning that u_x is above u_r should be smaller if you already know that u_x is likely above u_r . By a symmetric argument, with a negative signal, increasing u_x increases the magnitude of a negative boost.

Absent ordinal noise, panel (c) of Figure 2 illustrates these effects on the boost. Below u_r , an increase in u_x increases the magnitude of the negative boost. At u_r , there is a discontinuous jump as the ordinal signal switches from negative to positive. Above u_r , further increases reduce the magnitude of the positive boost.

With ordinal noise, the effect of an increase in u_x is smoother: the probability of receiving a positive signal increases continuously, rather than jumping discontinuously, as illustrated by panel (d) of Figure 2. In general, the maximal average boost occurs strictly above u_r : close to u_r , the DM is just as likely to receive a positive as a negative ordinal signal, so the average boost is close to zero; this is also clear in Figure 12 in Appendix G, which shows a heatmap of the full distribution of the mean posterior for a given value of $u_x - u_r$.

To obtain the effect of changes in u_x on mean beliefs, the effect of the boost must be combined with the effect on the reference independent portion of beliefs. The overall effect can be seen in the top panels of Figure 2: while average beliefs about u_x are strictly increasing in its true value, these beliefs are attenuated away from the reference point and excessively sensitive close to it, where the implicit change in gain-loss categorizations through the ordinal signal produces a contrast effect.

Comparative statics with respect to the reference point. The mean posterior in Eq. (4) has novel comparative statics implications regarding the role of the reference point. As with the option itself, two distinct—but interrelated—reference quantities enter the mean posterior: the belief about the reference point, \tilde{u}_r , which affects the size of the boost, and the true utility of the reference point, u_r , which affects the probability of which ordinal signal will be received.

Fixing the ordinal signal, an increase in \tilde{u}_r *increases* the evaluation of the option at hand, through the size of the boost. This is intuitive. Suppose the ordinal signal is positive: then, the belief about the utility of an option should be higher when you are told that it is better than an object with utility of one million than when told it is better than one with utility of zero. This produces an “assimilation effect:” a higher reference belief pushes evaluations of the option *up*. This is illustrated in the bottom panels of Figure 3.

On the other hand, fixing beliefs about the reference point \tilde{u}_r , the average posterior mean about u_x *decreases* in the true utility of the reference point u_r . This is because u_r determines which ordinal signal is observed, and a higher reference point makes it more likely that the comparative signal is negative, lowering beliefs on average. The top panels of Figure 3 illustrates how, varying u_r while fixing \tilde{u}_r , a “contrast effect” exists locally around the reference point.⁹

The overall effect of a change in the utility of the reference point u_r is determined by the relative importance of these two effects. This, in turn, is affected by how beliefs \tilde{u}_r depend on u_r . Depending on the parameterization of the problem, either channel could dominate. For example, if ordinal signals are noisy and \tilde{u}_r is insensitive to u_r (e.g., with a strong prior about the value of the reference point), the contrast effect dominates, and an increase in u_r decreases the evaluation of u_x . On the other hand, if ordinal signals are noiseless and \tilde{u}_r is very sensitive to u_r , then an increase in u_r that does not change the sign of $u_x - u_r$ will increase the evaluation of x as the assimilation effect will dominate.

Comparative statics with respect to uncertainty. Proposition 1 makes it clear that the reference-dependent component in our model also depends on (i) the degree of value uncertainty about

⁹If we add to the model a cardinal comparative signal (as discussed in Section 3.2), a global contrast effect obtains even in the absence of ordinal noise; see Appendix B.2.

the option, σ_x^2 , which depends on both prior (σ_p^2) and signal (σ_ϵ^2) uncertainty; (ii) uncertainty about the utility of the reference point, σ_r^2 ; and (iii) how noisy the ordinal information is, σ_o^2 . This highlights a central implication of our approach, and a key point of difference with classic models of reference dependence—that the effect of the reference point depends on uncertainty. On the other hand, uncertainty also affects the reference-independent component, as it changes the Bayesian shrinkage coefficient λ .

Figures 2.(b) and (c) and 3 illustrate how the magnitude of value uncertainty about the option modulates the effects of the reference-dependent boost. A key insight is that the boost (positive or negative) is more pronounced when x is less well-known. As a result, the degree of reference dependence—such as the strength of the contrast effect around the reference point—*increases* in value uncertainty. This is a central prediction of our model.

As discussed above, our model typically predicts overestimation of u_x when it is right above the reference point and underestimation when it is far above. Changes in undercertainty σ_x^2 have two effects on this over- and underestimation. First, as we have just seen, it will increase the boost. Second, the Bayesian attenuation becomes more pronounced. As a result, an increase in value uncertainty will generally *push up* the posterior mean close to the reference point, yet *push it down* far away from the reference point. That is, it amplifies both the contrast and attenuation effects. This is illustrated by the red and blue lines in the top panel of Figure 2.

The uncertainty in the knowledge about the reference point (σ_r^2) and in the ordinal signal (σ_o^2) affect the boost in exactly the same way. If the ordinal signal is unexpected ($\tilde{u}_x^s - \tilde{u}_r$ and o have different signs) or if there is at least some uncertain ex-ante ($|\tilde{u}_x^s - \tilde{u}_r|$ not too large), then the boost is decreasing in both; this is very intuitive, as greater uncertainty yields smaller updating.

More surprisingly, when the ordinal signal is confirmatory and widely expected ($|\tilde{u}_x^s - \tilde{u}_r|$ is large and o has the same sign), then the boost can be *increasing* in both σ_r^2 and σ_o^2 . As we explain in Appendix B.4, this is because, if \tilde{u}_x^s is a long way above \tilde{u}_r and there is no noise in the reference process, a positive signal is barely informative—the DM already knows that x is better than r . However, the less r is known, the more uncertainty there is about its values, and the more informative the signal is. A similar logic applies to the role of σ_o^2 .¹⁰

Formal statements of comparative statics. We conclude by stating formally the comparative statics discussed above; Appendix B.4 provides additional comparative static results for the reference dependent boost in Eq. (4), both conditional on and averaging across signal realization. We begin with the comparative statics with respect to the option and the uncertainty terms.

¹⁰When averaging over signal realizations, an increase in the ordinal noise σ_o^2 has the additional effect of increasing the probability that contradictory signals are received. This has the effect of ‘flattening out’ the average boost, as can be seen when comparing panels (c) and (d) of Figure 2.

Proposition 2. Assume $\tilde{u}_x = \tilde{u}_r = u_r$ and $\sigma_p^2, \sigma_\varepsilon^2 > 0$.

1. (Over and under-sensitivity wrt u_x .) $\mathbb{E}[\tilde{u}_x^{s,0}]$ is strictly increasing in u_x . Moreover, there exists $\bar{\sigma} > 0$:

(a) if $0 < \sigma_o < \bar{\sigma}$, there exist $0 < \varepsilon < \varepsilon'$ such that

$$\underbrace{\frac{\partial \mathbb{E}[\tilde{u}_x^{s,0}]}{\partial u_x} > 1 \quad \text{if } |u_x - u_r| < \varepsilon,}_{\text{Oversensitivity near RP}} \quad \underbrace{0 < \frac{\partial \mathbb{E}[\tilde{u}_x^{s,0}]}{\partial u_x} < 1 \quad \text{if } |u_x - u_r| > \varepsilon'}_{\text{Attenuation away from RP}}$$

(b) if $\sigma_o \geq \bar{\sigma}$, $0 < \frac{\partial}{\partial u_x} \mathbb{E}[\tilde{u}_x^{s,0}] < 1$.

(c) If $\sigma_o = 0$, $\mathbb{E}[\tilde{u}_x^{s,0}]$ is discontinuous at $u_x = u_r$, while we have $0 < \frac{\partial}{\partial u_x} \mathbb{E}[\tilde{u}_x^{s,0}] < 1$ (attenuation) everywhere else.

2. (Over- and under-estimation.) With the same $\bar{\sigma}$ as in Part 1, if $0 < \sigma_o < \bar{\sigma}$, there exists $\Delta > 0$ such that

$$\underbrace{\mathbb{E}[\tilde{u}_x^{s,0}] > u_x \quad \text{for } u_x \in (-\infty, u_r - \Delta) \cup (u_r, u_r + \Delta),}_{\text{Overestimation right above RP or much below}} \quad \underbrace{\mathbb{E}[\tilde{u}_x^{s,0}] < u_x \quad \text{for } u_x \in (u_r - \Delta, u_r) \cup (u_r + \Delta, +\infty).}_{\text{Underestimation right below RP or much above}}$$

If $\sigma_o \geq \bar{\sigma}$, $\mathbb{E}[\tilde{u}_x^{s,0}] < u_x$ if $u_x > u_r$, $\mathbb{E}[\tilde{u}_x^{s,0}] > u_x$ if $u_x < u_r$.

3. (Effects of uncertainty.) Assume $u_x > u_r$ (signs are the opposite when $u_x < u_r$). Then:

(a) there exists $\Delta_r > 0$ such that

$$\frac{\partial \mathbb{E}[\tilde{u}_x^{s,0}]}{\partial \sigma_r^2} < 0 \quad \text{if } 0 < u_x - u_r < \Delta_r, \quad \frac{\partial \mathbb{E}[\tilde{u}_x^{s,0}]}{\partial \sigma_r^2} > 0 \quad \text{if } u_x - u_r > \Delta_r;$$

(b) $\frac{\partial \mathbb{E}[\tilde{u}_x^{s,0}]}{\partial \sigma_p^2} > 0$ if $u_x > u_r$ and there exists $\Delta_\varepsilon \geq 0$ such that

$$\frac{\partial \mathbb{E}[\tilde{u}_x^{s,0}]}{\partial \sigma_\varepsilon^2} < 0 \quad \text{if } u_x - u_r > \Delta_\varepsilon;$$

Part (1) shows that—relative to the full-information benchmark—the DM's evaluation is too sensitive to variation in u_x close to the reference point, and attenuated far away from it, as long as ordinal noise is not too high; if it is, only attenuation is present. Part (2) says that, under the same conditions, there is overestimation of u_x when it is close to and above the reference point, and underestimation when much above it (signs are flipped when it is below). Part (3) says that

both over- and under-estimation increase in the signal uncertainty about u_x and decrease in the uncertainty about u_r .

Next, we formalize the main comparative statics with respect to the reference point.

Proposition 3. *Suppose $\sigma_p^2, \sigma_\epsilon^2, \sigma_0^2 > 0$.¹¹ Then,*

1. (Contrast Effect) Fixing \tilde{u}_r , $\frac{\partial \mathbb{E}[\tilde{u}_x^{s,o}]}{\partial u_r} < 0$.
2. (Assimilation Effect) Fixing u_r , $\frac{\partial \mathbb{E}[\tilde{u}_x^{s,o}]}{\partial \tilde{u}_r} > 0$.
3. If $\tilde{u}_r = \alpha u_r + (1 - \alpha)\bar{u}_r$, there exists $\alpha^* \in (0, 1)$ such that $\frac{\partial \mathbb{E}[\tilde{u}_x^{s,o}]}{\partial u_r} > 0 \Leftrightarrow \alpha > \alpha^*$.

Parts (1)–(2) say that the posterior mean exhibits a contrast effect with respect to the actual value of the reference point and an assimilation effect with respect to the beliefs about the reference point, both of which increase in value uncertainty. Part (3) clarifies that, if only the reference point (or its underlying attributes) is observable by the analyst, its observed net effect depends on how much the changes in the reference point are internalized by the decision-maker: the assimilation effect dominates if and only if changes are sufficiently assimilated in the beliefs, while the contrast effect dominates otherwise.

3.5 Applying the Model: Approximations

A difficulty in adopting the formulas in Proposition 1 in theoretical or empirical applications is that they involve Inverse Mills ratios, which, although relatively well understood, are hard to work with analytically. Fortunately, these admit a relatively simple approximation based on a first-order Taylor approximation around $\psi(0) = \sqrt{\frac{2}{\pi}}$.

$$\psi(t, o) \approx \psi^a(t, o) = \begin{cases} \psi_+^a(t) := \max \left\{ \sqrt{\frac{2}{\pi}} - \frac{2}{\pi}t ; 0 \right\} & \text{if } o = + \\ \psi_-^a(t) := -\max \left\{ \sqrt{\frac{2}{\pi}} + \frac{2}{\pi}t ; 0 \right\} & \text{if } o = -. \end{cases} \quad (7)$$

We also provide an approximation of the average boost across signal realizations, $\mathbb{E}[\tilde{u}_x^{s,o}]$. Let $t := \frac{\lambda u_x + (1-\lambda)\tilde{u}_x - \tilde{u}_r}{\sqrt{\sigma_x^2 + \sigma_r^2 + \sigma_0^2}}$. Then, the average posterior mean belief can be approximated by

¹¹When $\sigma_0^2 = 0$, Part (2) remains true as is, but $\frac{\partial \mathbb{E}[\tilde{u}_x^{s,o}]}{\partial u_r}$ in Parts (1) and (3) may not be defined because $\mathbb{E}[\tilde{u}_x^{s,o}]$ jumps discontinuously downwards as u_r reaches and surpasses u_x ; even in that case, all statements remain true (with weak inequality for Part 1 and Part 3) wherever $\frac{\partial \mathbb{E}[\tilde{u}_x^{s,o}]}{\partial u_r}$ is defined.

$$\mathbb{E}[\tilde{u}_x^{s,o}] \approx \underbrace{\lambda u_x + (1 - \lambda) \tilde{u}_x}_{\text{Bayesian shrinkage to prior; Attenuation}} + \underbrace{\frac{\sigma_x^2}{\sqrt{\sigma_x^2 + \sigma_r^2 + \sigma_o^2}}}_{\text{Weight of ref.-dep.}} \underbrace{\bar{\psi}^a(t, u_x, u_r, \sigma_o)}_{\text{Avg. Boost Approx.}}$$

$$\text{where } \bar{\psi}^a(t, u_x, u_r, \sigma_o) := \underbrace{\left[\Phi\left(\frac{u_x - u_r}{\sigma_o}\right) \psi_+^a(t) + \left(1 - \Phi\left(\frac{u_x - u_r}{\sigma_o}\right)\right) \psi_-^a(t) \right]}_{\text{Gain / loss boosts weighted by prob. of correct ordinal}}. \quad (8)$$

When the ordinal signal is noiseless, the last part simplifies to the gain-loss component in the illustrative example in Section 2. Derivations of both approximations can be found in Appendix B.5, while simulations showing their performance can be found in Appendices F.1 and F.2.

4 Applications of Reference Points as Information

4.1 Contrast, Assimilation and Differences with Prospect Theory

Differences with Prospect Theory. The mean perceived utility in our model resembles, in some ways, the structure of prospect theory’s value function: steeper around the reference point and attenuated away from it. At the same time, there are several crucial differences.

First and foremost, our model is entirely information-based. We make no assumptions about preferences and instead simply study the beliefs of a Bayesian DM who receives comparative signals with respect to reference points. As a result, our model makes the distinctive comparative-static prediction that people respond more to reference points when their value uncertainty is higher (see Figures 2 and 3). This is in line with Graeber et al. (2025), who show that market returns following earnings surprises are strongly inverse-S shaped around the reference point of no earnings surprise, and that this effect is substantially more pronounced for firms for which valuation uncertainty is higher. Related, Graeber & Enke (2026) show experimentally that reference point effects are more pronounced when a decision problem involves more uncertainty about how to map problem inputs into valuations and decisions.

Second, our model does not feature any gain-loss asymmetry. While below we show that it can generate several empirical patterns typically predicted by Prospect Theory through loss aversion, it does not, as stated, generate loss aversion itself (even in riskless choice) or the endowment effect. In Section 7 we discuss a natural extension of our model that does generate these predictions by introducing aversion to value uncertainty (caution).

Third, as we show next, our model explains various empirical regularities that are not explained by prospect theory, such as coherent arbitrariness, assimilation and decoy effects.

Contrast and assimilation. In Prospect Theory and its variants, a higher reference point always *decreases* the evaluation of an option. This captures the compelling intuition that if I make \$100 today, I am happier if I made \$10 rather than \$80 yesterday. This contrast effect with respect to the true attribute values is present also in our model. Yet, as can be seen in Figure 3, our model adds the prediction that evaluations increase in subjective beliefs about the utility of the reference point. For example, our model also predicts that if I make \$100 today, I am happier if I believe that my utility from making \$80 yesterday was 300 rather than 200.

In many cases, we would expect the true value of the reference point and its perceived utility to move together. Indeed, in most datasets, only one of these two is observable. However, a recent body of work both in economics and in perceptual psychology measures both and provides evidence that is consistent with our model’s predictions.

Graeber & Enke (2026) propose that comparison points produce assimilation effects when they consist of decisions / evaluations, yet contrast effects when they are given by exogenous problem attributes. They implement experiments that induce both a reference point and a reference evaluation. For instance, subjects are asked to state their willingness-to-accept (WTA) for completing 10 real effort tasks, while being told that a different subject stated a WTA of \$8 for completing 20 tasks. Through the lens of our model, the reference point u_r is (some function of) the number of effort tasks the other subject was offered, while the evaluation of the reference point, \tilde{u}_r , is given by (some function of) the other subject’s WTA. Graeber & Enke (2026) find that stated WTA increases in the other subject’s WTA and simultaneously decreases in the other subject’s offered workload, in line with the predictions of our model.

There is related evidence also in the perceptual domain in psychology. In experiments in which subjects sequentially estimate the magnitude of visual stimuli, a recent result is that subjects’ estimates decrease in the reference stimulus (e.g. the previous stimulus) but increase in the subject’s estimate of the reference stimulus (e.g. Moon and Kwon, 2022; Sadil et al., 2024). Thus, again, beliefs exhibit a contrast effect with respect to the true reference quantity (u_r) but increase in the evaluation of the reference quantity (\tilde{u}_r).

Finally, Jin et al. (2024) find that physicians’ evaluations of a patient increase in their evaluation of the previous patient (the reference evaluation) and—most importantly for our purposes—that this assimilation effect is stronger when medical uncertainty is higher or the physician is less experienced or more fatigued. This is exactly what our model predicts: that higher evaluative uncertainty (difficult case, less experience) increases the reliance on the reference evaluation.

4.2 Avoiding Losses

Consider two two-dimensional options, x and y , with values $u_{x,1} > u_{y,1} > u_{r,1}$ and $u_{y,2} > u_{x,2} > u_{r,2}$, i.e., all values are initially above their corresponding reference points. Further suppose that when the DM is asked to separately value the two options, these evaluations are equal. Suppose that the options are now manipulated by shifting all values down by the same constant (e.g., $u_{x',1} = u_{x,1} - c$ for $c > 0$) in such a way that one option drops below the reference point in one dimension and all other values remain above the reference point. That is, $u_{x,1} - c > u_{r,1} > u_{y,1} - c$, i.e. only option y falls below the reference point in dimension 1, and $u_{y,2} - c > u_{x,2} - c > u_{r,2}$, i.e. the shift does not affect the position of any values relative to the reference point in dimension 2. Then, a prominent empirical regularity in the literature on reference dependence is that option y is now valued less than x . That is, empirically, there appears to be a ‘loss penalty’.

Our model trivially generates this ‘avoiding losses’ effect as long as the ordinal signal is not too noisy. Intuitively, when all values are above the reference point, they receive boosts from positive ordinal signals. When one dimension of one option drops below the reference point, this option now receives a negative boost rather than a positive one. Because the posterior mean is substantially more sensitive to changes around the reference point than away from it (compare Figure 2), the change in the ordinal ranking between option and reference point leads to a larger drop in the evaluation. No gain-loss asymmetry is required for this ‘avoiding losses’ effect.

One application of this general insight is to the widely-studied empirical regularity that, in lottery choice, options appear to receive a disproportionate penalty when, in some state of the world, they drop below the reference point of a payout of zero. In particular, many experimental researchers have implemented ‘shifting’ manipulations of the type described above, where all payouts are shifted down by a constant such that only one of them drops below zero. To informally capture this, we rely on the same setup as above and now assume that the two dimensions of an option are given by two equally likely states of the world. If the DM values two lotteries equally when all four payouts are above zero, then when all payouts are shifted down by the same constant, in our model it must be that the DM now values the option that involves a loss less, purely because of the steep drop in perceived utility that results from the negative ordinal comparison.

4.3 Over- and Underreaction

Suppose the DM’s status quo consumption is r , and he now evaluates x . When will the DM over- or underreact to the difference between the two, relative to the true utility difference? More formally, assuming $x > r$, when is $\mathbb{E}_{s_x,0}[\mathbb{E}[u_x - u_r | s_x, 0]] > u_x - u_r$?

Proposition 2 implies that the DM will overreact to small changes provided that the ordinal signal is not too noisy, and that he will underreact to large changes. The first effect is driven by the boost from the comparative signal and the second one by Bayesian attenuation. The dual pattern of overreaction to small and underreaction to large changes is a recurring one in the literature (e.g., Augenblick et al., 2025; Ba et al., 2025; Graeber et al., 2025; Meier et al., 2025).

4.4 Coherent Arbitrariness

In an influential paper, Ariely et al. (2003) make the observation that experimental subjects' initial valuations of an option are fairly arbitrary, yet their subsequent decisions are coherent in the sense that higher prices lead to lower demand, or higher quality to higher willingness-to-pay. Ariely et al. (2003) discuss these results as follows. "Suppose that the subject indicates, for whatever reason, that she would be willing to purchase the average bottle [of wine] for \$25. If we were to ask her a moment later whether she would be willing to purchase the "rare" bottle [of wine] for \$25, the answer would obviously be "yes," because from her perspective this particular "choice problem" has been solved and its solution is known: if an average wine is worth at least \$25, then a rare wine must be worth more than \$25!"

We believe that our model provides a natural interpretation of these results and intuitions. First, in the model, initial evaluations of u_x are noisy (and regressive to a prior) and thus, to some extent, arbitrary.¹² Yet once the DM also observes an ordinal signal about how x compares to the previous option (the reference point), subsequent evaluations will be coherent in the sense that they will respect the ordinal ranking implied by the option attributes. In our model, highly noisy absolute evaluations and coherent responses to changes are entirely compatible with each other.

4.5 'Evaluability' in Joint and Separate Evaluation

A prominent result in research on judgment and decision making is that people's valuations for two different goods are strongly influenced by whether they are evaluated separately or jointly. According to the 'evaluability hypothesis' due to Hsee (1996), people's valuations are considerably more strongly influenced by difficult-to-value dimensions when two options are evaluated jointly rather than separately.

To illustrate, suppose an experimental subject is asked to state their WTP for two two-dimensional goods, where each option is strictly better on one dimension. These experiments

¹²Our model predicts a weak version of 'arbitrariness' (that initial valuations are uncertain and noisy), not necessarily a strong one (that initial valuations are determined by random anchors).

are designed such that one dimension is very difficult to value (e.g., the units are very unfamiliar) and one relatively easy to value. Then, the canonical result is that in separate evaluations the difference in WTP between the two options is largely determined by the easy-to-value dimension, while when both options are evaluated jointly, the difficult-to-value dimension has a substantial effect as well. This often produces reversals in the WTP ranking between the two options, such that the DM has a higher WTP for the option that is better on the easy-to-value dimension in separate evaluation, but a higher WTP for the option that is better on the difficult-to-value dimension in joint evaluation.

Our model naturally captures this phenomenon. In separate evaluations, the DM only receives cardinal absolute signals in each dimension, and because they are much more noisy in one dimension than the other (producing Bayesian shrinkage to the prior), differences in valuations across the two goods are largely driven by the easy-to-value (less noisy) dimension.

In joint evaluations, on the other hand, we assume that the DM evaluates each option by using the other option as a reference point, thus generating two comparative ordinal signals (one in each dimension). Crucially, however, the impact of these ordinal signals will have a larger impact on final valuations in the difficult-to-value dimension because—as shown in Proposition 15—the size of the boost that is generated by the ordinal signal increases in the DM’s interim value uncertainty. Intuitively, our DM does not know how to cardinaly value the difficult-to-evaluate dimension (meaning it has little impact in separate evaluations) but he does know which option is ordinaly better, which helps him understand relative rankings.

4.6 Within- vs. Between-Subjects Experiments

It is well-understood in experimental economics that within-subjects designs potentially produce confounds that are immediately related to contrast effects. In reviews of the literature, Charness et al. (2012) and List (2025) suggest that subjects may be more responsive to treatments in within-subjects designs because “subjects have a reference or comparison point when responding to the second question” (Charness et al., 2012). Indeed, Schäfer and Schwarz (2019) report that the average effect size in within-subject psychology experiments is significantly larger than that in between-subject ones.

Consider one of the leading examples in Charness et al. (2012), a willingness-to-pay elicitation. Suppose that in the within-subjects experiment, the DM first evaluates x and then y (by generating signals s_x and s_y), where WLOG $u_x < u_y$; suppose the DM has identical priors between them. Then, the average difference in utility evaluations in the between-subjects design

(without an ordinal signal) is

$$\mathbb{E}_{s_x, s_y} [\mathbb{E}[u_x | s_x] - \mathbb{E}[u_y | s_y]],$$

while that in the within-subjects design is

$$\mathbb{E}_{s_x, s_y, o} [\mathbb{E}[u_x | s_x] - \mathbb{E}[u_y | s_y, o]].$$

But the latter is smaller (more negative), because the ordinal signal boosts the evaluation of u_y away from u_r . (This is a direct consequence of Proposition 15.). Thus, our model predicts stronger within- than between-subject effects.

4.7 Income Targeting and Bunching

A canonical implication of reference dependence is bunching right at, or just above, a reference point, such as when workers' accumulated earnings bunch at an income target (Camerer et al., 1997; Abeler et al., 2011; DellaVigna et al., 2017, 2022) or when the performance of athletes bunches at salient goals or reference points (e.g., Pope and Schweitzer, 2011; Allen et al., 2017). We now illustrate that our model also leads to bunching, focusing on the standard example of labor supply; the same intuition holds more generally.

A DM chooses hours a to work at wage $w > 0$. Utility is $U(a, w) = u(aw) - a^2/2$, where $u(aw) = \beta aw$. In general, the DM may face value uncertainty in both the earnings and effort-cost dimensions. We here focus on value uncertainty in the money dimension. This could be interpreted either as saying that there is more utility uncertainty about what to do with the money earned, or that the reference point is more precise in the earnings dimension than in the effort dimension.¹³ Value uncertainty in the earnings dimension, $u(aw)$, can reflect uncertainty about β , the wage, or the mapping between them.

For each a , the DM has a prior about $u(aw)|a \sim \mathcal{N}(\overline{\beta w a}, \sigma_p^2)$ and receives cardinal signals, $s_a = u(aw) + \epsilon^a$, $\epsilon^a \sim \mathcal{N}(0, \sigma_\epsilon^2)$, where $\overline{\beta w}$ could be interpreted as average marginal benefit in the past. Note that, to be realistic, we assume the prior mean varies with the action so that, in the prior, working more hours delivers higher utility in the earnings dimension. The DM has a reference point given by earnings of r and knows that the utility of this earnings level is $u_r = \beta r$. For simplicity, we assume they receive a noiseless ordinal signal o_a about how each $u(aw)$ compares with the utility of the reference point. (Introducing ordinal noise leads to

¹³Our results would reverse if we assumed the opposite—value uncertainty or a precise reference point only in the effort-cost dimension—just as standard loss-aversion models would predict effort targeting rather than income targeting when loss aversion applies to effort but not to money.

similar results, though ‘bunching’ is less sharp.)

Given that our DM understands monotonicity between the reference point and his earnings, it is further realistic to assume that the DM also receives ordinal signals about the dimension-specific utility ranking between different actions: he should know that working six hours delivers higher earnings than working five. This means that, for each action (except boundary actions), the DM receives two additional dimension-specific ordinal signals—comparing with the closest higher and the closest lower actions; without these, independent error realizations may lead to non-monotone posteriors, which is unrealistic. However, computing the Bayesian posterior mean for this rich signal structure is analytically intractable due to the interdependent truncations introduced by the many ordinal signals across actions. Fortunately, the posterior mean about the utility of each action can be well approximated using an alternative tractable model.

A simpler problem. Consider the problem above, but suppose instead that the error term in the cardinal signals is perfectly correlated across actions— $s'_a = u(aw) + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$ —yet the DM computes posterior beliefs ignoring this correlation. Because the error term in the cardinal signal is correlated, in this version, posterior beliefs always respect monotonicity. Crucially, Appendix F.3 presents simulations to show that the distribution of posterior means is remarkably similar to that induced by the original model with all ordinal signals; we will therefore focus on this simpler version, which can be solved using standard tools.¹⁴ For a given realization of ϵ , the DM’s maximization problem is given by

$$\max_{a \in \mathbb{R}_+} \lambda(\beta wa + \epsilon) + (1 - \lambda)\overline{\beta wa} + \sigma_x \psi \left(\frac{\lambda(\beta wa + \epsilon) + (1 - \lambda)\overline{\beta wa} - \beta r}{\sigma_x}, o_a \right) - \frac{a^2}{2} \quad (9)$$

where $o_a = +$ if $aw \geq r$, and $o_a = -$ otherwise.¹⁵ Denote $a^*(w)$ its solution for any $w > 0$. This is a standard maximization problem except for the discontinuous, positive jump at $a = \frac{r}{w}$. It will be useful to compare its solution to the solution of the equivalent problem without reference point, call it $a^{\text{no rp}}(w)$.¹⁶

Proposition 4. *There exist $0 < w_1 < w_2 < w_3$ such that:*

¹⁴To see why this is a good approximation, note that the models are identical except that the former has i.i.d. noise terms but ordinal signals across actions, and the latter has perfectly correlated noise but no ordinal signals across actions. While the ordinal signals of the former model yield monotone posteriors, the latter model achieves them automatically. But then, calibrating the noise terms will lead to a very similar distribution of posteriors outside boundary actions—precisely what the simulation shows.

¹⁵This relies on the simplifying assumption that the DM receives a positive ordinal signal in the case of equality.

¹⁶That is, $a^{\text{no rp}}(w)$ is the solution to $\max_{a \in \mathbb{R}_+} \lambda(\beta wa + \epsilon) + (1 - \lambda)\overline{\beta wa} - \frac{a^2}{2}$, meaning $a^{\text{no rp}}(w) = \lambda\beta w + (1 - \lambda)\overline{\beta w}$.

1. (Bunching and downward-sloping labor supply). $a^*(w) = \frac{r}{w}$ and $\frac{\partial a^*(w)}{\partial w} = -\frac{r}{w^2} < 0$ for all $w \in (w_1, w_3)$.
2. (Bunching of earnings). $a^*(w)w$ is strictly increasing for $0 < w < w_1$, is discontinuous at w_1 , is constant and equal to r for $w \in (w_1, w_3)$, is discontinuous at w_3 , and is strictly increasing for $w > w_3$.
3. (Comparison with reference-free). $a^*(w) < a^{\text{no rp}}(w)$ for $0 < w < w_1$ and $w > w_2$, $a^*(w) > a^{\text{no rp}}(w)$ for $w_1 < w < w_2$.

Figure 4 illustrates these results. Without reference points, the optimal action is strictly increasing in w , with standard behavioral attenuation (Enke et al., 2024). Adding reference effects (i) decreases the expected value of earnings below r and (ii) increases it above that. This has several implications. When wages are low, the reference point reduces the perceived marginal benefit of earnings—recall that the boost is negative and decreasing as we approach the reference point from below, see Figure 2. Lower marginal benefits then imply a lower optimum given the convex cost of effort, and the DM *reduces* their effort compared to the case without reference points—a discouragement effect. This effect continues (and becomes stronger) until wages grow sufficiently high (w_1 in the proposition) that it is worth “pushing it” and reaching the reference point r : at that point, both action and earnings discontinuously jump to meet the earnings threshold as the DM seeks to avoid the negative boost. As wages further increase, however, optimal earnings stay put at r : as long as we are above the reference point, marginal benefits of effort drop and stay low. In turn, this means that a^* decreases as the wage increases (since wa^* stays constant). This continues until wages grow enough that increasing the action is again optimal—the action jumps discontinuously and converges (from below) to $a^{\text{no rp}}(w)$. While Figure 4 plots the actions and earnings of a single DM (with a given noise realization), averaging across multiple DMs with independent noise realizations would give smoother patterns.

5 A Theory of Endogenous Reference Points

Empirical studies have shown that many alternatives can serve as reference points, including the status quo, expectations, memory, or the outcomes of others. Our model offers a novel approach to predicting what the reference point will be in a given situation. Since we understand reference points as a source of information, it is natural to look at which reference point is *most informative*— i.e., reduces value uncertainty and thus leads to the best subsequent choice, in expectation. This provides a simple, tractable criterion for endogenizing the reference point, grounded in information theory.

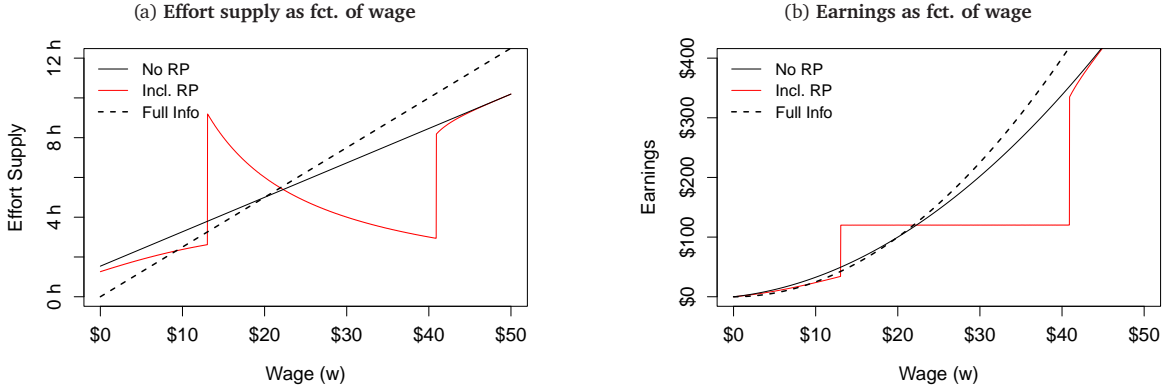


Figure 4: Effort supply and earnings as a function of the wage, with and without reference point. Parameters are $\bar{\beta} = 1$, $\beta = 1$, $\bar{w} = 20$, $\gamma = 4$, $\sigma_\epsilon = 80$, $\sigma_p = 120$, and reference earnings of \$120. Both panels show a single DM with noise realization $\epsilon = 0$.

We begin by making several observations about our approach. First, it may seem to demand an unrealistic level of sophistication to assume that people can “choose” the reference point. However, this need not be a conscious choice; rather, as is common in decision science, this selection may be an unconscious background process. Much like it is widely accepted that people often unconsciously bring to mind pertinent memories for each situation, they may also make use of a reference point that is most relevant for a given choice situation. For instance, when we evaluate the utility from a high-quality wine, we arguably compare it with other high-quality wines we’ve had in the past, rather than with the worst wine we’ve ever tried. Moreover, well-documented patterns in selective and associative memory may in fact implement precisely this kind of relevance-based retrieval mechanism.

Second, we formulate a model in which a different reference point can be selected for each choice situation. While theoretically convenient, in practice it is more plausible that the same reference point is used across a certain *class* of choice situations. Our approach could easily be extended in this way.

Third, we assume that any reference point can be used at zero cognitive cost. This abstracts from the possibility that some reference points could exogenously be made transparently available at low or zero cost (making them more likely to get adopted), for example when an experimenter induces a reference point.

Fourth, we restrict ourselves to a single reference point. Of course, the logic of reference points as information suggests that people may compare the current choice option with many past experiences. Without explicitly modeling it, we assume that it is infeasible to select multiple reference points. Relaxing this assumption represents an interesting potential future extension.

Fifth, we are agnostic about the nature of the mental construct that may serve as a refer-

ence point. As we show below, a DM may wish to adopt a reference point that aligns with the expected outcome they will face. Most naturally, past experiences and memories can play this role; however, reference points may also be formed through mental simulation or imagination, even in the absence of direct experience. For example, a DM may never have lived in Los Angeles but may nonetheless form an imagined representation of life there and use it as a benchmark when evaluating other cities. We thus allow for a broad class of potential reference points, with potentially infinite ordinal or reference point uncertainty to capture the idea that unexperienced or weakly imagined reference points are not very useful.

5.1 The Informational Value of a Reference Point

We first endogenize the reference point in a setting in which the DM is asked to value a single option. We extend to a choice between two options in Section 6.6. A DM knows they will face option x with utility $u_x \sim \mathcal{N}(\tilde{u}_x, \sigma_p^2)$. They wish to understand the value of reference point r , of which they know $u_r \sim \mathcal{N}(\tilde{u}_r, \sigma_{p,r}^2)$. They also know that, at the time of evaluation, they will receive (i) an ordinal signal o about u_x vs. u_r with noise σ_o^2 and (ii) cardinal signals, $s_x = u_x + \epsilon_x$ and $s_r = u_r + \epsilon_r$, with $\epsilon_x \sim \mathcal{N}(0, \sigma_{\epsilon_x}^2)$, $\epsilon_r \sim \mathcal{N}(0, \sigma_{\epsilon_r}^2)$.¹⁷

After observing all signals, the DM infers u_x and chooses action a to minimize the Mean Square Prediction Error (MSPE). We can think of a as the agent's estimate of u_x . We can then define the benefit of reference point r as the difference in the expected MSPE with and without the ordinal signal it generates, that is,

$$V(r) := \underbrace{\mathbb{E}_{s_x} [\min_{a \in \mathbb{R}} \mathbb{E}[(u_x - a)^2 | s_x]]}_{\text{Expected MSPE without } r} - \underbrace{\mathbb{E}_{s_x, s_r, o} [\min_{a \in \mathbb{R}} \mathbb{E}[(u_x - a)^2 | s_x, s_r, o]]}_{\text{Expected MSPE using } r}.$$

It turns out that this measure admits a particularly simple characterization and comparative statics. (We use the same notation as in previous sections: $\tilde{u}_x^{s_x}$ and σ_x^2 are the posterior mean and variance of the belief about u_x after observing s_x ; the equivalent is true for $\tilde{u}_r^{s_r}$ and σ_r^2 .)

Lemma 1. *The informational value of reference point r in evaluating x is*

$$V(r) = \frac{\sigma_x^4}{\sigma_x^2 + \sigma_r^2 + \sigma_o^2} \mathbb{E}_{s_x, s_r} \left[h \left(\frac{\tilde{u}_r^{s_r} - \tilde{u}_x^{s_x}}{\sqrt{\sigma_x^2 + \sigma_r^2 + \sigma_o^2}} \right) \right] \quad (10)$$

¹⁷We focus on one-dimensional objects for simplicity; extending to multiple dimensions is straightforward. By considering $\sigma_p^2, \sigma_{p,r}^2$ separately from $\sigma_{\epsilon_x}^2$ and $\sigma_{\epsilon_r}^2$, we allow a general treatment that allows for cases in which additional information is going to be revealed or not. Moreover, this allows for a general treatment of whether the evaluation of the reference point occurs before memories are accessed (high σ_x^2, σ_r^2 , low $\sigma_{\epsilon_x}^2, \sigma_{\epsilon_r}^2$), or after they are ($\sigma_{\epsilon_x}^2, \sigma_{\epsilon_r}^2 \rightarrow \infty$).

where $h(x) := \frac{\phi^2(x)}{\Phi(x)(1-\Phi(x))}$.

The Lemma shows that the informational benefit of a reference point is captured by a simple function which, unsurprisingly, is formally related to that of our boosts—indeed, one may notice that h is the product of the boosts in both directions. Note that when no further information is expected (that is, $\sigma_{\epsilon_x}^2, \sigma_{\epsilon_r}^2 \rightarrow \infty$), Lemma 1 further simplifies to

$$V(r) = \frac{\sigma_x^4}{\sigma_x^2 + \sigma_r^2 + \sigma_o^2} h\left(\frac{\tilde{u}_r - \tilde{u}_x}{\sqrt{\sigma_x^2 + \sigma_r^2 + \sigma_o^2}}\right). \quad (11)$$

The most informative reference point and the bias-variance trade-off. The simple functional form in Lemma 1 leads to a first comparative static on the information benefits of a reference point.

Proposition 5. *If $\sigma_x^2 > 0$, $V(r)$ is strictly decreasing in $|\tilde{u}_r - \tilde{u}_x|$ and in $\sigma_{\epsilon_r}^2$.*

This result provides a simple rule of thumb to assess the information value of reference points. First, reference points are more informative the more the DM knows about them—a small $\sigma_{\epsilon_r}^2$. This is immediate: the more *precise* the DM’s signal about r , the more informative any ordinal comparison is going to be. Second, reference points are more informative the closer they are to the *mean* of the distribution of u_x , that is, with a small *bias* $|\tilde{u}_r - \tilde{u}_x|$. To see why, note that when the reference point is centered exactly at \tilde{u}_x , we have an equal chance of either ordinal signal, and both ordinal signals are moderately informative. If, instead, \tilde{u}_r is far above \tilde{u}_x , the ordinal signal $u_x < u_r$ is almost guaranteed and (correspondingly) not very informative, while the ordinal signal $u_x > u_r$ would be very informative, but also very unlikely.¹⁸

Figure 5 visualizes these results. These two criteria in evaluating a reference point—how well-known and how well-centered it is—imply the existence of a bias-variance tradeoff. Yet when all potential reference points are equally well-known, the criteria also have a simple implication, which is that the expected outcome is the most informative reference point. This connects to theories in which the reference point is given by average outcomes or expectations, except that here it arises endogenously and purely for informational reasons.

Expected outcome vs. status quo. The tradeoff above has direct implications for when the status quo will be the most informative reference point. Consider a DM looking to buy a car. As a reference point, they can either use the car they currently own, the status quo (SQ). They know this car well ($\sigma_{p,rSQ}^2 = 0$), but it is likely to be significantly worse than the ones in the store

¹⁸Comparative statics with respect to $\sigma_{p,r}^2$ and σ_o^2 are also intuitive and discussed in Proposition 6 below.

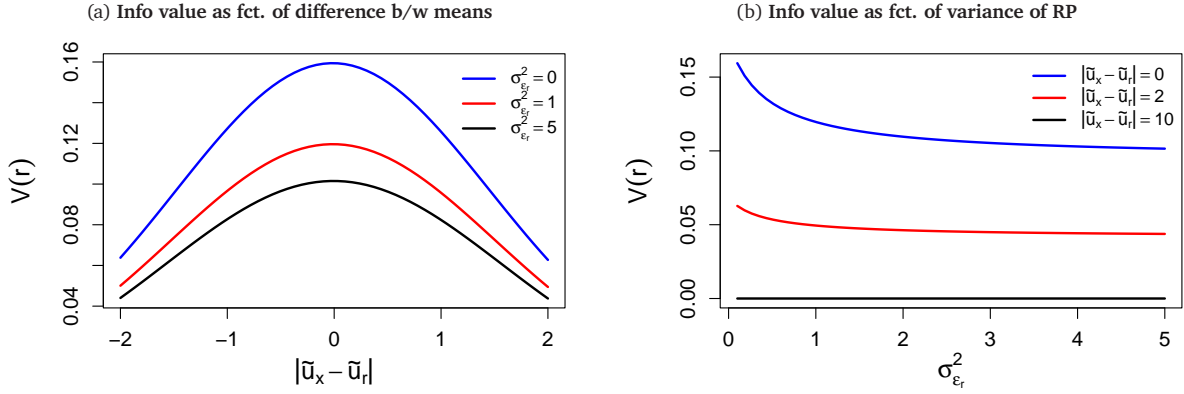


Figure 5: Comparative statics about the informational value of reference points. Left panel: $V(r)$ as a function of $|\tilde{u}_x - \tilde{u}_r|$. Right panel: $V(r)$ as a function of $\sigma_{\epsilon_r}^2$.

($\tilde{u}_{r_{SQ}}$ is well below \tilde{u}_x). Alternatively, he can use a car that his friend recently bought (A), which he knows less well ($\sigma_{p,r_A}^2 > 0$), but is more representative of the cars in the store ($\tilde{u}_{r_A} = \tilde{u}_x$).

The DM has a choice between a backward-looking reference point—the status quo, which is well known, but far away from what they expect to have to evaluate—or a forward-looking one—which is less well known but is selected to be close to the expected outcome. Our results give a precise rule for which is more informative. The optimal reference point is the status quo if it is not *too* far from what they expect to see in the store. There is an exact, unique threshold: Once $\tilde{u}_{r_{SQ}}$ and \tilde{u}_x are sufficiently far apart, the friend’s car A will be more informative. Similarly, the better the alternative reference point A is known (lower $\sigma_{\epsilon_{r_A}}^2$), the lower is this threshold. This can be formalized into a simple formula, a direct implication of Lemma 1. (For completeness, the Appendix includes a proof.)

Observation 1. Consider r_{SQ} and r_A such that r_{SQ} is biased ($b_{SQ} := |\tilde{u}_{r_{SQ}} - \tilde{u}_x| > 0$) but perfectly known ($\sigma_{p,r_{SQ}}^2 = 0$) while r_A is unbiased ($\tilde{u}_{r_A} = \tilde{u}_x$) but imperfectly known ($\sigma_{p,r_A}^2 > 0$). Suppose for simplicity that no further information is expected and that there is no ordinal noise for either ($\sigma_{\epsilon_x}^2, \sigma_{\epsilon_{r_A}}^2 \rightarrow \infty, \sigma_{o_{SQ}}^2 = \sigma_{o_A}^2 = 0$).¹⁹ Then, r_{SQ} is a better reference point than r_A if and only if its bias is small enough relative to the uncertainty about x , that is,

$$V(r_{SQ}) > V(r_A) \Leftrightarrow b_{SQ} < b^* \quad \text{where} \quad b^* := \sigma_x h^{-1} \left(\frac{2}{\pi} \frac{\sigma_x^2}{\sigma_x^2 + \sigma_{r_A}^2} \right) > 0.$$

¹⁹These results can be easily extended to the general case.

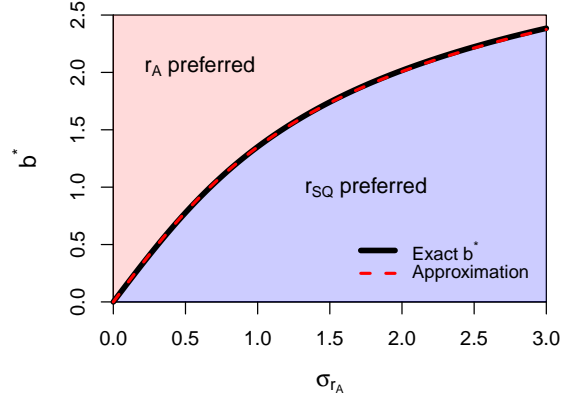


Figure 6: Bias-Variance tradeoff in most informative reference point. It assumes that r_{SQ} and r_A are such that r_{SQ} is biased ($b_{SQ} := |\tilde{u}_{r_{SQ}} - \tilde{u}_x| > 0$) but perfectly known ($\sigma_{p,r_{SQ}}^2 = 0$) while r_A is unbiased ($\tilde{u}_{r_A} = \tilde{u}_x$) but imperfectly known ($\sigma_{p,r_A}^2 > 0$). It also assumes $\sigma_p^2 = 1$, no further information expected ($\sigma_{\epsilon_x}^2, \sigma_{\epsilon_{r_A}}^2 \rightarrow \infty$), and no ordinal noise ($\sigma_{o_{SQ}}^2 = \sigma_{o_A}^2 = 0$).

Moreover, b^* can be well-approximated by²⁰

$$b^* \approx \frac{3}{2} \sigma_x \sqrt{\log \left(1 + \frac{5}{4} \cdot \frac{\sigma_{r_A}^2}{\sigma_x^2} \right)}.$$

Figure 6 plots the correct and approximated threshold.

Assuming that the status quo becomes better known over time, a direct implication of our model is that the “tenure” of an option as a status quo determines how likely it is to serve as a reference point. Central to this argument is the information value: a “just-adopted” status quo, that is less known, is also a less informative reference point.

Our results also speak to the experimental literature on reference points. A frequently made informal argument in this literature is that reference points (in particular expectations-based ones) need to “sink in” to affect behavior (see Heffetz, 2021, for a literature review). From our model’s perspective, the simple explanation is that reference points are adopted only when the DM has accumulated sufficient information about their utility values.

The role of ordinal noise. An additional comparative static result pertains to ordinal noise (and prior variance):

Proposition 6. *Suppose $\sigma_x^2 > 0$. Then, there exists $\bar{\delta}$ such that if $|\tilde{u}_r - \tilde{u}_x| < \bar{\delta}$, $V(r)$ is strictly decreasing in σ_o^2 and $\sigma_{p,r}^2$. If $|\tilde{u}_r - \tilde{u}_x|$ is large enough, however, $V(r)$ is strictly increasing in both.*

²⁰Numerical simulations suggest errors are typically below 1.1% in the relevant range when $\sigma_{r_A} \leq 3\sigma_x$.

As long as the bias is not too big, the DM prefers a reference point with smaller ordinal noise (and a smaller prior variance). This is intuitive, as less noise in the ordinal signal implies better information.²¹ This means that reference points will be particularly useful when they allow precise ordinal comparisons, for example, because they are nearby or because they are presented in similar formats and units. For these reasons, a DM may adopt a reference point that is easier to compare, even if it is less well-calibrated or less well-known.

5.2 Cues and Similarity-Based Memory Reference Points

A long literature in psychology formalizes memory through similarity-based recall (e.g. Tulving and Thomson, 1973; Shepard, 1987): essentially, the probability that a memory is recalled is proportional to its similarity to current circumstances. Building on this work, (e.g., Bordalo et al., 2020) model a DM whose reference point is given by similarity-based memory. Our model provides one potential microfoundation for this mechanism.

As noted above, when all potential reference points are equally well-known, the most informative one is given by the expected (average) outcome. Oftentimes, however, the DM may be unsure of the distribution of outcomes he faces and thus cannot directly use the expected utility outcome as a reference point. In this case, our DM tries to infer the average outcome from “cues” in the environment. For example, entering a new store, the DM may not know the average utility of clothes in this store, but he may observe contextual cues—how fancy is the store, are there mice on the floor? From these cues, the DM infers “what type of store” it is, and which reference point will be most informative to retrieve from memory. The idea is that the customer has different experiences stored in memory, and the cue provides information about which will one be most useful as a comparator. For instance, when the customer is served a glass of champagne upon entering a store, it is not particularly informative to retrieve as reference point the utility of the clothes from his college days.

Memory cues. Suppose the DM believes $u_x \sim N(\tilde{u}_x, \sigma_x^2)$, where the average outcome is uncertain itself, $\tilde{u}_x \sim N(\hat{u}_x, \sigma_\mu^2)$. Suppose that there is a cue c that is perfectly informative about the average outcome \tilde{u}_x (this is just for simplicity and easily generalizes). The analyst observes this cue, while the DM observes it with some noise—he receives $c_o = c + e$, where $e \sim N(0, \sigma_e^2)$. The DM chooses a reference point from the set R of possibilities.

²¹Notably, this holds only if the bias is not too large: when the bias is very large, then the DM may actually prefer some ordinal noise or prior variance about r . The reason is the same as that of the non-monotonicity of the boost w.r.t. σ_r^2 and σ_o^2 discussed in Section 3.4. When the bias is large the DM has little doubt about the ordinal signal they will get, meaning that they expect little information; however, adding some ordinal noise makes this expected ordinal signal *more* informative, because it lowers the (relative) chance of signals from values which are very close.

Imagine, for simplicity, that the DM has a wide set of potential reference points stored in memory. Given a cue c , what are the chances that the DM retrieves from memory the reference point $r = m$?

Corollary 1. *Suppose all potential reference points are available for a given noise level (that is, $\sigma_{p,r}^2$ and σ_e^2 are the same for all $r \in R$; and for each $z \in \mathbb{R}$, there exists $r \in R$ with $\tilde{u}_r = z$). Then, the probability density that m is the most informative reference point is*

$$\Pr(r = m|c) = \frac{a e^{-\tau(c-m)^2} \Pr(m)}{N} \quad (12)$$

where $a := \frac{1}{\sqrt{2\pi} \sigma_{cond}^2}$, $\tau := \frac{1}{2\sigma_{cond}^2}$, $\sigma_{cond}^2 := \frac{\sigma_x^2 \sigma_e^2}{\sigma_x^2 + \sigma_e^2}$ and N is a normalizing term.

This corollary follows immediately from our results.²² It says that the probability of a particular memory becoming the reference point increases in its similarity to the cue. The most informative reference point is given by the expected mean of the object to be evaluated, conditional on the cue received. The probability of a given memory being used is therefore the probability of a cue that gives rise to such expectations, which is decreasing in the distance between the memory and the true situation.

Remarkably, the functional form in Eq. (12) is precisely the rule commonly used for recall or memory-based reference points in psychology (Kahana, 2012) and economics (Bordalo et al., 2020). The only difference is that here this is derived as optimal for a decision maker who has value uncertainty and uses the informational value of reference points.²³

Summary. Viewing reference points through the lens of the most informative comparator for choice by reducing value uncertainty generates several natural predictions and links to earlier literature. We can summarize the results with four principles:

1. When the DM knows the distribution of outcomes and all potential reference points are equally well known, the most informative reference point is the one that aligns with the expected outcome (rational expectations).

²²To see why, notice that, if all potential reference points are available and have the same noises, then by Lemma 1, the agent will choose the reference point r with expected utility equal to the posterior about \tilde{u}_x ; that is, the probability that a given r is chosen is equal to the probability that the posterior about \tilde{u}_x is equal to r . By Bayes' rule, the probability that the posterior is equal to r when the cue is c is equal to the probability that the cue was c given that the posterior is equal to r times the prior of r , and normalized. But the probability that the cue is c given that the posterior is equal to r —thus that $\tilde{u} = c$ —is simply the density of the posterior, $\frac{1}{\sqrt{2\pi} \sigma_{cond}^2} \exp\left(-\frac{(c-m)^2}{2\sigma_{cond}^2}\right)$, giving us the formula above.

²³In this sense, our model builds on an earlier literature in psychology that emphasizes that similarity-based recall is sometimes optimal based on Bayesian principles (e.g., Anderson, 2013; Malmendier and Wachter, 2024).

2. When the DM knows the distribution of outcomes, yet some potential reference points are better known than others, a bias-variance tradeoff exists. If, for example, the status quo is known much better than the expected outcome, then it is the most informative reference point if it is not too biased.
3. In addition to the above, the DM will choose reference points that are easier to compare—for example, they are expressed in the same units, or in similar formats, or adjacent.
4. When the DM does not know the (mean of) the distribution of utility outcomes and all potential reference points are equally well known, the most informative reference point is determined by similarity-based recall.

Finally, notice how our results could, in principle, also be used to study costly reference points. For instance, it may be cognitively costly to retrieve temporally distant experiences from memory, or it may be costly to conduct the ordinal comparisons that the reference point enables. Our machinery can be applied to this question because Eq. (10) gives the informational value of a reference point, which can be compared with its cost.

6 Discrete Choice

Our analysis thus far focused on the case in which the DM evaluates a single option in the presence of a reference point. However, in practice, many contexts will potentially give rise to ordinal comparisons, potentially multiple ones. This may be the case with binary choice with or without a reference point, choices from more than two items, or multiple reference points—one of the advantages of thinking about reference dependence from an informational perspective is that it naturally allows for multiple reference points to play a role.

We now apply the same “technology” introduced above to study these settings, with a particular emphasis on binary choice in the presence of a reference point (or other comparison point). Throughout, we assume that the DM also receives ordinal signals on how available options compare in each dimension, in addition to the comparative signals with respect to a potential reference point. For instance, the DM not only knows that seven apples are better than the reference point of three apples, but also understands that seven apples are better than the alternative choice option of five apples.²⁴

²⁴This has an immediate implication: when there is no ordinal noise, the DM will never choose a dominated option. This is, of course, in line with empirical evidence, which shows that, when dimensions are easy to ordinally compare—e.g., involve numbers with clear rankings—the probability of choosing dominated options vanishes; however, this is less true when ordinal comparisons are less transparent (e.g., Choi et al., 2007; Nosofsky, 1986; Tversky, 1969), as predicted by our model when we add ordinal noise.

We proceed as follows. First, we introduce the formal results required to analyze situations with multiple ordinal signals—as would arise, for example, when an option is compared with both a reference point and another choice option. Second, we apply this general machinery to four special cases emphasized in the literature because of their associated robust empirical regularities:

1. A pure binary choice setting—or, with a reference point with infinite uncertainty—, where we show our model gives rise to the heuristic of “counting advantages” (Gigerenzer and Goldstein, 1996; Dawes, 1979).
2. A binary choice setting with an asymmetrically dominated comparator, giving rise to the asymmetric dominance decoy effect.
3. A binary choice setting comparing a reference point that is globally dominated vs. involves tradeoffs, giving rise to the Kahneman-Tversky improvements-vs.-tradeoffs effect.
4. A binary choice setting with a comparator that extends the range of options in the set, giving rise to range-based contrast effects.

Finally, we consider what the informationally-optimal reference points for the case of binary choice are.

6.1 Boosts with Multiple Comparators

When Comparators are Known. Suppose first that the DM gets ordinal signals with respect to multiple options, all of which are perfectly known, and the ordinal signal is noiseless. Let C be the (finite) set of comparators and focus on a specific dimension ($n = 1$); with multiple ones, the analysis simply applies to each dimension independently. If, for some $r \in C$, $u_r = u_x$, then posteriors are degenerate at u_r . If $u_r < u_x$ or $u_r > u_x$ for all $r \in C$, then the situation collapses to that in which the DM has a single comparator equal to the largest or smallest u_r , respectively; we can then simply use the formulas derived in the previous section. Otherwise, let $u^* := \arg \min_{r \in C | u_r > u_x} u_r$ and $u_* := \arg \max_{r \in C | u_r < u_x} u_r$. Then, the posterior mean about u_x is simply that of a two-sided truncated Normal,

$$\mathbb{E}[u_x | o, s] = \tilde{u}_x^s + \sigma_x \frac{\phi\left(\frac{u_* - \tilde{u}_x^s}{\sigma_x}\right) - \phi\left(\frac{u^* - \tilde{u}_x^s}{\sigma_x}\right)}{\Phi\left(\frac{u^* - \tilde{u}_x^s}{\sigma_x}\right) - \Phi\left(\frac{u_* - \tilde{u}_x^s}{\sigma_x}\right)}.$$

This “boost” can be positive or negative. If both u^* and u_* are below \tilde{u}_x , the boost is negative; if both are above \tilde{u}_x , the boost is positive. If $u^* > \tilde{u}_x > u_*$, the sign of the boost depends on whether \tilde{u}_x is closer to u^* or u_* : negative in the former case, positive in the latter.

Comparators that are not perfectly known. When the utilities of comparators are not perfectly known or when ordinal information is noisy, formulas become more complex. In Balakrishnan and Dean (2025), we derive the solution for two comparators; we include the precise formulae in the Appendix. Independent of the specific formulas, however, posteriors with multiple comparators follow several regularities. Two comparators that provide the same ordinal information (i.e., two positive or two negative boosts) reinforce each other, providing a bigger boost than either on its own; with conflicting information, the net boost depends on the mean distance between each comparator and the target and the relative variance of the two comparators. Moreover, the following proposition includes three additional regularities that will be particularly useful in what follows.

Proposition 7. *Let u_x, u_y, u_z be independent and Normally distributed. Then*

$$\begin{aligned} \mathbb{E}[u_x - u_y | u_x > u_y] &> \mathbb{E}[u_x - u_y | u_x > u_y > u_z] \\ \mathbb{E}[u_x - u_y | u_x > u_y] &> \mathbb{E}[u_x - u_y | u_z > u_x > u_y] \\ \mathbb{E}[u_x - u_y | u_x > u_y] &< \mathbb{E}[u_x - u_y | u_x > u_z > u_y]. \end{aligned}$$

Suppose that we are estimating the difference between u_x and u_y knowing that $u_x > u_y$. Then, the proposition shows how this estimate *shrinks* if we further learn that some u_z is above or below both ($u_x > u_y > u_z$ or $u_z > u_x > u_y$); intuitively, $u_x > u_y > u_z$ boosts u_y more than it does u_x , because u_y is necessarily closer to u_z , and we have already seen how this implies a bigger boost. On the other hand, this estimate grows if we learn that some u_z lies in between ($u_x > u_z > u_y$), as this pushes u_x upwards and u_y downwards. Notably, these results hold independently of the priors on these random variables, provided they are Normal and independent.

6.2 Binary Choice Without a Reference Point: Counting Advantages

The DM needs to choose between two options, x and y . The utilities are $U(x) = \sum u_{x,i}$ and $U(y) = \sum u_{y,i}$, the DM holds dimension-specific Normal priors for the utility in each dimension (which may or not coincide for the two options), and they receive a noisy cardinal signal $s_{a,i} \sim \mathcal{N}(u_{a,i}, \sigma_{\epsilon,a,i}^2)$ for each i and $a \in \{x, y\}$. The key novelty is that the DM also receives *ordinal* signals o_i that compare $u_{x,i}$ and $u_{y,i}$ for each i , with noise term $\sigma_{o,i}^2$. Denote $D_x \in \{1, \dots, n\}$ the

dimensions where x dominates and D_y those in which y dominates. We here assume that there is no reference point (or, equivalently, that the DM has infinite uncertainty about its value). Akin to (4), we have that for each dimension:

$$\mathbb{E}[u_{x,i}|o_i, s_{x,i}] = \sum_i \tilde{u}_{x,i}^{s_{x,i}} + \frac{\sigma_{x,i}^2}{\sqrt{\sigma_{x,i}^2 + \sigma_{y,i}^2 + \sigma_{o,i}^2}} \psi\left(\frac{\tilde{u}_{x,i}^{s_{x,i}} - \tilde{u}_{y,i}^{s_{y,i}}}{\sqrt{\sigma_{x,i}^2 + \sigma_{y,i}^2 + \sigma_{o,i}^2}}, o_i\right).$$

This gives the same formula as before, except that now the comparison is with another option, and the uncertainty about the reference point is replaced by that of the other option.

Consider the natural special case in which ordinal signals are noise-free and the DM has, in each dimension, the same prior about the two goods ($\tilde{u}_{x,i} = \tilde{u}_{y,i}$, $\sigma_{p,x,i}^2 = \sigma_{p,y,i}^2 =: \sigma_{p,i}^2 \forall i$) and the same noise variance ($\sigma_{\epsilon,x,i}^2 = \sigma_{\epsilon,y,i}^2 =: \sigma_{\epsilon,i}^2 \forall i$). Then, the posterior mean for the utility of x is

$$\mathbb{E}[u_{x,i}|s_i, o_i] = \underbrace{\sum_i \tilde{u}_{x,i}^{s_i}}_{\text{Belief after cardinal signals}} + \underbrace{\sum_{i \in D_x} \frac{\sigma_i}{\sqrt{2}} \psi\left(\frac{\lambda_i(s_{x,i} - s_{y,i})}{\sqrt{2}\sigma_i}, +\right)}_{\text{Positive comparisons}} + \underbrace{\sum_{i \in D_y} \frac{\sigma_i}{\sqrt{2}} \psi\left(\frac{\lambda_i(s_{x,i} - s_{y,i})}{\sqrt{2}\sigma_i}, -\right)}_{\text{Negative comparisons}},$$

where λ_i is the dimension-specific shrinkage weight and $\sigma_i := \frac{\sigma_{\epsilon,i}^2 \sigma_{p,i}^2}{\sigma_{\epsilon,i}^2 + \sigma_{p,i}^2}$ the dimension-specific posterior variance after the cardinal signal.

The evaluation of each option is given by a combination of the absolute cardinal assessment (comprising the cardinal signal and the prior) and a weighted sum of the boosts from the positive and negative pairwise comparisons in each dimension. The importance (or weight) of the boosts depends on the DM's uncertainty σ_i about a dimension.

Counting advantages. As value uncertainty grows, ordinal comparisons become more and more important. In the extreme case in which the DM can rely only on his prior and the ordinal comparisons, this leads to an even simpler formula based on counting advantages.²⁵

Observation 2. Suppose $\tilde{u}_{x,i} = \tilde{u}_{y,i}$ and $\sigma_{p,x,i}^2 = \sigma_{p,y,i}^2 =: \sigma_{p,i}^2 \forall i$ and $\sigma_{o,i}^2 = 0$. Then, for any $u_x, u_y \in \mathbb{R}^n$, there exists $\bar{\sigma}_\epsilon^2$ such that if $\sigma_{\epsilon,x_i}^2 = \sigma_{\epsilon,y_i}^2 > \bar{\sigma}_\epsilon^2$ for all i :

1. $\mathbb{E}[\mathbb{E}[\tilde{u}_x^{s_i}|s_i, o_i]] > \mathbb{E}[\mathbb{E}[\tilde{u}_y^{s_i}|s_i, o_i]] \Leftrightarrow \sum_{i \in D_x} \sigma_{p,i}^2 > \sum_{i \in D_y} \sigma_{p,i}^2$
2. If $\sigma_{p,i}^2 = \sigma_{p,j}^2 \forall i, j$, then $\mathbb{E}[\mathbb{E}[\tilde{u}_x^{s_i}|s_i, o_i]] > \mathbb{E}[\mathbb{E}[\tilde{u}_y^{s_i}|s_i, o_i]] \Leftrightarrow |D_x| > |D_y|$.

²⁵For completeness, the Appendix includes a complete proof. For an intuition, note that, as $\sigma_{\epsilon,x_i}^2 \rightarrow \infty$, Bayesian posteriors after cardinal signals will get arbitrarily close to the priors, meaning that $\lambda_i(s_{x,i} - s_{y,i}) \rightarrow 0$, $\psi\left(\frac{\lambda_i(s_{x,i} - s_{y,i})}{\sqrt{2}\sigma_i}, +\right) \rightarrow \sqrt{\frac{2}{\pi}}$, $\psi\left(\frac{\lambda_i(s_{x,i} - s_{y,i})}{\sqrt{2}\sigma_i}, -\right) \rightarrow -\sqrt{\frac{2}{\pi}}$, and $\sigma_i \rightarrow \sigma_p$. The observation then follows from Eq. (6.2).

When the noise in the cardinal signals is high enough, the DM simply “counts” the number of relative advantages and disadvantages of x and y , weighted by the prior uncertainty in that dimension. Intuitively, if priors are the same and no cardinal signals are given, then all the DM is left with is a collection of positive and negative boosts. If prior uncertainty is the same in each dimension, then the *optimal* choice of a fully rational Bayesian agent—not an approximation heuristic—boils down to counting which option has more advantages. When prior uncertainty varies across dimensions, this acts as a weight— if the DM knew that the variance in utilities in a given dimension was close to zero, for example, then he shouldn’t attribute much information to the fact that one option is better than the other in that dimension. This connects our model to a well-known finding in the heuristics literature: simple tallying strategies that count the number of favorable cues, assigning equal weight to each, often perform remarkably well in prediction and choice tasks (Dawes, 1979; Gigerenzer and Goldstein, 1996). Our framework provides a Bayesian foundation for this observation: tallying emerges as the *optimal* decision rule when cardinal signal quality is sufficiently low, rather than as an approximation heuristic adopted due to cognitive constraints.

Sensitivity to changes in ordinal rankings. An almost mechanical implication of our model is the high sensitivity to changes in ordinal ranks. Consider x, y with n dimensions, and suppose $u_{x,i} < u_{y,i}$ for some i . Because of Bayesian attenuation, the effect of increasing $u_{x,i}$ on the probability of choosing x is going to be limited; however, this probability will increase substantially if $u_{x,i}$ is increased enough to go above $u_{y,i}$. Just like the strong effects of surpassing the reference point we documented earlier in the paper, our model also predicts strong effects of ordinally surpassing other options in a given dimension.²⁶

6.3 Decoy Effects

We can use our model to study a classic empirical pattern in choice: the asymmetric dominance, or decoy effect, a classic violation of WARP; see Soltani et al. (2012) for a review. In it, there are two dimensions, and subjects are first asked to choose between x and y , where x is better on dimension 2 and y is better on dimension 1. In a second choice set, subjects are asked to choose between x , y , and a third option z , which is dominated by x in both dimensions but better than y in dimension 2. The empirical observation is that adding z increases the fraction

²⁶Consistent with this, empirical studies show strong discrete choice shifts when an ordinal dominance relation flips. These patterns are found in multi-attribute decisions (Campbell et al., 2007), laboratory inference tasks (Bröder, 2000) and risky choice (Brandstätter et al., 2006), showing limited sensitivity to small cardinal changes but large, discrete shifts when rankings change.

of subjects choosing x . In what follows, assume that DMs receive cardinal signals $s_{i,a}$ as we have been assuming thus far, as well as noiseless ordinal signals among available options; let s and o denote the vectors of all such signals, respectively. Our model naturally gives rise to this effect, as shown by the following result.

Corollary 2. *Suppose $n = 2$ and all ordinal signals are noiseless ($\sigma_{o,i}^2 = 0 \forall i$). Then,*

$$\mathbb{E}[U(x) - U(y) | u_{y,1} > u_{x,1}, u_{x,2} > u_{y,2}, s] < \mathbb{E}[U(x) - U(y) | u_{y,1} > u_{x,1} > u_{z,1}, u_{x,2} > u_{z,2} > u_{y,2}, s].$$

This result is an immediate corollary of Proposition 7. The addition of z boosts u_x relative to u_y in both dimensions: In dimension 1, u_z boosts both u_x and u_y upwards, but as u_x is closer to u_z than u_y , it receives the bigger boost. In dimension 2, u_z boosts u_x upwards and u_y downwards, once again increasing their difference. Thus, for every signal realization, the introduction of z strictly increases the assessment of x relative to y . Given that interim beliefs have full support, this means that the probability that x is chosen (when $\mathbb{E}[U(x) - U(y) | s, o] > 0$) strictly increases, in line with the evidence. (Also in line with the evidence, our model predicts that, because x dominates z , the latter is never chosen.)

Moreover, because our effect is based on information, it does not rely on z being choosable: the exact same effects would be at play if z is a observable but cannot be chosen. This is in line with the experimental evidence of the “phantom decoy” of Soltani et al. (2012).

Natenzon (2019) and subsequently Shubatt and Yang (2024) propose explanations for the asymmetric dominance effect that are conceptually very close. However, our approach does not rely on x being better than z in all dimensions for the effect to work. Consider an extension to the example above in which there is a third dimension in which x and y are equally valued, but z is superior to both. The introduction of z would still increase the valuation of x relative to y : The effect on the first two dimensions would be as before, while the effect on the third dimension would be the same on both x and y . Indeed, our results hold almost by-dimension: the addition of an option z that is closer to asymmetric dominance increases the value of one option over the other, as shown by the following corollary (which also immediately follows from Proposition 7).

Corollary 3. *Let $z'_i = z_i$ in all $i \neq j$, $x_j > z'_j > y_j$ while either $x_j > y_j > z_j$ or $z_j > x_j > y_j$. Then $\mathbb{E}[U(x) - U(y) | s, o] < \mathbb{E}[U(x) - U(y) | s, o']$ where o and o' indicate ordinal signals of options $\{x, y, z\}$ and $\{x, y, z'\}$ respectively.*

Compromise effect? A second classic violation of the independence of irrelevant alternatives is the compromise effect. Here, if x is better on dimension 2 while y is better on dimension 1, adding a third option z that is even better than x on dimension 1 and even worse on dimension

2 tends to push people towards choosing z . Our model is consistent with either the compromise effect or its opposite, depending on the exact placement of the decoy; see Appendix B.6 for a discussion.

6.4 Improvements versus Tradeoffs

A related regularity is that, starting from a reference point—such as a status quo—, individuals exhibit a preference for options that imply improvements rather than tradeoffs relative to the reference point (e.g. Tversky and Kahneman, 1991). For instance, suppose that a DM chooses between two two-dimensional options, x and y . There are two potential reference points, r and r' . From r option x implies improvements in both dimensions, whereas option y implies tradeoffs (one improvement and one loss). Suppose also that from r' the pattern reverses. Then, the empirical regularity is that the DM is more likely to choose x when the reference point is r than when it is r' .

We now show how this follows immediately from our model. Denote o the complete collection of ordinal signals when the reference point is r and o' that when the reference point is r' . The following is an immediate corollary of Proposition 7.

Corollary 4. *Suppose $n = 2$, $u_{x,1} > u_{r,1} > u_{y,1} > u_{r',1}$ and $u_{y,2} > u_{r',2} > u_{x,2} > u_{r,2}$. Then, $\mathbb{E}_s [\mathbb{E}[U(x) - U(y)|s, o]] > \mathbb{E}_s [\mathbb{E}[U(x) - U(y)|s, o']]$.*

For purely informational reasons, a Bayesian DM who receives ordinal signals gives an advantage to the option that dominates the reference point in both dimensions rather than the one that implies a tradeoff with the reference point. Intuitively, this is because dominance over the reference point entails a strictly positive boost in both dimensions. In contrast, a tradeoff entails a positive boost in one dimension and a negative boost in the other. This should also hold in the absence of gain-loss asymmetries in preferences, such as in perceptual tasks, where preferences play no role. Indeed, in line with these predictions, Trueblood (2015) experimentally shows that people exhibit the improvement-versus-tradeoffs effect also in magnitude estimation problems in which the reference point cannot affect preferences.

6.5 Range Effects

Our model also generates *range effects*: increasing the range of alternatives in one dimension affects the relative appeal of options that look good in that dimension—as if it modified the implicit weight given to that dimension.

To illustrate, consider three options, x , y , and z . Suppose x is better than y in dimension 1, that is, $u_{x,1} > u_{y,1}$. What happens to the belief about the relative value of x and y if we change the quality of z in dimension 1 ($u_{z,1}$) so that the *range* of utilities in dimension 1 grows? Standard rational choice predicts no effect, for we are changing the value of a third option. If, instead, it changes, we say that beliefs exhibit a *range effect on dimension 1 at $\{x, y, z\}$* . Following the literature, if increasing the range increases the relative value of x vs. y —equivalent to that dimension being overweighted—we call it a *range-contrast effect on dimension 1 at $\{x, y, z\}$* .²⁷

We now show that, with no ordinal noise, our model always exhibits a range-contrast effect. To see why, note that if changing $u_{z,1}$ changes the range, then either $u_{z,1} > u_{x,1} > u_{y,1}$ or $u_{x,1} > u_{y,1} > u_{z,1}$. Suppose the former is the case. In this situation, the fact that z is better in dimension 1 than x and y lowers the value of both—it acts as a ceiling. Since the value of x is higher, however, this ceiling is more binding for x than for y , lowering the beliefs about the utility difference—indeed, Proposition 7 gives $\mathbb{E}[u_{x,1} - u_{y,1} | u_{x,1} > u_{y,1}] > \mathbb{E}[u_{x,1} - u_{y,1} | u_{z,1} > u_{x,1} > u_{y,1}]$. Now suppose that $u_{z,1}$ grows. Then, this ceiling “lifts,” weakening its effects, thus increasing the belief in the utility difference between x and y in that dimension—precisely the range-contrast effect.²⁸ (The case in which $u_{z,1}$ is worse than both, and acts as a “floor,” is obtained with a symmetric intuition.)

The following proposition formalizes this discussion.

Proposition 8 (Range Effects). *Consider x, y and z such that $u_{x,m} > u_{y,m}$ for some $m \in \{1, \dots, n\}$, and $u_{z,m}$ is not between $u_{x,m}$ and $u_{y,m}$. Suppose we have the same prior and signal precisions in all dimensions ($\tilde{u}_{j,i}, \sigma_{p_{j,i}}^2 > 0, \sigma_{\epsilon_{j,i}}^2 > 0$ are the same for $j = x, y, z$ and $i = 1, \dots, n$) and no ordinal noise in any binary comparison ($\sigma_o^2 = 0$), and let s and o denote all cardinal and ordinal signals. Then beliefs exhibit a range-contrast effect on dimension m at $\{x, y, z\}$.*

See Bondi et al. (2025) for recent results and a survey of the experimental literature on range-based contrast effects. In general, the literature is divided on when range-based contrast or the opposite dominates, with the evidence suggesting that contrast dominates when the number of dimensions is not too small.

²⁷Formally, we say that beliefs exhibit a *range-contrast effect on dimension m at $\{x, y, z\}$* if the following holds. Let s and o denote all ordinal and cardinal signals. Then, for any $i, j \in x, y, z$ with if $u_{i,m} > u_{j,m}$, letting $l := \{x, y\} \setminus \{i, j\}$,

- (i) If $u_{i,m} > u_{i,m} > u_{j,m}$, then $\mathbb{E}[\mathbb{E}[U(i) - U(j) | s, o]]$ is strictly increasing in $u_{k,1}$;
- (ii) If $u_{i,1} > u_{j,1} > u_{l,1}$, then $\mathbb{E}[\mathbb{E}[U(i) - U(j) | s, o]]$ is strictly decreasing in $u_{k,1}$.

²⁸With ordinal noise, another force is at play, and opposite effects may take place. Appendix B.7 provides a complete characterization.

6.6 Endogenous Reference Point for Binary Choice

While in Section 5 we derived the most informative reference point for valuation, we can extend this analysis to the optimal reference point for binary choice. A decision maker knows they will have to choose one of $x, y \in \mathbb{R}^n$, where $u_{x,i} \sim N(\tilde{u}_{x,i}, \sigma_i^2)$ and $u_{y,i} \sim N(\tilde{u}_{y,i}, \sigma_i^2)$ for all i . (Note the same prior variance in each dimension.) Like in the rest of this section, the DM knows he will receive cardinal signals and dimension-by-dimension ordinal signals comparing the two options. In addition, he observes dimension-specific ordinal signals with respect to a reference point r , where for simplicity we assume that $u_{r,i}$ are known for all i and that ordinal signals are noise-free. What are the values of $u_{r,i}$ that maximize the expected value of this choice? Denote s the set of cardinal and o that of all ordinal signals. Then, the expected value of the choice is

$$U_{\max}(r) := \mathbb{E} [\max\{U(x), U(y)\} | s, o],$$

where we have emphasized the dependence on r . It turns out that the most informative r admits a simple characterization.

Proposition 9. *Let $u_{x,i} \sim N(\tilde{u}_{x,i}, \sigma_i^2)$ and $u_{y,i} \sim N(\tilde{u}_{y,i}, \sigma_i^2)$ for all i . Then $U_{\max}(r)$ is maximized by r such that $u_{r,i} = \frac{1}{2} (\tilde{u}_{x,i} + \tilde{u}_{y,i})$ for all i .*

This proposition shows that the most informative reference point for binary choice lies exactly halfway between the prior means of the two options, in each dimension.²⁹ Again, this is broadly consistent with the idea that reference points are given by average outcomes, as is (exogenously) assumed in multiple models of reference dependence but here derived from value uncertainty and reference points as information.

7 Caution

Our model thus far assumed that the DM simply chooses the option with the highest expected value without accounting for its uncertainty. However, it is well known that people are often *averse* to uncertainty—they may prefer an option with a lower expected value if it has much less uncertainty. A growing body of evidence suggests that people are indeed averse to their own value uncertainty (e.g., de Clippel et al., 2025), which can be modeled through *caution* (Cerreia-Vioglio et al., 2015, 2024) or *complexity aversion* (Mononen, 2025; Puri, 2025; Gabaix, 2025).

²⁹The intuition is the following. Since $\mathbb{E} [\max\{U(x), U(y)\}] = \frac{1}{2} (\mathbb{E}[U(x)] + \mathbb{E}[U(y)]) + \frac{1}{2} \mathbb{E}[|U(x) - U(y)|]$ and $\mathbb{E}[U(x)]$ and $\mathbb{E}[U(y)]$ are independent of r , then the expected value of the choice is maximized by the r that maximizes the Expectation of $|U(x) - U(y)|$. This is highest when, in each dimension, $u_{r,i}$ lies between $u_{x,i}$ and $u_{y,i}$: Proposition 7 shows that, in this case, expected values are pushed apart, while the opposite happens when $u_{r,i}$ is above or below both. Since this is most likely to happen when $u_{r,i} = \frac{1}{2} (\tilde{u}_{x,i} + \tilde{u}_{y,i})$, the proposition follows.

This attitude is both empirically and theoretically related to the broader notion of ambiguity aversion, of which it could be seen as a specific instance. We now extend our model to include such aversion to uncertainty, which we refer to as *caution*, and show how this has immediate implications for our model, such as generating the endowment effect.³⁰ For brevity, the following section contains only a sketch of the formal analysis and its implications for the endowment effect; these are developed in full in Appendix C.

Modeling Caution. We model aversion to uncertainty using the classical approach: adding a concave transformation, which yields aversion to uncertainty exactly like the concavity of a utility function leads to risk aversion in standard Expected Utility. We assume that the DM evaluates option x by integrating over the possible values of $U(x)$ a concave transformation of CARA form, that is, $-e^{-\alpha U(x)}/\alpha$ for $\alpha \neq 0$, where α captures the degree of aversion to uncertainty. In Appendix C, we prove that in this case, evaluations following the cardinal and ordinal signals in our model are given by:

$$\bar{V}(x|s, o) := \sum_i^n \text{CE}(u_{x,i}|s_i, o_i) \quad (13)$$

$$\text{CE}(u_{x,i}|s_i, o_i) = \underbrace{\tilde{u}_{x,i}^s}_{\text{Ref-ind.}} - \underbrace{\alpha \frac{\sigma_{x,i}^2}{2}}_{\text{Variance penalty}} + \underbrace{\bar{\psi} \left(\frac{\tilde{u}_{x,i}^s - \tilde{u}_{r,i}}{\sqrt{\sigma_{x,i}^2 + \sigma_{r,i}^2 + \sigma_{o,i}^2}}, \frac{\alpha \sigma_{x,i}^2}{\sqrt{\sigma_{x,i}^2 + \sigma_{r,i}^2 + \sigma_{o,i}^2}}, o_i \right)}_{\text{Caution-adjusted ref-dep. boost}}, \quad (14)$$

where the adjusted boost is given by

$$\bar{\psi}(m, \alpha k, o) := \begin{cases} \frac{1}{\alpha} (\ln \Phi(m) - \ln \Phi(m - \alpha k)) & \text{if } o = + \\ -\frac{1}{\alpha} (\ln \Phi(-m + \alpha k) - \ln \Phi(-m)) & \text{if } o = -. \end{cases}$$

This shows that preferences admit a very tractable representation, even if we add caution. First, it shows that we can represent aggregate preferences as the sum of the dimension-by-dimension *certainty equivalents* $\text{CE}(u_{x,i}|\mathcal{I})$. Moreover, the component of each dimension admits a simple functional form. The key difference to the mean evaluation without caution is the introduction of a variance penalty term—the DM dislikes value uncertainty. Moreover, the expression

³⁰In contrast to risk aversion (which is defined over prizes), caution can be understood as aversion to uncertainty about the utility (Cerreia-Vioglio et al., 2015, 2024). The fact that caution can generate the endowment effect should not be surprising, as this is precisely the content of Cerreia-Vioglio et al. (2024). Note, however, the setup of our paper is substantially different, as is the treatment of reference-dependence (since we do not have reference-dependent valuation but ordinal information), thus the results of this section do not follow from the previous literature.

for the reference-dependent boost changes a bit, though it maintains the core properties that we identified for mean beliefs above (e.g., it is strictly positive following a positive comparison). As intuitive, the value of each object is strictly *decreasing* in the uncertainty aversion parameter α .

The Endowment Effect and Status Quo Bias. Consider the classic evidence of the endowment effect. For a given object x of true utility $u_x \in \mathbb{R}_{++}$, the DM needs to compute: the price p_{WTP} they are willing to pay to acquire the object when they do not own it and the price p_{WTA} they are willing to accept to sell the object if they own it. We follow the assumptions we held throughout the paper: the agent has a prior $N(\mu_x, \sigma_{p,x}^2)$ for u_x , of $N(\mu_n, \sigma_{p,n}^2)$ for the utility u_n of not having it, and of $N(\mu_y, \sigma_{p,m}^2)$ on monetary exchanges of $\$y$; receives cardinal signals with variance σ_ϵ^2 on all unknowns; receives all ordinal signals in each dimension, with ordinal noise σ_o^2 .

As in the other parts of the paper, we assume that the DM has additional knowledge about the reference point, which we take to be the value of the endowment. Formally, we assume that the DM receives an additional cardinal signal about the value of the endowment with noise σ_ϵ^2 . As in the other parts of the paper, this additional information about the reference point has several potential interpretations. For instance, the DM may pay more attention to the reference object than to the alternative option, resulting in a more precise mental simulation of its utility value; this is rational, as being endowed with an object often implies that the agent may want to have a more accurate estimate of its value to decide whether to keep it. Similarly, being endowed with an object may trigger memories of similar objects, again resulting in a more precise evaluation of its utility value.

For transparency of the calculations, and to highlight which effects are due to our boosts and caution, in what follows we assume that all priors are well calibrated (that is, $\mu_x = u_x$, $\mu_r = u_r$, $\mu_y = y$), that the utility of not having the object is 0, that the true utility of money is the identity function, that the priors of owning or not owning an object have the same variance ($\sigma_{p,x}^2 = \sigma_{p,n}^2$), and that there is no ordinal noise ($\sigma_o^2 = 0$). (It is easy to see how our results hold more generally.)

Under these assumptions, a WTP-WTA gap follows straightforwardly from our model, using the formulas above.³¹ Indeed, computing p_{WTA} and p_{WTP} when cardinal signals are received at their expected value, accounting for all ordinal signals, yields (denoting σ_x^2 and $\hat{\sigma}_x^2$ the posterior

³¹Here, we only sketch the calculations under the assumption that there is no uncertainty about the value of money; Appendix C develops them in full also relaxing this assumption, and considers the price p_{Choice} that makes the agent indifferent between receiving the object or money, when they are endowed with neither.

variance about u_x when it is the endowment or not, respectively, after cardinal signals alone)

$$p_{\text{WTP}} = u_x - \frac{\alpha}{2} (\sigma_x^2 - \hat{\sigma}_x^2) + \frac{1}{\alpha} \ln \frac{\Phi\left(\frac{u_x + \alpha \hat{\sigma}_x^2}{\sqrt{\hat{\sigma}_x^2 + \sigma_x^2}}\right)}{\Phi\left(\frac{u_x - \alpha \sigma_x^2}{\sqrt{\hat{\sigma}_x^2 + \sigma_x^2}}\right)}, \quad p_{\text{WTA}} = u_x + \frac{\alpha}{2} (\sigma_x^2 - \hat{\sigma}_x^2) + \frac{1}{\alpha} \ln \frac{\Phi\left(\frac{u_x + \alpha \sigma_x^2}{\sqrt{\sigma_x^2 + \hat{\sigma}_x^2}}\right)}{\Phi\left(\frac{u_x - \alpha \hat{\sigma}_x^2}{\sqrt{\sigma_x^2 + \hat{\sigma}_x^2}}\right)},$$

giving us

$$\text{EE} := p_{\text{WTA}} - p_{\text{WTP}} = \alpha (\sigma_x^2 - \hat{\sigma}_x^2) + \frac{1}{\alpha} \ln \frac{\Phi\left(\frac{u_x + \alpha \sigma_x^2}{\sqrt{\sigma_x^2 + \hat{\sigma}_x^2}}\right) \Phi\left(\frac{u_x - \alpha \sigma_x^2}{\sqrt{\hat{\sigma}_x^2 + \sigma_x^2}}\right)}{\Phi\left(\frac{u_x + \alpha \hat{\sigma}_x^2}{\sqrt{\hat{\sigma}_x^2 + \sigma_x^2}}\right) \Phi\left(\frac{u_x - \alpha \hat{\sigma}_x^2}{\sqrt{\sigma_x^2 + \hat{\sigma}_x^2}}\right)}.$$

Intuitively, the comparisons between owning and not owning the mug generate “boosts” in several directions, but our DM dislikes value uncertainty and their utility evaluation of the reference bundle is more precise. Indeed, it is immediate to show that, as long as $\alpha > 0$, $\sigma_x^2 > 0$, and $\sigma_e^2 > 0$, the endowment effect above is strictly positive.

At the same time, note that our model also yields that the extent of the endowment effect varies with the extent of uncertainty about the value of the object—it is easy to see from the formula above that the endowment effect is strictly increasing in σ_x^2 ; Figure 7 in Appendix C plots it. Indeed, in the absence of value uncertainty ($\sigma_x = 0$), our model does not predict an endowment effect, reminiscent of the conjecture by Kahneman et al. (1990) that “The valuation ambiguity produced by . . . lack of commensurability is necessary . . . for both loss aversion and a buying-selling discrepancy.” The literature review in Cerreia-Vioglio et al. (2024) highlights that there is now a broad body of evidence to suggest that the endowment effect is highly susceptible to manipulations of (what we can think of as) value uncertainty, including experience or knowing the object’s market price.

Finally, we note that while the calculations above were applied to show how our model can generate the endowment effect, very similar ones show how it also generates status quo bias, following an identical logic.

8 Conclusion

This paper has argued that the combination of value uncertainty and reference points as information provides a framework for thinking about multiple phenomena at the core of behavioral economics. Building on the present paper, in Dean et al. (2026) we study how a similar mechanism—ordinal information from one’s own past actions—gives rise to an information-based model of habit formation, status quo bias and reasoning by analogy (“internal herding”).

We conclude by discussing limitations and potential extensions and applications.

Limitations. Our model has several limitations, and there are many behavioral economics phenomena and psychological mechanisms our model does not speak to. For instance, our model is silent on effects related to selective attention allocation such as salience (Bordalo et al., 2022) or rational inattention (Maćkowiak et al., 2023). Our model also does not allow for a role of memory cues that are entirely uninformative (or perceived to be entirely uninformative by the decision maker). Still in the same vein, we also do not capture categorization effects other than those that pertain to the gain-loss distinction (e.g. Mullainathan, 2002; Bordalo et al., 2024). We believe that it may be fruitful to combine our ‘cognitive’ approach to reference points with some of these other cognitive mechanisms. Finally, we emphasize that our approach does not rule out the existence of reference-dependent preferences, adaptation effects and/or loss aversion. For instance, it seems uncontroversial that lukewarm water feels more pleasant after cold than after hot water, a classic example of reference-dependent sensations. We believe that an interesting avenue for future research will be to study which of the effects of reference dependence reflect comparative information and which reflect preferences such as loss aversion.

Potential extensions and applications. We believe our model encourages several potential extensions, applications and empirical tests. A first potential extension involves social reference points. While introspection and evidence suggest that social reference points matter partly for preference-based reasons (e.g., emotions such as envy), it seems similarly intuitive that social reference points often contain information. For example, if I learn that a co-worker who is very similar to me makes twice as much money as I do, then this not only triggers envy in me but also suggests that I got screwed over by my boss, or that my marginal product is higher than I used to believe. Incorporating social reference points into our model might shed light on this.

A second useful extension might be to decision-making under risk and intertemporal choice. Here, the boost from the ordinal comparison may shed new light on familiar phenomena. For instance, when asked how much \$100 in 6 months is worth to the decision maker today, he may be uncertain (Enke et al., 2025; Chakraborty, 2021), yet through simple ordinal reasoning, he will know that his present value ought to be less than \$100, producing a downward boost in the evaluation of delayed payments. As a result of this ordinal information, a Bayesian decision maker *ought* to exhibit steeper discounting between today and tomorrow than between two future dates.

Third, because our information-theoretic model provides a novel way of thinking about reference points, it also offers a new perspective on the classic question of what the reference point is.

We believe that the broad idea that the reference point is “optimal,” in the sense that it reflects the most informative option from the perspective of reducing value uncertainty, can be fruitfully extended and applied in various different ways.

Finally, our model also encourages empirical applications and experimental tests. A natural first question is to what extent (and which) reference effects depend on value uncertainty. In the lab, experiments would need to separately vary value uncertainty about the option and the reference point. In field applications, one could study to what extent reference effects depend on experience or choice complexity (as some prior work has done).

References

- Abeler, Johannes, Armin Falk, Lorenz Goette, and David Huffman (2011) “Reference points and effort provision,” *American Economic Review*, 101 (2), 470–492.
- Agahi, Hamzeh (2015) “An elementary proof of the covariance inequality for Choquet integral,” *Statistics & Probability Letters*, 106, 173–175.
- Agranov, Marina and Pietro Ortoleva (2017) “Stochastic choice and preferences for randomization,” *Journal of Political Economy*, 125 (1), 40–68.
- Allen, Eric J, Patricia M Dechow, Devin G Pope, and George Wu (2017) “Reference-dependent preferences: Evidence from marathon runners,” *Management Science*, 63 (6), 1657–1672.
- Anderson, John R (2013) *The adaptive character of thought*: Psychology Press.
- Ariely, Dan, George Loewenstein, and Drazen Prelec (2003) ““Coherent arbitrariness”: Stable demand curves without stable preferences,” *The Quarterly journal of economics*, 118 (1), 73–106.
- Augenblick, Ned, Eben Lazarus, and Michael Thaler (2025) “Overinference from weak signals and underinference from strong signals,” *The Quarterly Journal of Economics*, 140 (1), 335–401.
- Ba, Cuimin, J Aislinn Bohren, and Alex Imas (2025) “Over-and underreaction to information,” mimeo University of Pennsylvania.
- Balakrishnan, Sivaraman and Mark Dean (2025) “Unknown,” mimeo Columbia University.
- Barberis, Nicholas C (2013) “Thirty years of prospect theory in economics: A review and assessment,” *Journal of economic perspectives*, 27 (1), 173–196.
- Bondi, Tommaso, Daniel Csaba, Evan Friedman, and Salvatore Nunnari (2025) “Range Effects in Economic Choice: The Role of Complexity,” Technical report, CESifo Working Paper.
- Bordalo, Pedro, Nicola Gennaioli, Giacomo Lanzani, and Andrei Shleifer (2024) “A Cognitive Theory of Reasoning and Choice,” NBER WP 33466.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2020) “Memory, attention, and choice,” *The Quarterly journal of economics*, 135 (3), 1399–1442.
- (2022) “Salience,” *Annual Review of Economics*, 14 (1), 521–544.
- Braida, Louis D and Nathaniel I Durlach (1972) “Intensity Perception. II. Resolution in One-Interval Paradigms,” *The Journal of the Acoustical Society of America*, 51 (2B), 483–502.
- Brandstätter, Eduard, Gerd Gigerenzer, and Ralph Hertwig (2006) “The priority heuristic: making choices without trade-offs,” *Psychological review*, 113 (2), 409.
- Bröder, Arndt (2000) “Assessing the empirical validity of the “take-the-best” heuristic as a model of human probabilistic inference.,” *Journal of Experimental Psychology: learning, memory, and Cognition*, 26 (5), 1332.

- Butler, David J and Graham C Loomes (2007) “Imprecision as an account of the preference reversal phenomenon,” *American Economic Review*, 97 (1), 277–297.
- Butler, David and Graham Loomes (2007) “Imprecision as an Account of the Preference Reversal Phenomenon,” *American Economic Review*, 97 (1), 277–297.
- (2011) “Imprecision as an account of violations of independence and betweenness,” *Journal of Economic Behavior & Organization*, 80 (3), 511–522.
- Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler (1997) “Labor supply of New York City cabdrivers: One day at a time,” *The Quarterly Journal of Economics*, 112 (2), 407–441.
- Campbell, Danny, W George Hutchinson, and Riccardo Scarpa (2007) “Lexicographic preferences in discrete choice experiments: Consequences on individual-specific willingness to pay estimates,” mimeo Queen’s University Belfast.
- Cerreia-Vioglio, Simone, David Dillenberger, and Pietro Ortoleva (2015) “Cautious Expected Utility and the Certainty Effect,” *Econometrica*, 83 (2), 693–728.
- (2024) “Caution and Reference Effects,” *Econometrica*, 92 (6), 2069–2103.
- Chakraborty, Anujit (2021) “Present bias,” *Econometrica*, 89 (4), 1921–1961.
- Charness, Gary, Uri Gneezy, and Michael A Kuhn (2012) “Experimental methods: Between-subject and within-subject design,” *Journal of economic behavior & organization*, 81 (1), 1–8.
- Choi, Syngjoo, Raymond Fisman, Douglas Gale, and Shachar Kariv (2007) “Consistency and heterogeneity of individual behavior under uncertainty,” *American economic review*, 97 (5), 1921–1938.
- de Clippel, Geoffroy, Paola Moscardiello, Pietro Ortoleva, and Kareen Rozen (2025) “Caution in the Face of Complexity,” mimeo, Princeton University.
- Cubitt, Robin P, Daniel Navarro-Martinez, and Chris Starmer (2015) “On preference imprecision,” *Journal of Risk and Uncertainty*, 50 (1), 1–34.
- Dawes, Robyn M (1979) “The robust beauty of improper linear models in decision making,” *American Psychologist*, 34 (7), 571–582.
- Dean, Mark, Benjamin Enke, Thomas Graeber, and Pietro Ortoleva (2026) “Internal Herding,” mimeo Columbia University.
- DellaVigna, Stefano, Jörg Heining, Johannes F Schmieder, and Simon Trenkle (2022) “Evidence on job search models from a survey of unemployed workers in germany,” *The Quarterly Journal of Economics*, 137 (2), 1181–1232.
- DellaVigna, Stefano, Attila Lindner, Balázs Reizer, and Johannes F Schmieder (2017) “Reference-dependent job search: Evidence from Hungary,” *The Quarterly Journal of Economics*, 132 (4), 1969–2018.

- Enke, Benjamin and Thomas Graeber (2023) “Cognitive uncertainty,” *The Quarterly Journal of Economics*, 138 (4), 2021–2067.
- Enke, Benjamin, Thomas Graeber, and Ryan Oprea (2025) “Complexity and time,” *Journal of the European Economic Association*, 23 (5), 1838–1867.
- Enke, Benjamin, Thomas Graeber, Ryan Oprea, and Jeffrey Yang (2024) “Behavioral Attenuation,” mimeo Harvard University.
- Frydman, Cary and Lawrence J Jin (2022) “Efficient coding and risky choice,” *The Quarterly Journal of Economics*, 137 (1), 161–213.
- Gabaix, Xavier (2014) “A sparsity-based model of bounded rationality,” *The Quarterly Journal of Economics*, 129 (4), 1661–1710.
- (2019) “Behavioral inattention,” in *Handbook of behavioral economics: Applications and foundations 1, 2*, 261–343: Elsevier.
- (2025) “A Theory of Complexity Aversion,” mimeo Harvard University.
- Garner, Wendell R (1953) “An informational analysis of absolute judgments of loudness.,” *Journal of experimental psychology*, 46 (5), 373.
- Gershman, Samuel (2021) *What makes us smart: The computational logic of human cognition*: Princeton University Press.
- Gigerenzer, Gerd and Daniel G Goldstein (1996) “Reasoning the fast and frugal way: models of bounded rationality,” *Psychological review*, 103 (4), 650.
- Gilboa, Itzhak and Massimo Marinacci (2011) “Ambiguity and the Bayesian paradigm,” mimeo Bocconi University.
- Gilboa, Itzhak and David Schmeidler (1989) “Maxmin expected utility with non-unique prior,” *Journal of Mathematical Economics*, 18, 141–153.
- Graeber, Thomas © Benjamin Enke (2026) “Comparisons,” mimeo Harvard University.
- Graeber, Thomas, Christopher Roth, and Marco Sammon (2025) “Coarse Categories Compete in a Complex World,” mimeo HBS.
- Griffiths, Thomas, Nick Chater, and Joshua B Tenenbaum (2024) *Bayesian models of cognition*: MIT Press.
- Gul, Faruk (1991) “A theory of disappointment aversion,” *Econometrica*, 59 (3), 667–686.
- Halevy, Yoram, David Walker-Jones, and Lanny Zrill (2023) *Difficult decisions*: University of Toronto, Department of Economics.
- Heckman, James J (1979) “Sample selection bias as a specification error,” *Econometrica: Journal of the econometric society*, 153–161.

- Heffetz, Ori (2021) “Are reference points merely lagged beliefs over probabilities?” *Journal of Economic Behavior & Organization*, 181, 252–269.
- Hollands, JG and Brian P Dyre (2000) “Bias in proportion judgments: the cyclical power model.,” *Psychological review*, 107 (3), 500.
- Hsee, Christopher K (1996) “The evaluability hypothesis: An explanation for preference reversals between joint and separate evaluations of alternatives,” *Organizational behavior and human decision processes*, 67 (3), 247–257.
- Jin, Lawrence, Rui Tang, Han Ye, Junjian Yi, and Songfa Zhong (2024) “Path dependency in physician decisions,” *Review of Economic Studies*, 91 (5), 2916–2953.
- Kahana, Michael Jacob (2012) *Foundations of human memory*: OUP USA.
- Kahneman, Daniel, Jack L Knetsch, and Richard H Thaler (1990) “Experimental tests of the endowment effect and the Coase theorem,” *Journal of political Economy*, 98 (6), 1325–1348.
- Kahneman, Daniel and Amos Tversky (1979) “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47 (2), 263–292.
- Khaw, Mel Win, Ziang Li, and Michael Woodford (2021) “Cognitive imprecision and small-stakes risk aversion,” *The review of economic studies*, 88 (4), 1979–2013.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005) “A smooth model of decision making under ambiguity,” *Econometrica*, 73 (6), 1849–1892.
- Kőszegi, Botond and Matthew Rabin (2006) “A Model of Reference-Dependent Preferences,” *Quarterly Journal of Economics*, 121 (4), 1133–1165.
- Laming, Donald (1984) “The relativity of ‘absolute’ judgements,” *British Journal of Mathematical and Statistical Psychology*, 37 (2), 152–183.
- (2009) “Weber’s law,” in *Inside Psychology: A science over 50 years*, 179–191: Oxford Univ. Press London, UK.
- Laming, Donald Richard John (1997) *The measurement of sensation* (30): Oxford University Press.
- List, John A (2025) “The Experimentalist Looks Within: Toward an Understanding of Within-Subject Experimental Designs,” Technical report, National Bureau of Economic Research.
- Loomes, Graham and Robert Sugden (1986) “Disappointment and Dynamic Consistency in Choice under Uncertainty,” *The Review of Economic Studies*, 53 (2), 271–282.
- Luce, R Duncan and David M Green (1974) “The response ratio hypothesis for magnitude estimation,” *Journal of Mathematical Psychology*, 11 (1), 1–14.
- Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt (2023) “Rational inattention: A review,” *Journal of Economic Literature*, 61 (1), 226–273.

- Malmendier, Ulrike and Jessica A Wachter (2024) “Memory of past experiences and economic decisions,” Available at SSRN 4013583.
- Masatlioglu, Yusufcan and Efe A. Ok (2005) “Rational Choice with Status-Quo Bias,” *Journal of Economic Theory*, 121, 1–29.
- (2014) “A Canonical Choice Model with Initial Endowment,” *Review of Economic Studies*, 81 (2), 851–883.
- Meier, Pascal Flurin, Raphael Flepp, and Egon Franck (2025) “Expectational reference points and belief formation: Field evidence from financial analysts,” *Journal of Economic Behavior & Organization*, 229, 106788.
- Mononen, Lasse (2025) “On preference for simplicity and probability weighting,” University of Bielefeld.
- Moon, Jongmin and Oh-Sang Kwon (2022) “Attractive and repulsive effects of sensory history concurrently shape visual perception,” *BMC biology*, 20 (1), 247.
- Mullainathan, Sendhil (2002) “Thinking through categories,” Technical report, Working Paper, Harvard University.
- Natenzon, Paulo (2019) “Random choice and learning,” *Journal of Political Economy*, 127 (1), 419–457.
- Nishimura, Hiroki and Efe A Ok (2021) “Preference Structures,” mimeo, NYU.
- Nosofsky, Robert M (1986) “Attention, similarity, and the identification–categorization relationship,” *Journal of experimental psychology: General*, 115 (1), 39.
- Oaksford, Mike and Nick Chater (2007) *Bayesian rationality: The probabilistic approach to human reasoning*: Oxford University Press, USA.
- Ok, Efe, Pietro Ortoleva, and Gil Riella (2011) “Revealed (P)Reference Theory,” Mimeo California Institute of Technology.
- Parducci, Allen (1965) “Category judgment: a range-frequency model,” *Psychological review*, 72 (6), 407.
- Pope, Devin G and Maurice E Schweitzer (2011) “Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes,” *American Economic Review*, 101 (1), 129–157.
- Puri, Indira (2025) “Simplicity and risk,” *The Journal of Finance*, 80 (2), 1029–1080.
- Sadil, Patrick, Rosemary A Cowell, and David E Huber (2024) “The push–pull of serial dependence effects: Attraction to the prior response and repulsion from the prior stimulus,” *Psychonomic Bulletin & Review*, 31 (1), 259–273.
- Sagi, Jacob (2006) “Anchored Preference Relations,” *Journal of Economic Theory*, 130, 283–295.
- Sampford, Michael R (1953) “Some inequalities on Mill’s ratio and related functions,” *The Annals of Mathematical Statistics*, 24 (1), 130–132.
- Schäfer, Thomas and Marcus A Schwarz (2019) “The meaningfulness of effect sizes in psychological research: Differences between sub-disciplines and the impact of potential biases,” *Frontiers in psychology*, 10, 813.

- Shepard, Roger N (1987) “Toward a universal law of generalization for psychological science,” *Science*, 237 (4820), 1317–1323.
- Shubatt, Cassidy and Jeffrey Yang (2024) “Tradeoffs and Comparison Complexity,” arXiv preprint arXiv:2401.17578.
- Soltani, Alireza, Benedetto De Martino, and Colin Camerer (2012) “A range-normalization model of context-dependent choice: a new model and evidence,” *PLoS computational biology*, 8 (7), e1002607.
- Stewart, Neil, Gordon DA Brown, and Nick Chater (2005) “Absolute identification by relative judgment.,” *Psychological review*, 112 (4), 881.
- Stewart, Neil, Nick Chater, and Gordon DA Brown (2006) “Decision by sampling,” *Cognitive psychology*, 53 (1), 1–26.
- Sugden, Robert (2003) “Reference-dependent subjective expected utility,” *Journal of economic theory*, 111 (2), 172–191.
- Trueblood, Jennifer S (2015) “Reference point effects in riskless choice without loss aversion.,” *Decision*, 2 (1), 13.
- Tulving, Endel and Donald M Thomson (1973) “Encoding specificity and retrieval processes in episodic memory.,” *Psychological review*, 80 (5), 352.
- Tversky, Amos (1969) “Intransitivity of preferences.,” *Psychological review*, 76 (1), 31.
- Tversky, Amos and Daniel Kahneman (1991) “Loss Aversion in Riskless Choice: A Reference-Dependent Model,” *Quarterly Journal of Economics*, 106 (4), 1039–1061.
- Vieider, Ferdinand M (2024) “Decisions under uncertainty as bayesian inference on choice options,” *Management Science*, 70 (12), 9014–9030.
- Villas-Boas, J Miguel (2024) “Toward an information-processing theory of loss aversion,” *Marketing Science*, 43 (3), 523–541.
- Woodford, Michael (2012) “Inattentive Valuation and Reference-Dependent Choice,” mimeo, Columbia University.
- (2020) “Modeling imprecision in perception, valuation, and choice,” *Annual Review of Economics*, 12 (1), 579–601.
- Wu, Keyu (2024) “Context-dependent perceptions and decision-making,” mimeo University of Zurich.