Decisions under Uncertainty as Bayesian Inference*

Ferdinand M. Vieider[1,2]

[1]*RISLαβ, Department of Economics, Ghent University*
[2]*RISLαβ Africa, AIRESS, University Mohammed VI Polytechnic,*
ferdinand.vieider@ugent.be

March 14, 2022

Abstract

I derive a model of decision-making under uncertainty based on the premise that choice stimuli are mentally encoded by noisy signals, which are optimally combined with pre-existing information about the environment to obtain actionable decision parameters. The model predicts decision patterns as described by prospect theory (PT), while also predicting PT violations. It sheds new light on puzzling phenomena, such as the overinsurance of modest risks. The model makes a number of empirical predictions, which I test on some prominent datasets originally collected to fit PT. The model predictions on parameter correlations and PT violations find substantial support in the data. The model furthermore consistently outperforms PT in terms of predictive ability.

Keywords: risk taking; prospect theory; noisy cognition; ambiguity attitudes

JEL-classification: C93; D03; D80; O12

1 Motivation

Risk attitudes play a central role in economic models. With the increasing availability of experimental and empirical data, subsequently more sophisticated models of decision-making under risk and uncertainty have thus been developed to

*I gratefully acknowledge financial support from BOF at Ghent University under the project "The role of noise in the determination of risk preferences". A previous version of this paper has been circulated under the title "Noisy Coding and Decisions under Uncertainty". I am indebted to Mohammed Abdellaoui, Aurélien Baillon, Ranoua Bouchouicha, Thomas Epper, Cary Frydman, Olivier L'Haridon, Tom Verguts, Peter P. Wakker, Michael Woodford, and Horst Zank for helpful comments and discussions. All errors remain my own.

describe aggregate decision patterns (Bernoulli, 1738/1954; Markowitz, 1952; Savage, 1954; Kahneman and Tversky, 1979; Gul, 1991; Tversky and Kahneman, 1992; Schmidt, Starmer and Sugden, 2008). Due to their deterministic nature, these models need to be combined with a separately chosen stochastic model of choice to be fit to actual data (Hey and Orme, 1994; Harless and Camerer, 1994; Apesteguia and Ballester, 2018). Recently, Khaw, Li and Woodford (2021) proposed a theoretical framework that markedly departs from this paradigm. Building on the influential Bayesian Brain paradigm in neuroscience (Knill and Pouget, 2004; Doya, Ishii, Pouget and Rao, 2007; Vilares and Kording, 2011), they proposed that numerical quantities are mentally encoded in a probabilistic fashion including some noise, and subsequently decoded by Bayesian updating with a mental prior. Applied to choices between risky wagers, Khaw et al. (2021) showed how such a process can explain small-stake risk aversion (Rabin, 2000), while at the same time accounting for stochastic choice patterns.

The promise of the framework proposed by Khaw et al. (2021) is that it may result in a unified view of individual decision-making, while reducing arbitrariness in modelling choices and in combinations of deterministic decision models with stochastic choice models. Here, I build on the numerical perception framework of Khaw et al. (2021) to develop a general model of choice under uncertainty. This allows me to test the model on existing datasets on choices under both risk (known probabilities) and uncertainty (unknown or vague probabilities).[1] Given that probabilities need to be inferred from inherently noisy signals about choice stimuli, the difference between risk and uncertainty indeed becomes one of degree, rather than being a categorical difference as under PT (Abdellaoui, Baillon, Placido and Wakker, 2011). The model can account for many empirical regularities, such as the coexistence of lottery play and insurance uptake (Friedman and Savage, 1948; Markowitz, 1952), or for systematic changes of risk attitudes across

[1]Since the details of the encoding-decoding process underlying the model is contingent on the choice stimuli being presented to a subject, different choice stimuli require different setups. I here focus on tradeoffs between sure amounts of money and non-degenerate binary prospects under both risk (known probabilities) and uncertainty(unknown or vague probabilities), since these are the types of tradeoffs used in the data I analyze. I willl discuss below how the model can be extended to more general comparisons of non-degenerate wagers.

likelihoods (Preston and Baratta, 1948; Kahneman and Tversky, 1979; Fehr-Duda and Epper, 2012). None of these phenomena can be rationalized within the risky choice framework as presented by Khaw et al. (2021), which predicts universal risk aversion for gain, and universal risk seeking for losses.

Starting from an optimal choice rule entailing expected value maximization, I describe the weights resulting from the optimal Bayesian combination of noisy signals about choice options with a mental prior indicating their likelihood in the environment. I refer to this process as *Bayesian Inference*—and hence to the model as the Bayesian Inference Model (*BIM*)—since it rests on inferences the mind makes about the outside world.[2] The weights governing the optimal combination of signal and prior entail systematic nonlinear distortions of both outcomes and likelihoods. The model can thus be seen as a *generative model for PT*, inasmuch as it predicts PT-like decision parameters. At the same time, the model also predicts systematic *violations* of PT. The latter affect primarily the strict separability between the probability and outcome dimensions postulated by prospect theory (Wakker, 2010; Abdellaoui et al., 2011; Dimmock, Kouwenberg and Wakker, 2015). These will result in likelihood-distortions being stake-dependent—a phenomenon which has indeed been widely documented in the literature (Hogarth and Einhorn, 1990; Fehr-Duda, Bruhin, Epper and Schubert, 2010; Bouchouicha and Vieider, 2017), but which thus far lacked a solid theoretical explanation. They will also result in outcome distortions being affected by the nature of the outcome-generating process—which I document in this paper using an original experiment comparing likelihood- and outcome distortions under risk and uncertainty.

Mental priors play an important role in the model I present here. The mental priors for the likelihood and outcome dimensions have a natural interpretation as the expectation of the likelihood-ratio and of the benefit-cost ratio, respectively. The neuroscientific literature has documented that mental priors quickly adapt to

[2]The terms *Bayesian Inference* and *Active Inference* are increasingly used in both neuroscience and artificial intelligence to refer to decisions based on noisy signals about the outside world, which need to be decoded in an optimal Bayesian way to result in actionable representations. See for instance Friston, Schwartenbeck, FitzGerald, Moutoussis, Behrens and Dolan (2013) and Friston, FitzGerald, Rigoli, Schwartenbeck and Pezzulo (2017).

the decision environment, and that they are specific to the given decision context (Adams, Graf and Ernst, 2004; Friston, 2005; Petzschner and Glasauer, 2011). Given these insights, the model predicts that parameter estimates gathered in laboratory experiments—where the costs and benefits from taking the wager tend to be equal on average by design—will perform poorly when applied to situations in which costs and benefits are highly asymmetric. Such calibration issues have, for instance, been highlighted in the case of insurance uptake, where typical PT-parameters have been shown to be ill-suited to account for the pervasive over-insurance of moderate risks (Sydnor, 2010). In empirical applications, this has led to the combination of probability distortions with loss aversion to capture e.g. deductible choices (Barseghyan, Molinari, O'Donoghue and Teitelbaum, 2013). In the BIM such choice patterns are indeed predicted by a context-specific prior, which is identified by a specific model parameter indicating pessimism for situations in which potential costs outweigh potential gains.

The BIM also makes a number of novel predictions about correlations between noise and model parameters we ought to observe in PT setups. I test these predictions on two large datasets collected to estimate PT parameters under risk—the data of Bruhin, Fehr-Duda and Epper (2010), which include three large experiments to identify PT parameters in Switzerland and China, and the data of L'Haridon and Vieider (2019), including 3000 subjects from 30 countries. Estimating individual-level parameters in a Bayesian random-parameter setup, I find the BIM parameters to be closely correlated with the parameters estimated from a PT model plus additive noise.[3] This serves to illustrate the model's generative nature. I then proceed to documenting systematic correlations between the PT parameters and noise, as well as amongst the PT parameters themselves. While puzzling from the point of view of PT—where the deterministic and stochastic

[3]Being deterministic in nature, PT *per se* is silent about errors. Random utility models with additive errors remain the most popular choice in empirical implementations, and they have been applied by both Bruhin et al. (2010) and L'Haridon and Vieider (2019) to the data I use for this comparison. While I mostly refer to additive noise models, it is important to note that my results carry over to more advanced error models, such as e.g. the contextual utility model of Wilcox (2011). I furthermore estimate random-parameter models throughout, so that my results are not confined to a pure random utility setup, and remain unaffected by recent critcisms of that model (Apesteguia and Ballester, 2018).

components are generally assumed to be orthogonal to each other (a 'white noise' assumption)[4]—these correlations are predicted by the BIM. Finally, I test the BIM against PT based on their predictive performance. The BIM significantly outperforms PT in all instances. PT's relatively poor predictive performance can be traced back to the overfitting of certain aspects of the data. The BIM avoids this issue through the integration of noise and decision parameters, which constrains the parameters in ways that give it an edge in terms of predictive performance.

The approach taken here builds on an influential theoretical paradigm in neuroscience, according to which the brain uses probabilistic mechanisms to encode perceptual information about the outside world, and decodes this information by Bayesian updating with a mental prior (Knill and Pouget, 2004; Doya et al., 2007; Vilares and Kording, 2011). Such an approach will be optimal whenever mental resources are scarce, and need to be distributed between a variety of different tasks. Behavioural features such as outcome- and likelihood-distortions result from deviations from an optimal choice rule that can be traced back to cognitive constraints imposed by limited cognitive resources. The noisy neural coding of choice stimuli will thus naturally result in discrimination difficulties for values that are observed relatively infrequently, an intuition that is shared with a variety of other models (Robson, 2001; Netzer, 2009; Woodford, 2012; Payzan-LeNestour and Woodford, 2021; Netzer, Robson, Steiner and Kocourek, 2021). Such discrimination difficulties for relatively infrequent stimuli, in turn, will result in systematic biases in those regions. Discounting signals indicating choice stimuli that are considered unlikely a priori will thus result in a sort of regression to the mean that can explain a number of behavioural regularities (Zhang and Maloney, 2012; Enke and Graeber, 2019). As such, the approach holds the promise of delivering a unifying foundation for many decision phenomena that are currently explained by separate models. Closely related models to the one here presented have been shown to

[4]Some stochastic models of choice under risk, such the the contextual utility model of Wilcox (2011), specifically allow for interactions between preferences and decision noise. Nevertheless, the choice of decision model and stochastic model in these setups remains largely independent, and to some extent, arbitrary. We will see below that allowing for such extensions of the error structure within a PT model does not affect the conclusions reached in the analysis.

account for seemingly hyperbolic or present-biased time discounting (Gabaix and Laibson, 2017; Vieider, 2021), further showcasing the generality of the approach.

This paper proceeds as follows. Section 2 introduces the theoretical framework. Section 3 presents the empirical analysis. I discuss the results in section 4, and section 5 concludes the paper.

2 The Noisy Coding Model

In this section, I lay out the theoretical model in several steps. I start by showing that coding uncertainty as log-likelihood ratios of complementary events constitutes a highly efficient form of representing and updating probabilistic information, thus resulting in an optimal choice rule under uncertainty. I then show how the noisy mental encoding of choices between wagers yields subjective distortions of the objective quantities in the optimal choice rule. Finally, I use the noisy coding model to derive predictions on decision patterns under risk and ambiguity.

2.1 Noisy coding of uncertainty as log-likelihood ratios

I start by showcasing the efficiency of coding probabilistic information in terms of likelihood ratios. Assume a decision maker (DM) needs to take a decision based on her assessment of the likelihood of an uncertain event e.[5] Assume that the DM does not know the true likelihood of e occurring, but observes a noisy signal, s, about this likelihood. Assume further that the DM knows the probability of the signal conditional on the event, $P[s|e]$, as well as the probability of the signal conditional on the complementary event $\tilde{e}$, $P[s|\tilde{e}]$. The DM can then use the

[5]For analytical tractability, I will generally work with discrete events (e.g., a ball extracted from an urn is blue), and associated probabilities $P[e]$. For continuous outcome variables (the stock market index increases by over 2% in a given year; cumulative rainfall during the agricultural planting season falls between 500mm and 700mm), this constitutes a slight abuse of notation for the more accurate $P[a < e < b] = \int_a^b f[e]de$, where f indicates the probability density function.

following equation to infer the likelihood of the event, conditional on the signal:

$$\frac{P[e|s]}{P[\tilde{e}|s]} = \frac{P[s|e]}{P[s|\tilde{e}]} \times \frac{\widehat{P}[e]}{\widehat{P}[\tilde{e}]}, \tag{1}$$

where the ratio $\frac{\widehat{P}[e]}{\widehat{P}[\tilde{e}]}$ indicates the prior likelihood ratio, which incorporates any knowledge the DM may have previously held about the likelihood of event e.[6]

This setup can be used to arrive at an optimal choice rule based on the relative costs and benefits of different actions. Take a wager offering x conditional on event e, but $y < x$ if $\tilde{e}$ obtains instead. Integrating the costs and benefits from taking the wager into equation 1, the wager will be preferred to a sure option, c, whenever:

$$\frac{P[e|s]}{P[\tilde{e}|s]} \times \frac{(x-c)}{(c-y)} > 1, \tag{2}$$

where $(x-c)$ indicates the benefit from taking the wager, and $(c-y)$ indicates the cost. This choice rule is optimal, being based on Bayesian updating and expected value maximization.[7]

From a computational point of view, it may be convenient to take the natural logarithm of the expression, thus transforming the choice rule from a product to a sum. This will also allow for the additive combination of several signals with the prior as in equation 1. Gold and Shadlen (2001) explicitly discuss the neurological realism of such a choice rule, and Gold and Shadlen (2002) show how a similar setup has been used by British code-breakers to decode the supposedly indecipherable Enigma Code during World War II. It is important to note that the optimality of the choice rule is unaffected by such an operation, since any monotonic transformation will leave this choice rule unaltered. This holds both ex ante *and* ex post. That is, the results presented below do not change in any substantive way if I use the choice rule in equation 2 instead (online appendix S1.2

[6]This follows from an application of Bayes rule, whereby $P[e|s] = \frac{P[s|e]*\widehat{P}[e]}{P[s]}$. The use of likelihood ratios means that $P[s]$ cancels out of the expression.

[7]This choice rule is used without loss of generality. That is, if the decision-maker instead has a concave utility function defined over lifetime wealth, the choice model can be augmented by such a preference without altering the qualitative conclusions that follow.

presents such an alternative derivation). The choice rule then becomes:

$$ln\left(\frac{P[e|s]}{P[\tilde{e}|s]}\right) > ln\left(\frac{c-y}{x-c}\right). \tag{3}$$

The wager should thus be chosen over the sure option whenever the log-likelihood ratio of event e and its complement $\tilde{e}$ exceeds the log cost-benefit ratio. While the choice rule can easily be generalized to multiple outcomes (Green, Swets et al., 1966; Gold and Shadlen, 2001), the subsequent mental transformations to be discussed shortly require a different setup building on neural network architectures. In what follows, I will thus focus on tradeoffs between binary wagers and sure outcomes. Extensions to the comparison of binary prospects with the same state space, such as in the popular task of Holt and Laury (2002), obtain trivially and are thus not discussed further.

The model I will present below is based on the premise that mental processing of uncertain wagers functions much like in the optimal decision rule above. According to an influential theoretical paradigm in neuroscience, the human brain acts as a Bayesian inference machine, continuously combining noisy signals about the environment with prior beliefs to come up with actionable decision parameters (Knill and Pouget, 2004; Doya et al., 2007; Vilares and Kording, 2011). Even numerically represented quantities, such as a monetary outcome x or an objectively given probability p, will be mentally represented by a noisy signal before entering the choice process. Noise will then arise especially when assessments of the relevant quantities are made quickly and intuitively, rather than being calculated based on an optimal choice rule (see Khaw et al., 2021, section 1, and Woodford, 2020, for extensive discussions).

2.2 Bayesian mental processing of noisy signals

I now present a step-by-step derivation of subjective choice parameters from the mental encoding-decoding process of choice stimuli. I will show that if the brain implements an optimal choice rule such as derived above, and if the stimuli are mentally encoded with some noise and subsequently decoded using a mental prior,

then this process will result in an actionable choice rule into which the encoding noise and the parameters characterizing the mental prior will enter as subjective parameters. An alternative derivation based on the unlogged choice rule in equation 2 is presented in Online Appendix S1.2, and does not alter the conclusions drawn below in any substantive way.

The mental encoding stage and the likelihood

Take a mental signal, r, encoding the characteristics of a given choice problem. This mental signal can be thought of as a neural firing rate, encoding the desirability of the choice stimuli in the brain. The mental signal r thus plays a role that is analogous to that of the signal s used above to illustrate the optimal choice rule. Just like s above, the mental signal will generally be noisy, and the noise itself may carry useful information about the accuracy of the signal.[8] Other than above, however, not only uncertain events are represented by mental signals, but so are nominally certain monetary outcomes, the exact nature of which also needs to be inferred by the mind from noisy signals. I will thus assume that there are two such mental signals, r_e and r_o, encoding the desirability of the likelihood ratio and the cost-benefit ratio, respectively.[9] For stimuli that are close together—roughly equal probabilities of e and $\tilde{e}$, or roughly equal costs and benefits—the signals r_e and r_o will be weak. If, however, the likelihood ratio or the cost-benefit ratio deviate significantly from 1, the relative signal will be strong. Costs $c-y$ and benefits $x-c$ from taking the wager are best conceived of as approximate comparisons, rather than the precise mathematical quantities entailed by the differences depicted, thus providing a further rationale for their noisy mental representations.

[8] More generally, encoding noise may be a function of the cognitive resources available for a given task, in which case we would expect systematic adaptation of the likelihood to the prior. The prior, in turn, would be itself learned from the signals about the outside world. In this paper, I focus on a context where the model parameters are given.

[9] The setup with two mental signals corresponds to a minimal setup in which the two choices can be compared. In particular, modelling costs and benefits through two separate mental signals will only result in a rescaling of decision noise, and is thus inconsequential from any practical point of view. I abstract from explicitly modelling the encoding of other aspects of the decision situation, such as the location on the screen of the risky wager, or the colors associated with winning and losing options, since they are not of central importance to the decision to be taken and do not enter the optimal choice rule derived above. Its has indeed been shown that numerical perception is independent from location and colour perception—see Dehaene (2011) for a book-length treatment.

Stimuli will have compressed mental representations for efficiency reasons, and logarithmic functions are typically used to model such mental representations (Dayan and Abbott, 2001; Petzschner, Glasauer and Stephan, 2015). It has indeed been shown that the activation of neurons thought responsible for the approximate mental representations of numbers is best fit by a normal distribution with as its argument the logarithm of the number being represented (Dehaene and Changeux, 1993; Dehaene, 2003). I thus assume that r_e and r_o are independently drawn from the following two distributions:

$$r_e \sim \mathcal{N}\left(ln\left(\frac{P[e]}{P[\tilde{e}]}\right), \nu_e^2\right) \quad , \quad r_o \sim \mathcal{N}\left(ln\left(\frac{c-y}{x-c}\right), \nu_o^2\right) \tag{4}$$

where ν_e and ν_o represent the encoding noise of the likelihood ratio and the cost-benefit ratio, conveying information on the uncertainty with which a given stimulus is perceived. Likelihoods of 0 or 1 are assumed to be perceived with certainty. The representation as the logarithm of the stimuli implies that the difference between two stimuli necessary for those stimuli to be reliably discriminated will be proportional to the magnitude of the stimuli themselves (Dayan and Abbott, 2001, ch. 3). This amounts to a so-called *just noticeable difference* in the stimulus m, Δm, so that $\frac{\Delta m}{m}$ is a constant. The neural encoding process can thus be conceived of as the driving factor behind a behavioural phenomenon that has long been known in psychophysics as the Weber-Fechner law (Fechner, 1860).

The mental decoding stage and the posterior expectations

The information provided by the noisy mental signals r_e and r_o needs to be decoded to be transformed into actionable quantities that can inform the decision process. This is due to the uncertainty in the mental representation of the stimuli, which makes it desirable to combine the encoded signal with prior information about what sort of stimuli are likely to be encountered in the given decision environment. It seems natural to let the mental priors used to this effect follow a normal distribution, since the choice of a conjugate prior distribution will minimize

the burden in terms of neural computations:

$$ln\left(\frac{P[e]}{P[\tilde{e}]}\right) \sim \mathcal{N}\left(ln\left(\frac{\widehat{P}[e]}{\widehat{P}[\tilde{e}]}\right), \sigma_e^2\right) \quad , \quad ln\left(\frac{c-y}{x-c}\right) \sim \mathcal{N}\left(ln\left(\frac{\widehat{k}}{\widehat{\tilde{b}}}\right), \sigma_o^2\right) \tag{5}$$

where $ln\left(\frac{\widehat{P}[e]}{\widehat{P}[\tilde{e}]}\right)$ is the mean of the likelihood ratio prior, $ln\left(\frac{\widehat{k}}{\widehat{\tilde{b}}}\right)$ the mean of the cost-benefit ratio prior, and σ_e and σ_o are the associated standard deviations of the the two priors.

Combining the likelihoods in equation 4 with the priors in equation 5 by Bayesian updating, and defining $\xi \triangleq \left(\frac{\widehat{P}[e]}{\widehat{P}[\tilde{e}]}\right)^{1-\gamma}$ and $\zeta \triangleq \left(\frac{\widehat{k}}{\widehat{\tilde{b}}}\right)^{(1-\alpha)}$, we obtain the following posterior expectations of the log-likelihood ratio and log-cost-benefit ratio, conditional on the noisy mental signals (see online appendix S1.1):

$$E\left[ln\left(\frac{P[e]}{P[\tilde{e}]}\right)|r_e\right] = \gamma \times r_e + ln(\xi) \quad , \quad E\left[ln\left(\frac{c-y}{x-c}\right)|r_o\right] = \alpha \times r_o + ln(\zeta), \tag{6}$$

where $\gamma \triangleq \frac{\sigma_e^2}{\sigma_e^2+\nu_e^2}$ and $\alpha \triangleq \frac{\sigma_o^2}{\sigma_o^2+\nu_o^2}$ thus constitute the weights assigned to the log-likelihood ratio and the log cost-benefit ratio of the stimulus, relative to the weights assigned to the means of the priors. The relative uncertainty associated with the mental signal and the prior, as captured by $\{\nu_e, \nu_o\}$ and $\{\sigma_e, \sigma_o\}$, will thus determine how much weight will be attributed to the likelihood versus the prior in the posterior distribution, which furnishes the actionable quantities on which the decision will be based. Note that as the coding noise parameters converge to 0, the signal will accurately reflect the stimuli, and the weights will converge to 1. For coding noise larger than 0 and weights strictly smaller than 1, however, the expectations of the log-likelihood ratio and the log-cost benefit ratio conditional on the signals will reflect a weighted average of signals and priors.

The actionable choice rule

The posterior means just derived can now be used to inform the decision process. In particular, we must amend the optimal choice rule described in equation 3 by

replacing the objective quantities with their mental representations:

$$E\left[ln\left(\frac{P[e]}{P[\tilde{e}]}\right)|r_e\right] > E\left[ln\left(\frac{c-y}{x-c}\right)|r_o\right], \tag{7}$$

indicating that the wager on event e will be accepted whenever the posterior expectation of the log-likelihood ratio exceeds the posterior expectation of the log cost-benefit ratio. Notice, once again, that the mental choice rule is here defined over the logged likelihood- and cost-benefit ratios simply because of the computational advantages this entails. That is, logging the choice rule does in no way determine the results I derived here, which can equally well be derived from an unlogged version of the choice rule above. Such an alternative derivation is shown in Online Appendix S1.2.

Substituting the posterior expectation in equation 6 into the choice rule in equation 7 and solving for the mental signals, we get:

$$\gamma \times r_e - \alpha \times r_o > ln(\delta)^{-1}, \tag{8}$$

where $\delta \triangleq \xi \times \zeta^{-1}$, and where $ln(\delta)^{-1}$ provides the threshold which the weighted difference of mental signals on the left-hand side needs to exceed in order for the wager to be accepted. The threshold parameter δ has a natural interpretation, since it is made up by the prior mean of the likelihood ratio, multiplied by the prior mean of the benefit-cost ratio (i.e., the inverse of the cost-benefit ratio)—the more favourable the prior expectation of the likelihood ratio and of the benefit to cost ratio, the more likely the DM will be to accept the wager.

The probabilistic choice rule

To derive an expression for the probability with which the wager will be chosen over the sure outcome that can be observed by an econometrician, we now need to obtain an expression free of the unobservable mental signals. To this end, we obtain the z-score of the weighted difference in mental signals in equation 8 by jointly distributing the two signals exploiting the known distributions in equation 4. Obtaining the z-score, and comparing it to the z-score of the threshold equation

8 (see online appendix S1.1), gives us the following probabilistic choice rule:

$$Pr[(x,e;y) \succ c] = \Phi\left(\frac{ln(\delta) + \gamma \times ln\left(\frac{P[e]}{P[\tilde{e}]}\right) - \alpha \times ln\left(\frac{c-y}{x-c}\right)}{\sqrt{\nu_e^2\gamma^2 + \nu_o^2\alpha^2}}\right), \tag{9}$$

where Φ is the standard normal cumulative distribution function.

The setup just derived readily generalizes to a situation where all outcomes are translated into losses. Assume that under event e the DM stands to lose x, or else lose $y < x$ under $\tilde{e}$, and that this scenario is compared to a sure loss of c. The gains and losses are now flipped relatively to the setup discussed above, so that $x - c$ constitutes the cost from taking the wager, and $c - y$ the potential benefit:

$$E\left[ln\left(\frac{P[e]}{P[\tilde{e}]}\right)|r_e\right] > E\left[ln\left(\frac{x-c}{c-y}\right)|r_o\right], \tag{10}$$

where all derivations follow the same steps as for gains. Such a setup will then naturally result in decreasing sensitivity towards both gains and losses.

A disproportionate dislike for losses when compared to monetarily equivalent gains, such as captured by loss aversion in PT, could arise from either of two different mechanisms. One, there may be a specific prior for the situation where the benefit is constituted by a gain, and the cost by a loss. It would indeed seem evolutionarily plausible that priors in such situations may have evolved to make people shy away from losses, which could affect their evolutionary fitness in a dynamic setting. Alternatively, loss aversion could be captured by asymmetries in the precision with which gains and losses are encoded, and in particular, by a higher coding precision for losses compared to gains (Tom, Fox, Trepel and Poldrack, 2007). Since disentangling these different explanations requires specific experimental tests, loss aversion is excluded from the empirical analysis below and left for future research.

2.3 The BIM generates PT functionals for binary wagers

I next discuss the implications of the model I have derived above, and show that the setup results in a stochastic version of prospect theory. Abstracting momentarily

from decision noise, we can write the point of indifference underlying the choice probability in equation 9 as follows[10]:

$$ln\left(\frac{c-y}{x-c}\right) = \alpha^{-1}\left[ln(\delta) + \gamma \times ln\left(\frac{P[e]}{P[\tilde{e}]}\right)\right]. \tag{11}$$

This represents the point of indifference between $(x, e; y)$ and c, where the likelihood dimension on the right-hand side—subjectively transformed by the mental representation parameters α, γ, and δ—is traded off against the outcome dimension on the left-hand side. Setting $P[\tilde{e}] = 1 - P[e]$, as assumed under prospect theory, the expression on the left can be transformed as follows:

$$ln\left(\frac{c-y}{x-c}\right) = ln\left(\frac{\frac{c-y}{x-y}}{1-\frac{c-y}{x-y}}\right) \triangleq ln\left(\frac{\pi(P[e])}{1-\pi(P[e])}\right), \tag{12}$$

where $\pi(P[e])$ is the solution of the equation $c = \pi(P[e])x + (1 - \pi(P[e]))y$, constituting a dual-EU representation of the choice problem (Yaari, 1987), for the decision weight $\pi(P[e])$. The right-hand side in equation 12 can thus be interpreted as the log of the ratio of the decision weights assigned to the high and to the low outcome in the wager. Substituting equation 12 into equation 11, and solving for $\pi(P[e])$, we obtain:

$$\pi(P[e]) = \frac{\delta^{1/\alpha}(P[e])^{\gamma/\alpha}}{\delta^{1/\alpha}\left(P[e]\right)^{\gamma/\alpha} + (1-P[e])^{\gamma/\alpha}}. \tag{13}$$

For $\alpha = 1$, this expression reduces to a probability-distortion function commonly used in the decision-making literature (Goldstein and Einhorn, 1987; Gonzalez and Wu, 1999; Bruhin et al., 2010). The general case of $\alpha \leq 1$ allows for outcome distortions in addition to probability distortions, and shows how probability weighting emerges from the tension between the two. The parameter γ mostly governs the slope of the function, capturing likelihood-sensitivity. The fact that γ decreases in the noisiness of the coding process receives support from findings

[10]Given that the model is stochastic, we can conceive of the point of indifference as the point at which a decision maker would choose either option 50% of the time on average over many trials.

showing that probabilistic sensitivity increases with cognitive ability (Choi, Kim, Lee, Lee et al., 2021). The parameter δ determines mostly the elevation of the function, and thus has an interpretation of optimism when used as a weighting of ranked gains, and of pessimism when applied to ranked losses.[11]

The mapping just presented shows how PT-like parameters naturally emerge from noisy mental representations of choice stimuli and their mental decoding by a prior distribution, indicating the likelihood of different stimuli in a given decision environment. A difference from PT is that outcome-distortions are defined over the costs and benefits associated with different events, rather than over single outcomes. Furthermore, the parameters α and γ are naturally restricted to fall between 0 and 1, whereas under PT they are only restricted to be positive. I will discuss the implications of these restrictions extensively in the empirical analysis below. Most importantly, however, the derivation of the PT-like parameters from the native BIM parameters results in a new interpretation of the PT parameters themselves. This is indeed the sense in which the BIM is a *generative* model for PT—it presents a causal account of how the PT parameters come about from a probabilistic mental representation of uncertain choice stimuli.

We can further derive substantive predictions about behaviour from the intuitions emerging from the model. Under the BIM, the parameter $\delta \triangleq \xi \times \zeta^{-1}$ has a natural interpretation as the mean of the prior. Assume that $\zeta = 1$, i.e. that the prior indicates costs and benefits that are equal on average (a realistic assumption for typical experimental setups). The constituent part of its likelihood-specific component ξ, $\psi \triangleq \widehat{P}[e]$, can then be shown to coincide with the fixed point of the probability-distortion function, i.e. the point where the function crosses the 45° line. Substituting ψ into the decision weight we obtain:

$$ln\left(\frac{\pi(\psi)}{1-\pi(\psi)}\right) = \gamma \times ln\left(\frac{\psi}{1-\psi}\right) + (1-\gamma) \times ln\left(\frac{\psi}{1-\psi}\right) = ln\left(\frac{\psi}{1-\psi}\right), \quad (14)$$

from which follows $\pi(\psi) = \psi$, and by extension, $\pi(P[e]) = P[e]$. In other

[11] Looking at this equation in isolation it may appear that for $\delta = 1$ the two parameters α and γ are not separately identified. This is, however, not the case, since α and γ also enter into the noise term. I will discuss the empirical identification of the model parameters in section 3.1.

words, the expectation of the mental prior coincides with the fixed point of the probability-distortion function, at which probabilities are perceived without subjective distortions.

Take now the general case in which $\zeta \neq 1$. The benefit-cost prior ζ^{-1} suggests that optimism ought to be particularly high in situations that present a large potential benefit with a small liklihood, as may be the case for lottery play. The inverse happens for large potential losses, where the parameter captures pessimism. This, in turn, ought to increase the likelihood of insurance uptake due to the large downside of the wager relative to the cost of the insurance itself. Such a pessimistic prior may then explain calibration issues that have been observed when PT parameters estimated in lab experiments are used to explain insurance decisions (Sydnor, 2010), which is why subsequent approaches have used a combination of probability distortions and loss aversion to explain such choices (Barseghyan et al., 2013). That is, the prior for lab experiments—where costs and benefits tend to be equal on average—should not be expected to be the same as in an insurance context, where the large potential downside of the loss ought to result in increased levels of pessimism.

The setup above can immediately be applied to decision-making under risk and ambiguity. The case of risk obtains directly by setting $P[e] = p$, with p an objectively known probability. The case of ambiguity obtains when subjects are asked to bet on Ellsberg-urns with unknown colour proportions (Ellsberg, 1961; Abdellaoui et al., 2011). Given that unknown odds are more difficult to encode than known odds (Petzschner et al., 2015), we should expect coding noise to increase as the information about the probabilities involved becomes more vague.[12] This, in turn, yields the prediction that $\gamma_a < \gamma_r$, where the subscripts a and r stand for 'ambiguity' and 'risk' respectively—a phenomenon known as *ambiguity-insensitivity*, and

[12]In keeping with the rest of the paper, I limit my discussion to the case of Ellsberg urns with unknown colour proportions. Findings may well deviate from the ones described here in other contexts and for different sources of uncertainty, since it is well-known that ambiguity attitudes over natural sources of uncertainty may vary depending on a DM's knowledge of the specific context or the DM's perceived competence in a given task (Heath and Tversky, 1991; Fox and Tversky, 1995; Abdellaoui et al., 2011). A good descriptive fit to choices under genuine uncertainty would thus require augmenting the decision model here presented with a learning model for the prior, which is beyond the scope of this paper.

which has been found to be near-universal (Abdellaoui et al., 2011; Dimmock et al., 2015; Trautmann and van de Kuilen, 2015; L'Haridon, Vieider, Aycinena, Bandur, Belianin, Cingl, Kothiyal and Martinsson, 2018). This prediction is further supported by the finding that time pressure, which presumably serves to augment coding noise, increases ambiguity-insensitivity (Baillon, Huang, Selim and Wakker, 2018). Another implication of increased encoding noise under ambiguity is that we would simultaneously expect an uptick in decision noise under ambiguity—a prediction that is born out by the increase in noisiness under ambiguity relative to risk documented by L'Haridon et al. (2018).

The prior mean, δ, may further be affected by changes in discriminability across decision situations. Since $\delta_a = \left(\frac{\psi_a}{1-\psi_a}\right)^{1-\gamma_a}$, the prior is directly impacted by the discriminability parameter γ_a unless $\psi_a = 0.5$. Furthermore, the fixed point of the function need not be the same as under risk, i.e. generally $\psi_a \neq \psi_r$. Ambiguity aversion—a dislike of unknown probabilities unrelated to their log-likelihood ratio—would then enter the model as a pessimistic prior, resulting in $\psi_a < \psi_r$. This may constitute a plausible assumption in conditions where the experimenter may be in a position to deceive subjects, as is the case when a colour choice is not allowed in urn choice problems, or more generally because of the impression generated by the inherent setup of the Ellsberg experiment that the experimenter is expressly withholding relevant information (Frisch and Baron, 1988; Fox and Tversky, 1995).

2.4 Violations of probability-outcome separability

It is now at the time to circle back to the noise term. Other than under typical stochastic implementations of PT, where the decision model and the noise model are combined in an *ad hoc* fashion, decision noise naturally emerges from the same mental encoding-decoding process as the other model parameters. Although conceptually distinct, encoding noise ν_e and ν_o cannot be distinguished in practice from observed choices. This follows from the joint distribution of r_e and r_o (see equation 18 in the online appendix), showing that this inseparability is a funda-

mental feature of the underlying mental coding process. I will thus henceforth write $\nu_e = \nu_o \triangleq \nu$, which results in a simplified decision error $\omega \triangleq \nu \times \sqrt{(\alpha^2 + \gamma^2)}$. Assuming that $\zeta = 1$—a realistic assumption for the experimental data I will examine below—the BIM has the same number of free parameters as a PT model plus additive noise.[13]

The decision noise parameter ω holds the key to understanding the BIM. PT is most often estimated in conjunction with an additive decision noise variable that is assumed to be orthogonal to the deterministic decision model (see section 3.1 for a short formal exposition of a typical PT setup). In the BIM, on the other hand, the noisiness of the choice process is intimately and inextricably linked to the other parameters of the model. Outcome-discriminability α and likelihood-discriminability γ directly enter the definition of decision noise ω, and through them their component parts ν, σ_o, and σ_p. This implies that the noisiness of the decision process will be impacted both by the uncertainty connected to the stimulus encoding, and the dispersion of the mental priors. Model parameters and decision noise thus move together, resulting in an inherently stochastic model. An exception to this rule concerns the mean of the prior. While optimism δ is affected through γ, given its definition $\delta = \left(\frac{\psi}{1-\psi}\right)^{1-\gamma}$, the native BIM parameter ψ making up the mean of the prior distribution remains unaffected. This will have important interpretative consequences, to which I will return in the discussion.

The inherent stochasticity of the BIM has a number of substantive implications for the decision-making patterns we would expect under risk and uncertainty. For one, parameters estimated in a PT setup, which neglects these intricate interrelations, should be correlated in a systematic way. This concerns primarily correlations between the noise parameter and likelihood- and outcome-sensitivity, with knock-on effects on other parameters. The neglect of these systematic corre-

[13]Focusing on prospects over gains only, the BIM has the native parameters ν, indicating encoding noise; ψ, indicating the fixed point of the probability-distortion function; and the two standard deviations of the likelihood and outcome priors, σ_e and σ_o. This corresponds to PT's 3 deterministic parameters $\widehat{\alpha}$ (utility curvature), $\widehat{\gamma}$ (likelihood-sensitivity), and $\widehat{\delta}$ (optimism), plus a noise parameter, call it $\widehat{\omega}$, where the 'hat' on the parameters serves to distinguish the native PT parameters from the equivalent BIM-derived parameters discussed in the text. These parameters are exactly replicated for losses, under the assumption that the error term under PT is allowed to be heteroscedastic across choice domains.

lations, in turn, may result in inferior predictive ability of the model. As a consequence, different model parameters estimated under PT will jointly be affected by the same observable characteristics of the decision-maker or the environment, with model parameters often moving in opposite directions in terms of revealed risk attitudes with explanatory variables, making an analysis of the correlates of risk attitudes based on PT difficult to interpret.

A further consequence of the inextricability of ν_e and ν_o is that no strict separability between the likelihood and outcome dimensions seems warranted, in opposition to what is assumed by PT. I illustrate the consequences with two examples. In first instance, we may expect that coding noise increases in the range of the stakes. This sort of adaptation may indeed be optimal in a context of scarce resources. Frydman and Jin (2022) present a model of *efficient coding*, in which the parameter of the likelihood function adapts to the prior exactly in this way. They furthermore provide experimental evidence that when outcome ranges increase, so does the noisiness in the encoding process. Such an increase in noisiness should then immediately be reflected in lower outcome-discriminability, resulting in apparent patterns of increasing relative risk aversion (Holt and Laury, 2002). Given the impossibility of separating the different dimensions of encoding noise, however, it will also result in apparent changes in likelihood-discriminability, thus impacting the probability dimension. Under PT, such effects will take the form of violations of probability-outcome separability, whereby changes in stakes ought to be reflected purely in utility curvature and leave probability-distortions unaffected. Such violations are well-documented in the literature (Hogarth and Einhorn, 1990; Fehr-Duda et al., 2010; Bouchouicha and Vieider, 2017).

The flip-side of this issue is observed when moving from risk to uncertainty or ambiguity. Given the increase in encoding noise of likelihood ratios we expect under ambiguity, the BIM predicts a lowered level of likelihood-discriminability. The latter, in turn, is expected to go hand-in-hand with a lowering of outcome-discriminability. While remaining undocumented to date, such a pattern would contradict PT, according to which ambiguity attitudes ought to be reflected purely

in probability weighting (Wakker, 2010; Abdellaoui et al., 2011; Dimmock et al., 2015).[14] Notice that the theoretical limiting case is one whereby coding noise for outcomes and probabilities is completely independent, such that $\nu_e \perp \nu_o$, in which case we would indeed expect outcome-discriminability to be unaffected by properties pertaining to the probability dimension, as predicted by PT. That case, however, appears conceptually unlikely due to the reasons laid out above.

3 Empirical evidence

In this section, I use several datasets to test the predictions and performance of the BIM. For space reasons, only essential details are provided in the main text, while additional details and results are provided in the online appendix.

3.1 Preliminaries to empirical estimations

I use Bayesian random parameter models to estimate the model parameters (Gelman and Hill, 2006; Gelman, Carlin, Stern, Dunson, Vehtari and Rubin, 2014b). This allows me to obtain individual-level estimates of all the parameters, while estimating the full variance-covariance matrix of the model parameters. This method is geared towards maximizing the *predictive* power of the model for new data, rather than towards optimizing its fit to existing data, to contain issues arising from overfitting. I conduct the analysis using Stan (Carpenter, Gelman, Hoffman, Lee, Goodrich, Betancourt, Brubaker, Guo, Li and Riddell, 2017), and I test model performance by means of leave-one-out cross-validation (LOO; Gelman, Hwang and Vehtari, 2014a; Vehtari, Gelman and Gabry, 2017). Models are thus compared on their out-of-sample predictive ability, rather than based on their fit to a given dataset.[15] Correlations discussed in the text refer to Spearman

[14] It is, of course, also in contradiction to models that capture ambiguity attitudes purely through utility curvature, such as the one proposed by Klibanoff, Marinacci and Mukerji (2005).

[15] The results obtained with the LOO criterion are similar to, but more stable than, those obtained using the Watanabe-Akaike information criterion (WAIC; Watanabe and Opper, 2010), which is a Bayesian generalization of the deviance information criterion. All model comparisons reported are unaffected by using the WAIC instead.

rank correlations on the estimated parameters. Any p-values reported are always two-sided. Figures may cut some outliers for better visual display.

I am interested in two sets of parameters. The first set consists in the BIM parameters, which consist of two sub-groups: 1) the native BIM parameters ν, ψ, σ_e, and σ_o; and 2) the BIM-derived PT parameters α, γ, δ, and ω. I refer to γ as *likelihood-discriminability* and to α as *outcome-discriminability*. The second set of parameters I am interested in are the native PT parameters resulting from a typical PT setup plus additive noise—the most frequently used stochastic implementation of PT, and the one that is used in the both papers from which I take the data. Under PT, a prospect $(x, p; y)$ will be chosen over a sure outcome c whenever

$$c < u^{-1}[w(p)u(x) + (1 - w(p))u(y)] + \epsilon, \tag{15}$$

where $w(p)$ is a probability-distortion function mapping probabilities into decision weights, and $u(.)$ is a utility function mapping outcomes into utilities. This model is applicable to both gains and losses. To allow for data fitting, an additive error term $\epsilon \sim \mathcal{N}(0, \widehat{\omega}^2)$ is appended to the deterministic part of the model. This random error is customarily interpreted to take the form of 'white noise', i.e. implicitly assumed to be orthogonal to the deterministic part of the model.[16] To ensure comparability with the BIM, I will use power utility throughout, so that $u(x) = x^{\widehat{\alpha}}$. I will use the Goldstein and Einhorn (1987) probability-distortion function for the same reason, with parameters $\widehat{\gamma}$ and $\widehat{\delta}$, where the 'hat' serves to distinguish the PT parameters from the equivalent BIM-generated parameters.[17] I will refer to $\widehat{\gamma}$ as *likelihood-sensitivity*, and to $\widehat{\alpha}$ as *outcome-sensitivity*, to distinguish them from the equivalent BIM-derived parameters.

[16] Many implementations allow for some sort of heteroscedasticity of the error term. For instance, both Bruhin et al. (2010) and L'Haridon and Vieider (2019) allow errors to be heteroscedastic across gains and losses, as well as proportional to the outcome range in the given prospect. Contextualizing the error by making it proportional to the *utility* range, as proposed by Wilcox (2011), does not affect any of the conclusions I draw below.

[17] Prelec (1998) presents an popular alternative 2-parameter weighting function that has been used inter alia by L'Haridon and Vieider (2019). Notice, however, that the two weighting functions provide a very similar fit except for prospects with extremely small or extremely large probabilities, which do not occur in the data I analyze.

Under risk, the BIM parameters are identified whenever the PT parameters are, thus requiring sufficient variation over both probabilities and outcomes. A crucial element for both models is the inclusion of non-zero lower outcomes, without which the model parameters are only unique up to a power. All data I use fulfil this criterion. Ambiguity attitudes are trickier to identify, given how the BIM makes different predictions from PT. This means that datasets collected to fit PT models of ambiguity attitudes are generally not suited to identify BIM parameters. I will thus limit my analysis to only one original dataset, which includes the required nonzero lower outcomes needed to identify outcome-distortions under ambiguity.

Noise deserves some extra attention in this context, given its centrality to the BIM setup. Khaw et al. (2021) identify noisiness by repeating identical choices multiple times. Noise, however, can also be identified in contexts where different choices are similar, though not identical, since such choices will result in inconsistencies and stochastic dominance violations across choice lists in regions where stimuli cannot be discriminated. Experiments imposing or nudging subjects towards consistency within choice lists may somewhat underestimate noise. Notice, however, that this would bias the tests I conduct against the BIM, so that my tests on data collected to fit PT can be considered conservative.

3.2 The BIM generates PT patterns

In now illustrate some stylized patterns for two-outcome prospects based on the data of Bruhin et al. (2010) and L'Haridon and Vieider (2019). The details of the econometric approach, as well as a host of additional results and graphs for the two papers, can be found in the online appendix. I start by illustrating the generative nature of the BIM. To this end, I use a few selected countries from the data of L'Haridon and Vieider (2019) to illustrate the correlations between the BIM-derived PT parameters, α, γ, δ, and ω and the corresponding native PT parameters, $\widehat{\alpha}$, $\widehat{\gamma}$, $\widehat{\delta}$, and $\widehat{\omega}$. The four countries selected are representative of the overall trends in the data, but allow for a more nuanced visualization. Results for other data—e.g. for other countries, for losses, or for the Bruhin et al. (2010)

data— are similar and can be found in the online appendices.

The correlations are displayed in figure 1. The correlation between decision noise in the two models, shown in panel 1(a), is extremely tight ($\rho = 0.947, p < 0.001$), as one might expect given that both are identified from the noisiness of choices. Given that ω contains in itself the model parameters α and γ, whereas $\widehat{\omega}$ is implemented as an independent dimension, we would however expect larger differences to emerge for those other parameters. These differences start surfacing in panel 1(b), showing the correlation of the optimism parameter emerging from the two models ($\rho = 0.711, p < 0.001$). The distribution of the parameter under the BIM is much narrower when compared to the equivalent parameter in the PT model (IQR of 1.25 for $\widehat{\delta}$ versus 0.336 for δ), and is a manifestation of the shrinking effect of the $1 - \gamma$ power applied to the prior mean in the BIM. The parameter correlations decrease further once we move to likelihood distortions, shown in panel 1(c) ($\rho = 0.662, p < 0.001$). Notice also that values of $\widehat{\gamma} > 1$ do not generally map into values of γ close to 1, but rather into intermediate values.

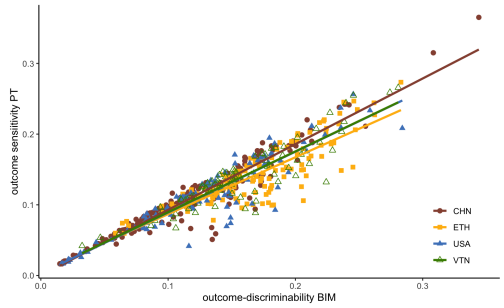

(a) Scatter plot of ω^+ and $\widehat{\omega}^+$

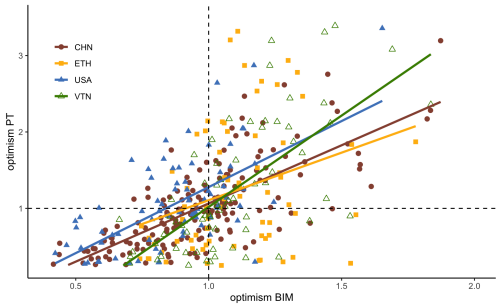

(b) Scatter plot of δ^+ and $\widehat{\delta}^+$

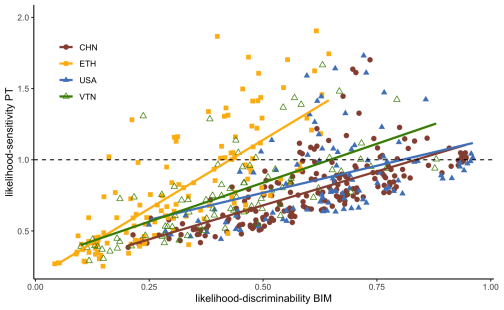

(c) Scatter plot of γ^+ and $\widehat{\gamma}^+$

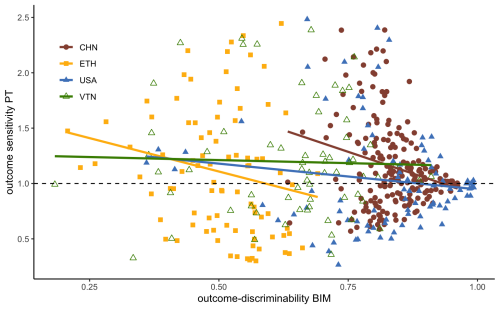

(d) Scatter plot of α and $\widehat{\alpha}^+$

Figure 1: BIM-generated parameters and PT parameters for L'Haridon and Vieider (2019)

Once we move to outcome-distortions, shown in panel 1(d), we no longer find any systematic correlation between the PT and BIM parameters ($\rho = -0.031, p = 0.522$). This absence of correlation is based on two main factors. The first is the different definition of outcome-distortions in the two models. Whereas PT applies outcome distortions to individual outcomes, the BIM results in distortions applied to the costs and benefits of a wager. The second reason concerns the coexistence of values of $\widehat{\alpha}$ larger and smaller than 1 under PT. We will see shortly that both are associated with noise. Indeed, values of the parameters falling relatively close to 1, which tend to be associated with low noise levels (see below), are very similar across the two models.

3.3 PT parameters show systematic inter-correlations

We have seen above that the BIM predicts systematic correlations in the PT parameters it generates, due to encoding noise entering the definition of both likelihood-discriminability and outcome-discriminability. The latter, in turn, my impact optimism/pessimism though their shrinking effect on those parameters. To the extent that PT parameters are indeed generated by noisy coding, we should thus expect to find systematic correlations between PT parameters. Figure 2 shows correlations between decision noise $\widehat{\omega}$ and likelihood-sensitivity $\widehat{\gamma}$ and outcome-sensitivity $\widehat{\alpha}$ for both gains and losses in the data of Bruhin et al. (2010) (the patterns in the data of L'Haridon and Vieider (2019) are very similar, and are illustrated in online appendix S2.3).

Panel 2(a) shows a scatter plot of the PT noise parameter, $\widehat{\omega}^{+}$, against the likelihood-sensitivity parameter, $\widehat{\gamma}^{+}$, for gains. The two parameters show a strong negative correlation ($\rho = -0.422, p < 0.001$), which is present in each of the 3 experiments. The results for losses, shown in panel 2(b), are very similar ($\rho = -0.389, p < 0.001$). For both gains and losses, a small group stands out that has virtually no noise and likelihood-sensitivity arbitrarily close to 1. These are the expected value maximizers detected in the mixture model of Bruhin et al. (2010), who most likely based their responses on precise calculation of expected values

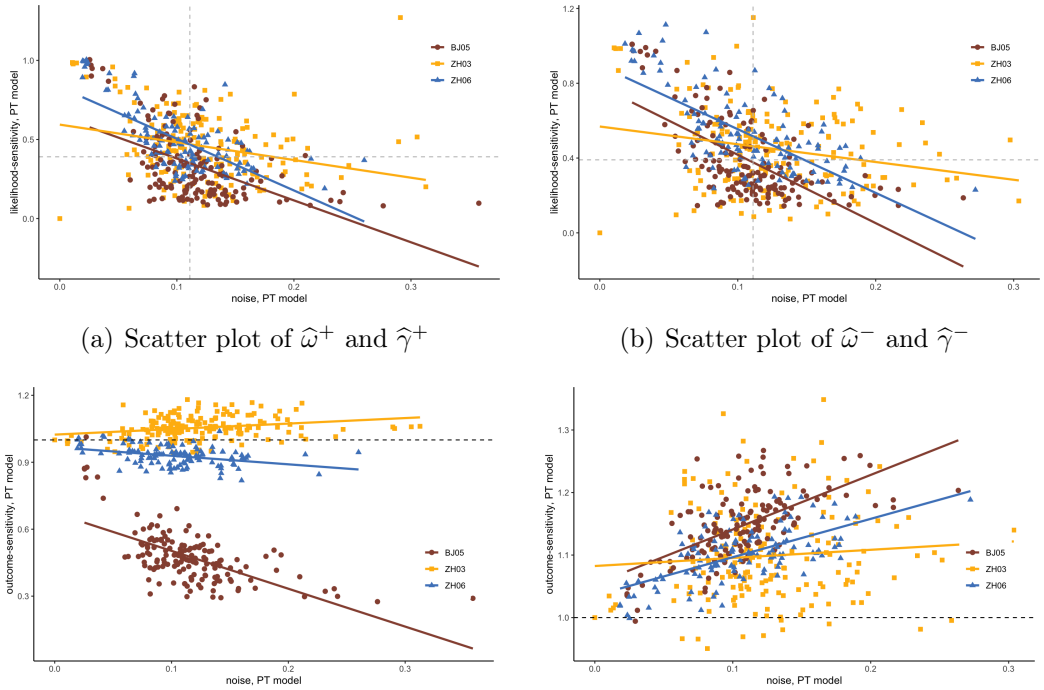

(a) Scatter plot of $\widehat{\omega}^{+}$ and $\widehat{\gamma}^{+}$

(b) Scatter plot of $\widehat{\omega}^{-}$ and $\widehat{\gamma}^{-}$

(c) Scatter plot of $\widehat{\omega}^{+}$ and $\widehat{\alpha}^{+}$

(d) Scatter plot of $\widehat{\omega}^{-}$ and $\widehat{\alpha}^{-}$

Figure 2: Scatter plot of PT parameters

The parameters have been obtained from the estimation of a PT model plus additive noise. The different colours and shapes represent the 3 experiments in Bruhin et al. (2010): ZH03 stands for Zurich 03; ZH6 for Zurich 06; and BJ05 for Beijing 05. The dashed lines indicate the median parameter values in panels 2(a) and 2(b), while indicating the salient point of 1 in panels 2(c) and 2(d).

rather than on quick and approximate judgments.

Panel 2(c) shows the correlations between the PT noise parameter, $\widehat{\omega}^{+}$, and the outcome sensitivity parameter, $\widehat{\alpha}^{+}$, for gains. Once again we witness systematic correlations, although the correlation is negative in the Beijing 05 data ($\rho = -0.566, p < 0.001$), as well as in the Zurich 06 data ($\rho = -0.357, p < 0.001$), but positive in the Zurich 03 data ($\rho = 0.279, p < 0.001$). Noise can thus be correlated with either excess outcome sensitivity or with insensitivity towards outcomes under PT. This intuition is further confirmed for losses, shown in panel 2(d). Here we witness a positive correlation in the aggregate data ($\rho = 0.258, p < 0.001$), as well as in the three individual experiments (ZH03: $\rho = 0.07, p = 0.38$; BJ05: $\rho = 0.612, p < 0.001$; ZH06: $\rho = 0.558, p < 0.001$). This is driven by concave utility for losses in all three experiments. Such patterns could, for instance, arise if noisy choice data are overfit. I further discuss this issue below.

The correlations just shown further have knock-on effects on correlations between the deterministic model parameters. In particular, likelihood-sensitivity $\widehat{\gamma}$ and outcome-sensitivity $\widehat{\alpha}$ will tend to be either positively or negatively correlated, depending on whether the correlation of $\widehat{\omega}$ with $\widehat{\alpha}$ is negative or positive. For instance, likelihood sensitivity and outcome-sensitivity for gains are positively correlated in the Beijing 05 data ($\rho = 0.463, p < 0.001$), as well as in the Zurich 06 data ($\rho = 0.753, p < 0.001$). They are, however, negatively correlated in the Zurich 03 data ($\rho = -0.254, p < 0.001$), given the positive correlation between decision noise and outcome-sensitivity. In terms of optimism δ^+ and pessimism δ^-, larger values of gamma tend to compress the distributions towards one (see online appendices for graphs and tests).

3.4 The BIM beats PT in predictive accuracy

A natural next question concerns the predictive performance of the two models. Table 1 reports test statistics for Bruhin et al. (2010), separately for their three experiments, and further separated for gains and losses. In addition to the additive noise model described above, the table also presents a test of the BMI against PT including a contextual noise term, whereby the error is made proportional to the utility difference of the prospect (Wilcox, 2011). The BIM outperforms PT in all 6 test cases for both additive noise and contextual noise, and the difference in performance is large. While there are yet other error specifications one could add to the PT model, this highlights precisely the arbitrariness in composition of decision model and error model discussed in the introduction, and is thus besides the point. The good performance of the BIM using the Bruhin et al. (2010) data is in no way exceptional. Table S2 in the online appendix provides tests of the predictive ability of the BIM versus PT using the data of L'Haridon and Vieider (2019). I conduct the tests country by country, and separately for gains and losses. Conditional on the design, this gives me 60 independent tests. Once again, the BIM easily outperforms PT in all 60 cases, with all test statistics two orders of magnitude larger than the associated standard errors.

Table 1: Evidence in favour of BIM over PT plus additive noise

	additive noise		contextual noise	
	gains	losses	gains	losses
Zurich 03	−14084.9	−13942.6	−14085.8	−13943.9
(SE)	(46.3)	(41.7)	(46.0)	(45.9)
Zurich 06	−7824.2	−7794.7	−7820.3	−7799.6
(SE)	(33.0)	(33.1)	(33.0)	(33.1)
Beijing 05	−4768.5	−4778.0	−4777.0	−4775.7
(SE)	(31.5)	(30.8)	(31.2)	(31.1)

Test results refer to ELPD (expected log pointwise predictive density) differences, and indicate the comparative performance of the two models (Gelman et al., 2014a; Vehtari et al., 2017). Negative differences constitute evidence in favour of the BIM over PT. Tests using WAIC yield virtually identical results, and are thus not reported for parsimony.

The superior predictive performance of the BIM as compared to PT may appear puzzling in the light of the fact that the BIM does not allow for discriminability parameters larger than 1, i.e $\alpha \leq 1$ and $\gamma \leq 1$. This seems due to the fact that the BIM avoids the overfitting of noisy data, which remain a problem in the PT setup notwithstanding the hierarchical model geared towards discounting noisy outliers. This fact is best illustrated with an example. Take subject 319 from L'Haridon and Vieider (2019). His PT parameters indicate extreme likelihood-sensitivity, $\widehat{\gamma} = 1.91$, joint to very low levels of optimism, at $\widehat{\delta} = 0.16$. The noisiness of his decisions is relatively high, at $\widehat{\omega} = 0.19$. Panel 3(a) shows the fit of this probability-distortion function to the raw data points, which are obtained under the assumption of linear utility. The fit appears to be terrible. At the same time, the fit of the function based on the BIM-derived PT parameters appears much more reasonable. The example from the data of Bruhin et al. (2010) shown in panel 3(b) reveals similar issues, albeit muted by the multiplicity of observations.

The underlying reason for the poor fit of the PT function can be traced back to the other parameters of the model, and particularly tends to go hand-in-hand with extreme values of outcome-sensitivity (the value for the subject in panel 3(a) is $\widehat{\alpha} = 3.74$; outcome-sensitivity for the subject in panel 3(b) is $\widehat{\alpha} = 1.05$). The outcome-discriminability parameter under the BIM tends to have much more reasonable values due to the constraints emerging from the definition of the parameter, and to the explicit modelling of parameter interactions. Taking into account the role of

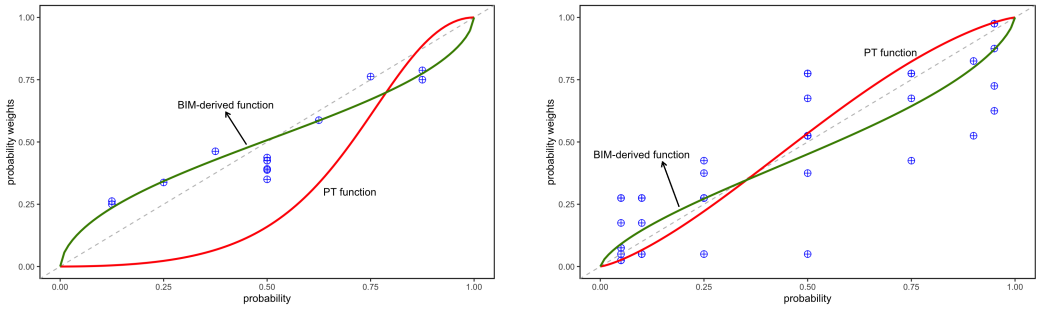

(a) Example L'Haridon and Vieider (2019) (b) Example Bruhin et al. (2010)

Figure 3: Examples of probability-distortion fit for PT vs the BIM

Fit of estimates probability-distortion functions derived from PT vs the BIM. Non-parametric data points are shown in blue, and are derived under the assumption that $u(x) = x$. Decision weights $\pi(p) = \frac{c-y}{x-y}$ are then plotted on the y-axis against the probability on the x-axis.

α in contributing to noise, this results in much better predictive performance, as well as improving the fit of the probability-distortion function to the data. Section S2.5 in the online appendix shows more example of individual-level data fit.

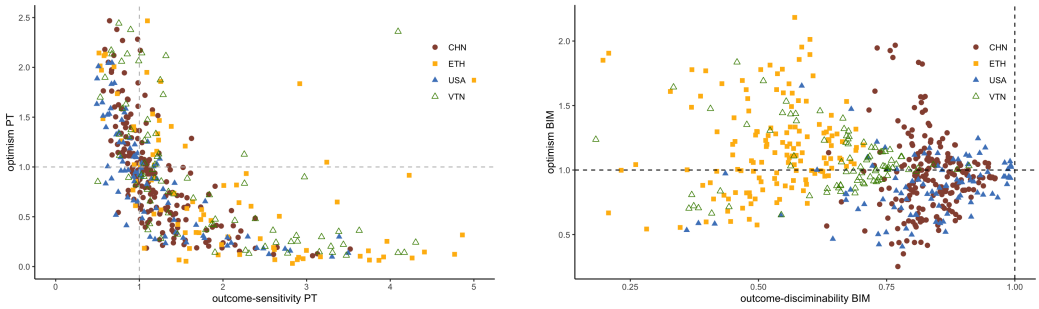

(a) Scatter plot of $\widehat{\alpha}^{+}$ and $\widehat{\delta}^{+}$ (b) Scatter plot of α^{+} and δ^{+}

Figure 4: Scatter plot of outcome-distortion and optimism for L'Haridon and Vieider (2019)

The pattern in the examples just shown is far from exceptional. Indeed, the poor fit of PT, and its even poorer performance in terms of predictive accuracy, is largely driven by the poor separability of some of its parameters, and their assumed orthogonality to noise. One particularly problematic dimension is the correlation between outcome-sensitivity $\widehat{\alpha}$ and optimism $\widehat{\delta}$ under PT. This correlation is shown in panel 4(a) of figure 4. It can be seen to be L-shaped, with large values of outcome-sensitivity corresponding to low levels of optimism, and vice versa. Since values of outcome-sensitivity larger than 1 are further associated with high noise levels, this pattern points at poor separability of the parameters in a traditional

PT setup (Zeisberger, Vrecko and Langer, 2012). The equivalent parameters α and δ in the BIM shown in panel 4(b) appear to be orthogonal to each other.

3.5 Probability-outcome separability under ambiguity

From the discussion above it follows immediately that the BIM predicts ambiguity-insensitivity, as well as allowing for ambiguity aversion. These patterns are well-documented in the literature, and can be fit well by PT. The BIM, however, also makes additional predictions. For one, we would expect coding noise, which I denote by ν_a to emphasize that it is specific to ambiguity, to be negatively correlated with ambiguity-insensitivity. In addition, we would expect coding noise to be substantially larger than under risk, so that $\nu_a > \nu_r$, where the subscript r indicates known probabilities.

Published data on ambiguity attitudes tend to be somewhat limited in that they have been collected explicitly to elicit parameters under the assumption that PT holds (or else for nonparametric analysis). That is, they typically lack the sort of stimuli that would allow one to assess whether $\widehat{\alpha}$ may be affected by ambiguity in addition to $\widehat{\gamma}$, since utility curvature for risk and ambiguity is assumed to be the same under PT due to the strict separation of the likelihood and outcome dimensions inherent in the model (Wakker, 2010; Abdellaoui et al., 2011; Dimmock et al., 2015). This strict separation does not hold in the BIM due to the effect of encoding noise on both γ and α, resulting in a prediction that both likelihood-discriminability *and* outcome-discriminability should be affected by ambiguity.

To test this prediction, I use an original dataset that contains a richer choice setup. The data contain observations for 47 subjects, and the data structure and experimental procedures closely follow those in L'Haridon et al. (2018). The stimuli are, however, richer in that the stimuli for risk are replicated exactly for ambiguity, including variation over outcomes and non-zero lower outcomes for both risk and ambiguity, thus allowing for the identification of utility curvature and outcome-discriminability under ambiguity as well as risk. Further details can be found in the online appendix.

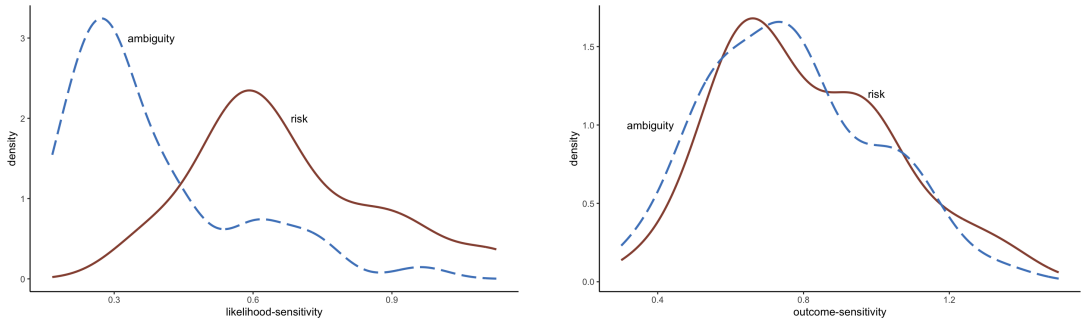

(a) Densities of $\widehat{\gamma}_r^+$ and $\widehat{\gamma}_a^+$

(b) Densities of $\widehat{\alpha}_r^+$ and $\widehat{\alpha}_a^+$

Figure 5: Densities of $\widehat{\gamma}$ and $\widehat{\alpha}$ for risk and ambiguity

I start by testing the performance of the different models against each other. The BIM easily outperforms the PT specification (elpd difference: -3521.4 , SE: 21.3). Figure 5 shows density plots for $\widehat{\gamma}$ and $\widehat{\alpha}$, for risk and ambiguity. The difference between risk and ambiguity for $\widehat{\gamma}$, shown in panel 5(a), is sizeable. This sort of lowering of likelihood-sensitivity under ambiguity relative to risk, termed ambiguity-insensitivity, is well-documented in the literature (Abdellaoui et al., 2011; Dimmock et al., 2015; Trautmann and van de Kuilen, 2015; L'Haridon et al., 2018). Enke and Graeber (2019) have documented empirically that this phenomenon goes hand-in-hand with a lowering of the confidence people declare to have in their choices, which is consistent with the account presented here. The difference between outcome sensitivity for risk and ambiguity, displayed in panel 5(b), is much less pronounced. Nevertheless, the distribution under ambiguity appears to be slightly shifted to the left. This is confirmed by a Mann-Whitney test, indicating that outcome-sensitivity is indeed reduced under ambiguity ($p = 0.012$, two-sided), thus resulting in a violation of PT.

Let us now take a look at the patterns emerging from the BIM, displayed in figure 6. Panel 6(a) shows a scatter plot of encoding noise for risk and ambiguity. The two are not significantly correlated. In accordance with the predictions discussed above, however, encoding noise for ambiguity is much larger than encoding noise for risk. This is in turn reflected in the other model parameters. In particular, likelihood-discriminability γ is much lower under ambiguity than under risk,

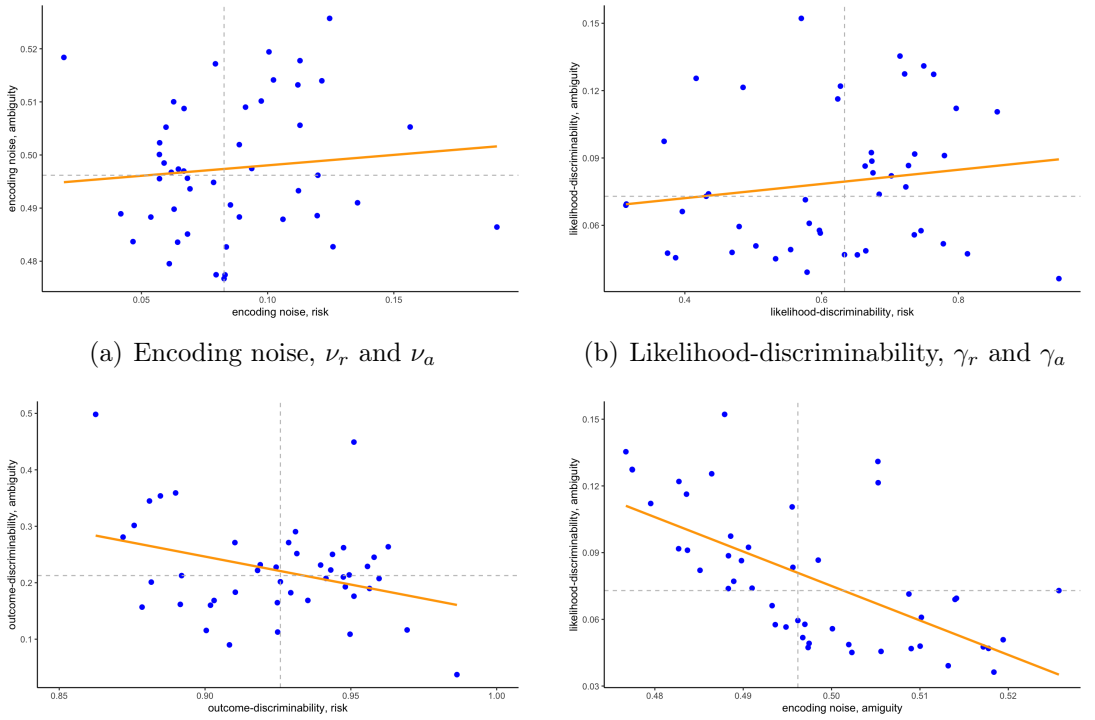

(a) Encoding noise, ν_r and ν_a

(b) Likelihood-discriminability, γ_r and γ_a

(c) Outcome-discriminability, α_r and α_a

(d) Noise ν_a and likelihood-discriminability γ_a

Figure 6: BIM parameters under risk and ambiguity

as is apparent from panel 6(b). Similar results obtain for outcome-discriminability α, shown in panel 6(c). Once again discriminability is much lower under ambiguity than it is under risk. Notice also that this difference is much more pronounced than it was under PT, which can be traced back to the different definition of outcome distortions and its interaction with encoding noise. Finally, panel 6(d) shows the correlation between encoding noise and likelihood discriminability under ambiguity. The two are strongly negatively correlated ($\rho = -0.711, p < 0.001$), as is predicted by the BIM.

4 Discussion

Traditional models of decision-making under risk and uncertainty represent behaviour by means of deterministic models applied to objectively perceived stimuli. These deterministic models are subsequently augmented by an independently chosen stochastic model to allow for choice parameters to be recovered from noisy

data. The Bayesian inference model presented in this paper turns this process on its head. Starting from the insight that the choice stimuli themselves may be encoded by noisy mental signals, I have shown how the optimal decoding of these signals by means of a mental prior may result in decisions that deviate systematically from the underlying optimal choice rule, and in particular, which may result in systematic outcome- and likelihood-distortions akin to those documented in the prospect theory literature. While predicting PT-like choice patterns, however, the model simultaneously predicts PT violations, in particular pertaining to the strict separation between the probabilistic and outcome dimensions under PT. It also suggests a novel interpretation of PT parameters, which may in turn yield better predictive performance, delivered for instance through the intuition that the mental prior may differ across different choice situations, which allows to better account for insurance and lottery play decisions than PT does based on standard parameter estimates (Sydnor, 2010).

Deviations from the normative benchmark of expected utility maximization such as predicted by the BIM are typically depicted as biases or rationality failures in economics. It is thus interesting to see how such patterns may result from a cognitive process which may be considered optimal. Indeed, the BIM builds on an optimal choice rule, which is based on accumulating information by means of Bayesian updating, and which entails expected value maximization. The deviations from the expected value benchmark purely arise from limited cognitive resources dedicated to a given decision problem, which will result in noise in the encoding of external stimuli, and hence regression to the mean of the mental prior. In the limit as coding noise goes to zero, we should indeed expected value maximization, implemented in a deterministic and noise-free fashion. This suggests a resource-saving rationale which may be highly adaptive, and which is reminiscent of concepts of bounded rationality and resource-saving through adaptation as originally discussed by Herbert Simon (1955; 1959).

Relative to prospect theory, the BIM seems especially promising when it comes to the analysis of empirical data. For one, I have documented the predictive ability

of the model to be vastly superior to the one of PT. The model also suggests that choice parameters may be malleable and amenable to manipulation, which could be exploited in policy making. When it comes to analyzing the correlates of behaviour under risk and uncertainty, the BIM furthermore provides a clear division of parameters pertaining to noisy cognition, as captured by the noisy mental encoding parameter, and those pertaining to the mental prior, which may incorporate more stable elements. This separation may help overcome issues that have hampered empirical analysis using PT, where several parameters usually co-vary in regression analysis, making an interpretation of their effects fraught with difficulties (see online appendix of L'Haridon and Vieider, 2019, for a number of examples).

A recent literature in neuroscience as well as economics has focused on how to endogenize the noisiness of the coding process (Heng, Woodford and Polania, 2020; Park and Pillow, 2020; Frydman and Jin, 2022). This has generally followed the *efficient coding* paradigm first proposed by Barlow et al. (1961). These models are, however, silent about the adaptation process of the prior, which is simply assumed to be fully adapted to the situation at hand. Many economic questions, however, seem to require an answer to how the prior itself adapts to a given environment, or how context-specific priors may interact across contexts. Effects of priming such as documented e.g. by Cohn, Engelmann, Fehr and Maréchal (2015) may be a manifestation of priors that are correlated across contexts. There may also be a longer-term component of priors, which is much less studied in neuroscience. Conceivably, this component may be influenced by long-term experiences (Malmendier and Nagel, 2011), shaped by conscious education efforts (Doepke and Zilibotti, 2014; 2017), forged by environmental adaptation (Di Falco and Vieider, 2022), or determined by intergenerational transmission (Galor and Michalopoulos, 2012; Bouchouicha and Vieider, 2019). Investigating the process under which the more stable components of the mental prior is formed ought to be a priority area for future research, as is the explicit modelling of its context-specificity.

The model proposed in this paper is more loosely linked to other models of

adaptive preferences. Most notable is the connection to the evolutionary models of Robson (2001) and Netzer (2009), who use a fitness-maximizing model predicated on limited discernibility of outcomes to derive a utility function that adapts to the local environment. A common element is that all these models incorporate the intuition of just noticeable differences in utility, although in the BIM this is a result of the compressed mental representation of stimuli, whereas it is a modelling assumption in the Robson-Netzer framework, where it results in utility taking the form of a step function. Netzer et al. (2021) present a model in which an agent receives noisy signals about different lottery arms. The agent may thereby decide to oversample some lottery arms, which could lead to the overweighting of small probability events. While some of the underlying intuitions are similar to those developed in this model, the focus is different. Netzer et al. (2021) primarily focus on the question of what happens when decision noise goes to zero, resulting in insights that are complementary to those presented in this paper.

The general framework based on the noisy neural coding of stimuli—and its extension to setups allowing for the continuing adaptation of the model parameters to changes in the environment—presents the promise of a unifying theory of individual choice behaviour. On the one hand, encoding noise plays a central role in the model, driving deviations from optimal behaviour, such as expected value maximization for small-stake risks. This results in behavioural regularities that are likely to impact decisions far beyond the particular setup used in this paper. For instance, we may well expect the presentation format of stimuli to impact choices, and the noisiness of stimulus encoding may be expected to increase systematically with the difficulty of the choice tasks. The precise implications of this insight deserve close attention, and are thus left for future work.

5 Conclusion

I proposed a model of decision-making under uncertainty resting on the noisy coding of external stimuli. These stimuli are subsequently optimally decoded using a mental prior to result in actionable decision inputs. I showed that this setup results

in a linear in log-odds specification of a popular probability-distortion function, whereby subjectively distorted likelihood ratios are traded off against a subjectively distorted cost-benefit ratio. The Bayesian inference model derived in this paper can thus be seen as a generative model of prospect theory preferences. At the same time, the model makes predictions that suggest that prospect theory will be violated in certain situations due to its strong separability precepts between outcome and probability representations. Taking the model to a variety of data sets, I have shown that the noisy coding model outperforms prospect theory significantly in terms of predictive performance. This superior performance of the noisy coding model can be traced back to its stochastic nature. While under prospect theory decision noise is added onto the deterministic model and is generally assumed to be orthogonal to the latter, in the noisy coding model decision noise and the core model parameters result from one and the same process, thus being inextricably linked. The generation of both decision model and stochastic assumptions from one and the same underlying process helps avoiding the arbitrariness in the composition of decision and stochastic models that has permeated the literature thus far. The interaction between noise and decision parameters predicted by the model appears to be a central feature of data on decisions under uncertainty, explaining the superior performance of the Bayesian inference model compared to prospect theory.

References

Abdellaoui, Mohammed, Aurélien Baillon, Lætitia Placido, and Peter P. Wakker (2011) 'The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation.' *American Economic Review* 101, 695–723

Adams, Wendy J, Erich W Graf, and Marc O Ernst (2004) 'Experience can change the'light-from-above'prior.' *Nature Neuroscience* 7(10), 1057–1058

Apesteguia, Jose, and Miguel A Ballester (2018) 'Monotone stochastic choice models: The case of risk and time preferences.' *Journal of Political Economy* 126(1), 74–106

Baillon, Aurélien, Zhenxing Huang, Asli Selim, and Peter P Wakker (2018) 'Measuring ambiguity attitudes for all (natural) events.' *Econometrica* 85, 1839–1858

Barlow, Horace B et al. (1961) 'Possible principles underlying the transformation of sensory messages.' *Sensory Communication*

Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua Teitelbaum (2013) 'The Nature of Risk Preferences: Evidence from Insurance Choices.' *American Economic Review* 103(6), 2499–2529

Bernoulli, Daniel (1738/1954) 'Exposition of a New Theory on the Measurement of Risk.' *Econometrica* 22(1), 23–36

Bouchouicha, Ranoua, and Ferdinand M. Vieider (2017) 'Accommodating stake effects under prospect theory.' *Journal of Risk and Uncertainty* 55(1), 1–28

_ (2019) 'Growth, Entrepreneurship, and Risk-Tolerance: A Risk-Income Paradox.' *Journal of Economic Growth* 24(3), 257–282

Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper (2010) 'Risk and Rationality: Uncovering Heterogeneity in Probability Distortion.' *Econometrica* 78(4), 1375–1412

Carpenter, Bob, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell (2017) 'Stan: A probabilistic programming language.' *Journal of Statistical Software* 76(1), 1–32

Choi, Syngjoo, Jeongbin Kim, Eungik Lee, Jungmin Lee et al. (2021) 'Probability

weighting and cognitive ability.' *Management Science, forthcoming*

Cohn, Alain, Jan Engelmann, Ernst Fehr, and Michel André Maréchal (2015) 'Evidence for countercyclical risk aversion: an experiment with financial professionals.' *The American Economic Review* 105(2), 860–885

Dayan, Peter, and Laurence F Abbott (2001) *Theoretical neuroscience: computational and mathematical modeling of neural systems* (Computational Neuroscience Series)

Dehaene, Stanislas (2003) 'The neural basis of the weber–fechner law: a logarithmic mental number line.' *Trends in cognitive sciences* 7(4), 145–147

_ (2011) *The number sense: How the mind creates mathematics* (OUP USA)

Dehaene, Stanislas, and Jean-Pierre Changeux (1993) 'Development of elementary numerical abilities: A neuronal model.' *Journal of cognitive neuroscience* 5(4), 390–407

Di Falco, Salvatore, and Ferdinand M. Vieider (2022) 'Environmental adaptation of risk preferences.' *Economic Journal, forthcoming*

Dimmock, Stephen G., Roy Kouwenberg, and Peter P. Wakker (2015) 'Ambiguity Attitudes in a Large Representative Sample.' *Management Science* 62(5), 1363–1380

Doepke, Matthias, and Fabrizio Zilibotti (2014) 'Culture, Entrepreneurship, and Growth.' In 'Handbook of Economic Growth,' vol. 2

_ (2017) 'Parenting with style: Altruism and paternalism in intergenerational preference transmission.' *Econometrica* 85(5), 1331–1371

Doya, Kenji, Shin Ishii, Alexandre Pouget, and Rajesh PN Rao (2007) *Bayesian brain: Probabilistic approaches to neural coding* (MIT press)

Ellsberg, Daniel (1961) 'Risk, Ambiguity and the Savage Axioms.' *Quarterly Journal of Economics* 75(4), 643–669

Enke, Benjamin, and Thomas Graeber (2019) 'Cognitive uncertainty.' Technical Report, National Bureau of Economic Research

Fechner, Gustav Theodor (1860) (Kessinger's Legacy Reprints)

Fehr-Duda, Helga, Adrian Bruhin, Thomas F. Epper, and Renate Schubert (2010)

'Rationality on the Rise: Why Relative Risk Aversion Increases with Stake Size.' *Journal of Risk and Uncertainty* 40(2), 147–180

Fehr-Duda, Helga, and Thomas Epper (2012) 'Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences.' *Annual Review of Economics* 4(1), 567–593

Fox, Craig R., and Amos Tversky (1995) 'Ambiguity Aversion and Comparative Ignorance.' *Quarterly Journal of Economics* 110(3), 585–603

Friedman, Milton, and L. J. Savage (1948) 'The Utility Analysis of Choices Involving Risk.' *Journal of Political Economy* 56(4), 279–304

Frisch, Deborah, and Jonathan Baron (1988) 'Ambiguity and rationality.' *Journal of Behavioral Decision Making* 1(3), 149–157

Friston, Karl (2005) 'A theory of cortical responses.' *Philosophical transactions of the Royal Society B: Biological sciences* 360(1456), 815–836

Friston, Karl, Philipp Schwartenbeck, Thomas FitzGerald, Michael Moutoussis, Tim Behrens, and Raymond J Dolan (2013) 'The anatomy of choice: active inference and agency.' *Frontiers in Human Neuroscience* 7, 598

Friston, Karl, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, and Giovanni Pezzulo (2017) 'Active inference: a process theory.' *Neural Computation* 29(1), 1–49

Frydman, Cary, and Lawrence J Jin (2022) 'Efficient coding and risky choice.' *Quarterly Journal of Economics* 136, 161–213

Gabaix, Xavier, and David Laibson (2017) 'Myopia and discounting.' Technical Report, National bureau of economic research

Galor, Oded, and Stelios Michalopoulos (2012) 'Evolution and the Growth Process: Natural Selection of Entrepreneurial Traits.' *Journal of Economic Theory* 147(2), 759–780

Gelman, Andrew, and Jennifer Hill (2006) *Data analysis using regression and multilevel/hierarchical models* (Cambridge university press)

Gelman, Andrew, Jessica Hwang, and Aki Vehtari (2014a) 'Understanding predictive information criteria for bayesian models.' *Statistics and computing*

24(6), 997–1016

Gelman, Andrew, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin (2014b) *Bayesian data analysis,* vol. 2 (CRC press Boca Raton, FL)

Gold, Joshua I, and Michael N Shadlen (2001) 'Neural computations that underlie decisions about sensory stimuli.' *Trends in cognitive sciences* 5(1), 10–16

_ (2002) 'Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward.' *Neuron* 36(2), 299–308

Goldstein, W, and H Einhorn (1987) 'Expression Theory and the Preference Reversal Phenomena.' *Psychological Review* 94, 236–254

Gonzalez, Richard, and George Wu (1999) 'On the Shape of the Probability Weighting Function.' *Cognitive Psychology* 38, 129–166

Green, David Marvin, John A Swets et al. (1966) *Signal detection theory and psychophysics,* vol. 1 (Wiley New York)

Gul, Faruk (1991) 'A Theory of Disappointment Aversion.' *Econometrica* 59(3), 667

Harless, David, and Colin F. Camerer (1994) 'The Predictive Utility of Generalized Expected Utility Theories.' *Econometrica* 62(6), 1251–1289

Heath, C, and A Tversky (1991) 'Preference and belief: Ambiguity and competence in choice under uncertainty. Journal of Risk and.' *Journal of Risk and Uncertainty* 4(1), 5–28

Heng, Joseph A, Michael Woodford, and Rafael Polania (2020) 'Efficient sampling and noisy decisions.' *Elife* 9, e54962

Hey, John D., and Chris Orme (1994) 'Investigating Generalizations of Expected Utility Theory Using Experimental Data.' *Econometrica* 62(6), 1291–1326

Hogarth, Robin M., and Hillel J. Einhorn (1990) 'Venture Theory: A Model of Decision Weights.' *Management Science* 36(7), 780–803

Holt, Charles A., and Susan K. Laury (2002) 'Risk Aversion and Incentive Effects.' *American Economic Review* 92(5), 1644–1655

Kahneman, Daniel, and Amos Tversky (1979) 'Prospect Theory: An Analysis of

Decision under Risk.' *Econometrica* 47(2), 263 – 291

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

Knill, David C, and Alexandre Pouget (2004) 'The bayesian brain: the role of uncertainty in neural coding and computation.' *TRENDS in Neurosciences* 27(12), 712–719

Körding, Konrad Paul, and Daniel M Wolpert (2004) 'The loss function of sensorimotor learning.' *Proceedings of the National Academy of Sciences* 101(26), 9839–9842

L'Haridon, Olivier, and Ferdinand M. Vieider (2019) 'All over the map: A worldwide comparison of risk preferences.' *Quantitative Economics* 10, 185–215

L'Haridon, Olivier, Ferdinand M. Vieider, Diego Aycinena, Agustinus Bandur, Alexis Belianin, Lubomir Cingl, Amit Kothiyal, and Peter Martinsson (2018) 'Off the charts: Massive unexplained heterogeneity in a global study of ambiguity attitudes.' *Review of Economics and Statistics* 100(4), 664–677

Malmendier, Ulrike, and Stefan Nagel (2011) 'Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?' *The Quarterly Journal of Economics* 126(1), 373–416

Markowitz, Harry (1952) 'The Utility of Wealth.' *Journal of Political Economy* 60(2), 151–158

McElreath, Richard (2016) *Statistical Rethinking: A Bayesian Course with Examples in R and Stan* (Academic Press)

Netzer, Nick (2009) 'Evolution of time preferences and attitudes toward risk.' *American Economic Review* 99(3), 937–55

Netzer, Nick, Arthur Robson, Jakub Steiner, and Pavel Kocourek (2021) 'Endogenous risk attitudes.' Working Paper

Park, Il Memming, and Jonathan W Pillow (2020) 'Bayesian efficient coding.'

BioRxiv p. 178418

Payzan-LeNestour, Elise, and Michael Woodford (2021) 'Outlier blindness: A neurobiological foundation for neglect of financial risk.' *Journal of Financial Economics, forthcoming*

Petzschner, Frederike H, and Stefan Glasauer (2011) 'Iterative bayesian estimation as an explanation for range and regression effects: a study on human path integration.' *Journal of Neuroscience* 31(47), 17220–17229

Petzschner, Frederike H, Stefan Glasauer, and Klaas E Stephan (2015) 'A bayesian perspective on magnitude estimation.' *Trends in cognitive sciences* 19(5), 285–293

Prelec, Drazen (1998) 'The Probability Weighting Function.' *Econometrica* 66, 497–527

Preston, Malcolm G., and Philip Baratta (1948) 'An Experimental Study of the Auction-Value of an Uncertain Outcome.' *The American Journal of Psychology* 61(2), 183

Rabin, Matthew (2000) 'Risk Aversion and Expecetd Utility Theory: A Calibration Theorem.' *Econometrica* 68, 1281–1292

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

Savage, Leonard J. (1954) *The Foundations of Statistics* (New York: Wiley)

Schmidt, Ulrich, Chris Starmer, and Robert Sugden (2008) 'Third-generation prospect theory.' *Journal of Risk and Uncertainty* 36(3), 203–223

Simon, Herbert A (1955) 'A behavioral model of rational choice.' *The Quarterly Journal of Economics* 69(1), 99–118

_ (1959) 'Theories of decision-making in economics and behavioral science.' *The American Economic Review* 49(3), 253–283

Stan Development Team (2017) 'RStan: the R interface to Stan.' R package version 2.17.2

Sydnor, Justin (2010) '(Over)insuring Modest Risks.' *American Economic Journal: Applied Economics* 2(4), 177–199

Tom, Sabrina M., Craig R. Fox, Christopher Trepel, and Russell A. Poldrack (2007) 'The Neural Basis of Loss Aversion in Decision-Making Under Risk.' *Science* 315(5811), 515–518

Trautmann, Stefan T., and Gijs van de Kuilen (2015) 'Ambiguity Attitudes.' In 'The Wiley Blackwell Handbook of Judgment and Decision Making' (Wiley Blackwell)

Tversky, Amos, and Daniel Kahneman (1992) 'Advances in Prospect Theory: Cumulative Representation of Uncertainty.' *Journal of Risk and Uncertainty* 5, 297–323

Vehtari, Aki, Andrew Gelman, and Jonah Gabry (2017) 'Practical bayesian model evaluation using leave-one-out cross-validation and waic.' *Statistics and computing* 27(5), 1413–1432

Vieider, Ferdinand M. (2021) 'Noisy coding of time delays and reward discounting.' Working Paper WP WP 21/1036 of Faculty of Economics and Business Administration, Ghent University, Belgium WP 21/1036, Ghent University, Faculty of Economics and Business Administration

Vilares, Iris, and Konrad Kording (2011) 'Bayesian models: the structure of the world, uncertainty, behavior, and the brain.' *Annals of the New York Academy of Sciences* 1224(1), 22

Wakker, Peter P. (2010) *Prospect Theory for Risk and Ambiguity* (Cambridge: Cambridge University Press)

Watanabe, Sumio, and Manfred Opper (2010) 'Asymptotic equivalence of bayes cross validation and widely applicable information criterion in singular learning theory.' *Journal of machine learning research* 11(12), 3571–3594

Wilcox, Nathaniel T. (2011) ''Stochastically more risk averse:' A contextual theory of stochastic discrete choice under risk.' *Journal of Econometrics* 162(1), 89–104

Woodford, Michael (2012) 'Prospect theory as efficient perceptual distortion.' *American Economic Review* 102(3), 41–46

_ (2020) 'Modeling imprecision in perception, valuation, and choice.' *Annual Review of Economics* 12, 579–601

Yaari, Menahem E. (1987) 'The Dual Theory of Choice under Risk.' *Econometrica* 55(1), 95–115

Zeisberger, Stefan, Dennis Vrecko, and Thomas Langer (2012) 'Measuring the time stability of Prospect Theory preferences.' *Theory and Decision* 72(3), 359–386

Zhang, Hang, and Laurence T Maloney (2012) 'Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition.' *Frontiers in neuroscience* 6, 1

ONLINE APPENDIX

Decisions under Uncertainty as Bayesian Inference

S1 Derivation of the noisy coding model

S1.1 Derivation details

In this section, I provide additional details about the derivations underlying the equations shown in the main text. Combining the likelihoods in equation 4 in the main text with the priors in equation 5 in the main text by Bayesian updating, we obtain the following posterior distributions conditional on the mental signals:

$$ln\left(\frac{P[e]}{P[\tilde{e}]}\right)|r_e \sim \mathcal{N}\left(\frac{\sigma_e^2}{\sigma_e^2+\nu_e^2} \times r_e + \frac{\nu_e^2}{\sigma_e^2+\nu_e^2} \times ln\left(\frac{\widehat{P}[e]}{\widehat{P}[\tilde{e}]}\right), \frac{\nu_e^2\sigma_e^2}{\nu_e^2+\sigma_e^2}\right) \tag{16}$$

$$ln\left(\frac{c-y}{x-c}\right)|r_o \sim \mathcal{N}\left(\frac{\sigma_o^2}{\sigma_o^2+\nu_o^2} \times r_o + \frac{\nu_o^2}{\sigma_o^2+\nu_o^2} \times ln\left(\frac{\widehat{k}}{\widehat{\tilde{b}}}\right), \frac{\nu_o^2\sigma_o^2}{\nu_o^2+\sigma_o^2}\right). \tag{17}$$

where we define $\gamma \triangleq \frac{\sigma_e^2}{\sigma_e^2+\nu_e^2}$ and $\alpha \triangleq \frac{\sigma_o^2}{\sigma_o^2+\nu_o^2}$. It follows that $\frac{\nu_e^2}{\sigma_e^2+\nu_e^2} = 1-\gamma$ and $\frac{\nu_o^2}{\sigma_o^2+\nu_o^2} = 1-\alpha$. Taking the expressions multiplying the logarithms to the exponent and further defining $\xi \triangleq \left(\frac{\widehat{P}[e]}{\widehat{P}[\tilde{e}]}\right)^{1-\gamma}$ and $\zeta \triangleq \left(\frac{\widehat{k}}{\widehat{\tilde{b}}}\right)^{1-\alpha}$ gives the posterior expectations shown in equation 6 in the main text.

We can then substitute the posterior expectations into the mental choice rule in equation 7 to obtain the threshold equation 8 in the main text. Given that r_e and r_o follow a normal distribution, their weighted difference in equation 8 will itself follow a normal distribution with an expectation equal to the weighted difference of the means of the distributions of r_e and r_o:

$$\gamma \times r_e - \alpha \times r_o \sim \mathcal{N}\left(\gamma \times ln\left(\frac{P[e]}{P[\tilde{e}]}\right) - \alpha \times ln\left(\frac{c-y}{x-c}\right), \omega^2\right), \tag{18}$$

where $\omega \triangleq \sqrt{\gamma^2\nu_e^2 + \alpha^2\nu_o^2}$ represents the standard deviation of the weighted dif-

ference of noisy mental signals. This yields the following z-score:

$$z = \frac{\gamma \times r_e - \alpha \times r_o - \left[\gamma \times ln\left(\frac{P[e]}{P[\tilde{e}]}\right) - \alpha \times ln\left(\frac{c-y}{x-c}\right)\right]}{\omega}, \tag{19}$$

which follows a standard normal distribution. Comparing this z-score to an equivalent z-score for equation 8 yields the probabilistic choice rule in equation 9.

S1.2 Derivation for unlogged choice rule

The derivations in the main text as well as the details in the last subsection were based on the logarithmic choice rule in equation 3. If we were to use the median of the posterior instead of its mean in the choice rule, then the derivations would be identical if we were to base them on the unlogged choice rule in equation 1 instead. Körding and Wolpert (2004) show that this is indeed optimal if the absolute value of the error is used as a loss function. Given that the derivation in the main text is based on the mean, however, the two choice rules will result in a slightly different interpretation of the prior mean. I here present the derivation based on the unlogged choice rule. We then have the following mental choice rule:

$$E\left[\frac{P[e]}{P[\tilde{e}]}|r_e\right] > E\left[\frac{c-y}{x-c}|r_o\right]. \tag{20}$$

If we want to use the posterior expectation in the choice rule instead of the median, as done here, then we need to obtain the posterior means of the unlogged quantities by taking the exponential of the expectations in equation 6 in the main text:

$$E\left[\left(\frac{P[e]}{P[\tilde{e}]}\right)|r_e\right] = e^{\gamma \times r_e + ln(\xi) + \frac{1}{2}\sigma_e^2} \quad , \quad E\left[\left(\frac{c-y}{x-c}\right)|r_o\right] = e^{\alpha \times r_o + ln(\zeta) + \frac{1}{2}\sigma_o^2}, \tag{21}$$

which are derived based on the properties of the mean of the log-normal distribution. We can now substitute these quantities into the following mental choice rule, to obtain:

$$e^{\gamma \times r_e + ln(\xi) + \frac{1}{2}\sigma_e^2} > e^{\alpha \times r_o + ln(\zeta) + \frac{1}{2}\sigma_o^2}. \tag{22}$$

Taking the logarithm of both sides, defining $\widehat{\xi} \triangleq e^{ln(\xi)+\frac{1}{2}\sigma_e^2}$ and $\widehat{\zeta} \triangleq e^{ln(\zeta)+\frac{1}{2}\sigma_o^2}$, we obtain:

$$\gamma \times r_e - \alpha \times r_o > ln(\widehat{\delta})^{-1}, \tag{23}$$

where $\widehat{\delta} \triangleq \widehat{\xi} \times \widehat{\zeta}^{-1}$. All further derivations proceed like above. The sole difference with the results presented in the main text thus flow from the difference in definition of $\widehat{\delta}$ versus δ. That is, the variance of the priors will enter into the definition of $\widehat{\delta}$, while it will not enter the definition of δ. In the main text, I argued that the logged choice rule appears more plausible than its unlogged version for computational reasons. The less intuitive formulation of the prior emerging from the unlogged choice rule may constitute a further argument for my preferred setup. That being said, the difference between the two setups is *quantitative* in nature, rather than *qualitative*, and it does not affect the main conclusions drawn in the paper.

S2 Additional results

This section presents additional details and results for the analysis of 2-outcome wagers under risk, both under PT and the BIM. The structure follows roughly the one followed in the main text. Notice that the basic econometric setup in Bruhin et al. (2010) and L'Haridon and Vieider (2019) are the same, as are the type of stimuli used in the two papers, which will allow me to discuss the econometric approach in one unified setup. Whereas Bruhin et al. (2010) applied such a setup to a finite mixture model, and L'Haridon and Vieider (2019) estimated a (frequentist version of a) random preference model, I will use the basic setup to estimate a Bayesian random preference (or hierarchical) model.

S2.1 Estimation of the Bayesian hierarchical model

I estimate individual-level parameters using Bayesian hierarchical models (Gelman and Hill, 2006; Gelman et al., 2014b; McElreath, 2016). Take a parameter vector $\boldsymbol{\theta}_i$, indicating individual-level parameters. This vector follows the following

distribution:

$$\boldsymbol{\theta}_i \sim \mathcal{N}(\overline{\boldsymbol{\theta}}, \Sigma), \tag{24}$$

where $\overline{\boldsymbol{\theta}}$ is a vector of hyperparameters containing the means of the individual-level parameters, and Σ is a variance-covariance matrix of the individual-level parameters. I estimate the parameters in Stan launched from Rstan (Stan Development Team, 2017). Estimations typically employ 4 chains with 2000 iterations per chain, of which 1000 are warmup iterations—the default settings of Stan (in some cases iterations may be increased because of efficiency concerns in the estimated parameters). I check convergence by examining the R-hat statistics, and by checking for divergent iterations. The estimation algorithm was developed using simulations, and making sure that I could recover individual-level parameters from those simulations. The model endogenously estimates the priors for the individual-level parameters from the aggregate data, resulting in partial pooling. This is indeed a central strength of the model, which tends to reduce issue with overfitting of few observations. The priors for the estimation of the aggregate-level means are chosen in such a way as to be mildly regularizing (McElreath, 2016). That is, their variance is chosen in such a way that all plausible parameters fall into a region attributed high likelihood, but narrow enough to nudge the simulation algorithm towards convergence. In any case, the datasets I use have sufficient data at the aggregate level for the priors chosen to have little or no impact on the final result.

Following the estimation approach used in both paper and the structure of the data, I use directly the density around the observed switching point to estimate the model. While this deviates somewhat from the discrete choice model presented in the main text for the BIM, using such a model instead does not affect the results in any way, since both datasets push participants towards single switching. This is indeed the sense in which certainty equivalents of the type used in the data are not specifically geared towards estimating the BIM, being rather driven by the deterministic conceptions underlying PT. In this sense, the test of the BIM against PT is conservative in nature, since the data tend to naturally favour PT.

For the PT model, the model thus takes the following form:

$$ce \sim \mathcal{N}(u^{-1}[w(p)u(x) + (1 - w(p))u(y)], \widehat{\omega}^2 \times |x - y|), \tag{25}$$

with $u(x) = x^{\widehat{\alpha}}$ and $w(p) = \frac{\widehat{\delta}p^{\widehat{\gamma}}}{\widehat{\delta}p^{\widehat{\gamma}}+(1-p)^{\widehat{\gamma}}}$, and where multiplying the variance $\widehat{\omega}$ by $|x - y|$ allows for heteroscedasticity across choice lists with different step sizes between choices, following the approach in the original papers from which I took the data. Contextualizing choices by letting the error be heteroscedastic across to the *utility difference*, $|u(x) - u(y)|$, such as proposed by Wilcox (2011), does not affect the conclusions I draw in any way.

The priors are chosen such as to be informative about the expected location of the model parameters, without imposing any undue restrictions on the data. This is typically referred to as *mildly regularizing* priors, and it helps convergence in the model. For instance, the prior chosen for $\widehat{\gamma}$ has a mean of 0.7 on the original scale, with 95% of the probability mass allocated to a range of $[0, 3.92]$. This can be expected to encompass most likely parameter values. Given furthermore the large quantity of data present at the aggregate level, the data can easily overpower the prior even for parameters falling outside this range. Making the prior more diffuse and shifting the mean to e.g. 1, does not affected the estimated parameters in any way, showing that the prior has only a minimal influence on the ultimate parameter estimates.

For the BIM I proceed similarly, but I transform the data to obtain a 'dual-EU decision weight', $dw \triangleq \frac{ce-y}{x-y}$, which is distributed as follows:

$$dw \sim \mathcal{N}\left(\frac{\delta^{\frac{1}{\alpha}}p^{\frac{\gamma}{\alpha}}}{\delta^{\frac{1}{\alpha}}p^{\frac{\gamma}{\alpha}} + (1-p)^{\frac{\gamma}{\alpha}}}, \nu \times \sqrt{\alpha^2 + \gamma^2}\right), \tag{26}$$

where the parameters are defined as in the model in the text, i.e. they are derived from the underlying parameters $\nu, \sigma_p, \sigma_o, \psi$. Given that much less is known about these parameters, I have generally chosen priors with a mean of 0.5 and a larger variance, in order not to preclude any but the most extreme estimates ex ante. As above, the data are sufficient to overcome this prior if necessary, and changing the

specifications of the prior to make it more diffuse does not affect the estimates.

S2.2 The BIM generates PT functionals: Additional results

In what follows, I present additonal results illustrating the generative nature of BIM for PT functionals. Figure S7 shows the correlations between the BIM-derived PT parameters and the true PT parameters for losses in the L'Haridon and Vieider (2019) data. The results for losses track those for gains closely.

I start again by discussing the correlation between decision noise in the BIM with the SD of the additive noise term in PT, shown in panel 7(a). The association between the two variables is extremely tight ($\rho = 0.909, p < 0.001$), as one might expect given the fact that both are identified from the noisiness of choices. Given that ω^- contains in itself the model parameters α^- and γ^-, whereas $\widehat{\omega}$ is implemented as an independent dimension, we would however expect larger differences to emerge on those parameters. These differences start surfacing in panel 7(b), showing the correlation of the optimism parameter emerging from the two models. At $\rho = 0.862, p < 0.001$, they are again highly correlated. Indeed, δ^- is only indirectly influenced by the decision noise encapsulated in ω^-. Nevertheless, this indirect influence surfaces in the narrower distribution the parameter has under the BIM compared to the PT model (IQR of 0.522 for $\widehat{\delta}^-$ versus 0.276 for δ^-), and is a manifestation of the shrinking power of $1 - \gamma^-$ in the BIM.

The parameter correlations barely decrease once we move to likelihood distortions, shown in panel 7(c), and remain highly significant ($\rho = 0.827, p < 0.001$). As already witnessed for gains, values of $\widehat{\gamma}^- > 1$ estimated using the PT specification do not necessarily map into values of γ^- close to 1. This reflects a central mechanism of the BIM at work. Values of likelihood-sensitivity larger than 1 under PT are typically associated with high levels of noise. Given the close association between decision noise and γ in the BIM, those choice patterns are best reflected by increased decision noise, which in turn results in values of $\gamma < 1$. It is this which drives the superior predictive performance of the BIM. Outcome distortions

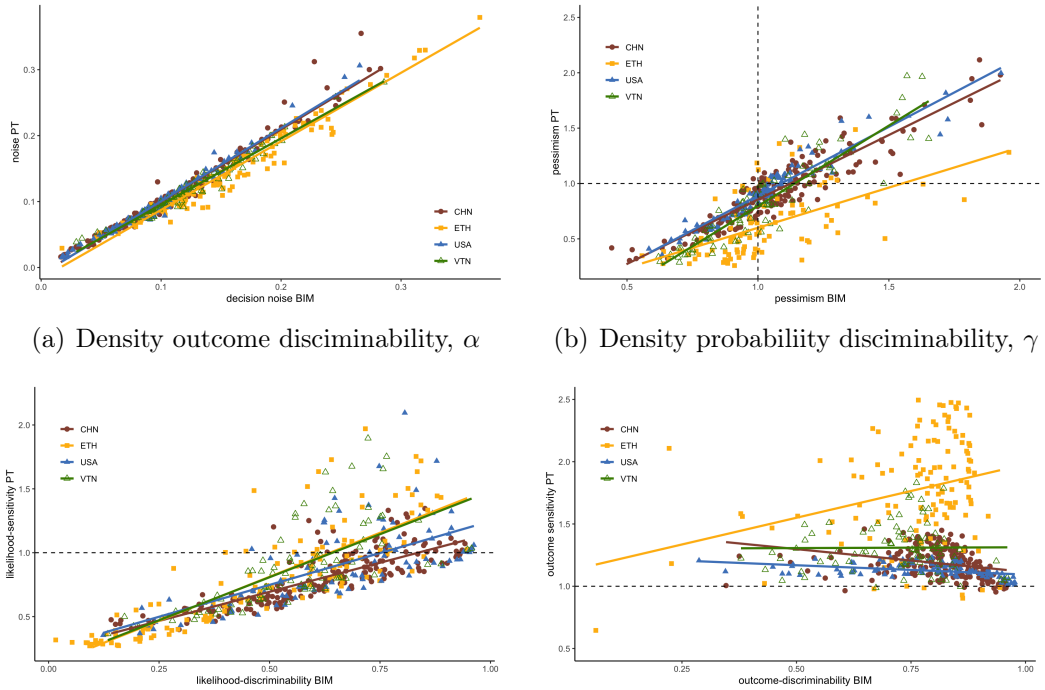

(a) Density outcome disciminability, α

(b) Density probabiliity disciminability, γ

(c) Density optimism

(d) Density decison noise

Figure S7: BIM-generated parameters and PT parameters for losses in L'Haridon and Vieider (2019)

are shown in panel 7(d), where as for gains we find no correlation between the PT and BIM parameter.

S2.3 PT parameter correlations: Additional results

This section adds some further results to the correlations amongst PT parameters presented in the main text. I will present additional results for both Bruhin et al. (2010) and L'Haridon and Vieider (2019), and for losses as well as for gains.

I start by presenting correlations between additional PT parameters in the Bruhin et al. (2010) data. Another interesting pattern emerges from the relation between likelihood-sensitivity and optimism for gains, shown in figure S8 panel 8(a), and between likelihood-sensitivity and pessimism for losses, shown in panel 8(b). In both cases, the values of $\widehat{\delta}$ are most dispersed for small values of $\widehat{\gamma}$, with the dispersion decreasing markedly as $\widehat{\gamma}$ increases, resulting in a funnel with the narrow part pointing to the right. Testing the correlation of absolute deviations of $\widehat{\delta}$ from

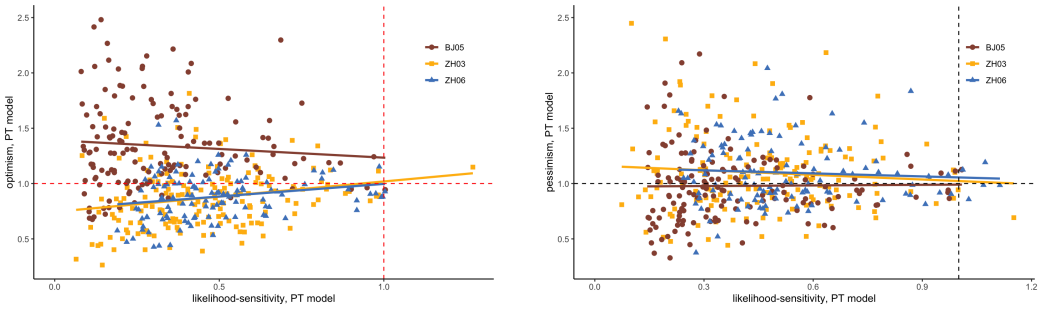

(a) Scatter plot of $\widehat{\gamma}^{+}$ and $\widehat{\delta}^{+}$ (b) Scatter plot of $\widehat{\gamma}^{-}$ and $\widehat{\delta}^{-}$

Figure S8: Scatter plot of PT parameters $\widehat{\gamma}$ and $\widehat{\delta}$
The parameters have been obtained from the estimation of a PT model plus additive noise. The different colours and shapes represent the 3 experiments in Bruhin et al. (2010): ZH03 stands for Zurich 03; ZH6 for Zurich 06; and BJ05 for Beijing 05.

1 with $\widehat{\gamma}$, I find highly significant effects for both gains ($\rho = -0.284, p < 0.001$) and losses ($\rho = -0.264, p < 0.001$). These patterns have no obvious explanation under PT. In the BIM, however, they are predicted by the definition of $\delta = \left(\frac{\psi}{1-\psi}\right)^{1-\gamma}$. That is, larger values of γ shift the attention from the prior to the likelihood, thus compressing δ towards 1. Panel 9(a) visualizes the correlations between $\widehat{\alpha}$ and $\widehat{\gamma}$ already discussed in the main text.

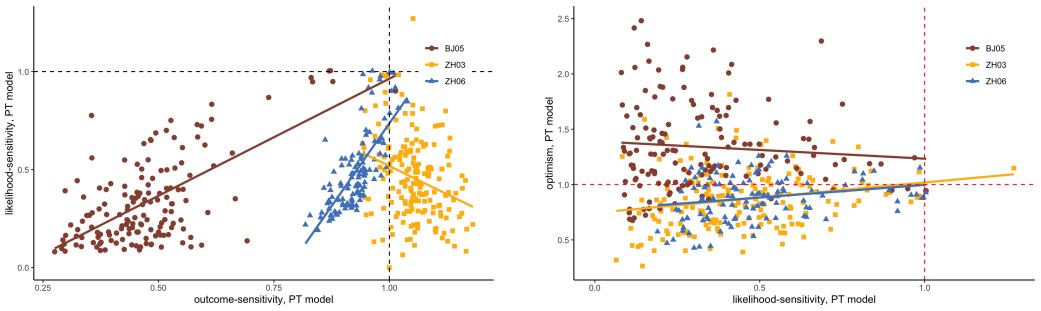

(a) Scatter plot of $\widehat{\alpha}$ against $\widehat{\gamma}$ (b) Scatter plot of $\widehat{\gamma}$ against $\widehat{\delta}$

Figure S9: Correlations between PT parameters in Bruhin et al. (2010)

Panel 9(b) further shows correlational patterns between $\widehat{\gamma}$ and $\widehat{\delta}$. According to the BIM, the results should depend on the initial value of ψ. In particular, the larger the value of $\widehat{\gamma}$, the closer the value of $\widehat{\delta}$ should be compressed towards 1. This is exactly what we observe. In experiments where we observe preponderantly values of $\widehat{\delta} < 1$, the distance to 1 decreases as $\widehat{\gamma}$ increases, thus resulting in a positive correlation between the two parameters. This is the case in the Zurich

03 experiment ($\rho = 0.262, p < 0.001$), as well as in the Zurich 06 experiment ($\rho = 0.369, p < 0.001$). In the Beijing 05 experiment, on the other hand, we observe very large values of $\widehat{\delta} > 1$. We may thus expect a negative relationship with $\widehat{\gamma}$. We fail to observe such a relationship in the data ($\rho = 0.045, p = 0.54$). This may be due to the fact that we have very few observations with large likelihood-sensitivity in that experiment.

I next document correlations amongst the PT parameters in the global data of L'Haridon and Vieider (2019), for all 30 countries and 3000 subjects. Different countries are not distinguished by colour or symbols as they were before, since there are not enough colours and symbols to render such a distinction intelligible. Correlations are tested on parameters demeaned at the country level, corresponding to a fixed effects specification.

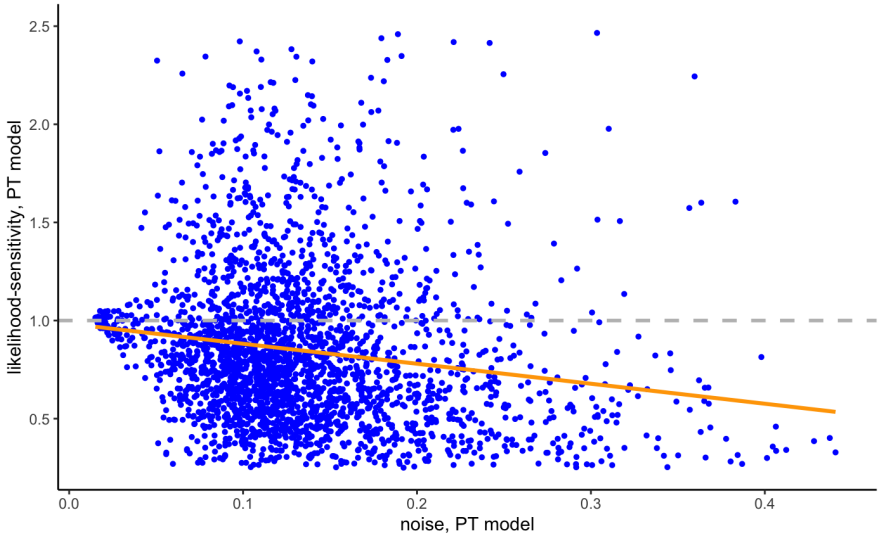

Figure S10: Correlatiion of $\widehat{\omega}^{+}$ and $\widehat{\gamma}^{+}$ in L'Haridon and Vieider (2019)

Figure S10 shows the correlation between noise and likelihood-sensitivity for gains under PT. We find the usual negative correlation already described for the risk and rationality data ($\rho = -0.2, p < 0.001$). At first, the negative correlation may appear somewhat weaker than witnessed for Bruhin et al. (2010). Closer examination reveals that this is due to the presence of a larger number of estimates of $\widehat{\gamma}^{+} > 1$. As already described for utility curvature in the main text, both negative and positive deviations tend to be correlated with noise for parameters

estimated in a PT context. Taking absolute deviations of likelihood sensitivity from 1, $|1 - \widehat{\gamma}^{+}|$, the correlation becomes indeed much stronger ($\rho = 0.386, p < 0.001$).

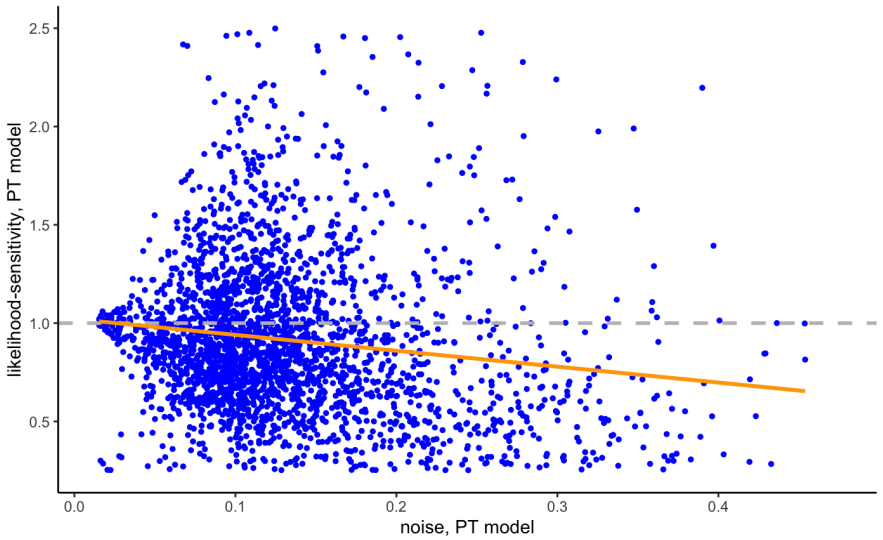

Figure S11: Correlatiion of $\widehat{\omega}^{-}$ and $\widehat{\gamma}^{-}$ in L'Haridon and Vieider (2019)

The patterns for losses are very similar, and are shown in figure S11. Once again, we observe a negative correlation in the global data ($\rho = -0.167, p < 0.001$), albeit one that is not as strong as we might have expected based on the Bruhin et al. (2010) results. This is once again driven by the presence of values of $\widehat{\gamma}^{-} > 1$, which are also associated with high noise levels. Looking at the correlations between noise and absolute deviations of the variable from 1, $|1 - \widehat{\gamma}^{-}|$, the correlation thus appears much stronger ($\rho = 0.445, p < 0.001$).

Figure S12 shows the global correlation between noise and outcome sensitivity for gains. The correlation in the global data is positive, reflecting the fact that the median outcome sensitivity parameter is positive at 1.125. The correlation is highly significant ($\rho = 0.141, p < 0.001$). Results for losses are shown in figure S13, where the same patterns appear to be even more accentuated ($\rho = 0.34, p < 0.001$).

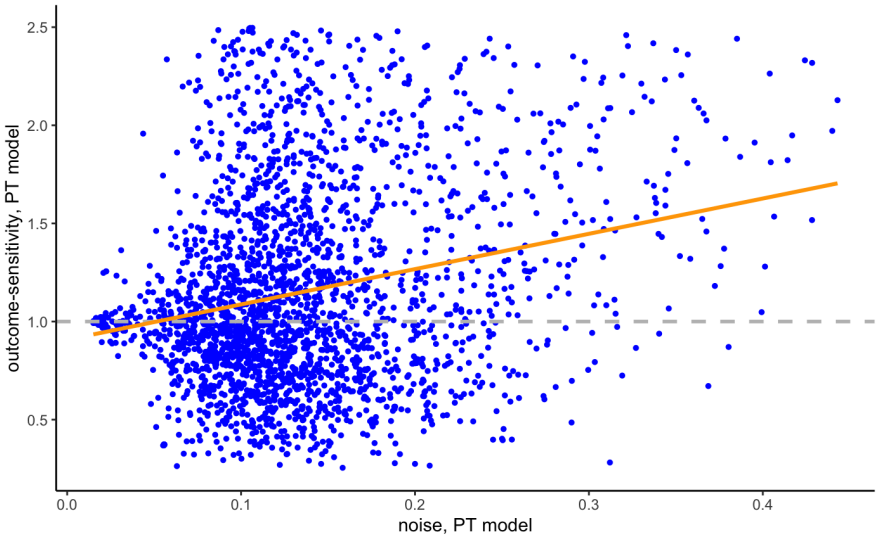

Figure S12: Correlatiion of $\widehat{\omega}^+$ and $\widehat{\alpha}^+$ in L'Haridon and Vieider (2019)

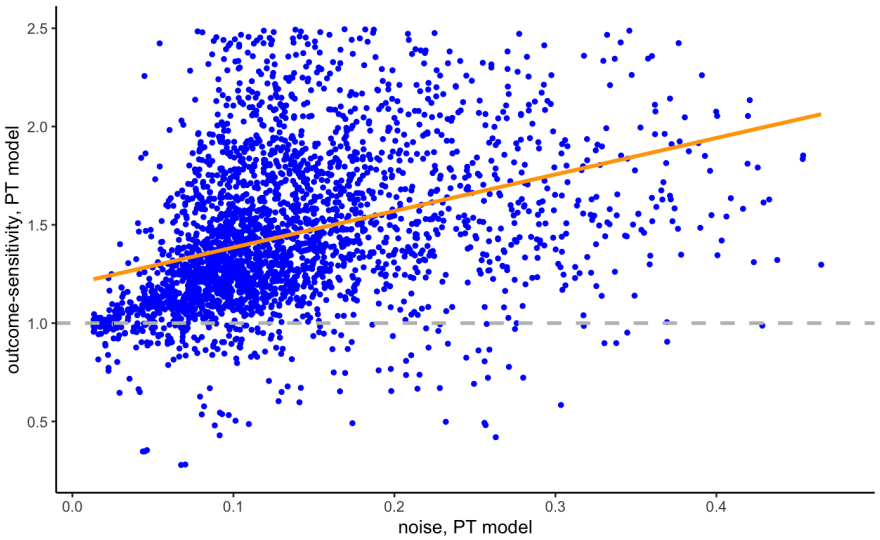

Figure S13: Correlatiion of $\widehat{\omega}^+$ and $\widehat{\alpha}^+$ in L'Haridon and Vieider (2019)

S2.4 Distribution of BIM and BIM-generated parameters

In this section, I show the distribution of the BIM parameters. This concerns both the native BIM parameters, and the BIM-derived PT parameters.

Let us now take a look at the native BIM parameters ν, ψ, σ_e and σ_o. Figure S14 shows the density functions for the 4 BIM parameters for gains for the Bruhin et al. (2010) data. The parameters are distributed over reasonable ranges, and resemble each other for the different experiments. The largest differences between

experiments emerge for the standard deviation of the outcome prior, σ_o, in panel 14(b) and for the fixed point of the probability-distortion function, ψ, in panel 14(d). The first has a relatively broad distribution in the Zurich 03 experiment. The second tends towards higher values in the Beijing 05 experiment.

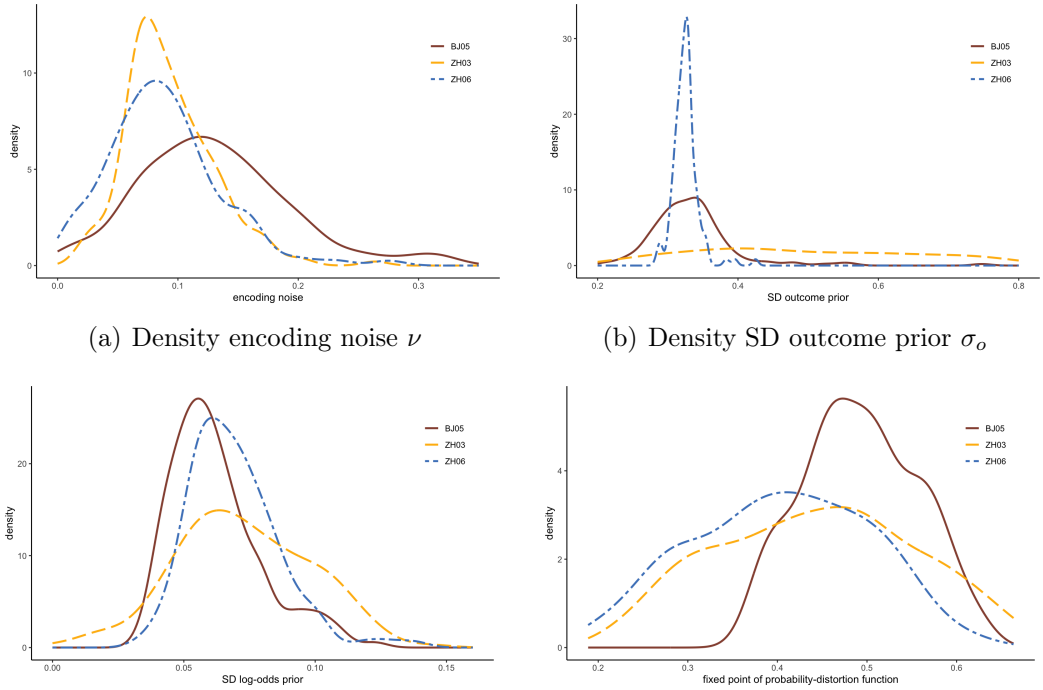

(a) Density encoding noise ν

(b) Density SD outcome prior σ_o

(c) Density SD probabillity prior σ_e

(d) Density fixed point probability-distortions ψ

Figure S14: Density functions of BIM parameters for Bruhin et al. (2010)

I next plot those same parameters for the data of L'Haridon and Vieider (2019), once again for gains and for the usual 4 countries. The parameter distributions can be seen in figure S15. Panel 15(a) shows systematic differences in encoding noise between countries, on top of the differences between individuals within each country. In particular, the US and China have relatively low encoding noise, while Ethiopia has very large encoding noise, with Vietnam falling in between. Panel 15(d) also shows systematic differences between countries, with Vietnam and Ethiopia having larger fixed points of the probability-distortion function on average. The between-country differences for the two scale parameters of the prior, shown in panel 15(b) and panel 15(c) are somewhat less pronounced.

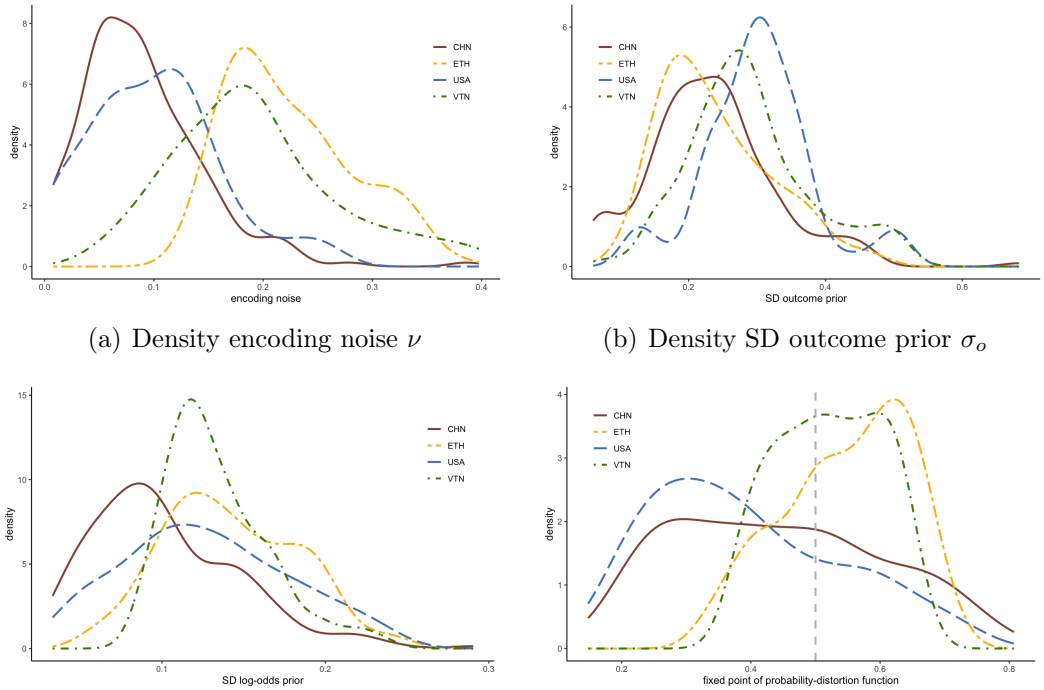

(a) Density encoding noise ν

(b) Density SD outcome prior σ_o

(c) Density SD probabillity prior σ_e

(d) Density fixed point probability-distortions ψ

Figure S15: Density functions of BIM parameters for L'Haridon and Vieider (2019)

Figure S16 shows the density functions for the 4 BIM parameters for losses estimated from the data of Bruhin et al. (2010). The parameters resemble each other for the different experiments. The largest differences between experiments emerge once more for the standard deviation of the outcome prior, σ_o^-, and for the fixed point of the probability-distortion function, ψ^-. The first has a relatively broad distribution in the two experiments in Zurich. The second tends towards higher values in the Beijing 05 experiment, indicating higher fixed points of the probability-distortion function, just as observed for gains. Encoding noise is somewhat higher in Zurich 06 and Beijing 05 when compared to Zurich 03.

I plot the BIM parameter distributions for losses in the data of L'Haridon and Vieider (2019) in figure S17 for the usual 4 countries. Panel 17(a) shows systematic differences in encoding noise between countries, on top the differences between individuals within each country. These differences are closely related to those observed for gains. In particular, the US and China have relatively low

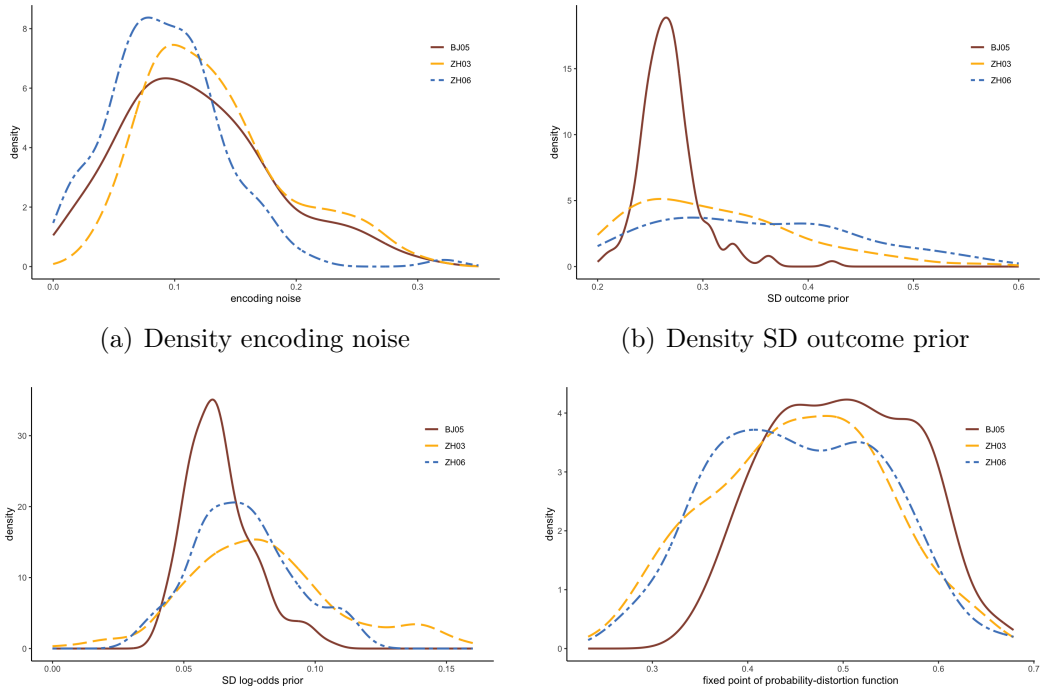

(a) Density encoding noise

(b) Density SD outcome prior

(c) Density SD probabillity prior

(d) Density fixed point probability-distortions

Figure S16: Density functions of BIM parameters for losses in Bruhin et al. (2010)

encoding noise, while Ethiopia and Vietnam have larger levels of encoding noise. Panel 17(d) also shows systematic differences between countries, with Vietnam and Ethiopia having somewhat fixed points of the probability-distortion function on average, thus inverting the pattern see for gains. The between-country differences for the two scale parameters of the prior are somewhat less pronounced, except for the somewhat diffuse distribution of the outcome prior variance in Ethiopia.

Let us now have a look at the distributions of the BIM-generated PT parameters, α, γ, δ, and ω. Figure S18 shows the density functions for gains in the data of Bruhin et al. (2010). The parameter values look reasonable, and fall into ranges we may expect to emerge from a PT model. Panel 18(a) and panel 18(b) show the largest differences between experiments, concerning outcome-discriminability α and likelihood-discriminability γ, respectively. Participants in Zurich in 2003 appear to be particularly discriminating towards outcomes, while Chinese participants in 2005 show particularly low discriminability in the likelihood dimension. Differences across experiments between other parameters tend to be less

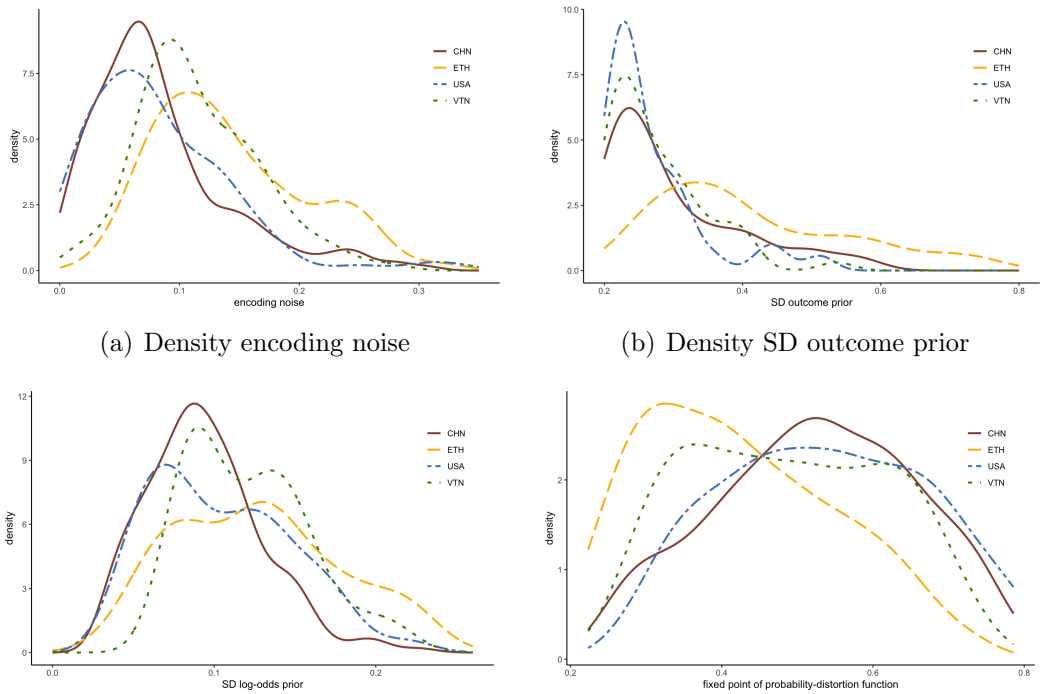

(a) Density encoding noise

(b) Density SD outcome prior

(c) Density SD probabillity prior

(d) Density fixed point probability-distortions

Figure S17: Density functions of BIM parameters for losses in L'Haridon and Vieider (2019)

pronounced, even though optimism δ tends to be larger in the Beijing 05 experiment than in the two experiments in Zurich. Decision noise ω is virtually indistinguishable across the three experiments.

Figure S19 shows the same distributions for the four countries extracted from the L'Haridon and Vieider (2019) data, once again for gains. Panel 19(a) shows the distribution of outcome-discriminability separately for the four countries. Discriminability is highest in the US and China, significantly lower in Vietnam, and lower even in Ethiopia. Very similar patterns are observed for likelihood-discriminability, displayed in panel 19(b). Again discriminability is highest in the US and China, with Vietnam and Ethiopia displaying significantly lower levels. This is hardly surprising, given the important role of encoding noise ν in the makeup of both parameters. Indeed, encoding noise is highly correlated with both α and γ ($\rho = -0.872$ and $\rho = -0.842$, respectively, with $p < 0.001$ in both cases), as one would expect based on the modelling setup and parameter definitions.

Similarly systematic patterns are shown in panel 19(c) for optimism, and in

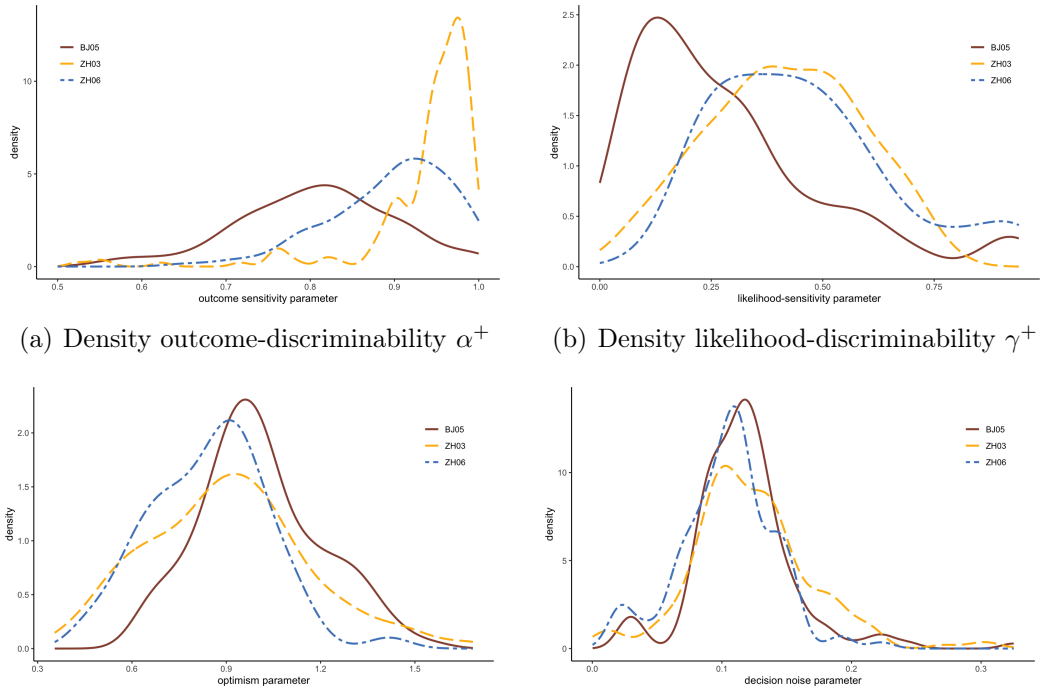

(a) Density outcome-discriminability α^+ (b) Density likelihood-discriminability γ^+

(c) Density optimism δ^+ (d) Density decison noise ω^+

Figure S18: Density functions of BIM-generated PT parameters for Bruhin et al. (2010)

panel 19(d) for decision noise. Optimism is lowest in the US data, closely followed by China. Vietnam and especially Ethiopia show considerably higher levels of optimism. The patterns documented for for these parameters are also reflected in differences in decision noise ω^+. In particular, decision noise is is markedly higher in Vietnam than it is in China or the US, with the US data exhibiting higher noise than the Chinese data. Decision noise is, however, highest in the Ethiopian data. The similarities in the differences across the parameters are indeed unsurprising, given that the parameters α and γ, which in turn are determined by ν, feed back into decision noise ω.

Figure S20 shows the distributions for the four countries extracted from the L'Haridon and Vieider (2019) data for losses. Panel 20(a) shows the distribution of outcome-discriminability separately for the four countries. Discriminability is highest in China, with the US having a broader distribution, i.e. more individual heterogeneity, while being lowest in Vietnam. Very similar patterns are observed for likelihood-discriminability, displayed in panel 20(b). Discriminability is highest

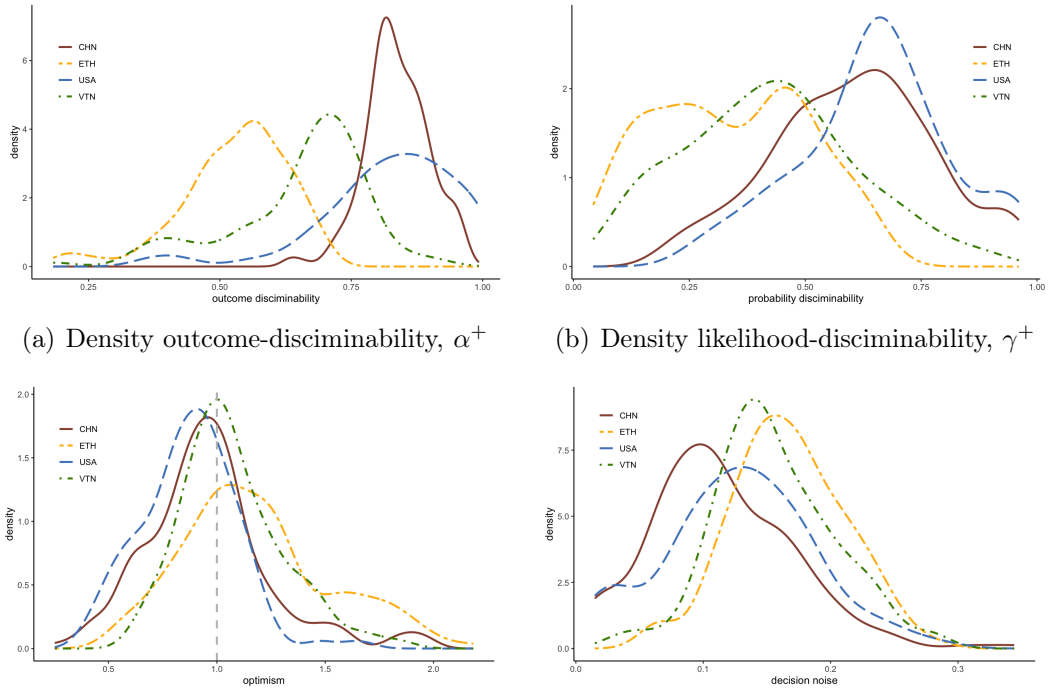

(a) Density outcome-disciminability, α^+ (b) Density likelihood-disciminability, γ^+

(c) Density optimism δ^+ (d) Density decison noise ω^+

Figure S19: Density functions of BIM-generated parameters for L'Haridon and Vieider (2019)

in the US and China, with Vietnam and especially Ethiopia displaying lower levels. The same systematic patterns are shown in panel 20(c) for pessimism, and in panel 20(d) for decision noise. Once again, the between-country differences in optimism are more muted than one might have expected from the differences in the fixed points of the function displayed in panel 17(d), which ones again happens due to the systematic differences in likelihood discriminability entering the definition of the latter. The same happens for decision noise ω^-. Although decision noise is higher in Vietnam and Ethiopia than it is in the US or China, the differences are less pronounced than those observed for encoding noise ν^-. This happens because the decision parameters α^- and γ^-, which are determined by ν^-, feed back into decision noise ω^-, compensating for the impact of ν^-.

S2.5 Testing BIM versus PT: Additional results

Table S2 shows tests of the predictive performance of the BIM model versus PT on the data of L'Haridon and Vieider (2019). Given that there are 30 countries,

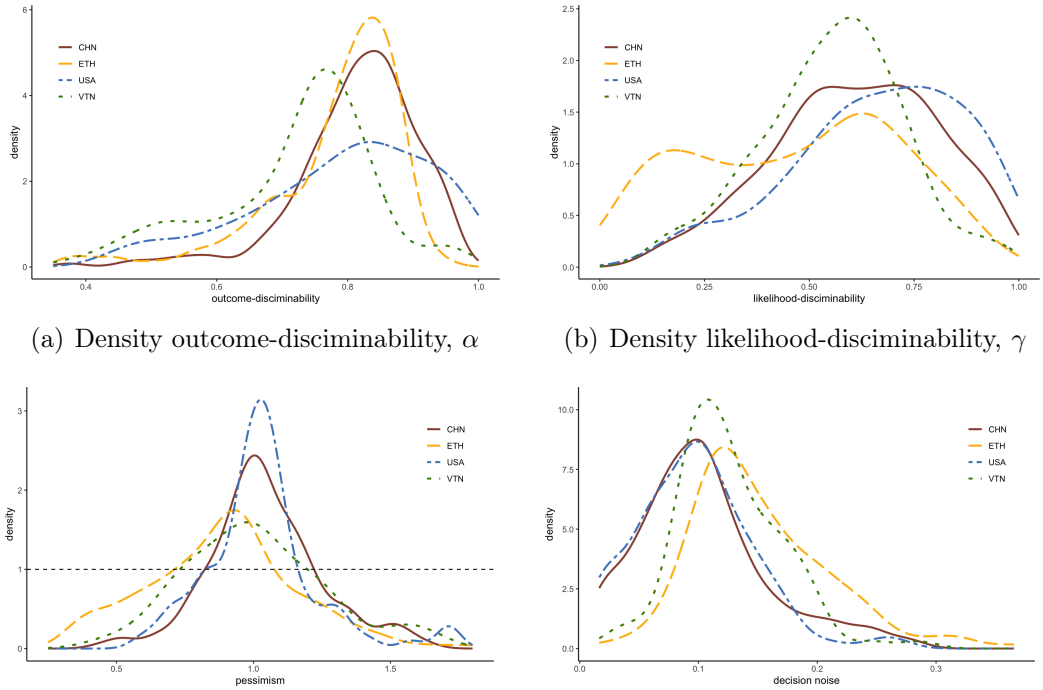

(a) Density outcome-disciminability, α

(b) Density likelihood-disciminability, γ

(c) Density optimism

(d) Density decison noise

Figure S20: Density functions of BIM-generated parameters for losses in L'Haridon and Vieider (2019)

and that I conduct the test separately on gains and losses, this gives me 60 tests, which are independent conditional on the experimental design. The BIM clearly outperforms PT in all 60 cases.

I illustrate the model-fit using some additional individual-level fit plots. I focus on individuals with $\widehat{\gamma} > 1.1$ in the data of L'Haridon and Vieider (2019), which are cases that are at first approximation difficult to account for by the BIM, which imposes $\gamma < 1$. Oftentimes, these are the observations that will have $\widehat{\alpha} \gg 1$ as well, as shown in the correlation graph in the main text. Rather than randomly selecting subjects, below I start by documenting the patterns in order for each individual displaying the described pattern, proceeding in order by subject number for the same 4 countries used as examples in the main text, unless stated otherwise.

Figure S21 shows the fit for the first 4 individuals. The pattern in panel 21(a) shows a pattern very similar to the one seen in the main text. Under PT, this subject displays oversensitivity at $\widehat{\gamma} = 1.45$, oversensitivity to outcomes at

Table S2: Model fit of BIM versus PT in L'Haridon and Vieider (2019)

	ELPD difference			ELPD difference	
country	gains	losses	country	gains	losses
Australia	−2374.5	−2147.3	Kyrgyzstan	−3783.6	−3425.8
(SE diff)	(13.3)	(12.1)	(SE diff)	(16.0)	(14.3)
Belgium	−3549.8	−3198.2	Malaysia	−2493.7	−2257.5
(SE diff)	(15.4)	(15.3)	(SE diff)	(13.5)	(12.1)
Brazil	−3277.0	−2949.3	Nicaragua	−4680.6	−4236.3
(SE diff)	(14.8)	(14.7)	(SE diff)	(17.7)	(16.2)
Cambodia	−3116.3	−2821.8	Nigeria	−7875.2	−7088.4
(SE diff)	(15.1)	(13.5)	(SE diff)	(23.1)	(20.7)
Chile	−3745.0	3361.8	Peru	−3705.2	−3346.2
(SE diff)	(15.8)	(16.7)	(SE diff)	(15.7)	(15.1)
China	−7952.8	−7191.2	Poland	−3468.2	−3128.5
(SE diff)	(23.2)	(21.6)	(SE diff)	(15.4)	(15.1)
Colombia	−4299.8	−3806.5	Russia	−2715.6	−2445.9
(SE diff)	(18.2)	(18.1)	(SE diff)	(15.3)	(14.9)
Costa Rica	−4134.9	−3725.4	Saudi Arabia	−2535.8	−2280.4
(SE diff)	(16.6)	(17.1)	(SE diff)	(18.0)	(13.3)
Czech Rep.	−3857.2	−3479.0	South Africa	−2769.4	−2499.1
(SE diff)	(16.2)	(15.7)	(SE diff)	(13.7)	(13.2)
Ethiopia	−5457.3	−4882.1	Spain	−3121.1	−2817.2
(SE diff)	(19.5)	(19.2)	(SE diff)	(14.4)	(13.9)
France	−3619.3	−3257.9	Thailand	−3078.2	−2764.3
(SE diff)	(16.5)	(16.4)	(SE diff)	(14.6)	(15.7)
Germany	−5067.6	−4574.1	Tunisia	−2872.4	−2603.4
(SE diff)	(18.4)	(18.1)	(SE diff)	(15.5)	(13.1)
Guatemala	−3264.9	−2958.4	UK	−3112.5	−2818.2
(SE diff)	(16.2)	(14.3)	(SE diff)	(15.2)	(13.7)
India	−3468.6	−3138.0	USA	−3782.7	−3385.5
(SE diff)	(15.7)	(14.4)	(SE diff)	(16.0)	(14.8)
Japan	−3274.1	−2956.4	Vietnam	−3387.6	−3064.2
(SE diff)	(14.8)	(14.4)	(SE diff)	(15.9)	(14.5)

The reported EPPD differences indicate the comparative performance of the two models (Gelman et al., 2014a; Vehtari et al., 2017). Negative differences constitute evidence in favour of the BIM over PT. All tests are significant at $p < 0.001$. Tests using WAIC yield similar results, and are thus not reported for parsimony.

$\widehat{\alpha} = 2.39$, and low optimism at $\widehat{\delta} = 0.49$, as well as moderately high levels of noise, at $\widehat{\omega} = 0.14$. The BIM-derived parameters, which provide a clearly better fit at least to the probability dimension, come in at $\gamma = 0.68$, $\alpha = 0.74$, and $\delta = 1.14$. Optimism in probability-distortions is thus substituted for the extreme distortion of outcomes, obtained under PT from relatively few observations.

The other comparisons show similarly clear evidence in favour of the BIM-derived functions over PT. The fit in panel 21(b) appears poor for both functions, but this derived from the relatively high levels of outcome-distortion under both

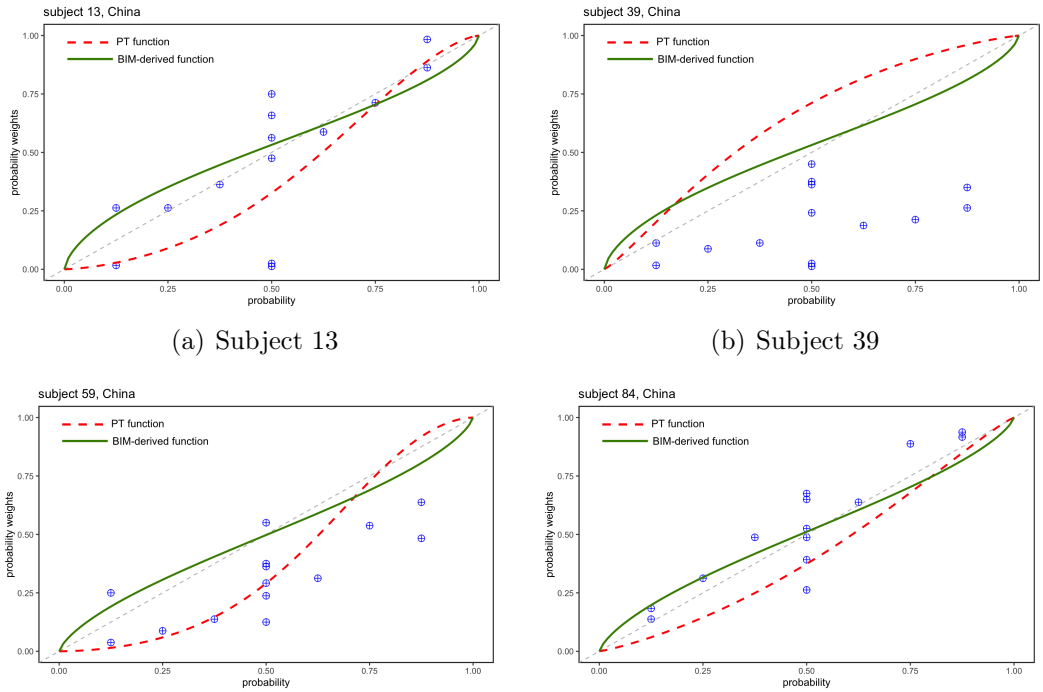

(a) Subject 13 (b) Subject 39

(c) Subject 59 (d) Subject 84

Figure S21: Individual-level data fit of PT vs BIM in L'Haridon and Vieider (2019)

models ($\alpha = 0.64$, with a virtually identical parameter under PT). Beyond this, the BIM-derived function seems to run at least roughly parallel to the data points, whereas the PT-parameter $\widehat{\gamma} = 1.17$ clearly over-estimates probabilistic sensitivity. A similar observation holds for 21(c), where the PT parameters would appear to fit the evidence better for $p \leq 0.5$, but where the strong S-shape also means that the fit for large probabilities is very poor, pointing at a slightly different manifestation of overfitting. By reproducing the qualitative patterns seen in the displayed probability patterns, and absorbing the level difference through outcome-distortions, the BIM-model again outperforms PT. In panel 21(d), the patter is crystal clear, and needs no further discussion. Figure S22 shows similar patterns for the next 4 subjects.

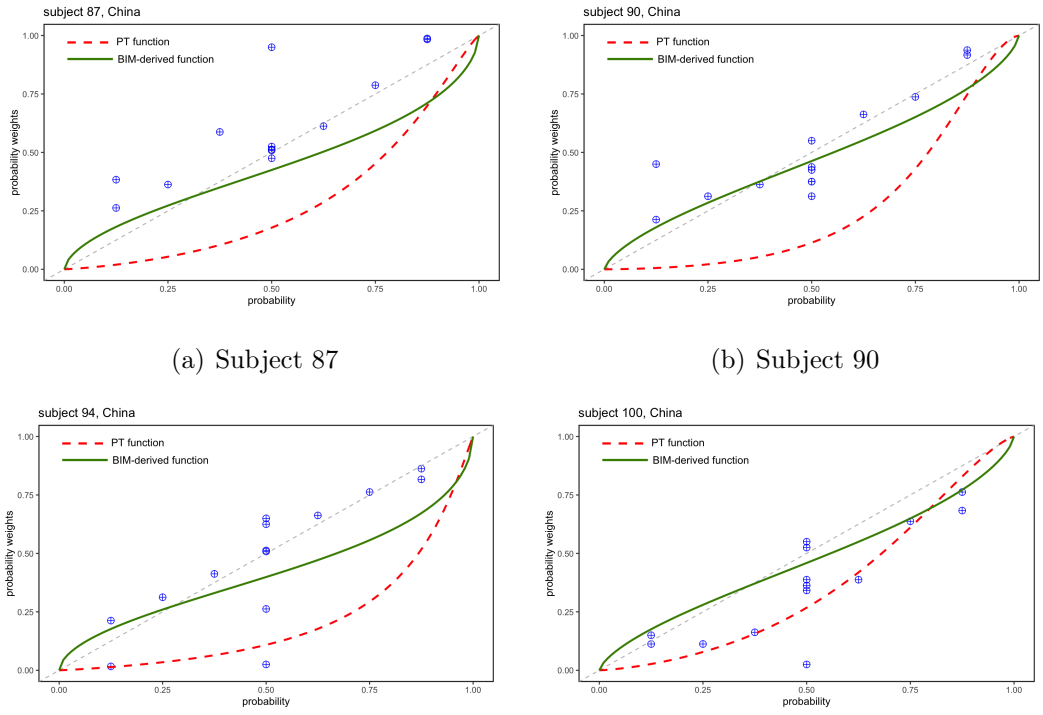

(a) Subject 87

(b) Subject 90

(c) Subject 94

(d) Subject 100

Figure S22: Individual-level data fit of PT vs BIM in L'Haridon and Vieider (2019)

S3 Separability violations under ambiguity

S3.1 Experimental details

Subjects. 48 subjects were recruited at the Melessa Lab at the University of Munich in June 2011. Only subjects who had participated in less than 3 experiments previously were invited. One subject was eliminated because she manifestly did not understand the task, and alternately chose only the sure amount or only the prospect. 38% of the subjects were male and the average age was 25 years. The experiment was run using paper and pencil.

Experimental tasks. I presented subjects with 56 different binary prospects (28 for gains, 26 for losses, and 2 mixed prospects over gains and losses). Subjects had to make a choice between these prospects and different sure amounts of money, bounded between the highest and the lowest amount in the prospect. Gains were

always presented first, and losses were administered from an endowment in a second part, the instructions for which were distributed once the first part was finished. Prospects were always kept in a fixed order. A pilot showed that this made the task less confusing for subjects, while no significant differences were found in certainty equivalents for different orders. Table S3 shows the prospects used in the usual notation $(x, p; y)$, where p indicates the probability of winning or losing x, and y obtains with a complementary probability $1 - p$.

Table S3: Decision tasks under risk and ambiguity

risky gains	uncertain gains	risky losses	uncertain losses
{0.5: 5; 0}	{0.5: 5; 0}	{0.5: –5; 0}	{0.5: –5; 0}
{0.5: 10; 0}	{0.5: 10; 0}	{0.5: –10; 0}	{0.5: –10; 0}
{0.5: 20; 0}	{0.5: 20; 0}	{0.5: –20; 0}	{0.5: –20; 0}
{0.5: 30; 0}	{0.5: 30; 0}	{0.5: –20; –5}	{0.5: –20; –5}
{0.5: 30; 10}	{0.5: 30; 10}	{0.5: –20; –10}	{0.5: –20; –10}
{0.5: 30; 20}	{0.5: 30; 20}	{0.125: – 20; 0}	{0.125: – 20; 0}
{0.125: 20; 0}	{0.125: 20; 0}	{0.125: –20; –10}	{0.125: –20; –10}
{0.125: 20; 10}	{0.125: 20; 10}	{0.25: –20; 0}	{0.25: –20; 0}
{0.25: 20; 0}	{0.25: 20; 0}	{0.385: –20; 0}	{0.385: –20; 0}
{0.385: 20; 0}	{0.385: 20; 0}	{0.625: –20; 0}	{0.625: –20; 0}
{0.625: 20; 0}	{0.625: 20; 0}	{0.75: –20; 0}	{0.75: –20; 0}
{0.75: 20; 0}	{0.75: 20; 0}	{0.875: –20; 0}	{0.875: –20; 0}
{0.875: 20; 0}	{0.875: 20; 0}	{0.875: –20; –10}	{0.875: –20; –10}
{0.875: 20; 10}	{0.875: 20; 10}	**mixed:** {0.5: 20; **–L**}	**mixed:** {0.5: 20; **–L**}

Prospects are displayed in the format $(p : x; y)$.

Notice how the exact same prospects were administered for risk (known probabilities) and uncertainty (unknown or vague probabilities). This will allow me to study ambiguity attitudes, i.e. the difference in behavior between uncertainty and risk. Preferences were elicited using choice lists, with sure amounts changing in equal steps between the extremes of the prospect.

Incentives. At the end of the game, one of the tasks was chosen for real play, and then one of the lines for which a choice had to be made in that task. This provides an incentive to reveal one's true valuation of a prospect, and is the standard way of incentivizing this sort of task. Subjects obtained a show-up fee of €4. The expected payoff for one hour of experiment was above €15.

Risk and uncertainty. Risk was implemented using an urn with 8 consecutively numbered balls. Uncertainty was also implemented using an urn with 8 balls, except that subjects were now told that, while the balls all had a number between

1 and 8, it was possible that some balls may recur repeatedly while others could be absent. The description as well as the visual display of the urns closely followed the design of Abdellaoui et al. (2011). The main differences were that I ran the experiment using paper and pencil instead of with computers; that I used numbers instead of colours in order to allow for black and white printing; and that I ran the experiment in sessions of 15-25 subjects instead of individually.

S3.2 The decision model and estimation code

The decision model closely follows the one for risk used above, but with all parameters doubled for ambiguity. Ambiguity parameters are coded as deviations of the equivalent parameters for risk, mostly to increase computational efficiency and without loss of generality. All estimations are executed directly using the density around the switching point.