Decisions under Uncertainty as Bayesian Inference on Choice Options*

Ferdinand M. Vieider[1]

[1] *RISL$\alpha\beta$, Department of Economics, Ghent University , ferdinand.vieider@ugent.be*

November 8, 2022

Abstract

Standard models of decision-making are deterministic. Inconsistencies in choices are accommodated by separate error models. The combination of decision model and error model, however, is arbitrary. Here, I derive a model of decision-making under uncertainty in which choice options are mentally encoded by noisy signals, which are optimally decoded by combination with pre-existing information. The model predicts a four-fold pattern of risk attitudes as documented in prospect theory. The model is, however, inherently stochastic, so that choices and noise are determined by the same underlying parameters. This results in several novel predictions, which I test in a series of experiments.

Keywords: risk taking; noisy cognition; prospect theory;

JEL-classification: C93; D03; D80; O12

1 Motivation

Take a wager paying a prize x if event e occurs, and $y < x$ under the complementary event $\tilde{e}$. Assume that a decision maker has to choose between this wager and

*I gratefully acknowledge financial support from the Research Foundation—Flanders (FWO) under the project "Causal Determinants of Preferences" (G008021N), and the special research fund (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, Benjamin Enke,Thomas Epper, Ernst Fehr, Cary Frydman, Olivier L'Haridon, Paulo Natenzon, Nick Netzer, Ryan Oprea, Pietro Ortoleva, Jakub Steiner, Tom Verguts, Peter Wakker, Michael Woodford, and Horst Zank for helpful comments and discussions. All errors remain my own.

a sure amount of money, c, where $x > c > y$. Traditionally, decision makers have been assumed to have deterministic utilities over outcomes, as well as being able to form subjective beliefs over the occurrence of event e (Savage, 1954). Choices of this type could then be used to recover the underlying belief and preference parameters of the decision maker. Empirical observations, however, have pointed to choices that are often inconsistent, even when the same choice options are repeated within relatively short time delays (Mosteller and Nogee, 1951; Tversky, 1969; Agranov and Ortoleva, 2017).

The dominant response to such challenges in work aimed at identifying preferences from choice data has been to append an additive error term to the deterministic choice model (Thurstone, 1927; Hey and Orme, 1994; Bruhin, Fehr-Duda and Epper, 2010). The combination of error model and deterministic choice model, however, is arbitrary, and provides many degrees of freedom to the econometric modeller. Buschena and Zilberman (2000) showed that inferences about the best-fitting choice model depend on the chosen error structure; and, vice versa, that conclusions about the most suitable error structure depend on the choice model adopted—a phenomenon that they describe as a 'path dependency problem for model selection' (p. 69). Conclusions reached based on any given combination of decision model and stochastic choice structure could thus be driven by arbitrary assumptions about the error structure (Alós-Ferrer, Fehr and Netzer, 2021).

I propose a model in which decision makers receive imperfect, noisy signals about the choice quantities. They then combine these noisy signals with mental priors describing the likelihood of different values in the environment to make inferences about the true choice options triggering the signals. The model is inherently stochastic, so that the decision model and stochastic choice model emerge organically from one and the same setup. The model predicts a four-fold pattern of risk attitudes much as documented in prospect theory (PT; Kahneman and Tversky, 1979; Tversky and Kahneman, 1992; Wakker, 2010), thus establishing micro-foundations for observed choice behaviour. At the same time, however, the model departs markedly from PT in other respects. In particular, the same parameters determine discriminability between choice options *and* decision noise. The model thus predicts that the 'white noise' assumption underlying typical PT implementations is not warranted. This is problematic inasmuch systematic correlations between arbitrary error structures and deterministic decision parameters may bias the inferences about preferences drawn from such models.[1]

A particularly stark challenge to standard models of decision making under risk arises from so-called *simplicity equivalents*. Oprea (2022) presents an experiment comparing a traditional PT setup measuring probability distortions by means of certainty equivalents to what he calls 'deterministic mirrors' of the same lotteries. Such mirrors consist of choice options that keep the complexity of the lottery representation, while entirely removing the underlying risk.[2] The choice problem is thus complex but deterministic, so that choices for the sure amount must represent a preference for simplicity, rather than reflecting attitudes towards risk. Remarkably, Oprea (2022) documents results for simplicity equivalents that closely track the results obtained under risk using certainty equivalents, including 'likelihood-insensitivity' in a context where likelihoods play no role.

The striking implication of these results is that the patterns we have observed in experiments investigating decision making under risk and uncertainty may not be about risk and uncertainty at all. This constitutes a fundamental challenge to models of decision making under risk and uncertainty, since the latter aim to describe underlying preferences about *risk*. The model I present here, on the other hand, is driven by approximate assessments of numerical quantities. Although I derive the model in the context of choices between uncertain options, one could equally well substitute deterministic mirrors for the uncertain events or likelihoods. The model treats choice stimuli as inherently uncertain, so that inferences need to be drawn about the true underlying choice options. This, in turn, implies that uncertainty about the choice options is a matter of degree. In other words, while models of decision making typically distinguish categorically between risk and ambiguity (Wakker, 2010; Abdellaoui, Baillon, Placido and Wakker, 2011), the

[1]Some recent papers have discussed errors in inferences on preference parameters resulting from the noise structure (Wilcox, 2011; Apesteguia and Ballester, 2018). While related to that line of research, the insights presented here rely on a model making precise predictions on the form of correlations we ought to observe in standard 'PT plus additive noise' setups.

[2]In practice, he uses 100 boxes containing different monetary amounts, with $0 < w < 100$ boxes containing a prize, and $100 - w$ containing nothing. To obtain CEs, he follows the standard procedure of paying one randomly extracted choice. Deterministic mirrors keep the same setup, but pay out the rewards for *all* the 100 boxes, scaled by 100 to keep the expected payoff constant.

Bayesian inference model I present places the two on a continuum of uncertainty about choice stimuli. The driving force underlying noisy inferences is given by the complexity of the choice environment, rather than by risk itself.

I present a number of empirical tests to assess the model predictions. The first type of tests aims at assessing novel predictions about patterns we should observe when estimating PT preference functionals on data that were truly generated by the Bayesian Inference Model (*BIM*) presented here. The first prediction concerns systematic correlations between the decision parameters and noise. Perhaps surprisingly, such correlations have not been systematically examined to date. Notice that correlations between noise and decision parameters violate the 'PT plus white noise' assumption, and may affect the very inferences about preference parameters that are the main interest of such empirical estimations. As showcased by Buschena and Zilberman (2000), this may not only affect preference parameters in quantitative terms, but can hit at the very heart of the debate on the descriptive validity of different decision models. Alós-Ferrer and Garagnani (2022) indeed showed that imposing a symmetric noise structure on data can produce inferences seemingly supporting a given decision model even if the underlying data have been generated based purely on a random choice algorithm.

The second test concerns PT's strict precept of probability-outcome separability. Violations of this precept under risk are well known, and take the form of probability weighting being affected by stake size (Hogarth and Einhorn, 1990; Fehr-Duda, Bruhin, Epper and Schubert, 2010; Bouchouicha and Vieider, 2017). The model I present further predicts a novel violation of probability-outcome separability under ambiguity. Knowledge about probabilities is thereby predicted to affect utility curvature, so that ambiguity will be reflected in utility curvature as well as in probability weighting. The driving force is once again complexity, since ambiguous probabilities are arguably more difficult to assess than known probabilities. To the best of my knowledge, this prediction remains untested. Here, I test this prediction using an original dataset that is rich enough to allow for the separate estimation of all the relevant quantities under both risk and ambiguity.

Finally, I present a direct test pitching PT against the BIM on binary choice data. While striking, the results presented by Oprea (2022) were obtained entirely based on a choice list design. If some subjects simply adopt a heuristic of switching towards the middle of any given list, this could yield inferences that mechanically resemble likelihood-distortions (Vieider, 2018). I thus compare binary choices for risk to binary choices for mirrors of the same tasks. The BIM makes no fundamental distinction between these two settings. PT, on the other hand, models the first setting as being driven by attitudes towards risk, whereas the second setting reduces to a comparison of sure outcomes (possibly augmented by some decision noise, which should however not exceed the noise observed in risky decisions). I find behaviour for mirrors to closely correspond to behaviour for risky lotteries. Predictive tests show that PT is indeed not able to handle such evidence well, whereas the BIM produces comparable parameter estimates in the two settings. Data on decision times further back the conclusion that decisions for risk and mirrors are driven by the same underlying factors.

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, Ishii, Pouget and Rao, 2007; Vilares and Kording, 2011). Such an approach will be optimal whenever decision settings are complex, and thus taxing for the constrained mental resources allocated to a problem. This serves to add computational realism to economic choices, which is often lacking in traditional models (Bossaerts, Yadav and Murawski, 2019).

The noisy neural coding of choice stimuli will 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; Steiner and Stewart, 2016; Netzer, Robson, Steiner and Kocourek, 2021). Such discrimination difficulties for relatively infrequent stimuli, in turn, will result in systematic biases in those regions. Discounting noisy signals indicating choice stimuli that are considered unlikely a priori will thus result in some form of regression to the mean that can explain a number of behavioural regularities (Zhang and Maloney, 2012). The degree of this phenomenon will, in turn, be governed by the relative confidence subjects have in the evidence, a phenomenon that has been documented empirically under the label of *cognitive undertainty* (Enke and Graeber, 2019, 2021).

The model is closely related to some recent contributions in economics. Notably, Khaw, Li and Woodford (2021) used a similar noisy numerical perception setup to explain small stake risk aversion (Rabin, 2000). My model differs from theirs in several respects. While they assume probabilities to be objectively perceived, I allow for the noisy perception of both probabilities *and* outcomes. This makes the model applicable to uncertainty as well as risk, and even to deterministic mirrors of risky choice setups (Oprea, 2022). It further makes it possible to account for insurance takeup and lottery play. Perhaps the most important difference resides in the role played by priors. Whereas the priors drop out from the model presented by Khaw et al. (2021), reducing observed behaviour to pure perception noise, they play a central role in the model I present.

The BIM is thus an inherently *adaptive* model of choice, predicting that decisions can quickly and efficiently adapt to the statistics of the environment to make the best possible use of constrained mental resources. Based on evidence from neuroscience, such a model should have two characteristics. One, priors should dynamically adapt to the choice environment. And two, priors should be hierarchical, so that salient events in one context may affect behaviour in other, related contexts. While a fully-fledged dynamic model of learning is beyond the scope of the present paper, the model allows a glimpse of how these elements could account for puzzling behaviour such as anti-cyclical risk attitudes (Cohn, Engelmann, Fehr and Maréchal, 2015), risk preferences depending on environmental circumstances (Di Falco and Vieider, 2022), and calibration issues when applying parameters measured in the lab to an insurance context (Sydnor, 2010).

The BIM is related to a number of recent models of choice behaviour. Gabaix and Laibson (2017) show that hyperbolic discounting can emerge from the noisy perception of future utilities by otherwise perfectly patient Bayesian decision makers. In a twin paper (Vieider, 2021), I discuss a model applying an identical numerical perception setup to time delays. Frydman and Jin (2022) inject an efficient coding setup into the model of Khaw et al. (2021), allowing them to endogenize coding noise as a function of environmental variability in outcomes.

They present evidence for this mechanism based on an experiment exogenously manipulating the variance of choice stimuli, while keeping their mean constant. They also document a correlation between risk aversion and noisiness in a numerical discrimination task, providing direct evidence for the central mechanism of noisy number perception. Natenzon (2019) uses a related Bayesian noisy perception setup to explain the attraction and compromise effects. Choices of seemingly dominated options are thus justified based on the optimal combination of noisy signals with prior information.

The results I present are further related to a large literature showcasing the instability of revealed preferences. Revealed preferences ought to be unstable to the extent that individual choices are based on noisy signals, and to the extent that priors change across time. Such instability has been documented extensively for inter-temporal correlations of risk attitudes—see Chuang and Schechter (2015) for a review and discussion. The same sort of instability has also been documented specifically for prospect theory parameters (Zeisberger, Vrecko and Langer, 2012).

The insights I present have potentially far-reaching policy implications. For one, the model suggests that revealed choice behaviour is not indicative of actual risk preferences, making it optimal for a welfare-maximizing policy maker to ignore such patterns. It should furthermore become possible to steer behaviour by simplifying the decision environment, providing training on how to efficiently handle complex decision situations, or by providing advice to overcome such complexities. Indeed, if what we think of as preferences is driven to a large extent by the noisy processing of complex decision situations and by changeable priors, then such decisions should be eminently malleable and subject to policy intervention.

This paper proceeds as follows. Section 2 derives the model, with section 2.3 building a 'mapping' between the BIM and PT. Section 3 presents the empirical tests. Section 4 presents a short discussion and concludes the paper.

2 The Bayesian Inference Model

In this section, I present the 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.[3] 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 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.[4]

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:

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

[3] 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.

[4] 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.

where $(x - c)$ indicates the benefit from taking the wager, and $(c - y)$ the cost. This choice rule is optimal, being based on Bayesian updating and expected value maximization. It is further used without loss of generality. If the DM has a concave utility function defined over lifetime wealth, the choice model can be augmented by such a preference without altering the conclusions that follow.

From a computational point of view, it is 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, which is a simpler operation than multiplication. Gold and Shadlen (2001) explicitly discuss the neural realism of such a choice rule. The optimality of the choice rule is unaffected by such an operation, since any monotonic transformation will leave this choice rule unaltered. Importantly, 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 (appendix S1.2 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 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. The choice rule can easily be generalized to the comparison of multiple outcomes. Appendix S1.3 presents a generalization of the model presented below to such a setting.

The model I will present below is based on the premise that mental processing of uncertain wagers functions much like in the optimal choice 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.

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 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 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 choice parameters. An alternative derivation based on the unlogged choice rule in equation 2 is presented in appendix S1.2, and does not alter the conclusions drawn below.

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.[5] 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.[6] 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

[5]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 exogenously given.

[6]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 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.

a further rationale for their noisy mental representations.[7]

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 for both animals and humans 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; Nieder and Miller, 2003; Piazza, Izard, Pinel, Le Bihan and Dehaene, 2004; Harvey, Klein, Petridou and Dumoulin, 2013).

Further support for the normality assumption comes from the central limit theorem, given that mental signals will be made up of the firing rates of thousands of neurons. The underlying rationale for the logarithmic coding can be traced back to the functional architecture of the neural system, where single neurons tend to have receptive fields that are tuned towards the detection of stimuli falling into a pre-determined range. In particular, Howard and Shankar (2018) show that logarithmic spacing of receptive fields is optimal inasmuch as it allows an organism to adapt to the statistics of the environment with maximum flexibility. Note that this conclusion holds independently of the distribution of actual stimuli in the environment, which is a priori unknown to the organism. I assume that r_e and r_o are independently drawn from the following distributions:

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

where ν is the common coding noise of the likelihood ratio and the cost-benefit ratio, conveying information on the uncertainty with which a given stimulus is

[7] In particular, transformation of single numerical quantities such as modelled by Khaw et al. (2021) could be included in the equation to provide a further rationale for the noisy comparison modelled here, so that my approach is complementary to theirs. Tversky (1969) discusses the superiority of such a comparative setup, inasmuch as it contains the transformation of single quantities within as a special case and facilitates otherwise taxing value judgments. I refrain from formally introducing such additional transformations for the sake of analytical tractability.

perceived.[8] Likelihoods of 0 or 1 are assumed to be perceived with certainty.[9] 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 seen 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. This, once again, serves to increase the biological realism of the model.[10] The priors take the following form:

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

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

[8] In what follows, I will assume the two signal to be independent of each other. Notice, however, that incorporting signal correlations is straightforward and may be beneficial in cases where such correlations may provide additional information about the choice options (see Natenzon, 2019, for a striking example).

[9] This assumption is justified by the observation that, in experimental choices as well as in real life, no probability information is typically provided for options that are certain. This, in turn, means that there is no option value to be encoded in the first place.

[10] In particular, obtaining precise posterior distributions in a conjugate setup requires only retaining 5 different parameters in memory in the current setting. Deriving even an approximation of the posterior in a fully general setup would require thousands of simulated points. The costs of assuming a normal distribution when this assumption does not strictly hold true may thus be relatively minor compared to the computational costs involved in abandoning this assumption.

Notice that the logit-normal distributions assumed[11] imply that the prior mean will be 0 for likelihoods and costs and benefits that are equal on average, and will tend towards plus and minus infinity as the differences between these quantities become extreme. Such a symmetric distribution around the mean thus provides a realistic representation of typical choice quantities.

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 $\zeta \triangleq \left(\frac{\widehat{k}}{\widehat{b}}\right)^{(1-\alpha)}$, we obtain the following posterior expectations of the log-likelihood ratio and log-cost benefit ratio, conditional on the noisy mental signals (see appendix S1.1):

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

where $\gamma \triangleq \frac{\sigma_e^2}{\sigma_e^2+\nu^2}$ and $\alpha \triangleq \frac{\sigma_o^2}{\sigma_o^2+\nu^2}$ constitute the weights assigned to the signals of 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 ν and $\{\sigma_e, \sigma_o\}$, will thus determine how much weight will be attributed to the signal versus the prior in the posterior mean. Note that as the coding noise parameter ν converges to 0, the signal will accurately reflect the stimuli, and the weights will converge to 1. For coding noise strictly 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 convex combination of signals and priors.

The actionable choice rule

The posterior expectations of the choice quantities just derived can now 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:

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

[11] A normal distribution of $ln\left(\frac{P[e]}{P[\tilde{e}]}\right)$ constitutes the canonical example of logit-normal distribution (Atchison and Shen, 1980). While the distribution of the log cost-benefit ratio falls farther from the literal interpretation of the logit-normal in terms of likelihood ratios, it follows the same *formalism* as long as we are only interested in the ratio, as is the case here.

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

$$Pr[(x, e; y) \succ c] = \Phi\left(\frac{\gamma \times ln\left(\frac{P[e]}{P[\tilde{e}]}\right) - \alpha \times ln\left(\frac{c-y}{x-c}\right) - ln(\delta)^{-1}}{\nu\sqrt{\gamma^2 + \alpha^2}}\right), \tag{9}$$

where Φ is the standard normal cumulative distribution function. Notice how the same parameters governing the discriminability between the choice options are also driving decision noise. All model parameters are indeed tightly intertwined.

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 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 different mechanisms. 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, and show that the setup results in stochastic micro-foundations for choice patterns as they have been documented in the PT literature. Note that this exercise serves purely to illustrate the mapping between the BIM and PT, and that the empirically implementable choice rule continues to be given by (9).

We can write the point of indifference in the nominator in (9) as follows:

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

This represents the point of indifference between $(x, e; y)$ and c, where the likelihood dimension on the right—subjectively transformed by the mental representation parameters α, γ, and δ—is traded off against the outcome dimension on the left. Setting $P[\tilde{e}] = 1 - P[e]$, as assumed in PT, 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 (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 (12) into (11) and solving for $\pi(P[e])$ we obtain:

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

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 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.

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 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. Such a comparative setting has a number of advantages over a outcomes being transformed individually. For one it is more general, and could accommodate transformations of single outcomes within it (Tversky, 1969). It will also facilitate the evaluation tasks, since outcomes are compared directly instead of being evaluated separately. The comparative setting thus seems well-suited for the approximate valuation entailed by the Bayesian inference model, since it allows to simplify the problem when some quantities are approximately equal.

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 where the function crosses the 45° line. Additionally assuming $\alpha = 1$ and 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, $\pi(P[e]) = P[e]$. The expectation of the mental prior thus 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 likelihood, such as lotteries. The inverse happens for large potential losses, where the same parameter captures pessimism. This 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 the lab are used to explain insurance decisions (Sydnor, 2010). 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 will 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 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*, which is well-documented (Abdellaoui et al., 2011; Dimmock, Kouwenberg and Wakker, 2015; Trautmann and van de Kuilen, 2015; L'Haridon, Vieider, Aycinena, Bandur, Belianin, Cingl, Kothiyal and Martinsson, 2018). Once again, this prediction is driven by the Bayesian aggregation of evidence and prior. With the evidence carrying less certainty under ambiguity, we should indeed expect regression to the mean of the prior to increase in strength. This prediction is further supported by the finding that time pressure, which presumably augments coding noise, increases ambiguity-insensitivity (Baillon, Huang, Selim and Wakker, 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 $\psi_a = 0.5$. Ambiguity aversion—a dislike of unknown probabilities unrelated to their log-likelihood ratio—could further 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. It may also hold more generally because of the impression inherent to the Ellsberg setting that the experimenter is expressly withholding relevant information (Frisch and Baron, 1988; Trautmann, Vieider and Wakker, 2008).

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. Outcome-discriminability α and likelihood-discriminability γ directly enter the definition of

[12] 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, 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 (Fox and Tversky, 1995; Abdellaoui et al., 2011).

decision noise, $\omega \triangleq \nu\sqrt{\alpha^2 + \gamma^2}$, and coding noise enters the definition of the discriminability parameters $\alpha \triangleq \frac{\sigma_o^2}{\sigma_o^2+\nu^2}$ and $\gamma \triangleq \frac{\sigma_e^2}{\sigma_e^2+\nu^2}$. This implies that the noisiness of the decision process will be impacted both by the uncertainty connected to the stimulus encoding, and by the dispersion of the mental priors. Model parameters and decision noise thus move together, resulting in an inherently stochastic model. This 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 correlations could result in biased inferences on the parameters of deterministic models if the noise process is mis-specified relative to the true data-generating process.

A further consequence of the role played by ν in the definitions of α and γ is that no strict separability between the likelihood and outcome dimensions seems warranted, in opposition to what is postulated by PT. We may thus expect coding noise to increase in stakes. Frydman and Jin (2022) present a model of *efficient coding*, in which the noisy coding parameter adapts to the prior variance exactly in this way. 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). The same change 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 indeed 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 coding noise 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).[13]

3 Empirical evidence

In this section, I empirically test the model in three different settings. First, I test the novel predictions about parameter correlations under PT in data originally collected to fit PT functions. Second, I test the novel prediction about probability-outcome separability under ambiguity in a PT setting. And third, I present a binary choice experiment randomly varying whether options are presented as risky lotteries, as traditionally done in the literature, or as 'deterministic mirrors' of those lotteries, and use the data to directly pitch the BIM against PT.

3.1 PT parameters show systematic correlations with noise

I start by examining parameter correlations in PT models. I use the rich individual-level data of Bruhin et al. (2010) collected in 3 different experiments, to estimate distributions of PT parameters (results based on the 30 country data of L'Haridon and Vieider (2019) are similar, see appendix S2.2). The basic econometric framework follows the one in the original paper (see appendix S2.1 for details). Other than in the original paper, I use Bayesian random parameter models to estimate the model parameters (Gelman, Carlin, Stern, Dunson, Vehtari and Rubin, 2014b). This method is geared towards maximizing the *predictive* power of the model for new data, rather than towards optimizing its fit to existing data.

The BIM predicts that the parameters governing sensitivity towards outcomes and likelihood ratios are governed by coding noise, jointly with the estimated prior variance. This results in a prediction that, when estimating a PT model, the noise term ought to be correlated with the decision parameters. In a typical PT setup plus noise, on the other hand, the additive noise term is assumed to take the form of 'white noise', and thus to be independent of the decision model itself. A violation of this assumption would be problematic, since any inferences on preference

[13] It is, of course, also in contradiction to models that capture ambiguity attitudes purely through utility curvature.

parameters one draws from the model may be sensitive to the specific assumptions about noise. I will test these predictions using Spearman rank correlations on the estimated parameters. Any p-values reported are always two-sided.

To ensure comparability with the model derived above, 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. These same functional forms were also used in the original paper.[14] 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, to which I refer as *discriminability* parameters. While I use the additive error form used in the original study, alternative error specifications, such as e.g. Wilcox's (2011) contextual error, do not affect the conclusions I draw.

Figure 1 shows correlations between decision noise and likelihood-sensitivity $\widehat{\gamma}$ for both gains and for losses. Note that the econometric model allows for heteroscedasticity in errors across outcome domains, and the errors for gain and losses are thus not the same. Panel 1(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 also present in each of the 3 individual experiments. The results for losses, shown in panel 1(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 rather than on quick and approximate judgments.

Figure 2 shows the correlations between the noise parameter and the outcome sensitivity parameter, $\widehat{\alpha}$, with panel 2(a) showing the results for gains. There is a negative correlation 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 a positive correlation in

[14] Prelec (1998) presents an popular alternative 2-parameter weighting function. 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.

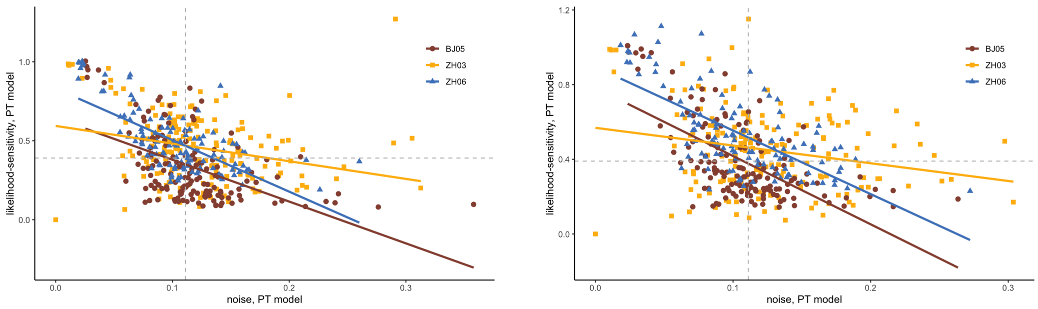

(a) Scatter plot of $\widehat{\omega}^{+}$ and $\widehat{\gamma}^{+}$ (b) Scatter plot of $\widehat{\omega}^{-}$ and $\widehat{\gamma}^{-}$

Figure 1: Scatter plot of PT parameters, likelihood-sensitivity
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. Solid lines indicate regression lines. The dashed lines indicate the median parameter values. Some outliers may be cut for better visual display.

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(b). 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.

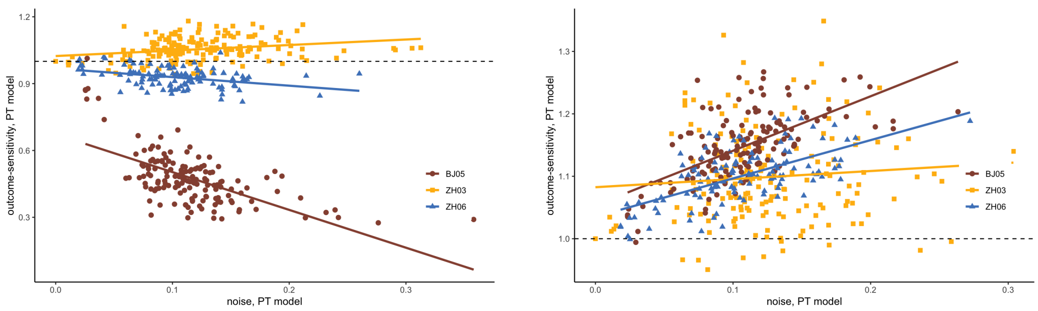

(a) Scatter plot of $\widehat{\omega}^{+}$ and $\widehat{\alpha}^{+}$ (b) Scatter plot of $\widehat{\omega}^{-}$ and $\widehat{\alpha}^{-}$

Figure 2: Scatter plot of PT parameters, outcome-sensitivity
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 reference parameter values of 1 (linear utility). Some outliers may be cut for better visual display.

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 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 noise and outcome-sensitivity. Larger values of $\widehat{\gamma}$ coincide with values of δ tending towards 1 (see appendix S2.2). These effects are not foreseen by PT, but line up exactly with the predictions emerging from the BIM.

3.2 Probability-outcome separability under ambiguity

The BIM predicts ambiguity-insensitivity based on the heightened level of difficulty in assessing unknown probabilities. Such patterns are 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. In addition, however, the model presented above predicts a novel violation of PT's probability-outcome separability, whereby utility curvature is also affected by ambiguity. This constitutes a violation of PT, 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). In this section, I present a test of that prediction in a PT setup.

Published data on ambiguity attitudes tend to be 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 $\widehat{\gamma}$. To test this prediction, I thus use an original dataset that contains a richer choice setup. The data contain observations for 47 subjects indicating certainty equivalents for both risky and ambiguous lotteries in a within-subject design. 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 under ambiguity as well as under risk (Fehr-Duda and Epper, 2012). The econometric approach closely follows the setup used above for the data of Bruhin et al. (2010). Further details can be found in appendix S3.

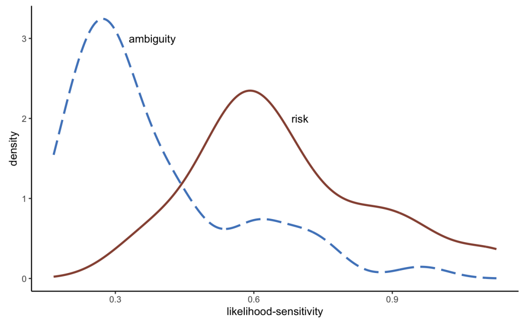

(a) Densities of $\widehat{\gamma}_r$ and $\widehat{\gamma}_a$

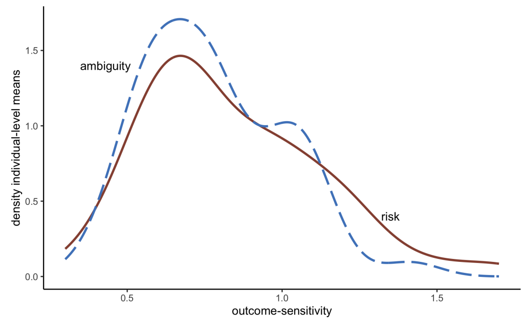

(b) Densities of $\widehat{\alpha}_r$ and $\widehat{\alpha}_a$

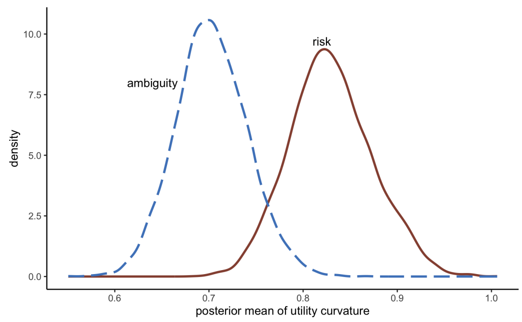

(c) Posterior densities of the means

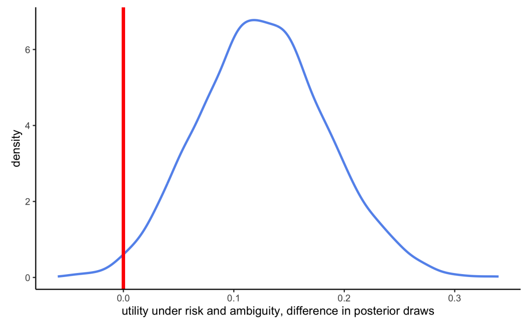

(d) Density difference of means

Figure 3: Likelihood- and outcome-sensitivity under risk and ambiguity
Densities of individual-level parameters (top) and of the mean parameters (bottom) of likelihood- and outcome-sensitivity. Panel 3(d) plots the difference in posterior draws between the two means, which is the quantity used in Bayesian tests of mean differences (Kruschke, 2014).

Figure 3 shows density plots for $\widehat{\gamma}$ and $\widehat{\alpha}$, for risk and ambiguity. The difference in individual-level estimates between risk and ambiguity for $\widehat{\gamma}$, shown in panel 3(a), is sizeable. This sort of lowering of likelihood-sensitivity under ambiguity relative to risk, termed ambiguity-insensitivity, is indeed a standard finding in the PT literature. The difference in individual-level parameters between outcome sensitivity for risk and ambiguity is displayed in panel 3(b). While it is less pronounced, the distribution under ambiguity is shifted to the left. This is confirmed by a Mann-Whitney test on the paired individual-level parameters, indicating that outcome-sensitivity is indeed reduced under ambiguity ($p = 0.002$, two-sided).

One can further test for the existence of differences using directly the posterior draws of the mean parameters for risk and ambiguity. Panel 3(c) shows the density of the posterior draws of the means of the two utility parameters for risk and ambiguity. The utility parameter for ambiguity is clearly smaller than the one for risk. To assess the statistical significance of this difference, however, we need to look directly at the *difference* in posterior draws, given that the draws are not independent (Gelman et al., 2014b; Kruschke, 2014). That difference is shown in panel 3(d). The overlap with 0 is minimal, and 0.99 % of the posterior probability mass falls above 0, indicating a significant difference in the means. This results in a violation of the probability-outcome separability precept of PT. It does, however, conform precisely to the prediction emerging from the Bayesian Inference Model.

3.3 Lottery choice versus deterministic mirrors

The tests shown so far aimed at testing predictions the BIM makes about what we should observe if we estimated the data in a PT setup, if the BIM is the true data-generating model. Here, I present a direct test of PT against the BIM. To this end, I deploy a variation of Oprea's (2022) deterministic mirrors. One potential shortcoming of the approach taken by Oprea (2022) is that the choice list format may be driving some of the results. That is, if some subjects pick switching points that mechanically fall towards the middle of the list, this could produce the same sort of patterns for risk and deterministic mirrors. I thus adapt the mirror setup to a rich binary choice setting to test the robustness of the effect. This furthermore allows me to field a number of additional tests, including structural estimations of the BIM for both risk and mirrors, out-of-sample predictive tests, and robustness tests using chronometric data.

The experiment was conducted in an introductory class of Behavioural Economics at Ghent University using a between-subject design. Students had heard about expected value and expected utility theory, but had not yet been introduced to behavioural models of decision making. Students were told to bring a laptop to class. 10 subjects were randomly selected to play one of their decisions for real money. After a short introduction mentioning the length of the experiment and the randomized incentive structure, the lecturer shared a link for the online experiment with the students. 133 students submitted complete responses in the allocated time.

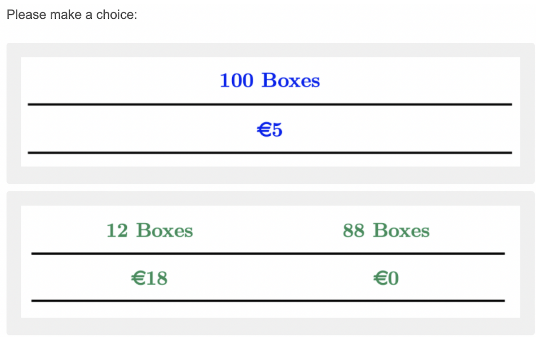

Figure 4: Screenshot of choice situation

Figure 4 shows a screenshot of the experiment. Subjects had to indicate their choice between two options. Each option had a number of boxes associated with monetary payoffs. Option A had some boxes associated with a higher payoff, and some with a lower payoff. The number of high outcome boxes could take the value 12, 32, 42, 48, 52, 58, 68, 88. The winning outcome ranged from €18 to €32, and the tasks included non-zero lower outcomes as well (see table 1 for a list). Option B always contained 100 boxes with the same outcome, which was made to vary in steps of €1 between the high and the low outcome in option A. The binary choice options were, however, completely randomized, as was the position of option A and option B on the screen. These stimuli were chosen to reflect typical choice stimuli used to investigate decision making under risk. Subjects faced 200 such binary choices in total, of which 20 were repeated choices.

The instructions provided the following information. First, all subjects saw an explanation of the general structure of the choice task. They were told that they had to decide between two options, and how different boxes could have different payoffs attached to them. Subsequently, subjects were randomly allocated to one of two treatments—a *random box* treatment, and an *average box* treatment. In the random box treatment, they learned that one random box from the chosen option would be opened to determine their payoff, following standard procedures

Table 1: Choice tasks and choice proportions by treatment

	characteristics			risk		mirrors		ranksum test
task	high	low	prob	risky	SD	complex	SD	p-value
1	18	0	0.12	0.14	0.35	0.14	0.34	0.672
2	18	0	0.32	0.22	0.42	0.24	0.43	0.614
3	18	0	0.68	0.52	0.50	0.55	0.50	0.142
4	18	0	0.88	0.70	0.46	0.72	0.45	0.301
5	27	5	0.48	0.39	0.49	0.40	0.49	0.704
6	27	5	0.52	0.49	0.50	0.53	0.50	0.287
7	32	5	0.42	0.39	0.49	0.38	0.49	0.718
8	32	0	0.58	0.42	0.49	0.46	0.50	0.334
9	43	18	0.48	0.42	0.49	0.45	0.50	0.410

Summary table of task characteristics and choice proportions. The column 'risky' indicates the choice proportion of the risky option A in the random box treatment. The column 'complex' indicates the choice proportion of the complex option A in the average box treatment. The column 'ranksum test' indicates p-values of Wilcoxon rank sum tests executed on choice proportions of Option A. The tests are based on 72 subjects assigned to the average box treatment, versus 61 subjects assigned to the random box treatment.

in the risky choice literature. In the average box treatment, on the other hand, they learned that they would be paid the amount in the average box. That is, the payoffs would be summed over the 100 boxes, and divided by 100. Examples served to clearly illustrate the choice situation. Subjects furthermore had to go through a number of comprehension questions that were meant to drive home the payoff mechanism (full instructions in appendix S4).

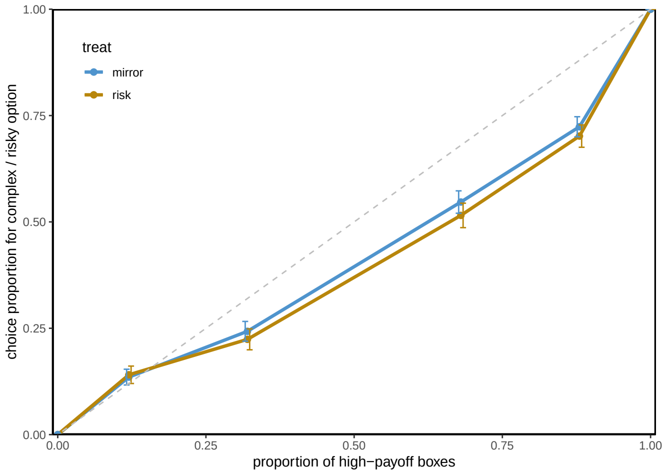

Figure 5: Choice proportions for risky or complex option
The graph plots the number of high-payoff boxes on the x-axis against choice proportions of the risky option (random box treatment) or the complex option (average box treatment). Given the equal steps in sure options, the graph can be interpreted as a (stochastic) variant of a dual-EU probability-distortion function.

Figure 5 shows the choice proportions for the risky option (random box treatment) or the complex option (average box treatment) for tasks providing different numbers of boxes associated with €18, or else 0. In the lottery treatment, the data reveal an inverse-S shape as typically found in the PT literature. Strikingly, the same pattern occurs in the mirror treatment, where no risk is present and choices are between deterministic options. Indeed, there is no difference between any of the choice proportions, and the same is true for the other choice tasks not shown in the graph (see table 1 for a complete list). Risk and choice complexity in a riskless setup thus produce virtually identical results.

Next, I estimate equation (9) on the complete data to examine the individual-level parameters. Figure 6, panel 6(a), shows a density plot of the individual-level likelihood-discriminability parameters (lottery treatment) and the individual-level 'complexity-discriminability' parameters (mirror treatment). Both distributions present a peak around 0.5, and look very similar. The lottery-discriminability has some probability mass to the left of the complexity-discriminability distribution, possibly indicating an effect of risk *on top* of the complexity effect, as also discussed by Oprea (2022). That effect, however, appears to be relatively small.

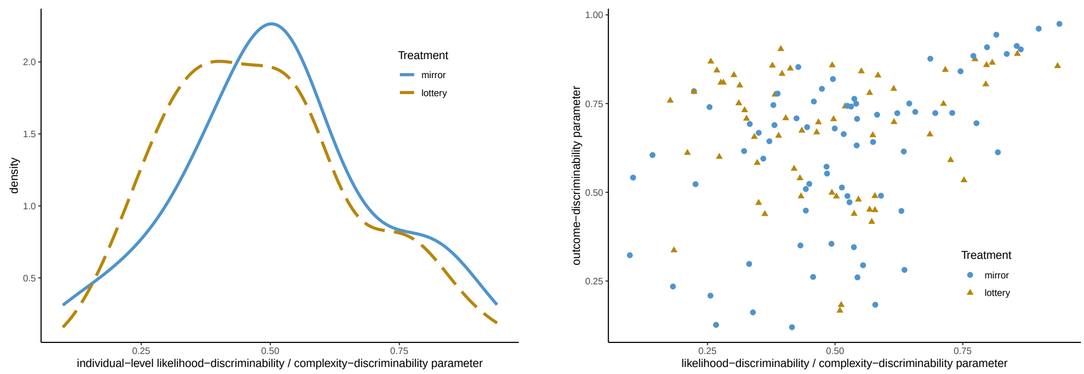

(a) Density of γ for lotteries and mirrors

(b) Scatter plot of γ and α

Figure 6: Individual-level parameters for lotteries and mirrors

Panel 6(b) shows a scatter plot of the parameters capturing sensitivity towards the number of boxes (i.e., likelihoods for risk and complexity for mirrors), and sensitivity towards outcomes. There is no indication that behaviour for mirrors might converge towards behaviour resembling value maximization (i.e., a deterministic choice for the option with the higher value). Remarkably, the results not only show equal distortions driven by the number of boxes across the two treatments, but also equal distortions of *outcomes*, providing direct evidence that outcome-discriminability is also driven by the complexity of the situation, rather than by any inherent attitude towards risk. This results in a clear violation of PT, and is indeed hard to reconcile with any deterministic model of risky choice.

The strength of this test relies on the insight that PT cannot account for these similarities across lotteries and deterministic mirrors. The BIM, however, can handle them easily since it models the perception of numerical quantities, rather than attitudes towards risk *per se*. This can be shown by directly comparing the predictive performance of the two models on the data. The BIM will then deploy one and the same modelling setup for both treatments, whereas PT will directly compare the utilities of the two deterministic outcomes in the mirror treatment, plus an additive error term just like the one used for risk. I test model performance based on their out-of-sample predictive ability (Gelman, Hwang and Vehtari, 2014a; Vehtari, Gelman and Gabry, 2017).[15]

Given the structure of the data, predictive tests can be executed in different ways. One way is to conduct leave-one-out sampling leaving out one choice at the time and predicting it based on the aggregate-level parameters estimated on the remaining data. Such a test yields a clear verdict in favour of the Bayesian Inference model (ELPD difference of -7653.5. with a standard error of 239.6). Arguably, however, predictive tests should take into account the hierarchical nature of the data, which comprise 200 different choices per individual. Such a true out-of-sample test indeed seems essential, given that imposing a given structure could yield spurious results in fitted data (Alós-Ferrer and Garagnani, 2022).[16] I thus conduct a stratified *leave-k-out* test, whereby a subset of $k\%$ of the choices of each individual are removed and predicted based on the remaining $(100-k)\%$ of the observations. This procedure is repeated for each subset of $k\%$ of the data. Notice that this is a harder test than grouping the data and predicting *across*

[15]The results obtained by means of out-of-sample cross-validation 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 Akaike information criterion (AIC). Using the latter instead does not affect the conclusions drawn in any way.

[16]In particular, Alós-Ferrer and Garagnani (2022) show that simply fitting a utility difference plus additive noise to simulated random choices can spuriously show a sigmoidal relationship between utility differences and strength of preference, which results purely from the stochastic choice model iself and has no grounding in the actual data.

subjects in the present setting, since it means that the comparison takes place within-subject and hence within-treatment. Setting $k = 10$, the results once again indicate that the BIM clearly outperforms PT (ELPD diff. of -753.8, $se = 51.9$).

Finally, let us take a look at decision times across the two treatments. Decision times are related to decision difficulty, and have been shown to be predictive of utility differences (Alós-Ferrer et al., 2021; Alós-Ferrer and Garagnani, 2022). Based on standard models of decision making under risk, we would expect decision times for lotteries and mirrors to look rather different. In particular, decision times in risky choice have been shown to be linked closely to differences in the expected utilities of choice options, whereas they tend to be only weakly linked to differences in expected values (Alós-Ferrer and Garagnani, 2022). For mirrors, on the other hand, we should expect decision time to peak around the point of equality in (expected) values, since the task boils down to picking the option with the higher deterministic value. According to the BIM, however, we should not see major differences between the treatments. In both cases, the BIM thus predicts decision time to follow an inverted-U shape, with its peak at the point of no discriminability, where the two choice options become difficult to tease apart.

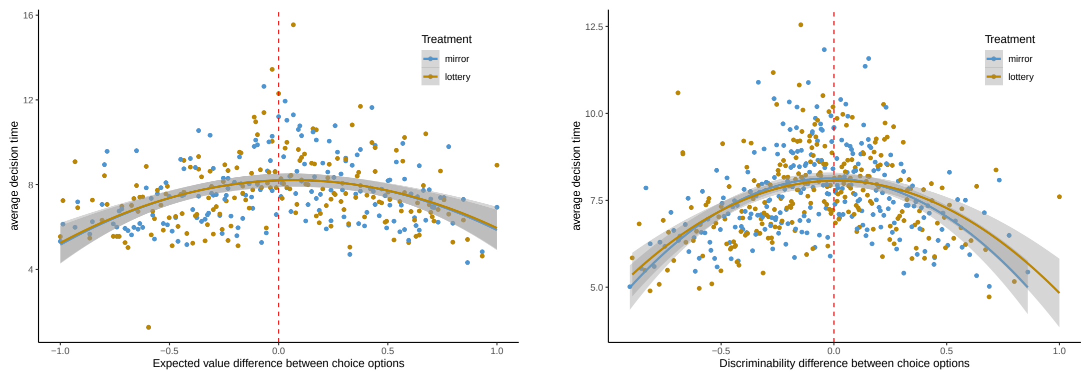

(a) Decision time and EV difference (b) Decision time and BIM discriminability

Figure 7: Average decision time as a function of EV difference and discriminability
Plot of expected value difference between the choice options (panel 7(a)) and discriminability between the choice options (panel 7(b)) against average decision time. Both differences in EV and differences in discriminability have been normalied to fall between -1 and 1. The discriminability plot has been obtained by means of a binning procedure applied to disciminability. The solid lines indicate a least squares fit of a second-degree polynomial.

Figure 7 plots average decision times against the difference in expected values (panel 7(a)) and against the discriminability given by the nominator of (9) (panel

7(b)). Decision times increase towards the point of equality in expected values, continue to increase slightly thereafter, and subsequently decline. Importantly, however, there is no difference whatsoever in this trend for lotteries and mirrors, as deterministic models of decision making under risk would suggest. The inverse-U shape as a function of discriminability in the Bayesian inference model shown in panel 7(b) is much more pronounced than the curve for the EV difference. The peak in decision times coincides with the point of zero discriminability, where both options appear equally good, and shows no difference between lotteries and mirrors. This further supports the conclusion that decisions in lotteries and mirrors are driven by the same processes, in contradictions to standard models of risky choice, but in line with the predictions of the Bayesian Inference Model.

4 Conclusion

Traditional models of decision-making under risk and uncertainty represent behaviour by means of deterministic models applied to objective stimuli. These deterministic models are augmented by an independently chosen stochastic model to allow for choice parameters to be recovered from noisy data. The combination between decision model and stochastic model of choice, however, is arbitrary. Such arbitrary combinations of decision models and stochastic choice architectures may bias inferences on the model parameters and even about the correct underlying model if the stochastic assumptions are not warranted.

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 providing micro-foundations for PT-like choice patterns, the model simultaneously predicts PT violations, in particular pertaining to the strict separation between the probabilistic and outcome dimensions under PT. Other than PT, the model is not specifically geared at describing attitudes towards risk and uncertainty, but rather captures the noisy perception of numerical quantities. The model is thus applicable not only to lottery choice, but also to the deterministic mirrors of those lotteries introduced by Oprea (2022).

The model proposed in this paper shares a common intuition with other models of adaptive behavior. Most notable is the connection to the evolutionary models of Robson (2001) and Netzer (2009), who propose 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.

There are also commonalities with other models of noisy perception. Steiner and Stewart (2016) present a model in which probabilities of binary wagers are subject to noisy perception, similar to the intuition developed here. Other than in the model presented here, however, outcomes are processed without uncertainty, and the model makes no predictions about decision noise. The sure option is the quantity driving probability distortions, rather than a learned prior summarizing the most likely stimuli in the environment. 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 probabilities. 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. Possibly the most far-reaching implications, however, concern the role of the prior, which amongst the recent contributions is unique to the model here presented. Understanding both short term determinants in the lab, and long-term determinants of the prior should thus be a priority topic for future research.

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

Agranov, Marina, and Pietro Ortoleva (2017) 'Stochastic choice and preferences for randomization.' *Journal of Political Economy* 125(1), 40–68

Alós-Ferrer, Carlos, and Michele Garagnani (2022) 'Strength of preference and decisions under risk.' *Journal of Risk and Uncertainty* pp. 1–21

Alós-Ferrer, Carlos, Ernst Fehr, and Nick Netzer (2021) 'Time will tell: Recovering preferences when choices are noisy.' *Journal of Political Economy* 129(6), 1828–1877

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

Atchison, J, and Sheng M Shen (1980) 'Logistic-normal distributions: Some properties and uses.' *Biometrika* 67(2), 261–272

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

Bossaerts, Peter, Nitin Yadav, and Carsten Murawski (2019) 'Uncertainty and computational complexity.' *Philosophical Transactions of the Royal Society B* 374(1766), 20180138

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

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

Buschena, David, and David Zilberman (2000) 'Generalized expected utility, heteroscedastic error, and path dependence in risky choice.' *Journal of Risk and Uncertainty* 20(1), 67–88

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*

Chuang, Yating, and Laura Schechter (2015) 'Stability of experimental and survey measures of risk, time, and social preferences: A review and some new results.' *Journal of Development Economics* 117, 151–170

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

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

_ (2021) 'Cognitive uncertainty in intertemporal choice.' Technical Report, Mimeo

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

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

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

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

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)

Harvey, Ben M, Barrie P Klein, Natalia Petridou, and Serge O Dumoulin (2013) 'Topographic representation of numerosity in the human parietal cortex.' *Science* 341(6150), 1123–1126

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

Howard, Marc W, and Karthik H Shankar (2018) 'Neural scaling laws for an uncertain world.' *Psychological review* 125(1), 47

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

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

Kruschke, John (2014) *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan* (Academic Press)

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

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

Mosteller, Frederick, and Philip Nogee (1951) 'An experimental measurement of utility.' *Journal of Political Economy* 59(5), 371–404

Natenzon, Paulo (2019) 'Random choice and learning.' *Journal of Political Economy* 127(1), 419–457

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

Nieder, Andreas, and Earl K Miller (2003) 'Coding of cognitive magnitude: Compressed scaling of numerical information in the primate prefrontal cortex.' *Neuron* 37(1), 149–157

Oprea, Ryan (2022) 'Simplicity equivalents.' *Working Paper*

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

Piazza, Manuela, Véronique Izard, Philippe Pinel, Denis Le Bihan, and Stanislas Dehaene (2004) 'Tuning curves for approximate numerosity in the human intraparietal sulcus.' *Neuron* 44(3), 547–555

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

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)

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

Steiner, Jakub, and Colin Stewart (2016) 'Perceiving prospects properly.' *American Economic Review* 106(7), 1601–31

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

Thurstone, Louis L (1927) 'Psychophysical analysis.' *The American journal of psychology* 38(3), 368–389

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

Trautmann, Stefan T., Ferdinand M. Vieider, and Peter P. Wakker (2008) 'Causes of ambiguity aversion: Known versus unknown preferences.' *Journal of Risk and Uncertainty* 36(3), 225–243

Tversky, Amos (1969) 'Intransitivity of preferences.' *Psychological review* 76(1), 31

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 (2018) 'Violence and risk preference: Experimental evidence from afghanistan: Comment.' *American Economic Review* 108(8), 2366–82

Vieider, Ferdinand M. (2021) 'Noisy coding of time delays and reward discounting.' Working Paper21/1036 FEB, Ghent University, Belgium

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

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 Bayesian Inference 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:

$$ln\left(\frac{P[e]}{P[\tilde{e}]}\right)\big|r_e \sim \mathcal{N}\left(\frac{\sigma_e^2}{\sigma_e^2+\nu^2}\times r_e + \frac{\nu^2}{\sigma_e^2+\nu^2}\times ln\left(\frac{\widehat{P}[e]}{\widehat{P}[\tilde{e}]}\right), \frac{\nu^2\sigma_e^2}{\nu^2+\sigma_e^2}\right) \qquad (15)$$

$$ln\left(\frac{c-y}{x-c}\right)\big|r_o \sim \mathcal{N}\left(\frac{\sigma_o^2}{\sigma_o^2+\nu^2}\times r_o + \frac{\nu^2}{\sigma_o^2+\nu^2}\times ln\left(\frac{\widehat{k}}{\widehat{b}}\right), \frac{\nu^2\sigma_o^2}{\nu^2+\sigma_o^2}\right). \qquad (16)$$

where we define $\gamma \triangleq \frac{\sigma_e^2}{\sigma_e^2+\nu^2}$ and $\alpha \triangleq \frac{\sigma_o^2}{\sigma_o^2+\nu^2}$. It follows that $\frac{\nu^2}{\sigma_e^2+\nu^2} = 1-\gamma$ and $\frac{\nu^2}{\sigma_o^2+\nu^2} = 1-\alpha$. 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 $\zeta \triangleq \left(\frac{\widehat{k}}{\widehat{b}}\right)^{1-\alpha}$ 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:

$$\gamma \times r_e - \alpha \times r_o \sim \mathcal{N}\left(\gamma \times ln\left(\frac{P[e]}{P[\tilde{e}]}\right) - \alpha \times ln\left(\frac{c-y}{x-c}\right), \omega^2\right), \qquad (17)$$

where $\omega \triangleq \sqrt{(\gamma^2+\alpha^2)\times\nu^2}$ represents the standard deviation of the weighted difference of noisy mental signals. This yields the following z-score:

$$z = \frac{\gamma \times r_e - \alpha \times r_o - \left[\gamma \times ln\left(\frac{P[e]}{P[\tilde{e}]}\right) - \alpha \times ln\left(\frac{c-y}{x-c}\right)\right]}{\omega}, \quad (18)$$

which follows a standard normal distribution. Subtracting this z-score from 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:

$$E\left[\frac{P[e]}{P[\tilde{e}]}|r_e\right] > E\left[\frac{c-y}{x-c}|r_o\right]. \quad (19)$$

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:

$$E\left[\left(\frac{P[e]}{P[\tilde{e}]}\right)|r_e\right] = e^{\gamma \times r_e + ln(\xi) + \frac{1}{2}\hat{\sigma}_e^2} \;,\; E\left[\left(\frac{c-y}{x-c}\right)|r_o\right] = e^{\alpha \times r_o + ln(\zeta) + \frac{1}{2}\hat{\sigma}_o^2}, \quad (20)$$

where $\hat{\sigma}_e^2 \triangleq \frac{\nu^2\sigma_e^2}{\nu^2+\sigma_e^2}$ and $\hat{\sigma}_o^2 \triangleq \frac{\nu^2\sigma_o^2}{\nu^2+\sigma_o^2}$ are the posterior variances. We can now substitute these quantities into the mental choice rule above, to obtain:

$$e^{\gamma \times r_e + ln(\xi) + \frac{1}{2}\hat{\sigma}_e^2} > e^{\alpha \times r_o + ln(\zeta) + \frac{1}{2}\hat{\sigma}_o^2} \quad (21)$$

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

$$\gamma \times r_e - \alpha \times r_o > ln(\widehat{\delta})^{-1}, \tag{22}$$

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 mean 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.

S1.3 Generalization to multiple outcomes

The model presented in the main text can easily be extended to the comparison of multi-outcome wagers, where it results in weighted, state-wise comparisons. Extensions of the model in the main text 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 not discussed further.

The choice rule in equation 3 can easily be generalized to multiple outcomes. To preserve mathematical tractability, let us assume that a wager $(x_1, e_1; x_2, e_2; ....; x_n, e_n)$, offering x_1 if event e_1 occurs, x_2 if event e_2 occurs, etc., is compared to a sure outcome, c. For convenience of exposition, I will assume that outcomes are ordered from highest to lowest, although this is not essential to the model itself. Comparisons of two non-degenerate wagers will result in state-wise comparisons, but is otherwise identical. The optimal choice rule according to which the risky wager will be chosen will then take the following form:

$$\frac{P[e_1]}{P[e_n]}\frac{(x_1 - c)}{(c - x_n)} + \frac{P[e_2]}{P[e_n]}\frac{(x_2 - c)}{(c - x_n)} + ... + \frac{P[e_{n-1}]}{P[e_n]}\frac{(x_{n-1} - c)}{(c - x_n)} > 1, \tag{23}$$

which tallies up the relative costs and benefits of the wager. The equation takes the form of an optimal choice rule used in signal detection theory (Green, Swets et al., 1966). Notice how the choice rule indeed incorporates the same two principles as the binary rule in the main text—optimal belief updating, and expected value maximization. While at first sight the rule may seem to compare everything to the worst-case scenario, the choice rule actually enshrines within all pairwise comparisons. To see this, let V_{ij} be the relative valuation to two arbitrary states of nature. We can then derive this valuation from the choice rule above as follows:

$$V_{ij} = \frac{V_{in}}{V_{jn}} = \frac{\frac{P[e_i]}{P[e_n]}\frac{(x_i-c)}{(c-x_n)}}{\frac{P[e_j]}{P[e_n]}\frac{(x_j-c)}{(c-x_n)}} = \frac{P[e_i]}{P[e_j]}\frac{(x_i-c)}{(x_j-c)}. \tag{24}$$

I postulate that each term in the comparison above is evaluated in parallel within a neural network, before being recomposed in the end. That implies that each neural network will generate signals for the log-likelihood ratio, as well as for the log-cost benefit ratio of the particular state-comparison it is encoding. For an arbitrary comparison, we will thus observe the following likelihoods

$$r_e^i \sim \mathcal{N}\left(ln\left(\frac{P[e_i]}{P[e_n]}\right), \nu^2\right) \quad , \quad r_o^i \sim \mathcal{N}\left(ln\left(\frac{\mathbb{1}(x_i-c)}{c-x_n}\right), \nu^2\right). \tag{25}$$

where $\mathbb{1} = 1$ if $x_i > c$, and $\mathbb{1} = -1$ if $x_i < c$. Notice that, other than in the binary case described in the main text, I now model the signal r_o as a signal for the *relative benefit-cost ratio*. Indeed, one can no longer simply refer to 'costs' and 'benefits' in the general case. For one, benefits will always be relative. What is more, while the difference $c - x_n$ will always be positive, the difference $x_i - c$ could be either positive or negative, thus constituting a 'relative benefit' or a 'relative cost', depending on the sign of the expression.

We can now follow the same steps as in the binary setup to obtain a threshold equation and a contribution to the overall evidence in favour of taking the wager from each of the single comparisons. That is, the comparisons are modelled as being executed in parallel in a neural network, so that each comparison will ultimately have a contribution equal to the following z-score:

$$z_i = \frac{\gamma ln\left(\frac{P[e_i]}{P[e_n]}\right) + \mathbb{1}\alpha ln\left(\frac{x_i - c}{c - x_n}\right) - ln(\delta)^{-1}}{\nu\sqrt{\alpha^2 + \gamma^2}} \tag{26}$$

The contributions of the different state-wise comparisons can then simply be combined into an overall contribution, since the sum of n standard normal random variables will itself follow a standard normal distribution. Notice that some of these contributions may add to the attractiveness of the wager, while other may detract from it, depending on the relative size of x_i and c.

While it may be interesting to augment such a model with an endogenous account of attention to different lottery arms, some reasonable intuitions already obtain from the stylized setup above. For one, lottery arms that have high benefits relative to costs and favourable odds, or vice versa, will stand out, whereas arms with roughly equal odds and similar costs and benefits will have their contribution drowned out by the noise, and will thus not contribute anything to the overall evaluation. Furthermore, once the contributions of the different lottery arms are combined, the noise is predicted to increase in a factor of $\sqrt{n}$. That is, we will see a progressive increase in the noisiness of responses as the number of outcomes in a wager (or, more generally, the number of states to be compared) increases, albeit at a decreasing rate. This corresponds closely to the intuition of what we would expect to happen as the choice task becomes more and more complex.

S2 Additional results on PT correlations

This section presents additional details and results for the analysis of 2-outcome wagers under risk in a PT setting.

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). I conduct the analysis using Stan (Carpenter, Gelman, Hoffman, Lee, Goodrich, Betancourt, Brubaker, Guo, Li and Riddell, 2017). Take a parameter vector $\boldsymbol{\theta}_i$, indicating individual-level parameters. This vector follows the following distribution:

$$\boldsymbol{\theta}_i \sim \mathcal{N}(\overline{\boldsymbol{\theta}}, \Sigma), \tag{27}$$

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. I checked convergence by examining the R-hat statistics, and by checking for divergent iterations. 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 Bruhin et al. (2010), I use the density around the observed switching point to estimate the model. The model takes the following form:

$$ce \sim \mathcal{N}(u^{-1}[w(p)u(x) + (1 - w(p))u(y)], \widehat{\omega}^2 \times |x - y|), \tag{28}$$

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 paper 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.

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. The full model in Stan is as follows (@to be added upon publication):

S2.2 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), as well as presenting equivalent results for the data of L'Haridon and Vieider (2019), for losses as well as for gains.

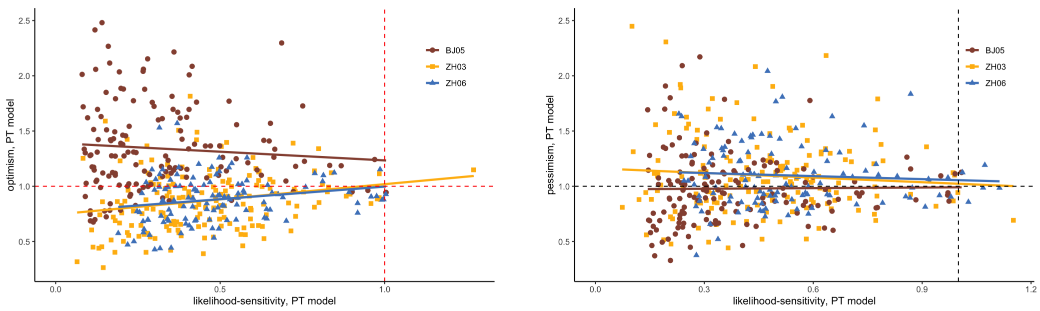

(a) Scatter plot of $\widehat{\gamma}^{+}$ and $\widehat{\delta}^{+}$

(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.

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 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.

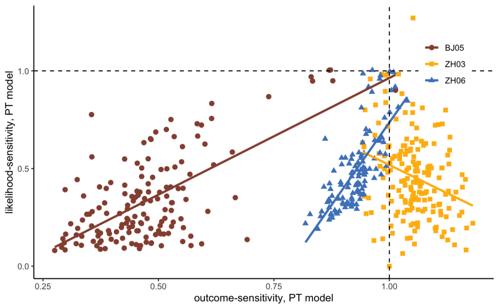

(a) Scatter plot of $\widehat{\alpha}$ against $\widehat{\gamma}$

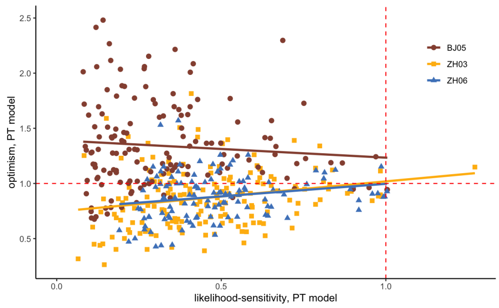

(b) Scatter plot of $\widehat{\gamma}$ against $\widehat{\delta}$

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.

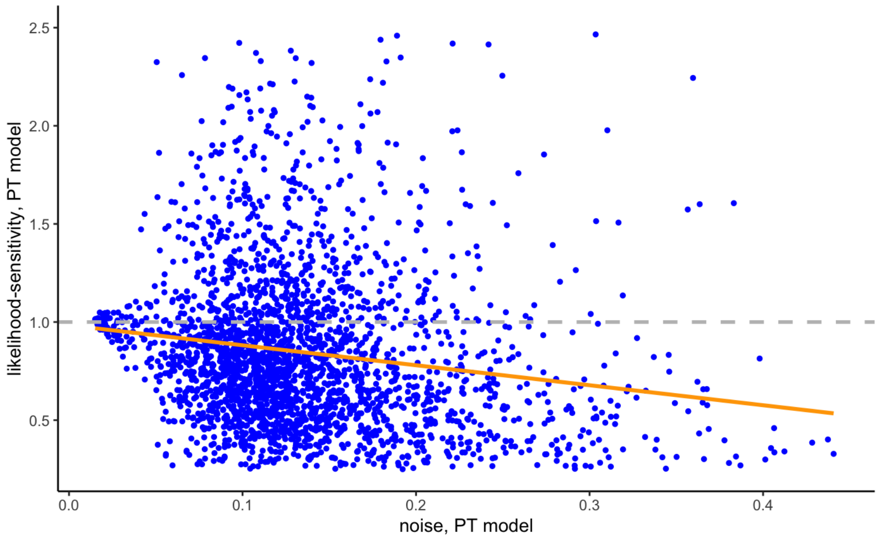

Figure S10: Correlatiion of $\widehat{\omega}^+$ and $\widehat{\gamma}^+$ in L'Haridon and Vieider (2019)

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. Correlations are tested on parameters demeaned at the country level, corresponding to a fixed effects specification. 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 ($\rho = 0.386, p < 0.001$).

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, $|1-\widehat{\gamma}^-|$,

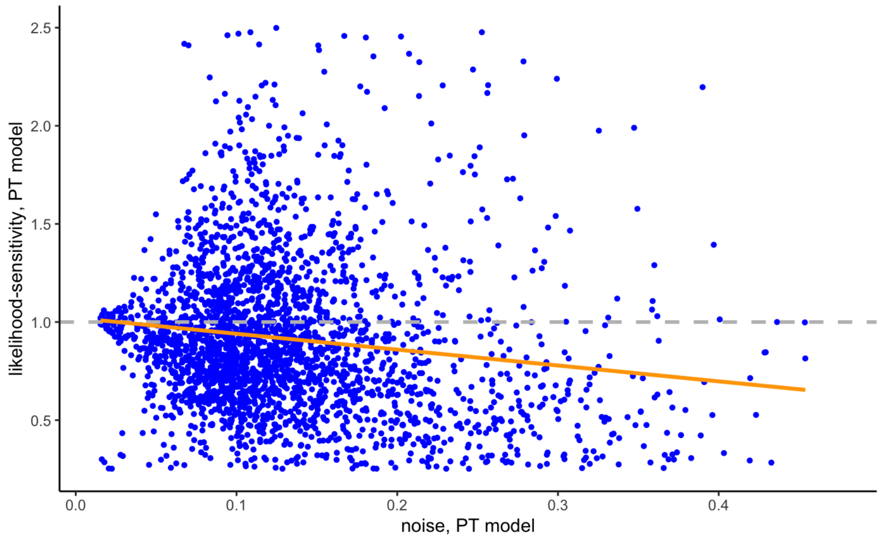

Figure S11: Correlatiion of $\widehat{\omega}^-$ and $\widehat{\gamma}^-$ in L'Haridon and Vieider (2019)

the correlation thus appears much stronger ($\rho = 0.445, p < 0.001$).

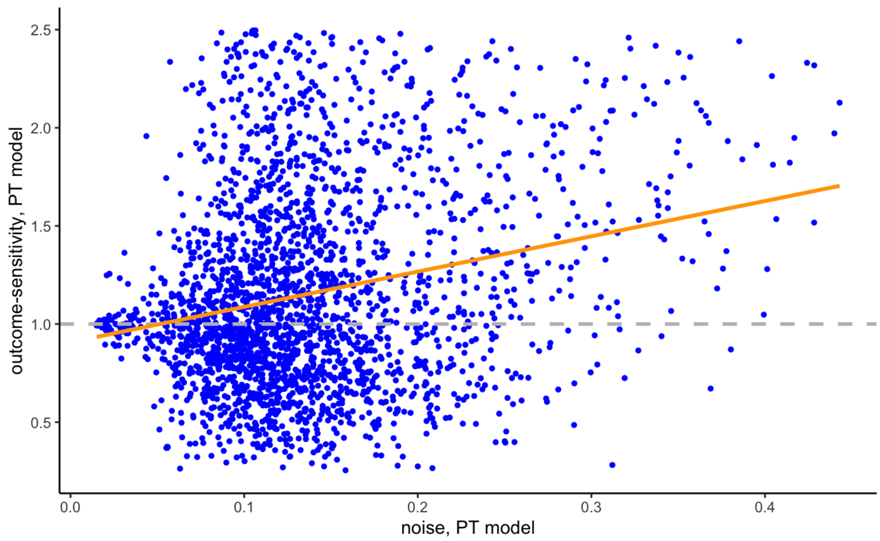

Figure S12: Correlatiion of $\widehat{\omega}^+$ and $\widehat{\alpha}^+$ in L'Haridon and Vieider (2019)

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 ($\rho = 0.34, p < 0.001$).

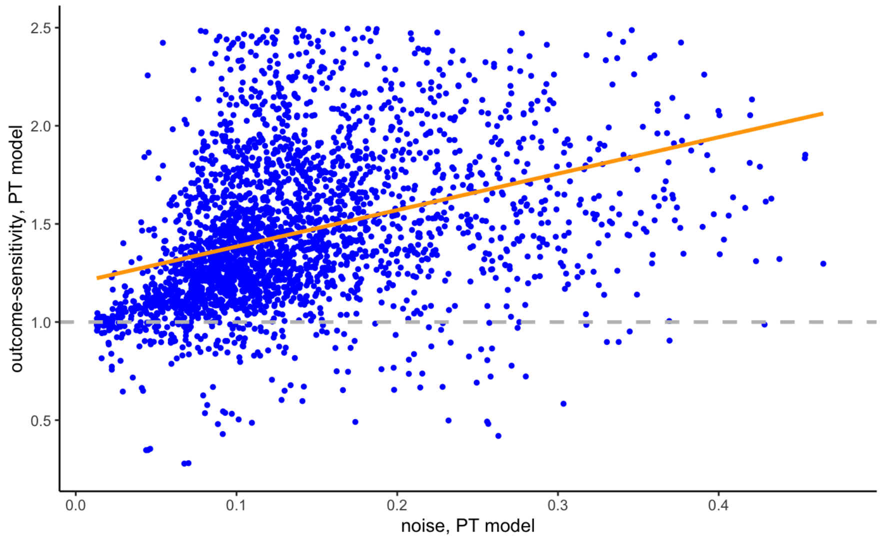

Figure S13: Correlatiion of $\widehat{\omega}^{+}$ and $\widehat{\alpha}^{+}$ in L'Haridon and Vieider (2019)

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 S2 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 S2: 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.

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. The model used for PT looks as follows (@add code upon publication):

S4 Instructions risky choice versus mirrors

First, all subjects saw the following general instructions:

Welcome to this experiment in decision making. Below, you will be asked to repeatedly choose between different options. In the end, **we will select 10 students to play one randomly chosen decision for real money**. Please read these instructions carefully, and pay close attention to the options being presented to you. Your final payoffs may depend on this.

In each task, you will be presented with 2 options. In both options, your payout will depend on **100 BOXES containing different monetary amounts**. The two sets of boxes will be displayed like in the example below. The counts of boxes are listed at the top, and below them you can see the monetary amounts contained in those boxes. In this example, the option at the top has 32 Boxes containing €18, and 68 Boxes containing €0 (nothing). The option at the bottom has 100 Boxes containing €6.

Please make a choice

32 Boxes	68 Boxes
€18	€0

100 Boxes
€6

Notice that the order of the boxes, the numbers of boxes, and associated amounts will change from task to task. **Please pay close attention to these dimensions before indicating your choice**. Please play careful attention to these instructions. After reading the instructions, you will be asked some comprehension questions. **You will only be allowed to proceed to the experiment if you correctly answer all comprehension questions**.

Subjects assigned to the RANDOM BOX treatment then saw the following instructions:

RANDOM BOX

I the upcoming tasks, **we will RANDOMLY DRAW one of the 100 boxes** from whichever option you have chosen to determine your payoff. All the boxes are equally likely to be drawn.

<u>Example</u>: Take a look at the example below. If you have selected the option at the top, we will make a **RANDOM DRAW** from a bag with numbers from 1 to 100. If the number drawn falls in the range 1–58, you get €32. If it falls in the range 59-100, you get € 0 (nothing). If you have selected the option at the bottom, you will simply be paid the sure amount of € 12, since all 100 boxes contain that same amount.

Please make a choice:

58 Boxes	42 Boxes
€32	€0

100 Boxes
€12

Subjects assigned to the AVERAGE BOX treatment instead saw the following instructions:

AVERAGE BOX

In the upcoming tasks, **we will pay you by calculating the AVERAGE amount of money across all 100 Boxes** from whichever option you have chosen. That is, we will add up the amount from each of the 100 boxes and divide the sum by 100.

<u>Example</u>: Take a look at the example below. If you have selected the option at the top, we will add (58 * €32 + 42 * €0)/100 = € 18.56. If you have selected the option at the bottom, you will simply be paid €12, since (€12 * 100)/100 = €12.

Please make a choice:

58 Boxes	42 Boxes
€32	€0

100 Boxes
€12

All subjects had to answer 6 comprehension questions as the ones shown below. The questions served the purpose of further emphasizing the treatment:

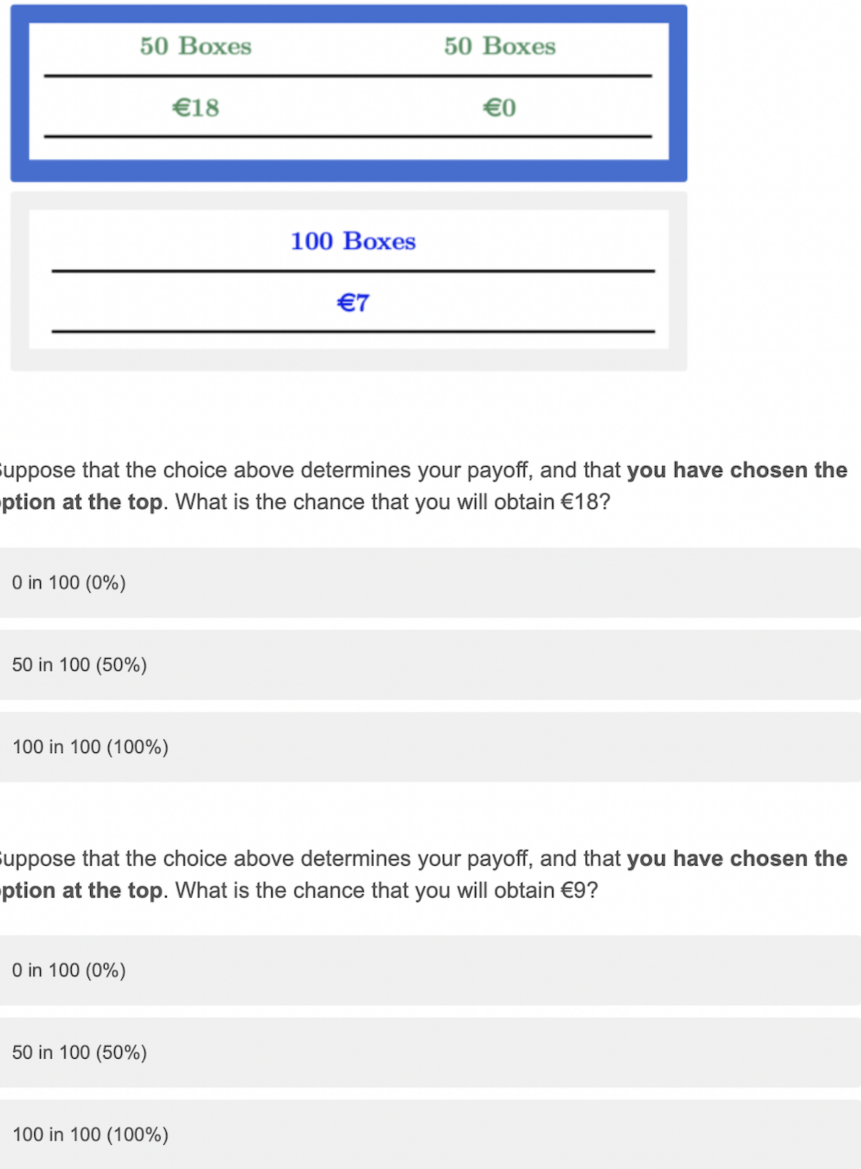

50 Boxes	50 Boxes
€18	€0

100 Boxes
€7

Suppose that the choice above determines your payoff, and that **you have chosen the option at the top**. What is the chance that you will obtain €18?

0 in 100 (0%)

50 in 100 (50%)

100 in 100 (100%)

Suppose that the choice above determines your payoff, and that **you have chosen the option at the top**. What is the chance that you will obtain €9?

0 in 100 (0%)

50 in 100 (50%)

100 in 100 (100%)