

WORKING PAPER

NOISY NEURAL CODING AND DECISIONS UNDER UNCERTAINTY

Ferdinand M. Vieider

August 2021
2021/1022

Noisy Neural Coding and Decisions under Uncertainty*

Ferdinand M. Vieider¹

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

August 20, 2021

Abstract

I derive a noisy neural coding model (*NCM*) and pit its performance against prospect theory plus additive noise (*PT*) using some prominent recent datasets collected to measure PT parameters. The NCM is based on the premise that choice patterns observed under uncertainty may originate from noisy perceptions of choice stimuli, which are optimally combined with mental priors to obtain actionable decision parameters. This contrast with PT, which models preferences as deterministic, but adds a noise term for empirical implementations. I show how the parameters emerging from the NCM naturally map into PT parameters. The NCM can thus be seen as a *generative model* for PT. At the same time, the NCM departs from PT in that it is inherently stochastic. This results in novel predictions about systematic correlations between PT parameters, as well as pointing to instances under which PT will be violated. Using Bayesian hierarchical models to fit the data, I find substantial support for these predictions. The NCM further consistently outperforms PT in terms of predictive ability. These results contribute to the nascent literature documenting the role played by imprecise cognition in economic decisions.

Keywords: risk taking; prospect theory; noisy cognition; efficient coding

JEL-classification: C93; D03; D80; O12

1 Motivation

Preferences over risk and uncertainty play a central role in economic models. The importance of having accurate descriptive models of decisions under risk and uncertainty has thus long been recognized. As the available evidence on decisions under risk increased, subsequently more sophisticated models of decision making have thus been proposed over the years ([Markowitz, 1952](#); [Kahneman and Tversky, 1979](#); [Loomes and Sugden, 1982](#); [1986](#); [Gul, 1991](#); [Tversky and Kahneman, 1992](#); [Bordalo, Gennaioli and Shleifer, 2012](#)).

*I gratefully acknowledge financial support from BOF at Ghent University under the project “The role of noise in the determination of risk preferences”. I am indebted to Mohammed Abdellaoui, Aurélien Baillon, Ranoua Bouchouicha, Cary Frydman, Olivier L’Haridon, Peter P. Wakker, Michael Woodford, and Horst Zank for helpful comments and discussions. All errors remain my own.

Two general features stand out in this feedback process between data and models. First, the models were generally created inductively by introducing features that could account for the median or average revealed choice patterns. Second, given the deterministic nature of the models, they need to be augmented by a stochastic component to fit empirical data. Much debate has indeed been generated by the question of whether observed behavioural deviations from any given model may be driven by systematic deviations from the choice model itself, or rather by noisy decision patterns pertaining to the stochastic component (Hey and Orme, 1994; Harless and Camerer, 1994; Loomes, 2005).

Recent contributions instead propose that the choice patterns we observe under risk may be the product of noise in the perception of the choice stimuli. According to an influential theoretical paradigm in neuroscience, 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). Mental representations of outside stimuli thus combine noisy information about choice options with mental priors indicating their likelihood in a given environment. Applying this intuition to choice between risky options, Khaw, Li and Woodford (2020) showed that the noisy perception of risky stimuli will result in small stakes risk aversion, which would otherwise appear paradoxical (Rabin, 2000). Frydman and Jin (2021) document that a given increase in payoffs has a greater effect on risky choice in a low-volatility environment than in a high-volatility environment, while at the same time making choice patterns less random. Further adding a numerical perception task under certainty, they showed the noisy perception parameters to be significantly correlated across the two tasks. Taken together, these results suggest that noisy perception is indeed driving choice patterns observed under risk.

In this paper, I estimate a noisy neural coding model (*NCM*) based on several prominent experimental datasets originally collected to estimate prospect theory (*PT*; Tversky and Kahneman, 1992). I then pit the the *NCM* against *PT*—widely considered to be the leading descriptive model of decision making under risk and uncertainty (Starmer, 2000; Wakker, 2010; Barberis, 2013)¹—and test them against each other in terms of their relative predictive performance. The use of data collected to measure *PT* preference

¹For prospects with at most two non-zero outcomes such as used in the datasets I analyze, nearly all existing non-EU theories are special cases of *PT* (Wakker, 2010, section 7.11). This includes, *inter alia*, original prospect theory (Kahneman and Tversky, 1979), rank-dependent utility (Quiggin, 1982), and disappointment aversion (Gul, 1991).

functions thereby makes for a conservative test of the NCM’s performance, given that the data are not specifically geared towards detecting the noisy choice patterns which play a central role in the NCM (see discussion on identification below).

As a basis for the empirical analysis, I start by deriving a noisy coding model that generalizes the setup proposed by [Khaw et al. \(2020\)](#), so that it can simultaneously capture distorted perceptions of likelihoods and outcomes. Starting from an optimal decision rule entailing expected value maximization, I describe the weights resulting from its neural implementation combining noisy perceptions of choice options with a mental prior by Bayesian updating. I then show how these weights result in mental representations that entail nonlinear distortions of both outcomes and likelihoods, and how they map into the parameters of a popular probability-distortion function first proposed by [Goldstein and Einhorn \(1987\)](#), and frequently used in the PT literature ([Gonzalez and Wu, 1999](#); [Bruhin, Fehr-Duda and Epper, 2010](#)).²

The NCM can thus be seen as a *generative model for PT*, inasmuch as it naturally results in PT-like decision parameters. In particular, both *likelihood-insensitivity*—the stylized fact whereby people are insufficiently sensitive to changes in probabilities in intermediate ranges, while attributing excessive weight to certainty and possibility ([Tversky and Wakker, 1995](#); [Fehr-Duda and Epper, 2012](#))—and decreasing sensitivity towards outcomes—the stylized fact whereby people increasingly discount outcomes as they depart from a reference point—can be traced back to noise in the encoding of the choice stimuli. These features of behaviour, which can be thought of as deviations from an optimal choice rule based on expected value maximization, can thus be traced back to cognitive constraints imposed by limited cognitive resources. The noisy neural coding of choice parameters will thus naturally result in discrimination difficulties for values that are observed relatively infrequently ([Netzer, 2009](#); [Woodford, 2012](#); [Payzan-LeNestour and Woodford, 2021](#)), resulting in systematic biases in those regions. While generating PT-like parameters, the NCM thus goes beyond PT by exploring the origin of the parameters, rather than focusing on the best functional fit to revealed choice behaviour.

The NCM also departs from PT in other ways. Randomness in observed choice be-

²[Khaw et al. \(2020\)](#) focus mainly on a setup in which outcomes are perceived with noise, while probabilities are perceived objectively. They do, however, also discuss a dual setup involving the noisy perception of probabilities, and resulting in the same probability distortion function I derive in this paper, albeit with some additional parameter restrictions. To pitch the model against PT, however, a joint model of the two dimensions turns out to be crucial, as is the inclusion of non-zero lower outcomes. I will come back to the differences between the two models in the discussion.

haviour emerges naturally from the model, and is intimately and inextricably linked to the model parameters, thus questioning the implicit orthogonality between a deterministic decision model and noise underlying empirical implementations of PT. This inherent stochasticity also results in substantive predictions about the model parameters themselves. For one, the NCM predicts that different parameters estimated using PT should show systematic correlations. The NCM also predicts that the outcome and probability dimensions should interact, thereby resulting in violations of PT’s strict separability precepts. Such violations of probability-outcome separability are indeed well-documented in decisions under risk (Hogarth and Einhorn, 1990; Fehr-Duda, Bruhin, Epper and Schubert, 2010; Bouchouicha and Vieider, 2017). The NCM predicts that separability should also be violated for decisions under *ambiguity*—a context where the precise decision probabilities are unknown or vague (Ellsberg, 1961)—so that ambiguity attitudes ought to be reflected in both likelihood-distortions and outcome-distortions, rather than purely in the likelihood dimension as predicted by PT (Wakker, 2010; Abdellaoui, Baillon, Placido and Wakker, 2011; Dimmock, Kouwenberg and Wakker, 2015).

Having derived the model and its predictions, I set out to test them on two large datasets collected in order to estimate PT parameters under risk—the data of Bruhin et al. (2010), which include three large experiments to identify PT parameters in Switzerland and China, and the data of L’Haridon and Vieider (2019), including 3000 subjects from 30 countries. I also test some specific predictions based on an original dataset on decisions under ambiguity implemented using Ellsberg urns.³ I analyze the data using a Bayesian hierarchical model (Gelman and Hill, 2006; Gelman, Carlin, Stern, Dunson, Vehtari and Rubin, 2014b; McElreath, 2016), and I test model performance using cross-validation methods (Gelman, Hwang and Vehtari, 2014a; Vehtari, Gelman and Gabry, 2017). These methods are explicitly devised to maximize out-of-sample predictive ability, thus avoiding potential issues arising from overfitting existing data.

I start by documenting systematic correlations between the parameters estimated from a PT model. In particular, I show that the PT parameters pertaining to the deterministic part of the model—far from being orthogonal to the noise component as traditionally assumed—show strong and systematic correlations with the estimated standard deviation of the additive noise term. This, in turn, results in knock-on effects resulting

³While the model I present is fully general and applicable also to natural sources of uncertainty, taking the model to the data requires additional information that is typically not available for decisions under naturally occurring uncertainty. I will discuss empirical identification after presenting the model.

in correlations between the model parameters themselves. While somewhat puzzling from a PT perspective, these correlations closely follow the predictions of the NCM. I then proceed to documenting correlations between the NCM-derived PT parameters and the parameters obtained from a traditional PT setup. While being closely correlated at times, the parameters of the two models show systematic departures in other instances, showcasing the commonalities and differences in the two approaches.

I then test the NCM and PT against each other in terms of their predictive performance. Implementing separate tests by experiment and for gains versus losses gives me 66 independent tests, conditional on the experimental designs. The NCM significantly outperforms PT in all 66 cases. Zooming in on parameter estimates for individual subjects, I show how PT overfits certain aspects of the data, thus resulting in poor predictive performance. This phenomenon, in turn, can be traced back to the correlations between the PT parameters mentioned above, which result in their poor separability based on relatively few individual-level data points. The NCM, on the other hand, avoids this issue due to the interaction between decision noise and the model parameters themselves, which serves to constrain the parameters in ways that give it an edge in terms of both data fit and predictive performance.

Finally, I test the PT prediction pertaining to the strict separation of the probability and outcome dimensions using an original dataset containing choices under both known and vague probabilities. Estimating a PT model, I show the typical pattern of likelihood-insensitivity increasing under ambiguity relative to risk ([Abdellaoui et al., 2011](#); [Trautmann and van de Kuilen, 2015](#)). At the same time, however, I also find outcome-sensitivity to decrease under ambiguity, violating the strict separability precept of PT. These patterns are further strengthened when estimating the NCM, under which they are predicted due to an increase in perception noise when probabilities are vague or unknown. This shows that the NCM cannot only account for typical PT patterns, but it can also predict under what circumstances PT will be violated.

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

2 The Noisy Coding Model

In this section, I lay out the theoretical model in several steps. I start by showing that coding uncertainty as log-likelihood ratios of complementary events constitutes a highly efficient form of representing and updating probabilistic information, thus resulting in an optimal choice rule under uncertainty. I then present a step-by-step derivation of a Bayesian mental encoding-decoding model, showing how the noisy mental encoding of choices between uncertain prospects 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. In particular, I show that the subjective distortions from the noisy encoding-decoding model map into the parameters of a popular probability-distortion function used in the PT literature, which are traded off against outcome-distortions akin to those modelled through utility functions.

2.1 Efficient coding of uncertainty as log-likelihood ratios

I start by showcasing the efficiency of coding probabilistic information in terms of likelihood ratios, or odds. Take a generic future event, e . For analytical tractability, I will work with discrete events (e.g., a ball extracted from an urn is blue). I will use the shorthand $P[e]$ to indicate the subjective likelihood of a generic event e .⁴ The probability of the complementary event, \tilde{e} , is indicated by $P[\tilde{e}]$. The model applies to choices between binary prospects $(x, e; y)$, where the monetary amount x obtains under the occurrence of event e , and $0 \leq y < x$ under the complementary event \tilde{e} .

Assume a decision maker (DM) needs to take a decision based on her assessment of the likelihood of an uncertain event e . Assume further that the DM does not know the true likelihood of e occurring, but observes a noisy signal, s , about this likelihood. For instance, we could think of a report on the health of the economy as a noisy signal about the likelihood that the stock market index will increase over the following week. The situation of interest is one where the DM knows the probability of the signal, conditional on the event occurring, $P[s|e]$, as well as the probability of the signal conditional on the event not occurring, $P[s|\tilde{e}]$. The DM can then use the following equation to infer the

⁴For 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.

likelihood of the event occurring, conditional on the signal:

$$\frac{P[e|s]}{P[\tilde{e}|s]} = \frac{P[s|e]}{P[s|\tilde{e}]} \times \frac{\hat{P}[e]}{\hat{P}[\tilde{e}]}, \quad (1)$$

where the ratio $\frac{\hat{P}[e]}{\hat{P}[\tilde{e}]}$ indicates the prior likelihood ratio (*PLR*), which incorporates any knowledge the DM may have previously held about the likelihood of event e .⁵

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 materializing, but $y < x$ if \tilde{e} obtains instead. The DM can weigh this wager against a sure option, c (replacing c with a non-degenerate binary prospect it straightforward). The costs and benefits from taking the wager can be integrated into equation 1. This results in an optimal choice rule whereby the DM should take the wager whenever:

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

where $(x - c)$ indicates the benefit from taking the wager under the contingency that e occurs, and $(c - y)$ indicates the cost in case of the occurrence of \tilde{e} . Since any monotonic transformation will leave this choice rule unaltered, we can take the natural logarithm of both sides. Further exploiting the equality in equation 1 and rearranging, we obtain:

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

The wager should thus be chosen over the sure option whenever the log-likelihood ratio of event e and its complement \tilde{e} (*LLR*) exceeds the log cost-benefit ratio (*LCBR*). The choice rule is indeed optimal, as information is aggregated using Bayesian updating and the various dimensions are traded off linearly against each-other. [Gold and Shadlen \(2002\)](#) showed how an analogous decision rule has been employed successfully by British codebreakers during WWII to break the supposedly unbreakable Enigma code employed by the German navy. [Gold and Shadlen \(2001\)](#) used reasoning based on neural processing to argue that the representation based on log-likelihood ratios underlying equation 3 results naturally when the truth of one hypothesis (say, event e will obtain) is signalled by one pool of neurons, and the truth of the alternative hypothesis (event \tilde{e} will obtain)

⁵This follows from an application of Bayes rule, whereby $P[e|s] = \frac{P[s|e] * \hat{P}[e]}{P[s]}$. The use of likelihood ratios means that $P[s]$ cancels out of the expression.

is signalled by a different pool of anti-neurons. Such a setup has the advantage of being invariant to changes in levels of neural activity.

The model I will present below is based on the premise that mental processing of uncertain stimuli or prospects functions much like the optimal decision rule just described. According to an influential theoretical paradigm in neuroscience, the human brain acts as a Bayesian calculation engine, 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 potentially noisy signal before entering the choice process. Noise will then arise especially when assessments of the relevant quantities are made quickly and intuitively (see Khaw et al., 2020, section 1, and Woodford, 2020, for extensive discussions). I now proceed to discussing how such actionable choice parameters are derived by the noisy mental encoding of uncertain choice stimuli and their decoding by combination with a mental prior.

2.2 Bayesian mental processing of noisy signals

I now present a step-by-step derivation of subjective choice parameters from the mental encoding-decoding process of choice stimuli. I will thus show that if the brain implements an optimal choice rule such as the one derived above, and if the stimuli are mentally encoded with some noise and subsequently decoded using a mental prior, then this process will result in an actionable choice rule into which the encoding noise and the parameters characterizing the mental prior will enter as subjective parameters. The derivations follow those of Khaw et al. (2020), but they are more general in that both outcomes and probabilities are simultaneously perceived with noise, and in that I allow for non-zero lower outcomes. In the next section, I show how this subjective choice rule derived from the mental representation of the stimuli maps into a tradeoff between a probability-distortion function and outcome distortions, as modelled under PT. Readers not interested in formal derivations may thus want to skip ahead to section 2.3.

The mental encoding stage and the likelihood

Take a mental signal, r , encoding the characteristics of a given choice problem. This mental signal will take the form of neural spike or firing rates, 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. I will assume that there are two such mental signals, r_e and r_o , encoding the desirability of the LLR and the LCBR, respectively.

These mental signals will generally be noisy, possibly because the noise itself may carry useful information about the accuracy of the signal.⁶ Assume further without loss of generality that each signal is based on *one* neuron, encoding the likelihood of the winning event or the benefit from winning, and one *anti*-neuron, encoding the likelihood of not winning or the associated costs. For stimuli that are close together—roughly equal probabilities of e and \tilde{e} , or roughly equal costs and benefits—the noisy firing rates of the neuron and anti-neuron will be virtually indistinguishable, resulting in low activations of r_e and r_o . If, however, the LLR or the LCBR deviate significantly from 0, the relative signal will exhibit high spike rates. This phenomenon—known as *discriminability* (cfr. [Dayan and Abbott, 2001](#), ch. 3)—will depend both on the difference in the mean spike count rate, and on the noisiness in the distribution.⁷

Neuronal firing rates react most strongly during rapid change, whereas reactions in a stable environment are much more subdued. Stimuli will thus have a compressed mental representation, and logarithmic or power functions are typically used to model such mental representations ([Dayan and Abbott, 2001](#), ch. 1; [Petzschner, Glasauer and Stephan, 2015](#)). I thus assume that r_e follows a normal distribution with as its mean the LLR, and that r_o follows a normal distribution with the LCBR as its mean:

$$r_e \sim \mathcal{N}\left(\ln\left(\frac{P[e]}{P[\tilde{e}]}\right), \nu_e^2\right) \quad (4)$$

$$r_o \sim \mathcal{N}\left(\ln\left(\frac{c-y}{x-c}\right), \nu_o^2\right), \quad (5)$$

where ν_e and ν_o represent the encoding noise of the likelihood ratio and the cost-benefit

⁶More generally, encoding noise may be a function of the cognitive resources available for a given task. Following the model of [Heng, Woodford and Polania \(2020\)](#), [Frydman and Jin \(2021\)](#) show formally that in a context where the likelihood adapts to the prior, such noisy encoding is indeed *efficient* according to several performance criteria, in the sense that it optimally exploits scarce cognitive resources (for further discussion of efficient coding, see [Ganguli and Simoncelli, 2014](#); [Summerfield and Tsetsos, 2015](#); [Wei and Stocker, 2015](#); [Polania, Woodford and Ruff, 2019](#)).

⁷Take the neuron pair encoding the cost-benefit ratio in r_o , and designate the neuron as r_o^n and the anti-neuron as r_o^a . Discriminability is then formally defined as $d = \frac{r_o^n - r_o^a}{\sigma_o}$, which measures the distance in the means of the activation rates of the two neurons in units of their common standard deviation. The same concept of discriminability is used in machine learning, where it serves to assign objects to distinct analytical categories (e.g. [Murphy, 2012](#)).

ratio, respectively. The standard deviations ν_e and ν_o thus convey information on the uncertainty with which a given stimulus is perceived.

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. This amounts to a so-called *just noticeable difference* in the stimulus m , Δm , so that $\frac{\Delta m}{m}$ is a constant. The neural encoding process can thus be conceived of as the driving factor behind a behavioural phenomenon that has long been known in psychophysics as the Weber-Fechner law (Fechner, 1860).

The mental decoding stage and the posterior

The information provided by the noisy mental signals r_e and r_o will need to be decoded based on prior information 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 as captured by ν_e and ν_o , which makes it desirable to combine the perceptual information 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 log-normal distribution, since the choice of a conjugate prior distribution will minimize the burden in terms of neural computations, so that:

$$\ln \left(\frac{P[e]}{P[\tilde{e}]} \right) \sim \mathcal{N} \left(\ln \left(\frac{\hat{P}[e]}{\hat{P}[\tilde{e}]} \right), \sigma_e^2 \right), \quad (6)$$

$$\ln \left(\frac{c - y}{x - c} \right) \sim \mathcal{N} (\ln(\mu_o), \sigma_o^2), \quad (7)$$

where $\frac{\hat{P}[e]}{\hat{P}[\tilde{e}]}$ is the mean of the prior likelihood ratio, μ_o the mean of the prior cost-benefit ratio, and σ_e and σ_o are the associated standard deviations of the priors. The latter fulfil a role analogous to the one of ν_e and ν_o in the encoding stage—they quantify the confidence associated with the expected values of the likelihood and cost-benefit ratios.

Combining the likelihoods in equations 4 and 5 with the priors in equations 6 and 7 by Bayesian updating, we obtain the following posterior distributions by the multiplication of the Gaussian likelihood with the conjugate Gaussian prior:

$$\ln \left(\frac{P[e]}{P[\tilde{e}]} \right) | r_e \sim \mathcal{N} \left(\frac{\sigma_e^2}{\sigma_e^2 + \nu_e^2} \times r_e + \frac{\nu_e^2}{\sigma_e^2 + \nu_e^2} \times \ln \left(\frac{\hat{P}[e]}{\hat{P}[\tilde{e}]} \right), \frac{\nu_e^2 \sigma_e^2}{\nu_e^2 + \sigma_e^2} \right) \quad (8)$$

$$\ln\left(\frac{c-y}{x-c}\right) | r_o \sim \mathcal{N}\left(\frac{\sigma_o^2}{\sigma_o^2 + \nu_o^2} \times r_o + \frac{\nu_o^2}{\sigma_o^2 + \nu_o^2} \times \ln(\mu_o), \frac{\nu_o^2 \sigma_o^2}{\nu_o^2 + \sigma_o^2}\right). \quad (9)$$

Let us define $\gamma \equiv \frac{\sigma_e^2}{\sigma_e^2 + \nu_e^2}$ and $\alpha \equiv \frac{\sigma_o^2}{\sigma_o^2 + \nu_o^2}$. The two variables α and γ are the weights assigned to the log-likelihood ratio and the log cost-benefit ratio of the stimulus, relative to the weights assigned to the prior means, since $\frac{\nu_e^2}{\sigma_e^2 + \nu_e^2} = 1 - \gamma$ and $\frac{\nu_o^2}{\sigma_o^2 + \nu_o^2} = 1 - \alpha$. The relative uncertainty associated with the stimulus and the prior, as captured by $\{\nu_e, \nu_o\}$ and $\{\sigma_e, \sigma_o\}$, will determine how much weight will be attributed to the likelihood versus the prior in the posterior distribution, which furnishes the actionable quantities on which the decision will be based.

Let us further define $\delta \equiv \left(\frac{\hat{P}[e]}{\hat{P}[\tilde{e}]}\right)^{1-\gamma}$. By substituting these variables into the posterior means in equations 8 and 9, we obtain the following posterior expectations of the log-likelihood ratio and log-cost-benefit ratio, conditional on the noisy mental signals:

$$E\left[\ln\left(\frac{P[e]}{P[\tilde{e}]}\right) | r_e\right] = \gamma \times r_e + \ln(\delta) \quad (10)$$

$$E\left[\ln\left(\frac{c-y}{x-c}\right) | r_o\right] = \alpha \times r_o, \quad (11)$$

where I assumed that $\ln(\mu_o) = 0$. This assumption corresponds to a prior indicating that the average ratio of costs to benefits is equal to 1. This echoes the assumption of [Khaw et al. \(2020\)](#) that the prior of different outcomes is distributed with the same mean. While keeping the number of free parameters contained, this assumption is not fundamental: allowing for $\ln(\mu_o) \neq 0$ simply adds a constant to the model.

The actionable choice rule

The posterior means just derived can now be used to inform the decision process. In particular, we must amend the optimal choice rule described in equation 3 by replacing the objective quantities with their noisy mental representations:

$$E\left[\ln\left(\frac{P[e]}{P[\tilde{e}]}\right) | r_e\right] > E\left[\ln\left(\frac{c-y}{x-c}\right) | r_o\right], \quad (12)$$

which indicates that the wager on event e will be accepted whenever the expectation of the log-likelihood ratio of the posterior exceeds the expectation of the log of the cost-benefit ratio. Substituting the posterior means in equations 10 and 11 into the choice rule in equation 12 and solving for the mental signals, we get:

$$\gamma \times r_e - \alpha \times r_o > \ln(\delta)^{-1}, \quad (13)$$

where the right-hand side 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. Given that r_e and r_o follow a normal distribution, their weighted difference in equation 13 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 :

$$\gamma \times r_e - \alpha \times r_o \sim \mathcal{N} \left(\gamma \times \ln \left(\frac{P[e]}{P[\bar{e}]} \right) - \alpha \times \ln \left(\frac{c-y}{x-c} \right), \omega^2 \right), \quad (14)$$

where $\omega \equiv \sqrt{\gamma^2 \nu_e^2 + \alpha^2 \nu_o^2 + 2\alpha\gamma\nu_e\nu_o\rho_{eo}}$ represents the standard deviation of the weighted difference of noisy mental signals under the assumption that the two noisy perception variables r_e and r_o exhibit a correlation equal to ρ_{eo} .

The probabilistic choice rule

To derive the probability with which the wager will be chosen over the sure outcome, we now need to obtain an expression free of the unobservable mental signals. To this end, we obtain the z-score of the weighted difference in mental signals:

$$z = \frac{\gamma \times r_e - \alpha \times r_o - \left[\gamma \times \ln \left(\frac{P[e]}{P[\bar{e}]} \right) - \alpha \times \ln \left(\frac{c-y}{x-c} \right) \right]}{\omega}, \quad (15)$$

which follows a standard normal distribution. The wager will then be chosen whenever this z-score exceeds the z-score based on the threshold equation 13:

$$Pr[(x, e; y) \succ c] = \Phi \left(\frac{\ln \left(\frac{P[e]}{P[\bar{e}]} \right) - \gamma^{-1} \left[\alpha \times \ln \left(\frac{c-y}{x-c} \right) - \ln(\delta) \right]}{\omega} \right), \quad (16)$$

where Φ is the standard normal cumulative distribution function.

I briefly discuss two special cases of this equation. If probabilities are perceived without noise, then γ and $\ln(\delta)$ drop out of the equation, and the standard deviation simplifies to $\omega = \nu_o$. If we further let $y \equiv 0$, we obtain a special case of the model that corresponds to the noisy perception of outcomes used by [Khaw et al. \(2020\)](#) to explain the [Rabin \(2000\)](#) paradox. For the general case where $y \neq 0$, outcome-discriminability is applied to costs and benefits, rather than to single outcomes. If instead we assume

a dual scenario under which outcomes are encoded without noise while probabilities are distorted, then α drops out of the equation and the standard deviation simplifies to ν_e .

Losses and loss aversion

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

$$E \left[\ln \left(\frac{P[e]}{P[\tilde{e}]} \right) | r_e \right] > E \left[\ln \left(\frac{x - c}{c - y} \right) | r_o \right], \quad (17)$$

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 (Woodford, 2012).

The case of mixed gambles, on the other hand, requires some special consideration. Studying neuronal activation rates for decisions involving gains and losses, Tom, Fox, Trepel and Poldrack (2007) have shown that losses are encoded as decreases in neuronal activity that are on average about twice as strong as the increases caused by equivalent monetary gains. They further showed that measures of ‘neural loss aversion’ correlated tightly with behavioural measures of loss aversion obtained from choices of the same subject (the correlation between neural loss aversion and behavioural loss aversion they report is 0.85). Such a hard-wired mechanism suggests that loss aversion ought to be generated directly by the noisy encoding of losses contrasted to gains.⁸ Assume a decision involving a loss of y that is compared to a gain of x , with a riskless outside option of c (this is usually set to 0 in experiments). I propose the following encoding process:

$$r_o^\pm \sim \mathcal{N} \left(\ln \left(\frac{\exp(c - y)^\beta}{x - c} \right), \nu_o^2 \right) \quad (18)$$

where the \pm indicates a signal specific to mixed prospects. The encoding for the benefit now follows the same distribution as discussed before. The distribution for the costs, now consisting of monetary losses, takes the form of a power law, which converges to the distribution for gains in the limit as $\beta \rightarrow 0$. For values of $\beta > 0$, however, this formulation captures the insight of Tom et al. (2007) that people tend to exhibit increased sensitivity

⁸Khaw et al. (2020) discuss a different mechanism whereby loss aversion may enter through the prior. I will contrast the two approaches in the discussion.

towards losses relative to gains, and that behavioural loss aversion is predicted by this increased sensitivity. This sort of distributions has been proposed for neural encoding by [Stevens \(1970\)](#), and is another form used to represent muted neuronal responses in steady state ([Dayan and Abbott, 2001](#); [Billock and Havig, 2018](#)).

2.3 Decisions under risk and ambiguity

Mental representation parameters map into PT parameters

I next discuss the implications of the model I have derived above in terms of PT. Abstracting momentarily from decision noise ω , we can write the point of indifference underlying the choice probability in equation 16 as follows:

$$\alpha^{-1} \left[\ln(\delta) + \gamma \times \ln \left(\frac{P[e]}{P[\bar{e}]} \right) \right] = \ln \left(\frac{c - y}{x - c} \right). \quad (19)$$

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

$$\ln \left(\frac{c - y}{x - c} \right) = \ln \left(\frac{\frac{c-y}{x-y}}{1 - \frac{c-y}{x-y}} \right) \equiv \ln \left(\frac{\pi(P[e])}{1 - \pi(P[e])} \right), \quad (20)$$

where $\pi(P[e])$ is the solution of the equation $c = \pi(P[e])x + (1 - \pi(P[e]))y$, constituting a dual-EU representation of the choice problem ([Yaari, 1987](#)), for the decision weight $\pi(P[e])$. The right-hand side in equation 20 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 prospect. Substituting equation 20 into equation 19, and solving for $\pi(P[e])$, we obtain:

$$\pi(P[e]) = \frac{\delta^{1/\alpha} (P[e])^{\gamma/\alpha}}{\delta^{1/\alpha} (P[e])^{\gamma/\alpha} + (1 - P[e])^{\gamma/\alpha}}. \quad (21)$$

For $\alpha = 1$, this expression reduces to a popular probability-distortion function first proposed by [Goldstein and Einhorn \(1987\)](#), and commonly used in the decision-making literature ([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, and is interpreted as a parameter of likelihood-sensitivity

(Gonzalez and Wu, 1999). 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 the weighting of ranked losses.

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

Decisions under risk and ambiguity

The setup just derived can be immediately 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), noisy encoding implies that we would expect the discrimination parameter to decline as the information about the probabilities involved becomes more vague.⁹ This is a natural consequence of the relative weights associated to the stimulus encoding in equation 8, whereby an increase in the uncertainty with which the stimulus is perceived, captured by ν_e , will lead to a shifting of the weight from the stimulus to prior.

This, in turn, yields the prediction that $\gamma_a < \gamma_r$, where the subscripts a and r stand for ‘ambiguity’ and ‘risk’ respectively—a phenomenon known as *ambiguity-insensitivity*,

⁹In keeping with the rest of the paper, I limit my discussion to the case of Ellsberg urns with unknown color proportions. Findings may well deviate from the ones described here in other contexts and for different sources of uncertainty, since it is well-known that ambiguity attitudes over natural sources of uncertainty may vary depending on a DM’s knowledge of the specific context or the DM’s perceived competence in a given task (Heath and Tversky, 1991; Fox and Tversky, 1995; Abdellaoui et al., 2011).

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

We can also derive predictions about the elevation parameter δ . Under the NCM, the parameter has a natural interpretation as the mean of the prior. Its constituent part, $\psi \equiv \widehat{P}[e]$, can be shown to coincide with the fixed point of the probability-distortion function, i.e. the point where the function crosses the 45° line.¹⁰ In other words, the expectation of the mental prior distribution of an event e occurring coincides with the fixed point of the probability-distortion function, at which probabilities are perceived objectively and without subjective distortions.

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

The NCM is inherently stochastic

Having discussed a deterministic choice rule to derive the PT-functionals, it is now at the time to circle back to the decision noise parameter ω . Other than under PT,

¹⁰To see this, substitute ψ into the decision weight to get:

$$\ln\left(\frac{\pi(\psi)}{1-\pi(\psi)}\right) = \gamma \times \ln\left(\frac{\psi}{1-\psi}\right) + (1-\gamma) \times \ln\left(\frac{\psi}{1-\psi}\right) = \ln\left(\frac{\psi}{1-\psi}\right), \quad (22)$$

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

where noise is added on to the deterministic part of the model and assumed to be orthogonal to it, the decision noise parameter of the NCM naturally emerges from the same mental encoding-decoding process as the other model parameters. This is indeed the direct consequence of the mental representations taking a probabilistic form. As I will document extensively below, this feature is central to the model. First, however, I need to discuss the relationship between the different encoding noise components.

Although conceptually distinct, encoding noise ν_e and ν_o cannot be distinguished in practice from observed choices unless one is willing to make strong additional assumptions about the underlying process. This follows from the joint distribution of r_e and r_o as modelled in equation 14, showing that this inseparability is a feature of the underlying mental coding process. It is further supported by extensive simulations, reflecting once again that the tradeoff between the probabilistic and outcome dimensions is central to decisions under uncertainty. The practical issues making the separation of the two components impossible are indeed such as to suggest a more substantive . I will thus henceforth write $\nu_e = \nu_o \equiv \nu$, which results in a simplified decision error $\omega = \nu \times (\alpha + \gamma)$. As a consequence, the NCM has the same number of free parameters as a PT model.¹¹

The decision noise parameter $\omega = \nu \times (\alpha + \gamma)$ holds the key to understanding the NCM. In PT, the corresponding decision noise variable is modelled as being orthogonal to the deterministic decision model (see section 3.1 for a short formal exposition of a typical PT setup). In the NCM, on the other hand, the noisiness of the choice process is intimately and inextricably linked to the other parameters of the model. Outcome-discriminability α and likelihood-discriminability γ directly enter the definition of decision noise ω , and through them their component parts σ_o and σ_p (on top of ν). That implies that the noisiness of the decision process will be impacted both by the uncertainty connected to the stimulus encoding, and the uncertainty in the mental priors. Model parameters and decision noise thus move together, resulting in an inherently stochastic model. An exception to this rule concerns the mean of the prior. While optimism δ is affected through γ , given its definition $\delta = \left(\frac{\psi}{1-\psi}\right)^{1-\gamma}$, the native NCM parameter ψ making

¹¹Focusing on prospects over gains only, the NCM has the native parameters ν , indicating encoding noise; ψ , indicating the fixed point of the probability-distortion function; and the two standard deviations of the likelihood and outcome priors, σ_e and σ_o . This corresponds to PT's 3 deterministic parameters $\hat{\alpha}$ (utility curvature), $\hat{\gamma}$ (likelihood-sensitivity), and $\hat{\delta}$ (optimism), plus a noise parameter, call it $\hat{\omega}$, where the 'hat' on the parameters serves to distinguish the native PT parameters from the equivalent NCM-derived parameters discussed in the text. These parameters are exactly replicated for losses, under the assumption that the error term under PT is allowed to be heteroscedastic across choice domains. One additional parameter is then required to capture loss aversion under both the NCM (β) and PT ($\hat{\lambda}$).

up the mean of the prior distribution remains unaffected. This will have important interpretative consequences, to which I will return in the discussion.

The inherent stochasticity of the NCM has a number of substantive implications for the decision-making patterns we would expect under risk and ambiguity. 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, although there are likely to be also knock-on effects on other parameters. The neglect of these systematic relationships, in turn, will result in inferior predictive ability of the model. As a consequence, different model parameters estimated under PT will jointly be affected by the same observable characteristics of the decision maker or the environment, with model parameters often moving in opposite directions with the observable characteristics of DMs in terms of revealed risk attitudes, making an analysis of the correlates of risk attitudes based on PT fraught with difficulties.

Probability-outcome separability and PT violations

A further consequence of the inextricability of ν_e and ν_o is that no strict separability between the likelihood and outcome dimension seems warranted, in opposition to what is assumed in PT. I illustrate the consequences with two examples.

In first instance, we may expect that coding noise increases in the range of the stakes. This sort of adaptation may indeed be optimal in a context of scarce resources. [Frydman and Jin \(2021\)](#) present a model of *efficient coding*, in which the parameter of the likelihood function adapts to the prior exactly in this way. They furthermore provide experimental evidence that when outcome ranges increase, so does the noisiness in the encoding process. Such an increase in noisiness should then immediately be reflected in lower outcome-discriminability, resulting in apparent patterns of increasing relative risk aversion ([Holt and Laury, 2002](#)). Given the impossibility of separating the different dimensions of encoding noise, however, it will also result in apparent changes in likelihood-discriminability. Under PT, such effects will then take the form of violations of the principle 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 ambiguity. Given the increase in encoding noise we expect under ambiguity, the NCM predicts a lowered level of likelihood-discriminability. The latter, in turn, may 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).¹² Notice that the theoretical limiting case is one whereby coding noise for outcomes and probabilities is completely independent, such that $\nu_e \perp \nu_o$, in which case we would indeed expect outcome-discriminability to be unaffected by properties pertaining to the probability dimension, as predicted by PT. That case, however, appears conceptually unlikely due to the reasons laid out above.

3 Empirical evidence

In this section, I use several datasets to test the predictions and performance of the NCM against those of PT. I use two large datasets used to estimate PT parameters in decision under risk—Bruhin et al. (2010), who present three large experiments from two countries, and L’Haridon and Vieider (2019), who present data for close to 3000 subjects from across 30 countries. I further use one original dataset on decisions under ambiguity. I only use this one dataset for decisions under ambiguity, and only ambiguity using Ellsberg urns, because existing datasets on ambiguity attitudes typically lack some stimuli required to fully identify the NCM parameters (see below for details). I also exclude loss aversion, since the specific model of loss aversion proposed above does not lend itself to being estimated using existing data, seen how it requires stake variations in mixed prospects due to the nonlinearities involved. Since this is furthermore not the only possible representation of loss aversion within the NCM (see discussion), this topic is best left for future research. I only provide essential details about the estimations in the main text. Extensive further details on the analysis, estimation scripts, and additional results can be found in the supplementary materials.

¹²It is, of course, also in contradiction to models that capture ambiguity attitudes purely through utility curvature, such as the one proposed by Klibanoff, Marinacci and Mukerji (2005).

3.1 Preliminaries to empirical estimations

I use Bayesian hierarchical models to estimate the model parameters (Gelman and Hill, 2006; Gelman et al., 2014b). This allows me to obtain individual-level estimates of all the parameters jointly with the aggregate-level estimates in one and the same framework. This method is geared towards maximizing the *predictive* power of the model for new data, rather than towards optimizing its fit to existing data. This serves to avoid issues of overfitting that occur when employing maximum likelihood procedures on individual-level data. I conduct the estimations using Stan (Carpenter, Gelman, Hoffman, Lee, Goodrich, Betancourt, Brubaker, Guo, Li and Riddell, 2017). Detailed explanations can be found in the supplementary materials.

I am interested in two sets of parameters. The first set consists in the NCM parameters, which consist of two sub-groups: 1) the native NCM parameters ν , ψ , σ_e , and σ_o ; and 2) the NCM-derived PT parameters α , γ , δ , and ω . I subscript the parameters by a when referring to ambiguity and by r when referring to risk, unless the context makes subscripting unnecessary. I superscript them by $+$ or $-$ to indicate gains versus losses wherever this distinction is relevant. I refer to γ as *likelihood-discriminability* and to α as *outcome-discriminability* to distinguish them from the equivalent PT parameters.

The second set of parameters I am interested in are the native PT parameters resulting from a typical PT setup plus additive noise. Under PT, a prospect $(x, p; y)$ will be chosen deterministically over a sure outcome c whenever

$$c < u^{-1}[w(p)u(x) + (1 - w(p))u(y)], \quad (23)$$

where $w(p)$ is a probability-distortion function mapping probabilities into decision weights, and $u(\cdot)$ is a utility function mapping outcomes into utilities. This model is applicable to both gains and losses. To allow for data fitting, an additive error term is then appended to the deterministic part above, so that the risky option is chosen whenever

$$c < u^{-1}[w(p)u(x) + (1 - w(p))u(y)] + \epsilon, \quad (24)$$

where $\epsilon \sim \mathcal{N}(0, \hat{\omega}^2)$ is a random error term supposed to take the form of ‘white noise’, i.e. implicitly assumed to be orthogonal to the deterministic part of the model. Deriving a Probit choice rule in keeping with what I did for the NCM above, we obtain:

$$Pr[(x, p; y) \succ c] = \Phi \left(\frac{u^{-1}[w(p)u(x) + (1 - w(p))u(y)] - c}{\hat{\omega}} \right) \quad (25)$$

where $\hat{\omega}$ is the standard deviation of the error term. To ensure comparability with the NCM, I will use power utility throughout, so that $u(x) = x^{\hat{\alpha}}$. I will use the [Goldstein and Einhorn \(1987\)](#) probability-distortion function for the same reason, with parameters $\hat{\gamma}$ and $\hat{\delta}$, where the ‘hat’ serves to distinguish the PT parameters from the equivalent NCM-generated parameters.¹³ I will refer to $\hat{\gamma}$ as *likelihood-sensitivity*, and to $\hat{\alpha}$ as *outcome-sensitivity*, to distinguish them from the equivalent NCM-derived parameters.

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

Noise deserves some extra attention in this context, given its centrality to the NCM setup. [Khaw et al. \(2020\)](#) identify noisiness by repeating identical choices multiple times. Noise, however, can also be identified in contexts where different choices are similar, though not identical, since such choices will result in inconsistencies and stochastic dominance violations across choice lists in regions where stimuli cannot be discriminated. Experiments imposing or nudging subjects towards consistency within choice lists—something that applies to all the data I use—may somewhat underestimate noise. Notice, however, that this would bias the tests I conduct against the NCM, so that my tests can be considered conservative. Testing the models against each-other on data collected to fit PT may thus indeed be considered a strength of the tests I present.

I test model performance by means of leave-one-out cross-validation (*LOO*; [Gelman et al., 2014a](#); [Vehtari et al., 2017](#)). Models are thus compared on their out-of-sample predictive ability, rather than based on their fit to a given dataset. The results obtained

¹³[Prelec \(1998\)](#) presents an popular alternative 2-parameter weighting function that has been used inter alia by [L’Haridon and Vieider \(2019\)](#). Notice, however, that the two weighting functions provide a very similar fit except for prospects with extremely small or large probabilities, which do not occur in the data I analyze.

with the LOO criterion are similar to, but more stable than, those obtained using the Watanabe-Akaike information criterion (WAIC; [Watanabe and Opper, 2010](#)), which is a Bayesian generalization of the deviance information criterion. All model comparisons reported are unaffected by using the WAIC instead. Correlations discussed in the text refer to Spearman rank correlations on the estimated parameters. Any p-values reported are always two-sided. Graphical displays may cut some outliers for better visual display.

3.2 Decisions under risk

Parameter correlations under PT

I start by documenting correlations between parameters estimated from a PT setup plus additive noise. Under PT, additive errors are supposed to take the form of ‘white noise’ assumed to be orthogonal to the decision parameters. The NCM, on the other hand, predicts that decision noise will be systematically correlated with the model parameters. I start by analyzing this issue in the data of [Bruhin et al. \(2010\)](#). The data are analyzed using an econometric approach analogous to the one used in the original paper, but in a random parameter setup rather than using a finite mixture model. I thus use directly the density around the switching point for the estimation, including heteroscedasticity across outcome ranges in the PT setup. Errors are allowed to vary across gains and losses (further details in section [S2](#) of the supplementary materials).

Panel [1\(a\)](#) of figure [1](#) shows a scatter plot of the PT noise parameter, $\hat{\omega}^+$, against the likelihood-sensitivity parameter, $\hat{\gamma}^+$, for gains. The two parameters show a strong negative correlation ($\rho = -0.422, p < 0.001$). This negative correlation is present for each of the 3 experiments (ZH03: $\rho = -0.311, p < 0.001$, ZH06: $\rho = -0.629, p < 0.001$, BJ05: $\rho = -0.512, p < 0.001$). The results for losses, shown in panel [1\(b\)](#), are very similar. Again, the negative correlation is highly significant in the aggregate data ($\rho = -0.389, p < 0.001$), as well as in the 3 individual experiments (ZH03: $\rho = -0.179, p = 0.016$; BJ05: $\rho = -0.587, p < 0.001$; ZH06: $\rho = -0.633, p < 0.001$). For both gains and losses, a small group stands out that has virtually no noise and likelihood-sensitivity arbitrarily close to 1. These are the expected value maximizers detected in the mixture model of [Bruhin et al. \(2010\)](#), who most likely based their responses on precise calculation of prospect values rather than on quick and approximate judgments.

Panel [1\(c\)](#) shows the correlations between the PT noise parameter, $\hat{\omega}^+$, and the

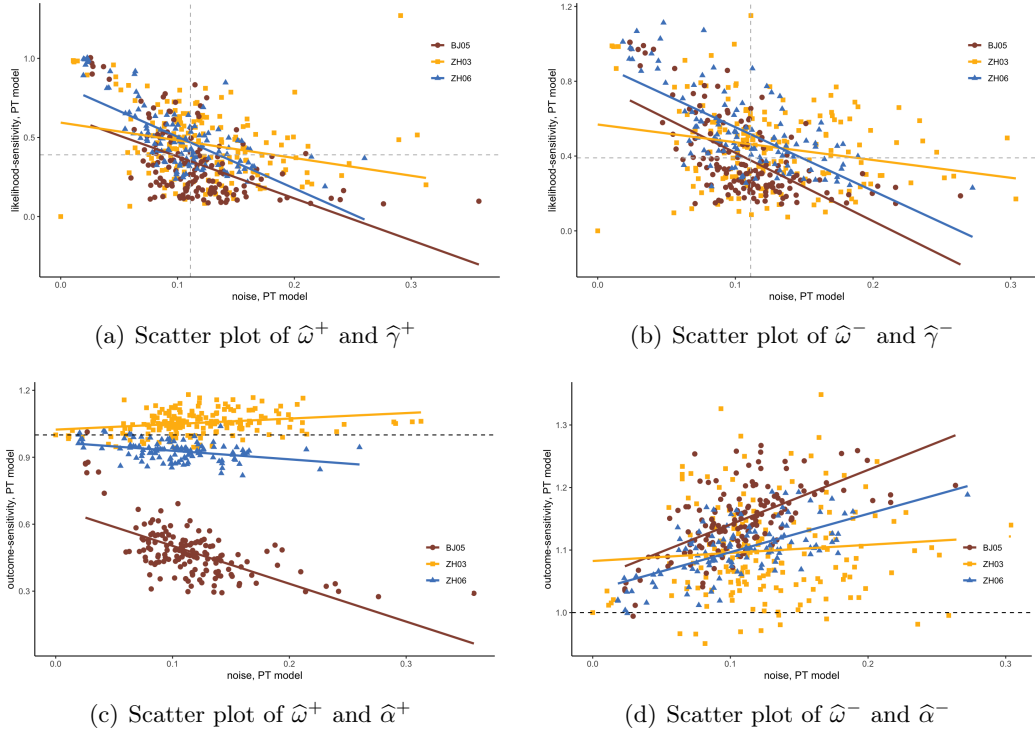


Figure 1: Scatter plot of PT parameters

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

outcome sensitivity parameter, $\hat{\alpha}^+$, for gains. Once again we witness systematic correlations, although the overall picture is complicated by the fact that the correlation goes in opposite directions in the three experiments. In particular, the correlation is strongly negative in the Beijing 05 data ($\rho = -0.566, p < 0.001$), as well as in the Zurich 06 data ($\rho = -0.357, p < 0.001$). It is, however, significantly positive in the Zurich 03 data ($\rho = 0.279, p < 0.001$). Indeed, the average power utility parameter reported by Bruhin et al. (2010) for the Zurich 03 experiment is larger than 1, which appears to be driving this behaviour. In other words, for parameters which may plausibly fall to either side of 1, as is the case for outcome-sensitivity, noise can be correlated with either excess sensitivity or insensitivity under PT. This intuition is further confirmed for losses, shown in panel 1(d). Here we witness a positive correlation in the aggregate data ($\rho = 0.258, p < 0.001$), as well as in the three individual experiments (ZH03: $\rho = 0.07, p = 0.38$; BJ05: $\rho = 0.612, p < 0.001$; ZH06: $\rho = 0.558, p < 0.001$). This is driven by concave utility for losses in all three experiments. I will discuss this issue at

some length further below in the context of the NCM.

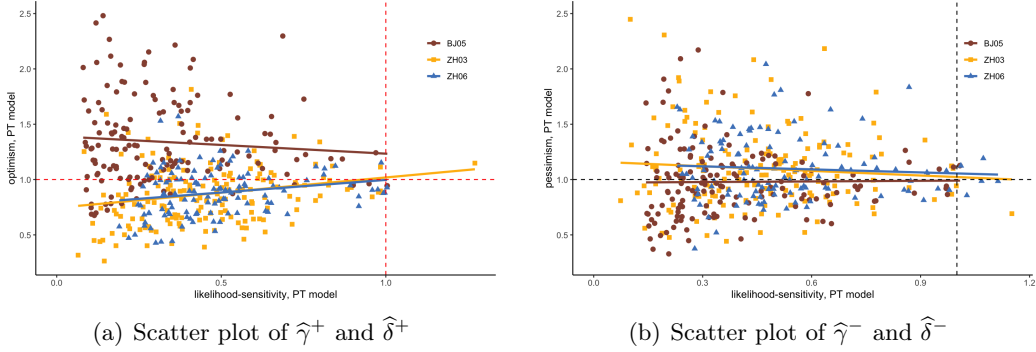


Figure 2: Scatter plot of PT parameters $\hat{\gamma}$ and $\hat{\delta}$

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.

The correlations just shown further have knock-on effects on correlations between the deterministic model parameters. For instance, likelihood-sensitivity $\hat{\gamma}$ and outcome-sensitivity $\hat{\alpha}$ will tend to be either positively or negatively correlated, depending on whether the correlation of $\hat{\omega}$ with $\hat{\alpha}$ is negative or positive. For instance, likelihood sensitivity and outcome-sensitivity for gains are positively correlated in the Beijing 05 data ($\rho = 0.463, p < 0.001$), as well as in the Zurich 06 data ($\rho = 0.753, p < 0.001$). They are, however, negatively correlated in the Zurich 03 data ($\rho = -0.254, p < 0.001$), given the positive correlation between decision noise and likelihood-sensitivity.

Another interesting pattern emerges from the relation between likelihood-sensitivity and optimism for gains, shown in figure 2 panel 2(a), and between likelihood-sensitivity and pessimism for losses, shown in panel 2(b). In both cases, the values of $\hat{\delta}$ are most dispersed for small values of $\hat{\gamma}$, with the dispersion decreasing markedly as $\hat{\gamma}$ increases, resulting in a funnel with the narrow part pointing to the right. Testing the correlation of absolute deviations of $\hat{\delta}$ from 1 with $\hat{\gamma}$, 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 NCM, 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.

The patterns in the L’Haridon and Vieider (2019) data are very similar to those just shown. I thus only show the correlation between decision noise and likelihood-sensitivity for gains, with additional results in section S3. To keep the display intelligible, I represent




Figure 3: Scatter plot of $\hat{\omega}^+$ and $\hat{\gamma}^+$ for L’Haridon and Vieider (2019)

Correlation plot of decision noise and likelihood-sensitivity in the global data of L’Haridon and Vieider (2019). The parameters have been obtained using a hierarchical model with two levels: individuals are nested within countries, which are nested within the global mean. The dashed line indicates perfect likelihood-sensitivity, $\hat{\gamma}^+ = 1$.

the data for four countries—China, the USA, Vietnam, and Ethiopia. These countries span the whole range of parameter values documented in the original paper. Graphs for the entire global dataset are shown in the supplementary materials. Figure 3 shows the correlation between $\hat{\omega}^+$ and $\hat{\gamma}^+$. As already seen for the Bruhin et al. (2010) data, the two parameters sport a significantly negative correlation ($\rho = -0.187, p < 0.001$). The correlation is somewhat weakened by some values of $\hat{\gamma}^+ > 1$, which also tend to be associated with higher levels of noise. And indeed, the correlation is much stronger when looking at absolute deviations of $\hat{\gamma}$ from 1 ($\rho = 0.492, p < 0.001$). This fact will become important when discussing the predictive ability of the different models father below.

Correlations between PT and NCM parameters

To zoom in on the commonalities and differences between PT and the NCM, it is useful to directly compare the native PT parameters, $\hat{\alpha}$, $\hat{\gamma}$, $\hat{\delta}$, and $\hat{\omega}$, with the equivalent NCM-derived PT parameters, α , γ , δ , and ω . I illustrate the correlations based only on the L’Haridon and Vieider (2019) data for gains. Results for losses and for other countries are similar (see section S3), as are the equivalent correlations in the Bruhin et al. (2010) data. Section S3.3 shows graphs of the distribution of the NCM-derived PT parameters.

The correlations are displayed in figure 4. I start from the correlation between decision noise in the two models, shown in panel 4(a). The association between the two

variables is extremely tight ($\rho = 0.947, p < 0.001$), as one might expect given that both are identified from the noisiness of choices. Given that ω contains in itself the model parameters α and γ , whereas $\hat{\omega}$ is implemented as an independent dimension, we would however expect larger differences to emerge for those other parameters. These differences start surfacing in panel 4(b), showing the correlation of the optimism parameter emerging from the two models. At $\rho = 0.711, p < 0.001$, they nevertheless are again highly correlated. Indeed, δ is only indirectly influenced by the decision noise encapsulated in ω . This indirect influence surfaces in the narrower distribution the parameter has under the NCM compared to the PT model (IQR of 1.25 for $\hat{\delta}$ versus 0.336 for δ), and is a manifestation of the shrinking power of $1 - \gamma$ in the NCM.

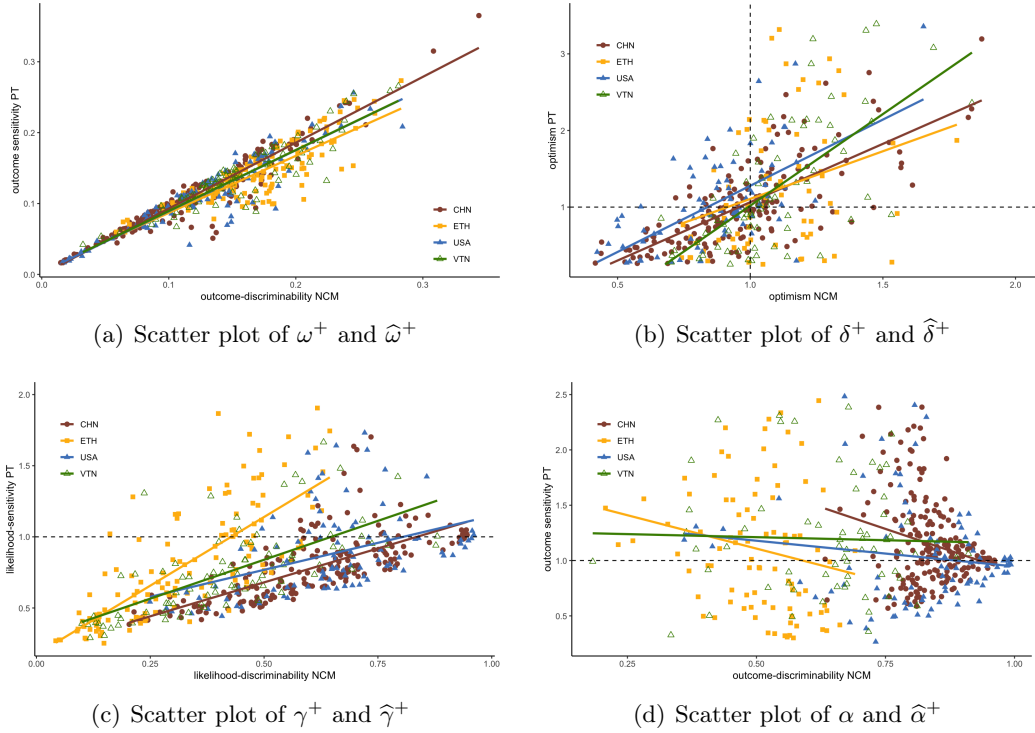


Figure 4: NCM-generated parameters and PT parameters for L'Haridon and Vieider (2019)

The parameter correlations decrease further once we move to likelihood distortions, shown in panel 4(c), although remaining highly significant ($\rho = 0.662, p < 0.001$). This shows the effect of the explicit interaction between decision noise and the model parameters. Notice also that values of $\hat{\gamma} > 1$ estimated using the PT specification do not generally map into values of γ close to 1, but rather result in intermediate values.

These effects become even more extreme for outcome-distortions, shown in panel 4(d),

where we find no correlation between the PT and NCM parameters ($\rho = -0.031, p = 0.522$). This absence of correlation is based on two main factors. The first is simply the different definition of outcome-distortions in the two models. Whereas PT applies outcome distortions to individual outcomes, the NCM results in distortions applied to the costs and benefits of a wager. The second reason concerns the coexistence of values of $\hat{\alpha}$ larger and smaller than 1 under PT. We have seen that both are associated with noise, implying some degree of random variation. Indeed, values of the parameters falling relatively close to 1, and associated with low noise levels, are very similar across the two models. Outcome-sensitivity under PT is usually identified from stimuli with relatively modest stake ranges. This issue is further aggravated by the observation that few stimuli of the crucial type with non-zero lower outcomes are typically included in experiments—only 5 out of the 14 stimuli for risky gains in the data of [L’Haridon and Vieider \(2019\)](#) fulfil this criterion—which creates a risk of overfitting noisy data.

The distribution of native NCM parameters

Let us now take a look at the native NCM parameters ν, ψ, σ_e and σ_o . Figure 5 shows the density functions for the 4 NCM parameters for gains for the [Bruhin et al. \(2010\)](#) data. The parameters are distributed over reasonable ranges, and resemble each other for the different experiments. The largest differences between experiments emerge for the standard deviation of the outcome prior, σ_o , in panel 5(b) and for the fixed point of the probability-distortion function, ψ , in panel 5(d). The first has a relatively broad distribution in the Zurich 03 experiment. The second tends towards higher values in the Beijing 05 experiment.

I next plot those same parameters for the data of [L’Haridon and Vieider \(2019\)](#), once again for gains and for the usual 4 countries. The parameter distributions can be seen in figure 6. Panel 6(a) shows systematic differences in encoding noise between countries, on top of the differences between individuals within each country. In particular, the US and China have relatively low encoding noise, while Ethiopia has very large encoding noise, with Vietnam falling in between. Panel 6(d) also shows systematic differences between countries, with Vietnam and Ethiopia having larger fixed points of the probability-distortion function on average. The between-country differences for the two scale parameters of the prior, shown in panel 6(b) and panel 6(c) are somewhat less pronounced.

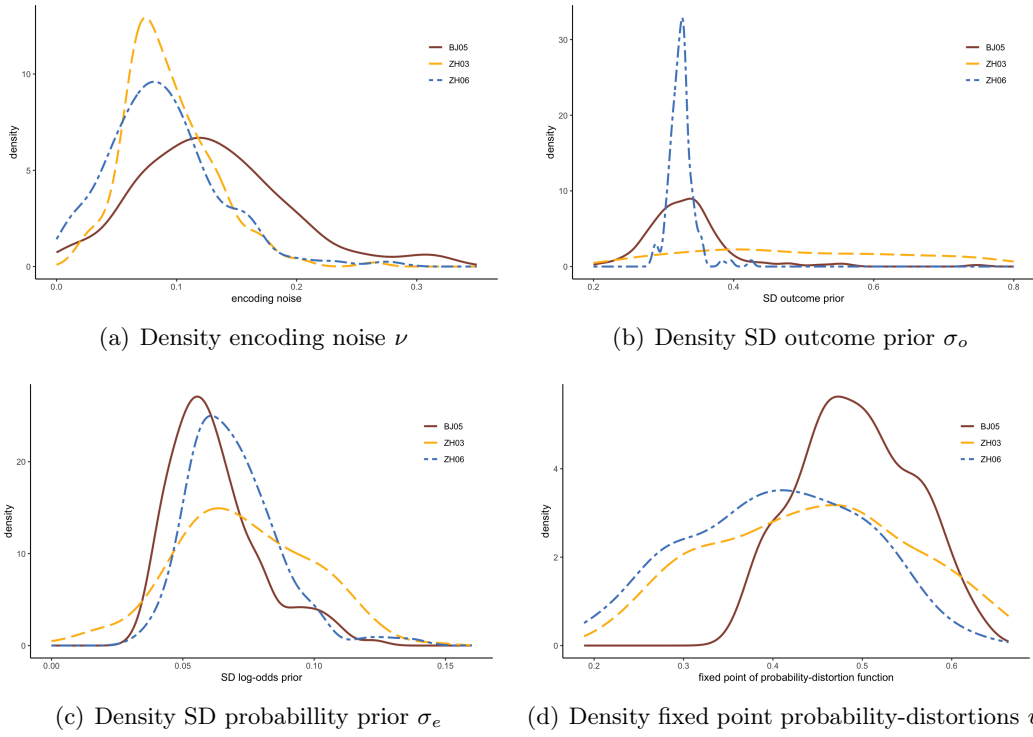


Figure 5: Density functions of NCM parameters for Bruhin et al. (2010)

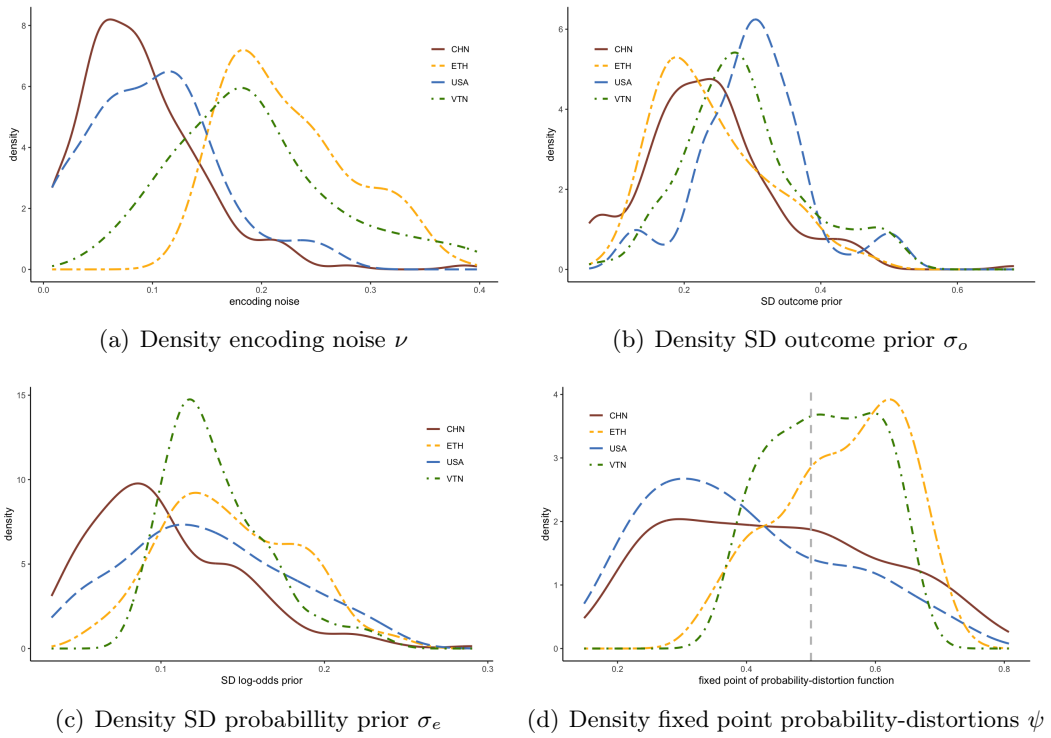


Figure 6: Density functions of NCM parameters for L'Haridon and Vieider (2019)

The NCM has higher predictive accuracy

Now that we know the distribution of the parameters, as well as the relation between the PT parameters and their NCM-generated counterparts, a natural next question concerns the predictive performance of the two models. Table 1 reports test statistics for the data of Bruhin et al. (2010), separately for their three experiments, and further separated for gains and losses. The NCM outperforms PT in all 6 test cases, and the difference in performance is large.

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

	gains	losses		gains	losses		gains	losses
Zurich 03	-16316.5	-16237.2	Zurich 06	-8745.5	-8654.3	Beijing 05	-5566.5	-5485.0
(SE)	(143.6)	(122.6)	(SE)	(137.9)	(88.1)	(SE)	(82.9)	(64.4)

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

The good performance of the NCM using the Bruhin et al. (2010) data is in no way exceptional. Table 2 provides tests of the predictive ability of the NCM versus PT using the data of L’Haridon and Vieider (2019). I conduct the tests country by country, and separately for gains and losses. Conditional on the design, this gives me 60 independent tests. Once again, the NCM easily outperforms PT in all 60 cases, with all tests significant at $p < 0.001$.

The superior predictive performance of the NCM as compared to PT may appear puzzling in the light of the fact that the NCM does not allow for discriminability values larger than 1, i.e $\alpha \leq 1$ and $\gamma \leq 1$. This is purely due to the fact that the NCM avoids the overfitting of noisy data, which remain a problem in the PT setup notwithstanding the hierarchical model geared towards discounting noisy outliers. This fact is best illustrated with an example. Take subject 319 in the global data of L’Haridon and Vieider (2019), from Ethiopia. His PT parameters indicate extreme likelihood-sensitivity, $\hat{\gamma} = 1.91$, joint to very low levels of optimism, at $\hat{\delta} = 0.16$. The noisiness of his decisions is relatively high, at $\hat{\omega} = 0.19$. Figure 7 shows the fit of this probability probability-distortion function to the raw data points, which are obtained under the assumption of linear utility. The fit appears to be terrible. At the same time, the fit of the function based on the NCM-derived PT parameters—with values $\gamma = 0.62$ and $\delta = 1.03$ —appears much more reasonable.

Table 2: Model fit of NCM versus PT

country	LEPD difference		country	LEPD difference	
	gains	losses		gains	losses
Australia	-2412.3	-2229.0	Kyrgyzstan	-3870.7	-3558.3
Belgium	-3602.2	-3311.4	Malaysia	-2556.9	-2341.7
Brazil	-3330.5	-2970.0	Nicaragua	-4855.3	-4385.9
Cambodia	-3222.9	-2957.3	Nigeria	-8151.8	-7393.2
Chile	-3812.2	-3475.3	Peru	-3831.7	-3471.3
China	-8125.3	-7443.3	Poland	-3566.2	-3235.7
Colombia	-4433.1	-3970.1	Russia	-2750.6	-2521.4
Costa Rica	-4225.3	-3824.1	Saudi Arabia	-2598.5	-2380.5
Czech Rep.	-3952.7	-3589.2	South Africa	-2824.2	-2581.3
Ethiopia	-5626.2	-5083.1	Spain	-3181.8	-2932.3
France	-3674.9	-3348.3	Thailand	-3172.8	-2785.0
Germany	-5155.8	-4708.7	Tunisia	-2961.8	-2701.8
Guatemala	-3370.6	-3042.8	UK	-3208.4	-2947.9
India	-3576.9	-3263.9	USA	-3844.4	-3488.4
Japan	-3381.2	-3018.4	Vietnam	-3457.4	-3130.4

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

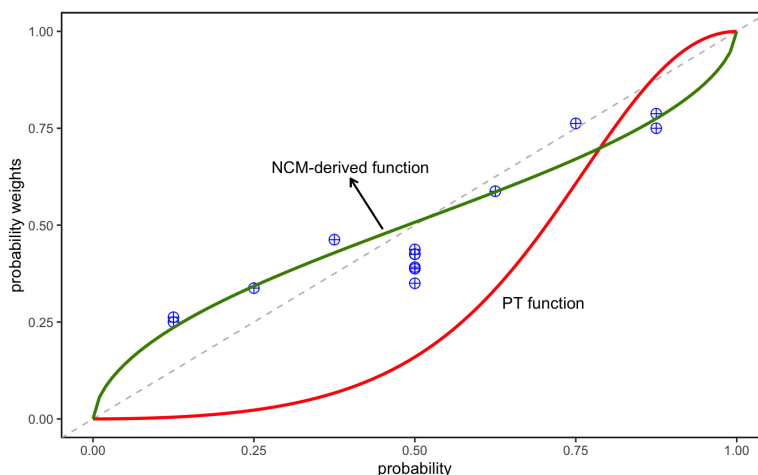


Figure 7: Example of probability-distortion fit for PT vs the NCM

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

The underlying reason for this bad fit of the PT function can be traced back to the other parameters of the model. At $\hat{\alpha} = 3.74$, outcome-sensitivity seems extreme. This extreme convexity of the utility function is derived mostly from the 5 data points at $p = 0.5$, which include 3 prospects with non-zero lower outcomes (two points at $p = 0.125$ and $p = 0.875$ also contain non-zero lower outcomes). Even upon casual inspection, these data points can be seen to be outliers from the overall trend, yet they

have a disproportionate effect on all the model parameters estimated under PT. The outcome-discriminability parameter under the NCM, on the other hand, has a much more reasonable value of $\alpha = 0.67$. Taking into account the role of α in contributing to noise, this results in much better predictive performance, as well as improving the fit of the probability-distortion function to the data. Section S4 in the supplementary materials shows many more example of individual-level data fit.

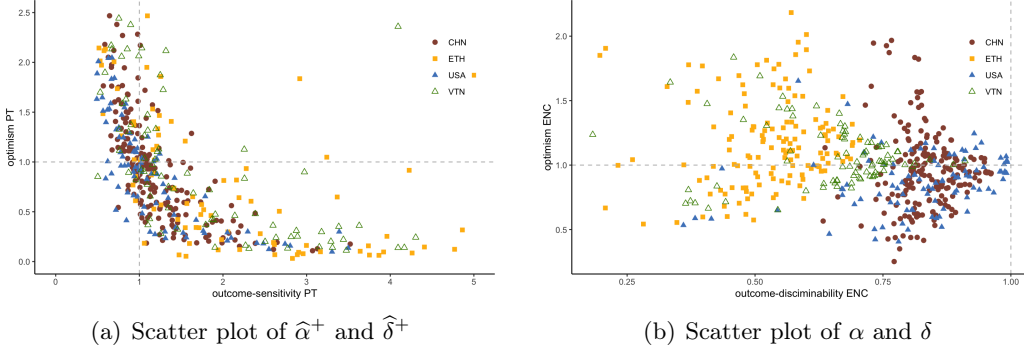


Figure 8: Scatter plot of outcome-distortion and optimism parameters for L’Haridon and Vieider (2019)

The pattern in the example just shown is far from exceptional. Indeed, the poor fit of PT, and its even poorer performance in terms of predictive accuracy, is largely driven by the poor separability of some of its parameters, and their assumed orthogonality to noise. While I have already documented some of these patterns above, it is nevertheless instructive to take a look at one particularly problematic dimension—the correlation between outcome-sensitivity $\hat{\alpha}$ and optimism $\hat{\delta}$ under PT. This correlation is shown in panel 8(a) of figure 8. It can be seen to be L-shaped, with large values of outcome-sensitivity corresponding to low levels of optimism, and vice versa. Since values of outcome-sensitivity larger than 1 are further associated with noise, this pattern points at poor separability of the parameters in a traditional PT setup. Zeisberger, Vrecko and Langer (2012) discuss this issue at some length for a 1-parameter specification of the probability-distortion function, with the increased degrees of freedom in the 2-parameter version likely to make this issue much worse. Panel 8(b) shows the correlation between the equivalent parameters α and δ derived from the NCM for comparison. In this setup, the two parameters indeed appear to be orthogonal to each other.

Is coding efficient?

My analysis so far was based on the noisy coding model I have derived. An interesting

additional question, however, concerns whether behaviour may also conform to *efficient coding* predictions, according to which the noisiness in the likelihood function would optimally adapt to the prior given constraints on cognitive resources. To this end, I use the prediction formally derived by [Frydman and Jin \(2021\)](#), who show that under an efficient coding setup the encoding noise ought to covary with the standard deviation of the prior. While any such correlation in the data I use would not be causal—there is no random variation in the choice stimuli such as the one used by [Frydman and Jin \(2021\)](#), and the variance in the stimuli is fixed—such a correlation would nevertheless be indicative of the coding noise being efficient, since different *perceptions* of volatility in the environment also ought to result in different encoding noise.

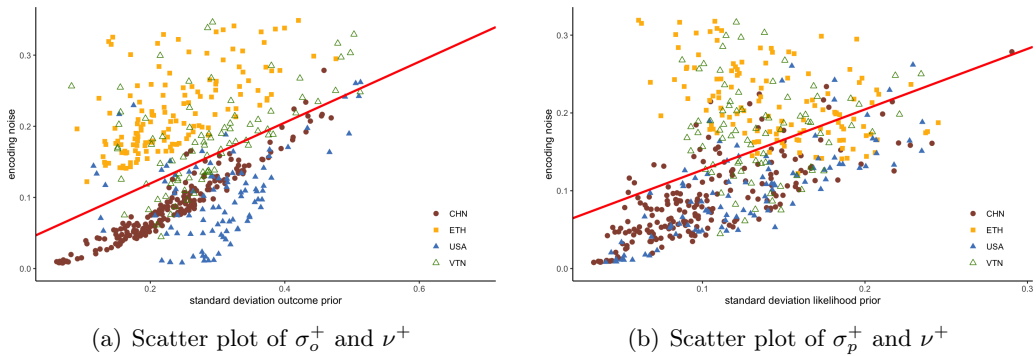


Figure 9: Scatter plot of prior SDs against encoding noise [L’Haridon and Vieider \(2019\)](#)

Figure 9 shows correlation between the standard deviations of the outcome prior (panel 9(a)) and of the likelihood prior (panel 9(b)) with encoding noise. For parsimony, I only show the patterns for the usual four countries in the [L’Haridon and Vieider \(2019\)](#) data, and only for gains. Patterns for losses, the global data, or the data of [Bruhin et al. \(2010\)](#) are qualitatively similar. We observe a clear positive correlation between the standard deviation of the outcome prior and encoding noise ($\rho = 0.410, p < 0.001$), as well as between the standard deviation of the likelihood prior and encoding noise ($\rho = 0.513, p < 0.001$). This provides suggestive evidence that the coding noise on which the observed choice patterns are based is indeed *efficient*.

3.3 Ellsberg urns and separability violations

From the discussion above it follows immediately that the NCM predicts ambiguity-insensitivity, as well as allowing for ambiguity aversion. These patterns are well-documented

in the literature, and can be fit well by PT. The NCM, however, also makes additional predictions. For one, we would expect coding noise, which I denote by ν_a to emphasize that it is specific to ambiguity, to be negatively correlated with ambiguity-insensitivity. In addition, we would expect coding noise to be substantially larger than under risk, so that $\nu_a > \nu_r$, where the subscript r indicates known probabilities.

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

To test this prediction, I use an original dataset that contains a richer choice setup. The data contain observations for 47 subjects, and the data structure and experimental procedures closely follow those in L’Haridon et al. (2018). The stimuli are, however, richer in that the stimuli for risk are replicated exactly for ambiguity, including variation over outcomes and non-zero lower outcomes for both risk and ambiguity, thus allowing for the identification of utility curvature and outcome-discriminability under ambiguity as well as risk. Further details about the experiment and estimations are reported in section S5 of the supplementary materials.

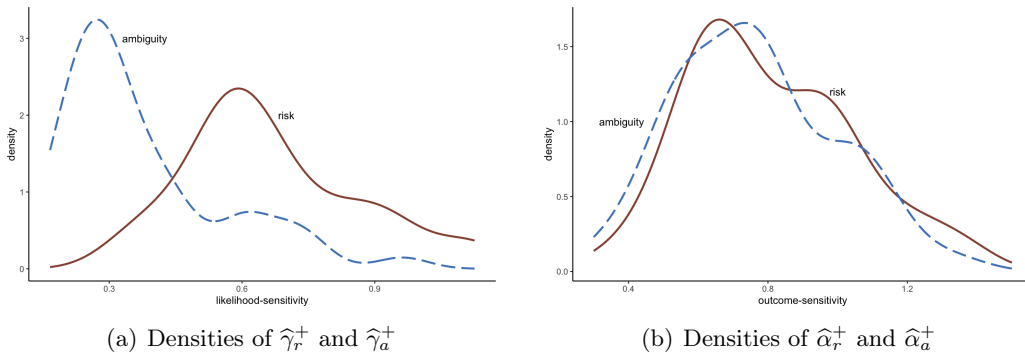


Figure 10: Densities of $\hat{\gamma}$ and $\hat{\alpha}$ for risk and ambiguity

I start by testing the performance of the different models against each-other. The NCM easily outperforms the PT specification (elpd difference: -3521.4 , SE: 21.3). Figure 10 shows density plots for $\hat{\gamma}$ and $\hat{\alpha}$, for risk and ambiguity. The difference between risk and ambiguity for $\hat{\gamma}$, shown in panel 10(a), is sizeable. This sort of lowering of likelihood-sensitivity under ambiguity relative to risk, termed ambiguity-insensitivity, is well-documented in the literature (Abdellaoui et al., 2011; Dimmock et al., 2015; Trautmann and van de Kuilen, 2015; L’Haridon et al., 2018). The difference between outcome sensitivity for risk and ambiguity, displayed in panel 10(b), is much less pronounced. Nevertheless, the distribution under ambiguity appears to be slightly shifted to the left. This is confirmed by a Mann-Whitney test, indicating that outcome-sensitivity is indeed reduced under ambiguity ($p = 0.012$), thus resulting in a violation of PT.

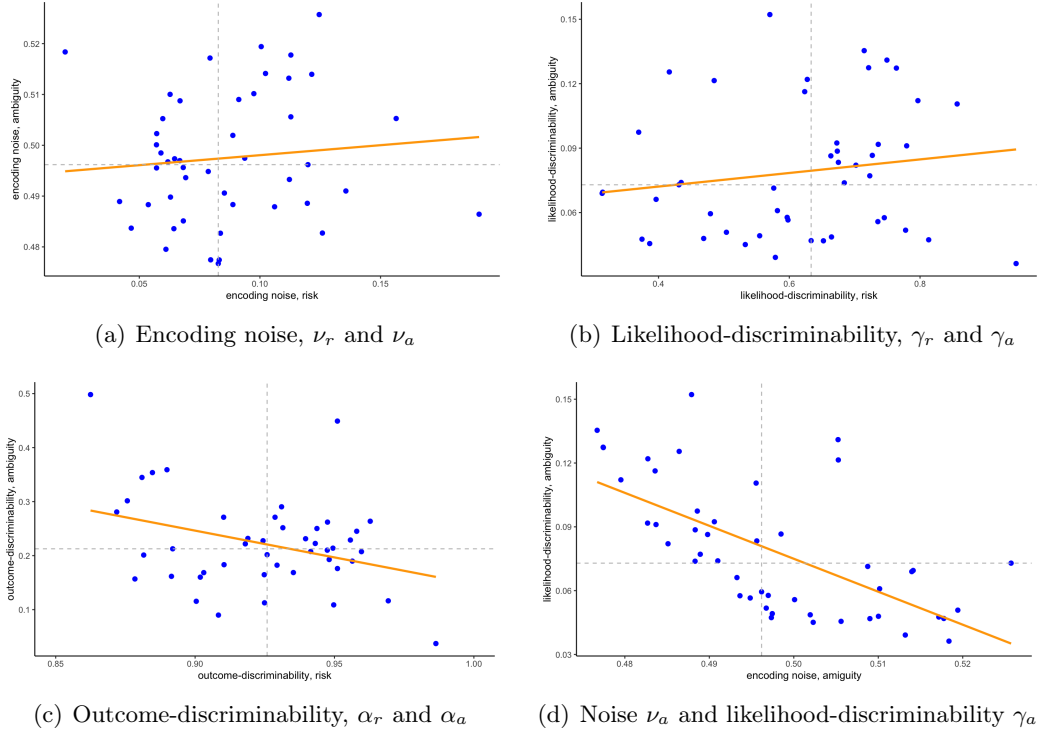


Figure 11: NCM parameters under risk and ambiguity

Let us now take a look at the patterns emerging from the NCM, displayed in figure 11. Panel 11(a) shows a scatter plot of encoding noise for risk and ambiguity. The two are not significantly correlated. In accordance with the predictions discussed above, however, encoding noise for ambiguity is much larger than encoding noise for risk. This is in turn reflected in the other model parameters. In particular, likelihood-

discriminability γ is much lower under ambiguity than under risk, as is apparent from panel 11(b). Similar results obtain for outcome-discriminability α , shown in panel 11(c). Once again discriminability is much lower under ambiguity than it is under risk. Notice also that this difference is much more pronounced than it was under PT, which can be traced back to the different definition of outcome distortions and its interaction with encoding noise. Finally, panel 11(d) shows the correlation between encoding noise and likelihood discriminability under ambiguity. The two are strongly negatively correlated ($\rho = -0.711, p < 0.001$), as is predicted by the NCM.

4 Discussion

Given the superior performance of the NCM compared to PT, an important question concerns the relative status of the two models. The NCM can be seen as a generative model for PT, given that PT functionals naturally emerge from the noisy neural coding of stimuli depicting choices under uncertainty and their mental decoding. In this sense, the NCM provides causal underpinnings for PT. This means in first instance that theoretical models based on PT remain valid, seen how the qualitative PT patterns are reaffirmed by the NCM. At the same time, however, the superior predictive performance of the NCM and the violations of the strict separability precepts of PT emerging from the NCM suggest that when it comes to data fitting and prediction, working directly with the NCM may be preferable, even when the ultimate goal is to interpret PT parameters, which can easily be derived from the native NCM setup.

In empirical analysis it may furthermore be desirable to directly work with the native NCM parameters. The latter provide a clear division of parameters pertaining to noisy cognition, as captured by the noisy mental encoding parameter, and those pertaining to the mental prior. This separation may help overcome issues that have hampered empirical analysis using PT, where several parameters usually co-vary in regression analysis, making an interpretation of their effects fraught with difficulties (see online appendix of L'Haridon and Vieider, 2019, for a number of examples). The cleaner separation of cognitive components, issues pertaining to the mental prior (the location of the fixed point of the likelihood-distortion function), and the confidence in this prior (the standard deviations associated with the prior means) ought to allow for a reconsideration of the effects of observable characteristics on behaviour under risk.

The model I have presented builds on the model of [Khaw et al. \(2020\)](#), and generalizes it along several dimensions. The latter used a noisy coding model to account for apparent risk aversion for small monetary stakes as typically used in experiments, while assuming perfect accuracy in the mental representation of probabilities, which thus enter their main equations linearly. While they also showed that a probability-distortion function can be derived from their setup, they showed this in a dual setup, and their result only holds in the limit as $\nu_p, \sigma_p \rightarrow 0$ and for a fixed ratio $\frac{\nu_p}{\sigma_p} \equiv \kappa$. In contrast, the derivation I presented here applies to uncertainty as well as risk, imposes no limits on the parameters, is based on the simultaneous derivation of outcome-distortions and probability-distortions, and is applicable to prospects with non-zero lower outcomes. These generalizations were indeed crucial for the data fitting exercise I conducted in this paper.

[Khaw et al. \(2020\)](#) discuss how loss aversion can be produced by a different prior mean for losses versus gains, or by differences in the prior variance. Loss aversion could then arise if either the mean or the variance of the prior, or both, are larger for losses than for gains. However, this assumption would need to hold across a variety of environments, some of them characterized by rapid economic growth, given the near-universal evidence for loss aversion ([Rieger, Wang and Hens, 2017](#); [L'Haridon and Vieider, 2019](#); [Brown, Imai, Vieider and Camerer, 2019](#)). It may thus seem unlikely that losses would be considered to be systematically larger than gains. This explanation would also seem to contradict rapid adaptation to an experimental environment as documented by [Frydman and Jin \(2021\)](#), since the average losses in experiments are usually chosen to be smaller than the average gains to limit the financial exposure of experimental subjects.

Fortunately, the two accounts of loss aversion can be distinguished experimentally. If loss aversion is driven by differential priors for gains and losses, it ought to be possible to systematically manipulate such priors in an experiment. If, on the other hand, loss aversion enters through a different encoding process as modelled in this paper, then loss aversion is predicted to increase in stake size due to the different degree of nonlinearity of the encoding functions for gains and losses. [Ert and Erev \(2013\)](#) present experimental evidence suggesting that this may indeed be the case. In their study 2, they show that a large majority of subjects shy away from 50-50 bets offering the possibility to gain 100 Shekels or else lose 100 Shekels. When the stakes are reduced to gains and losses of 10 Shekels, however, this aversion to mixed fair bets disappears. In their study 1a, they further compare the acceptance of a 50-50 bet over a gain or loss of 10 Shekels to a

condition involving the same actual payoffs, but where the nominal payoffs are presented in an experimental currency that inflates the amounts 100-fold. Once again, they find no evidence for loss aversion in the low-stake condition. However, they do find evidence for loss aversion for high nominal stakes, even though the actual stakes have not changed. Together, this provides evidence not only that loss aversion increases in stake size, but that the perception of nominal numerical amounts may play a role in this process.

A more philosophical difference between my approach and the one of [Khaw et al. \(2020\)](#) concerns the origin of preferences themselves. [Khaw et al. \(2020\)](#) mostly focus on ‘apparent preferences’ emerging from noise in perception. My own derivations, on the other hand, suggest that systematic differences between individuals may well exist beyond these cognitive components, even in a context where the optimal decision rule forming the backbone of the mental representations is based on expected value maximization.¹⁴ That indeed emerges from the model, since the prior mean, ψ , is unaffected by issues surrounding noisy cognition. The clear separation between the cognitive components of the model and the mean of the prior, however, promises to shed fresh light on a number of empirical correlations that have been documented in the literature on risk taking. It also means that a link is restored between the investigation of the correlates of risk attitudes and the modelling of decisions under risk, which was much less obvious given the traditional interpretation of PT parameters, so that the modelling of preferences and the investigation of the correlates of risk attitudes have largely proceeded in parallel ([Friedman, Isaac, James and Sunder, 2014](#)).

This distinction, of course, raises important questions about what may determine the prior itself. The neuroscience literature conceptually distinguishes two parts of the mental prior. One can quickly adapt to a given situation, and is thus subject to experimental manipulation and study. Effects of priming, such as documented e.g. by [Cohn, Engelmann, Fehr and Maréchal \(2015\)](#), may be a manifestation of such an effect. The other consists in a more stable component, which is much less studied in neuroscience. Conceivably, this component may be influenced by long-term experiences ([Malmendier and Nagel, 2011](#)), shaped by conscious education efforts ([Doepke and Zilibotti, 2014; 2017](#)), forged by environmental adaptation ([Di Falco and Vieider, 2021](#)), or more generally be determined by intergenerational transmission ([Galor and Michalopoulos, 2012;](#)

¹⁴It is further possible that the optimal decision rule may itself deviate from expected value maximization, e.g. when stakes get very large. [Khaw et al. \(2020\)](#) discuss this possibility in their footnote 21, and the same caveat applies to the setup used in this paper.

[Bouchouicha and Vieider, 2019](#)). Investigating the process under which the more stable components of the mental prior is formed ought to be a priority area for future research.

The model proposed in this paper is more loosely linked to other models predicting that preferences may adapt to the environment. Most notable is the connection to the evolutionary models of [Robson \(2001\)](#) and [Netzer \(2009\)](#), who use a fitness-maximizing model predicated on limited discernibility of outcomes to derive a Bernoulli utility function that adapts to the local environment. A common element is that both models incorporate the intuition of just noticeable differences in utility, although in the NCM this is a result of the compressed mental representation of stimuli, while it is a starting point in the Robson-Netzer framework, where it results in utility taking the form of a step function. Beyond that, the Robson-Netzer model addresses mostly the outcome dimension, whereas the NCM derived in this paper addresses both the outcome and the likelihood dimension, and provides an integrated mechanism of stochastic choice.

I have applied the model derived in this paper to data involving risk, or known probabilities, and data involving ambiguity or vague probabilities, obtained using Ellsberg urns. One of the reasons for this was data availability—data rich enough to identify the NCM are not available for natural uncertainty, due to the more restrictive assumptions under which they have been obtained. An additional reason concerns the way in which signals are translated into subjective probability judgements—an issue we could eschew for risk and ambiguity implemented through Ellsberg urns. PT papers typically measure these subjective probabilities in some form, and then use them as an input for the decision model ([Abdellaoui et al., 2011](#); [Baillon, Bleichrodt, Keskin, L’Haridon and Li, 2018a](#)), or circumvent the measurement of subjective probabilities entirely ([Baillon et al., 2018b](#)). Given the signal-detection approach to the problem underlying the NCM, however, it seems desirable to complement the model with a learning process based on the same neural mechanisms as the ones at work in the model itself.

The general framework based on the noisy neural coding of stimuli—and its extension to efficient coding schemes that allow for the endogenization of the likelihood function—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 risky choices. 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. These implications deserve considerable attention, and are thus left for future work.

5 Conclusion

I proposed a model of decision making under uncertainty resting on the noisy coding of external stimuli. These stimuli are subsequently optimally decoded using a mental prior to result in actionable decision inputs. I showed that this setup results in a linear in log-odds specification of a popular probability-distortion function, whereby subjectively distorted likelihood ratios are traded off against a subjectively distorted cost-benefit ratio. The noisy neural coding model derived in this paper can thus be seen as a generative model of prospect theory preferences. At the same time, the model makes predictions that suggest that prospect theory will be violated in certain situations due to its strong separability precepts between outcome and probability representations. Taking the model to a variety of data sets, I have shown that the noisy coding model outperforms prospect theory significantly in terms of predictive performance. This superior performance of the noisy coding model can be traced back to its stochastic nature. While under prospect theory decision noise is added onto the deterministic model and is generally assumed to be orthogonal to the latter, in the noisy coding model decision noise and the core model parameters result from one and the same process, thus being inextricably linked. This interaction between noise and decision parameters appears to be a central feature of data on decisions under uncertainty, explaining the superior performance of the noisy coding model.

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
- Baillon, Aurélien, Han Bleichrodt, Umut Keskin, Olivier L’Haridon, and Chen Li (2018a) ‘The effect of learning on ambiguity attitudes.’ *Management Science* 64(5), 2181–2198
- Baillon, Aurélien, Zhenxing Huang, Asli Selim, and Peter P Wakker (2018b) ‘Measuring ambiguity attitudes for all (natural) events.’ *Econometrica*
- Barberis, Nicholas C. (2013) ‘Thirty Years of Prospect Theory in Economics: A Review and Assessment.’ *Journal of Economic Perspectives* 27(1), 173–96
- Billock, Vincent A, and Paul R Havig (2018) ‘A simple power law governs many sensory amplifications and multisensory enhancements.’ *Scientific reports* 8(1), 1–7
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2012) ‘Salience theory of choice under risk.’ *The Quarterly Journal of Economics* 127(3), 1243–1285
- Bouchouicha, Ranoua, and Ferdinand M. Vieider (2017) ‘Accommodating stake effects under prospect theory.’ *Journal of Risk and Uncertainty* 55(1), 1–28
- (2019) ‘Growth, Entrepreneurship, and Risk-Tolerance: A Risk-Income Paradox.’ *Journal of Economic Growth* 24(3), 257–282
- Brown, Alexander L., Taisuke Imai, Ferdinand M. Vieider, and Colin Camerer (2019) ‘Meta-analysis of empirical estimates of loss aversion.’ *Working Paper, MetaArXiv*
- Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper (2010) ‘Risk and Rationality: Uncovering Heterogeneity in Probability Distortion.’ *Econometrica* 78(4), 1375–1412
- Carpenter, Bob, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell (2017) ‘Stan: A probabilistic programming language.’ *Journal of Statistical Software*
- 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)
- Di Falco, Salvatore, and Ferdinand M. Vieider (2021) ‘Environmental adaptation of risk preferences.’ *Mimeo*

- Dimmock, Stephen G., Roy Kouwenberg, and Peter P. Wakker (2015) ‘Ambiguity Attitudes in a Large Representative Sample.’ *Management Science* 62(5), 1363–1380
- Doepke, Matthias, and Fabrizio Zilibotti (2014) ‘Culture, Entrepreneurship, and Growth.’ In ‘Handbook of Economic Growth,’ vol. 2
- (2017) ‘Parenting with style: Altruism and paternalism in intergenerational preference transmission.’ *Econometrica* 85(5), 1331–1371
- Doya, Kenji, Shin Ishii, Alexandre Pouget, and Rajesh PN Rao (2007) *Bayesian brain: Probabilistic approaches to neural coding* (MIT press)
- Ellsberg, Daniel (1961) ‘Risk, Ambiguity and the Savage Axioms.’ *Quarterly Journal of Economics* 75(4), 643–669
- Ert, Eyal, and Ido Erev (2013) ‘On the descriptive value of loss aversion in decisions under risk: Six clarifications.’ *Judgment and Decision Making* 8(3), 214–235
- Fechner, Gustav Theodor (1860) ‘Elements of psychophysics, 1860.’
- Fehr-Duda, Helga, Adrian Bruhin, Thomas F. Epper, and Renate Schubert (2010) ‘Rationality on the Rise: Why Relative Risk Aversion Increases with Stake Size.’ *Journal of Risk and Uncertainty* 40(2), 147–180
- Fehr-Duda, Helga, and Thomas Epper (2012) ‘Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences.’ *Annual Review of Economics* 4(1), 567–593
- Fox, Craig R., and Amos Tversky (1995) ‘Ambiguity Aversion and Comparative Ignorance.’ *Quarterly Journal of Economics* 110(3), 585–603
- Friedman, Daniel, R. Mark Isaac, Duncan James, and Shyam Sunder (2014) *Risky Curves: On the Empirical Failure of Expected Utility* (London and New York: Routledge)
- Frisch, Deborah, and Jonathan Baron (1988) ‘Ambiguity and rationality.’ *Journal of Behavioral Decision Making* 1(3), 149–157
- Frydman, Cary, and Lawrence J Jin (2021) ‘Efficient coding and risky choice.’ *Quarterly Journal of Economics*, forthcoming
- Galor, Oded, and Stelios Michalopoulos (2012) ‘Evolution and the Growth Process: Natural Selection of Entrepreneurial Traits.’ *Journal of Economic Theory* 147(2), 759–780
- Ganguli, Deep, and Eero P Simoncelli (2014) ‘Efficient sensory encoding and bayesian

- inference with heterogeneous neural populations.’ *Neural computation* 26(10), 2103–2134
- Gelman, Andrew, and Jennifer Hill (2006) *Data analysis using regression and multi-level/hierarchical models* (Cambridge university press)
- Gelman, Andrew, Jessica Hwang, and Aki Vehtari (2014a) ‘Understanding predictive information criteria for bayesian models.’ *Statistics and computing* 24(6), 997–1016
- Gelman, Andrew, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin (2014b) *Bayesian data analysis*, vol. 2 (CRC press Boca Raton, FL)
- Gold, Joshua I, and Michael N Shadlen (2001) ‘Neural computations that underlie decisions about sensory stimuli.’ *Trends in cognitive sciences* 5(1), 10–16
- (2002) ‘Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward.’ *Neuron* 36(2), 299–308
- Goldstein, W, and H Einhorn (1987) ‘Expression Theory and the Preference Reversal Phenomena.’ *Psychological Review* 94, 236–254
- Gonzalez, Richard, and George Wu (1999) ‘On the Shape of the Probability Weighting Function.’ *Cognitive Psychology* 38, 129–166
- Gul, Faruk (1991) ‘A Theory of Disappointment Aversion.’ *Econometrica* 59(3), 667
- Harless, David, and Colin F. Camerer (1994) ‘The Predictive Utility of Generalized Expected Utility Theories.’ *Econometrica* 62(6), 1251–1289
- Heath, C, and A Tversky (1991) ‘Preference and belief: Ambiguity and competence in choice under uncertainty. Journal of Risk and.’ *Journal of Risk and Uncertainty* 4(1), 5–28
- Heng, Joseph A, Michael Woodford, and Rafael Polania (2020) ‘Efficient sampling and noisy decisions.’ *Elife* 9, e54962
- Hey, John D., and Chris Orme (1994) ‘Investigating Generalizations of Expected Utility Theory Using Experimental Data.’ *Econometrica* 62(6), 1291–1326
- Hogarth, Robin M., and Hillel J. Einhorn (1990) ‘Venture Theory: A Model of Decision Weights.’ *Management Science* 36(7), 780–803
- Holt, Charles A., and Susan K. Laury (2002) ‘Risk Aversion and Incentive Effects.’ *American Economic Review* 92(5), 1644–1655
- Kahneman, Daniel, and Amos Tversky (1979) ‘Prospect Theory: An Analysis of Decision under Risk.’ *Econometrica* 47(2), 263 – 291

- Khaw, Mel Win, Ziang Li, and Michael Woodford (2020) ‘Cognitive imprecision and small-stakes risk aversion.’ *The Review of Economic Studies*, forthcoming
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005) ‘A Smooth Model of Decision Making under Ambiguity.’ *Econometrica* 73(6), 1849–1892
- Knill, David C, and Alexandre Pouget (2004) ‘The bayesian brain: the role of uncertainty in neural coding and computation.’ *TRENDS in Neurosciences* 27(12), 712–719
- 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
- Loomes, Graham (2005) ‘Modelling the Stochastic Component of Behaviour in Experiments: Some Issues for the Interpretation of Data.’ *Experimental Economics* 8(4), 301–323
- Loomes, Graham, and Robert Sugden (1982) ‘Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty.’ *The Economic Journal* 92(368), 805–824
- (1986) ‘Disappointment and Dynamic Consistency in Choice under Uncertainty.’ *The Review of Economic Studies* 53(2), 271–282
- Malmendier, Ulrike, and Stefan Nagel (2011) ‘Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?’ *The Quarterly Journal of Economics* 126(1), 373–416
- Markowitz, Harry (1952) ‘The Utility of Wealth.’ *Journal of Political Economy* 60(2), 151–158
- McElreath, Richard (2016) *Statistical Rethinking: A Bayesian Course with Examples in R and Stan* (Academic Press)
- Murphy, Kevin P (2012) *Machine learning: a probabilistic perspective* (MIT press)
- Netzer, Nick (2009) ‘Evolution of time preferences and attitudes toward risk.’ *American Economic Review* 99(3), 937–55
- Payzan-LeNestour, Elise, and Michael Woodford (2021) ‘Outlier blindness: A neurobiological foundation for neglect of financial risk.’ *Journal of Financial Economics*
- Petzschner, Frederike H, Stefan Glasauer, and Klaas E Stephan (2015) ‘A bayesian per-

- spective on magnitude estimation.’ *Trends in cognitive sciences* 19(5), 285–293
- Polania, Rafael, Michael Woodford, and Christian C Ruff (2019) ‘Efficient coding of subjective value.’ *Nature neuroscience* 22(1), 134–142
- Prelec, Drazen (1998) ‘The Probability Weighting Function.’ *Econometrica* 66, 497–527
- Quiggin, John (1982) ‘A theory of anticipated utility.’ *Journal of Economic Behavior & Organization* 3(4), 323–343
- Rabin, Matthew (2000) ‘Risk Aversion and Expected Utility Theory: A Calibration Theorem.’ *Econometrica* 68, 1281–1292
- Rieger, Marc Oliver, Mei Wang, and Thorsten Hens (2017) ‘Estimating cumulative prospect theory parameters from an international survey.’ *Theory and Decision* 82(4), 567–596
- Robson, Arthur J (2001) ‘The biological basis of economic behavior.’ *Journal of Economic Literature* 39(1), 11–33
- Stan Development Team (2017) ‘RStan: the R interface to Stan.’ R package version 2.17.2
- Starmer, Chris (2000) ‘Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk.’ *Journal of Economic Literature* 38, 332–382
- Stevens, So S (1970) ‘Neural events and the psychophysical law.’ *Science* 170(3962), 1043–1050
- Summerfield, Christopher, and Konstantinos Tsetsos (2015) ‘Do humans make good decisions?’ *Trends in cognitive sciences* 19(1), 27–34
- Tom, Sabrina M., Craig R. Fox, Christopher Trepel, and Russell A. Poldrack (2007) ‘The Neural Basis of Loss Aversion in Decision-Making Under Risk.’ *Science* 315(5811), 515–518
- Trautmann, Stefan T., and Gijs van de Kuilen (2015) ‘Ambiguity Attitudes.’ In ‘The Wiley Blackwell Handbook of Judgment and Decision Making’ (Wiley Blackwell)
- Tversky, Amos, and Daniel Kahneman (1992) ‘Advances in Prospect Theory: Cumulative Representation of Uncertainty.’ *Journal of Risk and Uncertainty* 5, 297–323
- Tversky, Amos, and Peter P. Wakker (1995) ‘Risk Attitudes and Decision Weights.’ *Econometrica* 63(6), 1255–1280
- 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

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*

Wei, Xue-Xin, and Alan A Stocker (2015) ‘A bayesian observer model constrained by efficient coding can explain ‘anti-bayesian’ percepts.’ *Nature neuroscience* 18(10), 1509–1517

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

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

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

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

SUPPLEMENTARY MATERIALS (For online publication)

Efficient Neural Coding and Decisions under Uncertainty

S1 Estimation details

S1.1 The Bayesian hierarchical model

I estimate individual-level parameters using Bayesian hierarchical models (Gelman and Hill, 2006; Gelman et al., 2014b; McElreath, 2016). Take a parameter vector θ_i , indicating individual-level parameters. This vector follows the following distribution:

$$\theta_i \sim \mathcal{N}(\bar{\theta}, \Sigma), \quad (26)$$

where $\bar{\theta}$ 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 check convergence by examining the R-hat statistics, and checking for divergent iterations. The estimation algorithm was developed using simulations, and making sure that I could recover individual-level parameters from those simulations. Details about the models and the priors used are included below for the individual datasets.

S2 Additional results for Bruhin et al. (2010)

S2.1 Estimation details

The estimations use directly the density around the observed switching point, following the econometric approach in the original paper. For the PT model, this takes the following form:

$$ce \sim \mathcal{N}(u^{-1}[w(p)u(x) + (1 - w(p))u(y)], \hat{\omega}^2 \times |x - y|), \quad (27)$$

with $u(x) = x^{\hat{\alpha}}$ and $w(p) = \frac{\hat{\delta}p^{\hat{\gamma}}}{\hat{\delta}p^{\hat{\gamma}} + (1-p)^{\hat{\gamma}}}$, and where multiplying the variance $\hat{\omega}$ by $|x - y|$ allows for heteroscedasticity across choice lists with different step sizes between choices,

following the approach in the original paper.

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 $\hat{\gamma}$ 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 individual 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 affect the estimated parameters in any way, showing that the prior has only a minimal influence on the ultimate parameter estimates.

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

$$dw \sim \mathcal{N}\left(\frac{\delta_{\alpha}^{\frac{1}{2}} p_{\alpha}^{\frac{\gamma}{\alpha}}}{\delta_{\alpha}^{\frac{1}{2}} p_{\alpha}^{\frac{\gamma}{\alpha}} + (1-p)^{\frac{\gamma}{\alpha}}}, \omega^2\right), \quad (28)$$

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

S2.2 Additional results for gains

This section adds some further results to the analysis of the data for gains presented in the main text. I start by presenting correlations between other parameters. Panel 12(a) visualizes the correlations between $\hat{\alpha}$ and $\hat{\gamma}$ already discussed in the main text.

Panel 12(b) further shows correlational patterns between $\hat{\gamma}$ and $\hat{\delta}$. According to the NCM, the results should depend on the initial value of ψ . In particular, the larger the value of $\hat{\gamma}$, the closer the value of $\hat{\delta}$ should be compressed towards 1. This is exactly what we observe. In experiments where we observe preponderantly values of $\hat{\delta} < 1$, the distance to 1 decreases as $\hat{\gamma}$ increases, thus resulting in a positive correlation between the two parameters. This is the case in the Zurich 03 experiment ($\rho = 0.262, p < 0.001$),

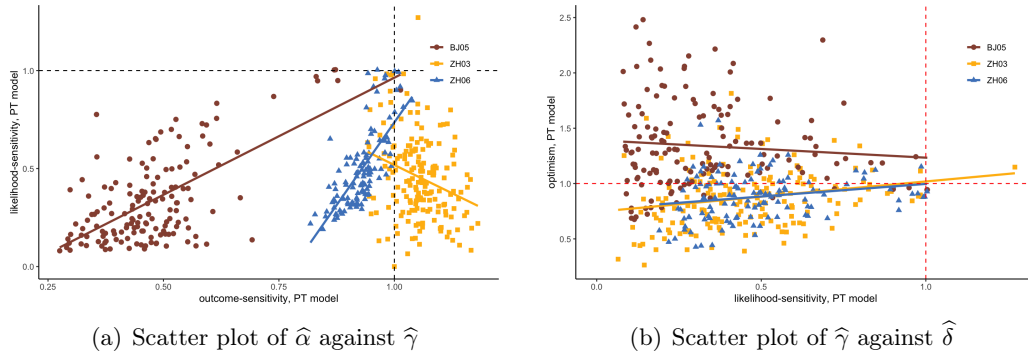


Figure S12: Correlations between PT parameters in Bruhin et al. (2010)

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 $\hat{\delta} > 1$. We may thus expect a negative relationship with $\hat{\gamma}$. 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.

S2.3 Results for losses

In this section, I replicate the results reported in the main text for Bruhin et al. (2010) for losses. Since several results for losses are already included in the main text, this mainly concerns the densities of the NCM parameters. Figure S13 thus shows the density functions for the 4 NCM parameters estimated from the data of Bruhin et al. (2010). The parameters resemble each other for the different experiments. The largest differences between experiments emerge once more for the standard deviation of the outcome prior, σ_o^- , and for the fixed point of the probability-distortion function, ψ^- . The first has a relatively broad distribution in the two experiments in Zurich. The second tends towards higher values in the Beijing 05 experiment, indicating higher fixed points of the probability-distortion function, just as observed for gains. Encoding noise is somewhat higher in Zurich 06 and Beijing 05 when compared to Zurich 03.

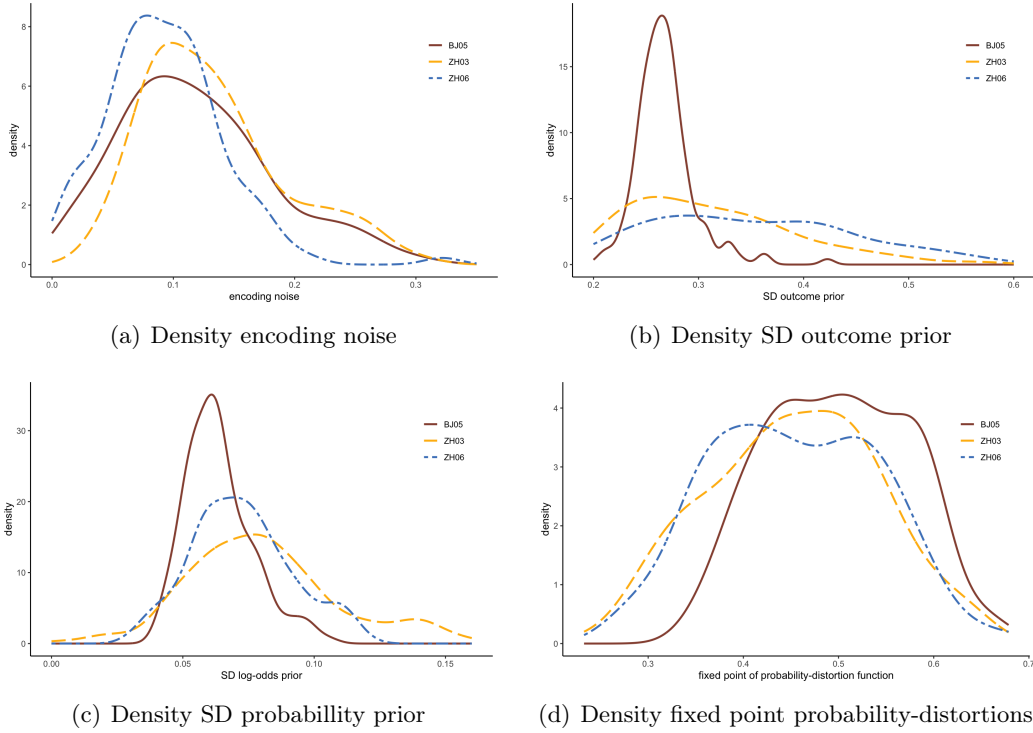


Figure S13: Density functions of NCM parameters for losses in Bruhin et al. (2010)

S3 Additional results for L'Haridon and Vieider (2019)

S3.1 Estimation details

The estimations follow the same econometric approach as those for Bruhin et al. (2010). The details reported below are thus identical to those reported for the latter paper above, and are reproduced here purely for convenience.

The estimations use directly the density around the observed switching point, following the econometric approach in the original paper. For the PT model, this takes the following form:

$$ce \sim \mathcal{N}(u^{-1}[w(p)u(x) + (1 - w(p))u(y)], \hat{\omega}^2 \times |high - low|), \quad (29)$$

with $u(x) = x^{\hat{\alpha}}$ and $w(p) = \frac{\hat{\delta}p^{\hat{\gamma}}}{\hat{\delta}p^{\hat{\gamma}} + (1-p)^{\hat{\gamma}}}$, and where multiplying the variance $\hat{\omega}$ by $|x - y|$ allows for heteroscedasticity across choice lists with different step sizes between choices, following the approach in the original paper.

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 $\hat{\gamma}$ 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 individual 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 affect the estimated parameters in any way, showing that the prior has only a minimal influence on the ultimate parameter estimates.

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

$$dw \sim \mathcal{N}\left(\frac{\delta_{\alpha}^{\frac{1}{\alpha}} p_{\alpha}^{\frac{\gamma}{\alpha}}}{\delta_{\alpha}^{\frac{1}{\alpha}} p_{\alpha}^{\frac{\gamma}{\alpha}} + (1-p)^{\frac{\gamma}{\alpha}}}, \omega^2\right), \quad (30)$$

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

S3.2 Correlations of PT parameters

I start from documenting correlations in the global parameters, for all 30 countries and 3000 subjects. Different countries are not distinguished, since there are not enough colours and symbols to render such a distinction intelligible. Correlations are tested on parameters demeaned at the country level, corresponding to a fixed effects specification.

Figure S14 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 $\hat{\gamma}^+ > 1$. 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, $|1 - \hat{\gamma}^+|$, the correlation becomes




Figure S14: Correlation of $\hat{\omega}^+$ and $\hat{\gamma}^+$ in L'Haridon and Vieider (2019)

indeed much stronger ($\rho = 0.386, p < 0.001$).




Figure S15: Correlation of $\hat{\omega}^-$ and $\hat{\gamma}^-$ in L'Haridon and Vieider (2019)

The patterns for losses are very similar, and are shown in figure S15. 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 $\hat{\gamma}^- > 1$, which are also associated with high noise levels. Looking at the correlations between noise and absolute deviations of the variable from 1, $|1 - \hat{\gamma}^-|$, the correlation thus appears much stronger ($\rho = 0.445, p < 0.001$).




Figure S16: Correlatiion of $\hat{\omega}^+$ and $\hat{\alpha}^+$ in L'Haridon and Vieider (2019)

Figure S16 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 S17, where the same patterns appear to be even more accentuated ($\rho = 0.34, p < 0.001$).




Figure S17: Correlatiion of $\hat{\omega}^+$ and $\hat{\alpha}^+$ in L'Haridon and Vieider (2019)

S3.3 NCM-derived PT parameters

Let us now have a look at the distributions of the NCM-generated PT parameters, α , γ , δ , and ω . Figure S18 shows the density functions for gains in the data of Bruhin et al. (2010). The parameter values look reasonable, and fall into ranges we may expect to emerge from a PT model. Panel 18(a) and panel 18(b) show the largest differences between experiments, concerning outcome-discriminability α and likelihood-discriminability γ , respectively. Participants in Zurich in 2003 appear to be particularly discriminating towards outcomes, while Chinese participants in 2005 show particularly low discriminability in the likelihood dimension. Differences across experiments between other parameters tend to be less pronounced, even though optimism δ tends to be larger in the Beijing 05 experiment than in the two experiments in Zurich. Decision noise ω is virtually indistinguishable across the three experiments.

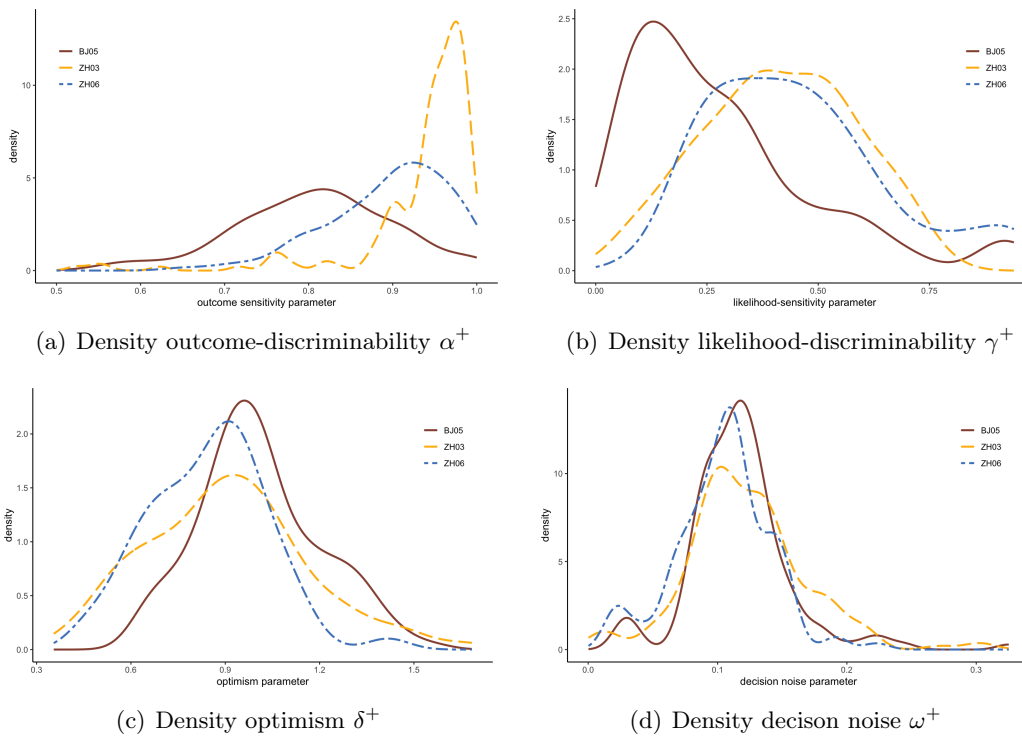


Figure S18: Density functions of NCM-generated PT parameters for Bruhin et al. (2010)

Figure S19 shows the same distributions for the four countries extracted from the L’Haridon and Vieider (2019) data, once again for gains. Panel 19(a) shows the distribution of outcome-discriminability separately for the four countries. Discriminability is highest in the US and China, significantly lower in Vietnam, and lower even in

Ethiopia. Very similar patterns are observed for likelihood-discriminability, displayed in panel 19(b). Again discriminability is highest in the US and China, with Vietnam and Ethiopia displaying significantly lower levels. This is hardly surprising, given the important role of encoding noise ν in the makeup of both parameters. Indeed, encoding noise is highly correlated with both α and γ ($\rho = -0.872$ and $\rho = -0.842$, respectively, with $p < 0.001$ in both cases), as one would expect based on the modelling setup and parameter definitions.

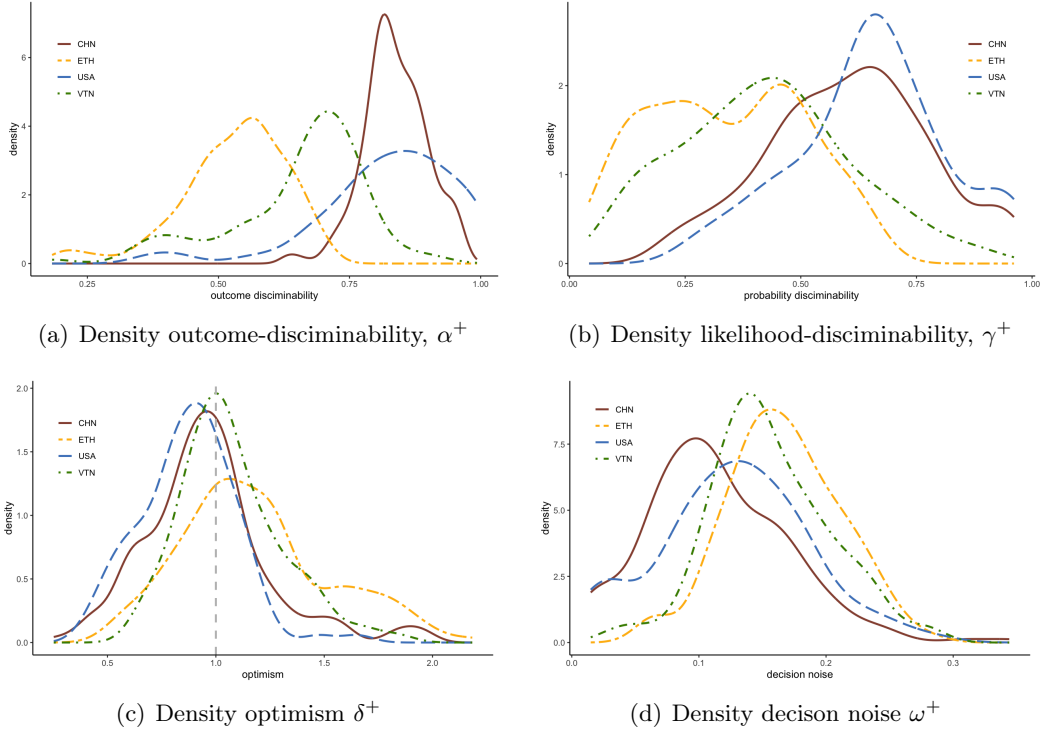


Figure S19: Density functions of NCM-generated parameters for L'Haridon and Vieider (2019)

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

S3.4 Correlations between NCM parameters

Finally, let us take a look at the correlations between the coding noise parameter ν , and the derived parameters α and γ . Figure S20 shows the correlations in the data of Bruhin et al. (2010). I show the correlations for gains only—the patterns for losses are similar, and can be found in the supplementary materials in section ???. The outcome discriminability parameter α shows a clear negative association with encoding noise in the overall data ($\rho = -0.472$). Whereas the pattern is striking in the Beijing 50 and Zurich 06 experiments, it is absent in the Zurich 03 experiment. The correlation between encoding noise and likelihood discriminability, shown in panel 20(b), is even more striking ($\rho = -0.798, p < 0.001$), and present across all three experiments. These strong negative correlations indeed form the cornerstone of the NCM, determining how much weight is being put on the likelihood relative to the prior.

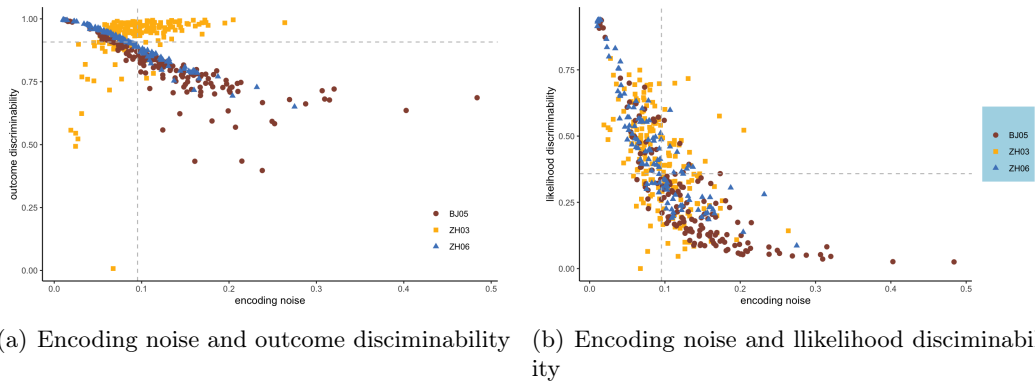


Figure S20: Correlations of NCM parameters in the data of Bruhin et al. (2010)

Figure S21 shows an equivalent scatter plot for the data of L'Haridon and Vieider (2019). The display follows the now-familiar pattern: Both outcome-discriminability and likelihood-discriminability show a strong negative correlation with encoding noise.

S3.5 Some results for losses

I plot the NCM parameter distributions for losses in the data of L'Haridon and Vieider (2019) in figure S22 for the usual 4 countries. Panel 22(a) shows systematic differences in encoding noise between countries, on top of the differences between individuals within each country. These differences are closely related to those observed for gains. In particular, the US and China have relatively low encoding noise, while Ethiopia and

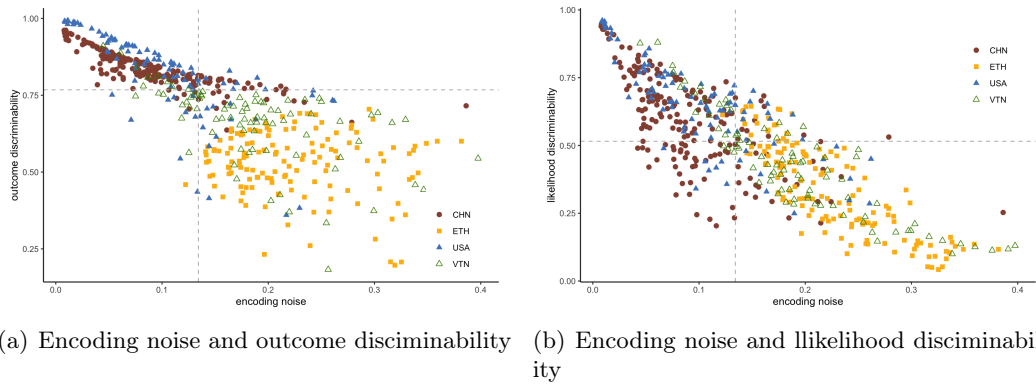


Figure S21: Correlations of NCM parameters in the data of L'Haridon and Vieider (2019)

Vietnam have larger levels of encoding noise. Panel 22(d) also shows systematic differences between countries, with Vietnam and Ethiopia having somewhat fixed points of the probability-distortion function on average, thus inverting the pattern seen for gains. The between-country differences for the two scale parameters of the prior are somewhat less pronounced, except for the somewhat diffuse distribution of the outcome prior variance in Ethiopia.

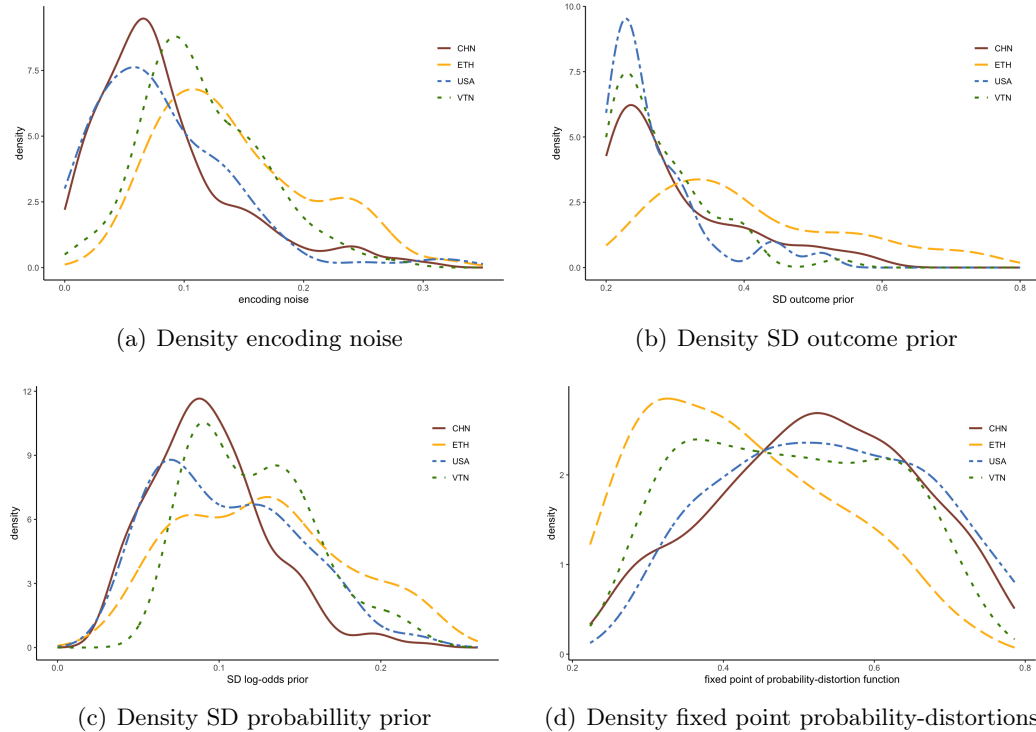


Figure S22: Density functions of NCM parameters for losses in L'Haridon and Vieider (2019)

To better interpret these patterns, it is useful to have a look at the NCM-derived parameters α^- , γ^- , δ^- , and ω^- . Figure S23 shows the distributions for the four countries extracted from the L’Haridon and Vieider (2019) data. Panel 23(a) shows the distribution of outcome-discriminability separately for the four countries. Discriminability is highest in China, with the US having a broader distribution, i.e. more individual heterogeneity, while being lowest in Vietnam. Very similar patterns are observed for likelihood-discriminability, displayed in panel 23(b). Discriminability is highest in the US and China, with Vietnam and especially Ethiopia displaying lower levels. The same systematic patterns are shown in panel 23(c) for pessimism, and in panel 23(d) for decision noise. Once again, the between-country differences in optimism are more muted than one might have expected from the differences in the fixed points of the function displayed in panel 22(d), which ones again happens due to the systematic differences in likelihood discriminability entering the definition of the latter. The same happens for decision noise ω^- . Although decision noise is higher in Vietnam and Ethiopia than it is in the US or China, the differences are less pronounced than those observed for encoding noise ν^- . This happens because the decision parameters α^- and γ^- , which are determined by ν^- , feed back into decision noise ω^- , compensating for the impact of ν^- .

Finally, let us take a look at the correlations between the NCM-derived PT parameters and the true PT parameters, which I illustrated based on the L’Haridon and Vieider (2019) data, and which are displayed in figure S24. The results for losses track those for gains closely.

I start by discussing the correlation between decision noise in the NCM with the SD of the additive noise term in PT, shown in panel 24(a). The association between the two variables is extremely tight ($\rho = 0.909, p < 0.001$), as one might expect given the fact that both are identified from the noisiness of choices. Given that ω^- contains in itself the model parameters α^- and γ^- , whereas $\hat{\omega}$ is implemented as an independent dimension, we would however expect larger differences to emerge on those parameters. These differences start surfacing in panel 24(b), showing the correlation of the optimism parameter emerging from the two models. At $\rho = 0.862, p < 0.001$, they are again highly correlated. Indeed, δ^- is only indirectly influenced by the decision noise encapsulated in ω^- . Nevertheless, this indirect influence surfaces in the narrower distribution the parameter has under the NCM compared to the PT model (IQR of 0.522 for $\hat{\delta}^-$ versus 0.276 for δ^-), and is a manifestation of the shrinking power of $1 - \gamma^-$ in the NCM.

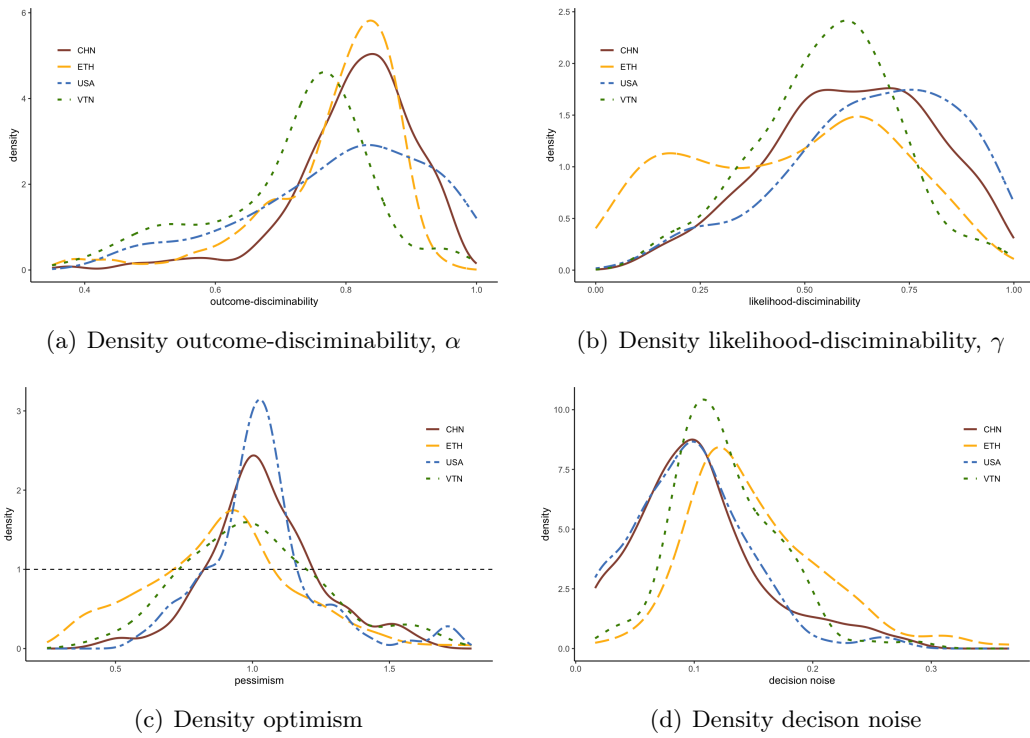


Figure S23: Density functions of NCM-generated parameters for losses in L'Haridon and Vieider (2019)

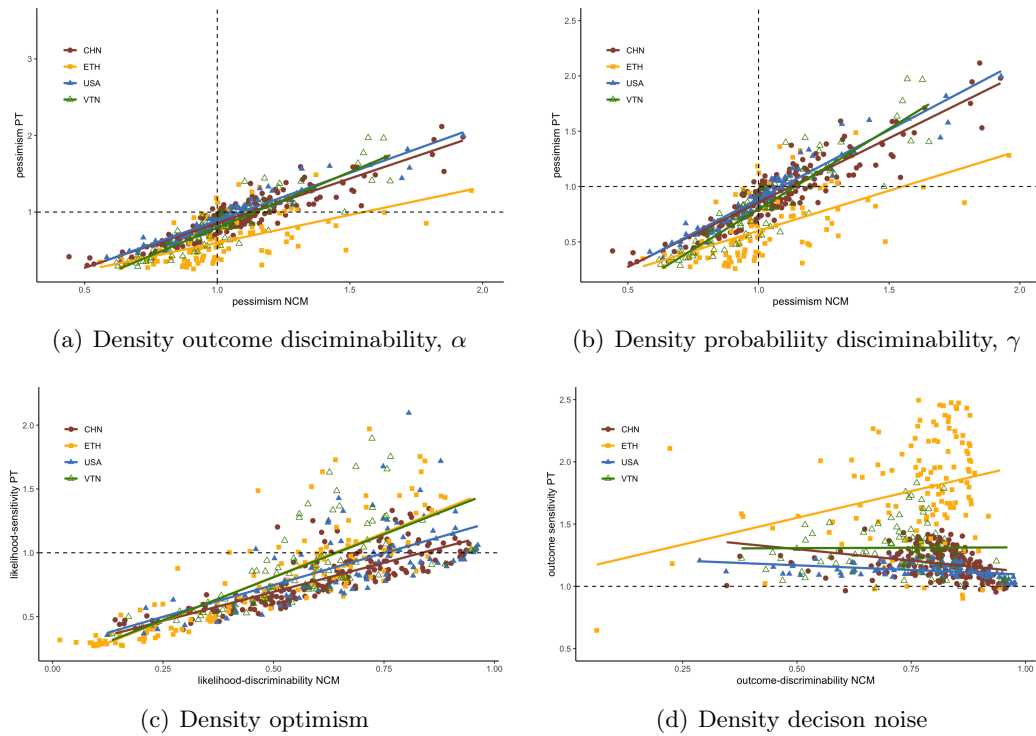


Figure S24: NCM-generated parameters and PT parameters for losses in L'Haridon and Vieider (2019)

The parameter correlations barely decrease once we move to likelihood distortions, shown in panel 24(c), and remain highly significant ($\rho = 0.827, p < 0.001$). As already witnessed for gains, values of $\hat{\gamma}^- > 1$ estimated using the PT specification do not necessarily map into values of γ^- close to 1. This reflects a central mechanism of the NCM at work. Values of likelihood-sensitivity larger than 1 under PT are typically associated with high levels of noise. Given the close association between decision noise and γ in the NCM, those choice patterns are best reflected by increased decision noise, which in turn results in values of $\gamma < 1$. It is this which drives the superior predictive performance of the NCM. Outcome distortions are shown in panel 24(d), where as for gains we find no correlation between the PT and NCM parameter.

S4 Fit of PT vs NCM to individual data

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

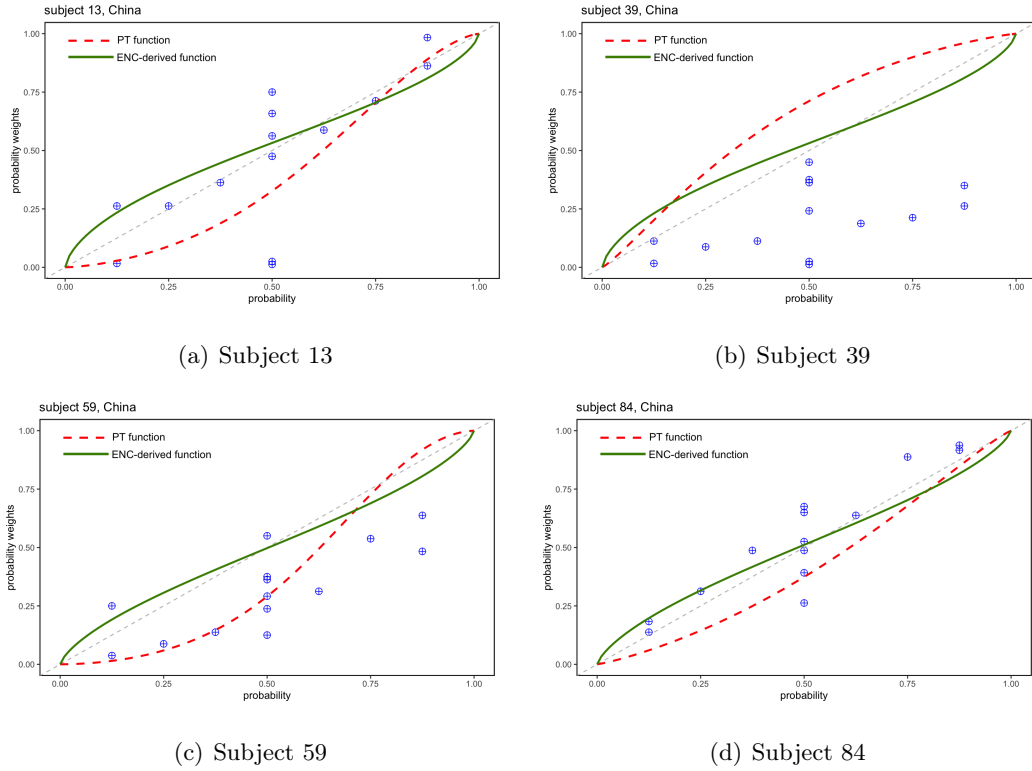


Figure S25: Individual-level data fit of PT vs NCM in L’Haridon and Vieider (2019)

Figure S25 shows the fit for the first 4 individuals. The pattern in panel 25(a) shows a pattern very similar to the one seen in the main text. Under PT, this subject displays oversensitivity at $\hat{\gamma} = 1.45$, oversensitivity to outcomes at $\hat{\alpha} = 2.39$, and low optimism at $\hat{\delta} = 0.49$, as well as moderately high levels of noise, at $\hat{\omega} = 0.14$. The NCM-derived parameters, which provide a clearly better fit at least to the probability dimension, come in at $\gamma = 0.68$, $\alpha = 0.74$, and $\delta = 1.14$. Optimism in probability-distortions is thus substituted for the extreme distortion of outcomes, obtained under PT from relatively few observations.

The other comparisons show similarly clear evidence in favour of the NCM-derived functions over PT. The fit in panel 25(b) appears poor for both functions, but this derived from the relatively high levels of outcome-distortion under both models ($\alpha = 0.64$, with a virtually identical parameter under PT). Beyond this, the NCM-derived function seems to run at least roughly parallel to the data points, whereas the PT-parameter $\hat{\gamma} = 1.17$ clearly over-estimates probabilistic sensitivity. A similar observation holds for 25(c), where the PT parameters would appear to fit the evidence better for $p \leq 0.5$, but where

the strong S-shape also means that the fit for large probabilities is very poor, pointing at a slightly different manifestation of overfitting. By reproducing the qualitative patterns seen in the displayed probability patterns, and absorbing the level difference through outcome-distortions, the NCM-model again outperforms PT. In panel 25(d), the pattern is crystal clear, and needs no further discussion. Figure S26 shows similar patterns for the next 4 subjects.




Figure S26: Individual-level data fit of PT vs NCM in L'Haridon and Vieider (2019)

S5 Ambiguity in the original data

S5.1 Experimental details

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

Experimental tasks. I presented subjects with 56 different binary prospects (28 for gains, 26 for losses, and 2 mixed prospects over gains and losses). Subjects had to make a choice between these prospects and different sure amounts of money, bounded between the highest and the lowest amount in the prospect. Gains were always presented first, and losses were administered from an endowment in a second part, the instructions for which were distributed once the first part was finished. Prospects were always kept in a fixed order. A pilot showed that this made the task less confusing for subjects, while no significant differences were found in certainty equivalents for different orders. Table S3 shows the prospects used in the usual notation $(x, p; y)$, where p indicates the probability of winning or losing x , and y obtains with a complementary probability $1 - p$.

Table S3: Decision tasks under risk and ambiguity

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

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

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

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

Risk and uncertainty. Risk was implemented using an urn with 8 consecutively numbered balls. Uncertainty was also implemented using an urn with 8 balls, except that

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

S5.2 The decision model and estimation code

The decision model closely follows the one for risk used in section [S3](#), 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.