Loss-Sensitivity versus Loss-Aversion*

Ranoua Bouchouicha[1], Yuchi Li[1], and Ferdinand M. Vieider[1]

[1] *RISL$\alpha\beta$, Department of Economics, Ghent University*

February 12, 2025

Abstract

Using a representative sample of the UK ($N = 1000$), we document risk taking in mixed gain-loss choices to be strongly stake-dependent: while our respondents are risk seeking for mean-preserving spreads around zero for moderate stakes up to £10, they become increasingly risk averse as stakes increase. Such patterns are predicted by recent models of adaptive behaviour based on 'noisy sampling' and 'noisy cognition'. We test the two adaptive models against each other and against traditional accounts based on diagnostic treatments for which they make opposite predictions. Prospect theory cannot account for the adaptive patterns we document. The evidence further supports the noisy cognition model over the sampling-based account. The reason is that the former is grounded in an optimization framework, whereas the latter is based on psychological intuition, which we show to be at the origin of the predictive differences between the two models.

Broadly stated, the task is to replace the global rationality of economic man with a kind of rational behavior that is compatible with the access to information and the computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist.

Herbert Simon (1955), p. 99

1 Motivation

Loss aversion—the observation that people tend to dislike losses more than they like monetarily equivalent gains—is one of the key concepts in behavioural eco-

*We are indebted to Mohammed Abdellaoui, Olivier L'Haridon and Horst Zank for helpful comments and discussions. All errors remain our own. This research was supported by the Research Foundation—Flanders under the project "Causal Determinants of Preferences" (G008021N). It received IRB approval from the Ethics review board of the FEB Ghent.

nomics. Prospect theory (PT) identifies loss aversion with a kink in the utility function at the origin (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992; Köbberling and Wakker, 2005). Hundreds of studies have measured the PT parameter governing this kink, and Brown, Imai, Vieider and Camerer (2024) put its meta-analytic average at just below 2. Recent studies have, however, cast doubt on the universality of loss aversion. Ert and Erev (2013) have argued that loss-attitudes are stake-dependent, and that stake-dependence is incompatible with how loss aversion is modelled in PT (though Bleichrodt and L'Haridon, 2023, dispute both these claims). Others have argued that loss aversion is an artifact of questionable measurement techniques, and that more appropriate measures produce loss neutrality (Yechiam and Hochman, 2013; Yechiam and Zeif, 2025). Using a large-scale sample representative of the US, Chapman, Snowberg, Wang and Camerer (2024) document wide-spread gain seeking (the opposite of loss aversion) in the American general population.

Here, we re-examine the concept of loss aversion in the light of recent models providing a *generative* account of the origin of choices over mixed-outcome gain-loss wagers. Other than preference-based accounts such as PT, these models ascribe observed choices to cognitive frictions in the decision process. Such cognitive frictions are overcome by leveraging information about the distribution of gains and losses in the environment. These models deliver distinctive predictions about how choice behaviour should change after exposing decision-makers to different distributions of gains and losses. Here, we exploit differences in predictions about environmental adaptation between the models to test these models against eachother, as well as against PT. As we argue in this paper, our results support a concept we call loss-*sensitivity*, which is distinct from loss-aversion, and which depicts attitudes towards losses as systematically modulated by stake size.

Risk taking in mixed gain-loss wagers is stake-dependent. We start by conducting an experiment with a representative sample of the population of the United Kingdom ($N = 1000$). We present subjects with forced binary choices between 50-50 gambles with a gain G and a loss L and the status quo of 0. We vary G and L to range between £2 and £30, and combine every gain with every loss. For mean-preserving spreads around 0 in which $G = L = x$, more than 50% of respondents prefer the risky option over the status quo up to £10 inclusive.[1] We thus replicate the finding of extensive risk seeking for mixed wagers by Chapman et al. (2024) for a different representative subject population (UK versus US). At the same time, we qualify that finding: while our subjects are on average risk seeking up to including £10, they turn to increasing risk aversion for larger amounts. Risk aversion thus systematically increases in the size of the stakes.

Stake-dependence implies *loss-sensitivity*. Any model that aspires to account for the patterns described above will need to incorporate heightened sensitivity towards losses relative to gains—a concept we refer to as *loss-sensitivity*. PT exhibits loss-sensitivity whenever utility for losses is 'steeper' than for gains. Such a pattern is indeed supported by the majority of empirical studies estimating PT parameters (see Imai, Nunnari, Vieider and Wu, 2025). Under commonly used behavioural definitions of loss aversion, which model loss attitudes as independent of probability weighting, the small-stake risk seeking we document further requires *gain-seeking*, in the sense of a loss aversion parameter $\lambda < 1$ (i.e., the opposite of loss aversion), which is less commonly observed. As we show in Online Appendix F, this condition continues to hold under more general definitions of loss aversion when one considers the universe of existing PT estimates. While PT can thus in principle organize the decision patterns we document, doing so requires parameter combinations that are not standard based on the accumulated empirical evidence.

An interesting question thus concerns where such parameter combinations may originate. PT is silent on this question, since it treats the parameters governing choices as exogenous. To answer this question, we thus identify two adaptive

[1] We use mean-preserving spreads to characterize risk taking behaviour nonparametrically, since they allow for a clear behavioural categorization simply by looking at the proportions of responses in favour of the mean-preserving spread over the sure option of 0 for equal expected value.

models of choice providing accounts of the cognitive origins of choice behaviour in mixed gain-loss wagers. The first is Decision by Sampling (*DbS*; Stewart, Chater and Brown, 2006). The second is the more recent Noisy Cognition Model (*NCM*) of Khaw, Li and Woodford (2021). Both provide cognitive micro-foundations for behaviour such as described in PT. Other than PT, however, both DbS and the NCM depict observed choices as arising from cognitive frictions affecting the decision process. The shortcomings arising from such cognitive frictions are, in turn, addressed by means of an adaptive mechanism exploiting regularities in the environment. Interestingly, however, the two models make different—and even opposite—predictions for specific choice situations. This presents an opportunity to test the models against each other, as well as against PT.

In DbS, noise arises in the evaluation of outcomes, given the psychological difficulty of mapping outcomes onto an absolute and unchanging utility scale (see Robson, 2001a, and Netzer, 2009, for different models where noise also arises at the stage of value attribution). To overcome such difficulties, utility is constructed from the comparison of the given outcome to be evaluated to a handful of outcomes sampled from memory. The utility of the outcome is then determined by its ordinal rank among the sampled outcomes. Stewart et al. (2006) use UK bank account transactions to show that few relatively large credits (monetary gains) are typically spent for shopping and bills, giving rise to many relatively small debits (monetary losses). The rank order of a given gain x will thus typically be lower when compared to typical credits than the rank order of a monetarily equivalent loss x compared to typical debits. This results in heightened sensitivity to losses relative to gains, and hence in a dislike of fair spreads around 0. Such loss-sensitivity is the key mechanism in DbS capturing attitudes toward mixed gain-loss wagers.

The NCM postulates that observed choice regularities arise from the noisy perception of outcomes. Noise thus arises at the perceptual stage, rather than at the stage of value-attribution as in DbS. The problem for the decision maker (the *mind*, given the neural motivation of the model) consists in inferring the true choice quantities from noisy signals about those quantities. The accuracy of these inferences can be maximized by 'decoding' the noisy signals by combination with a Bayesian prior, which summarizes the decision maker's expectations about the distribution of gains and losses in the choice environment. Loss-sensitivity obtains if subjects are more "attentive" to losses than they are to gains (i.e., if the noisy signal for losses is more accurate on average than the one for gains). The mean of the Bayesian prior furthermore determines whether a given outcome x will be over- or underestimated, thus naturally allowing for the co-existence of risk seeking and risk aversion over different stake sizes.

Empirical tests of predictions about choice adaptation. DbS and the NCM both postulate adaptive mechanisms exploiting information about the environment to overcome cognitive frictions. Notwithstanding this communality, the two models make distinctive predictions on how adaptation will impact observed choices in specific situations. To test these predictions, in experiment II we expose subjects to choices involving small versus large *gains* or small versus large *losses*. We do so in an initial adaptation phase containing pure gain or loss choices, which are followed by a common test set of mixed gain-loss choices. Other than existing tests of DbS, which have manipulated gains and losses *jointly* in the mixed outcome domain (Walasek and Stewart, 2015), this design allows us to separately test the predictions of DbS and the NCM, and to pitch them against each other, while keeping the common test stimuli separate from the adaptation stimuli.

The gain adaptation treatments allow us to pitch the two adaptive models against each other: In DbS, exposing subjects to larger gains ought to *decrease* the utility attributed to a given gain x, since it will have a lower rank order compared to a case where gains are expected to be smaller. This ought to result in an increase in risk aversion over mixed gambles. The NCM makes the exact opposite prediction: increasing the average gains in the environment ought to shift the mean of the prior upwards, thus increasing risk taking in mixed gain-loss gambles.

PT does not model adaptation in its standard formulations. It can, however, account for changes in choices by summoning changes in the reference point, which gives it additional explanatory power (see Kőszegi and Rabin, 2007, for a model of adaptation). One difficulty is that the literature is split e.g. on whether subjects are loss averse or gain seeking (see below for a review). This makes it almost impossible to make clear *ex ante* predictions on how choice behaviour should change following shifts in the reference point. One can nevertheless derive a partial restriction on PT's predictions upon reference point changes: if one observes changes in risk taking in a given direction when the reference point shifts upwards, then one ought to observe changes in the *opposite* direction when the reference point shifts downwards. Changes in risk taking in the *same* direction when comparing small to large losses and small to large gains would thus reject adaptation of the reference point as an explanation of our results, and hence PT.

Choice adaptation supports Noisy Cognition. We find that subjects exposed to large gains in the adaptation phase display systematically higher levels of risk taking in subsequent mixed gain-loss choices than subjects exposed to small gains. This result supports the noisy cognition model over Decision-by-Sampling, which predicts the exact opposite pattern. We furthermore find that subjects exposed to large losses in the adaptation phase also display more risk taking over mixed gambles than subjects exposed to small losses. While this finding is consistent with both DbS and (a specific case of) the NCM (see section 3.2), it contradicts PT. The finding of choices changing in the same direction after exposure to large gains or to large losses cannot be organized by reference-point adaptation in PT, no matter what the underlying parameter combinations may be. The NCM of Khaw et al. (2021) thus emerges as the clear winner from our tests.

Commonalities and differences between the models. This raises the question of where the difference in predictions between DbS and the NCM may originate, given the similar role the models attribute to environmental adaptation. We can trace the difference back to the Bayesian inference mechanism underlying the NCM. In DbS, the differences in experienced losses and gains directly result in the attribution of different utilities to a given gain and loss x, which arise from its ordinal ranking in comparison to typical gains or losses. In the model of Khaw et al. (2021), however, such a decision rule would be sub-optimal: inferences can be systematically improved by 'decoding' the noisy signal with a prior summarizing the learned characteristics of the environment.[2] This explains the difference in predictions: the NCM is based on a constrained optimization setup (with the constraint deriving from perceptual noise), whereas DbS is motivated based on psychological principles.

The optimizing machinery underlying the NCM is a feature that is clearly desirable from the point of view of evolutionary engineering. After all, if some noise is unavoidable due to biological information-processing constraints, then evolutionary pressures ought to have favoured mechanisms that optimally deal with such noise. Specifically for mixed gambles, this opens another interesting perspective. From an evolutionary point of view, the maximand is not per-period consumption, but rather evolutionary fitness (number of surviving offspring). As Robson (2001a) and Netzer (2009) powerfully argue, the impossibility of directly maximizing such a long-term goal would have led nature to create a mapping between consumption utility and fitness utility, which allows for the maximization of more proximate outcomes. In this article, we argue that the stylized neural mechanisms we discuss could be nature's solution to this design challenge. In particular, we argue that attention to gains and losses ought to be asymmetric with more "attention" being given to losses. *Loss-sensitivity* can then be conceptualized as the neural mechanism devised by nature to solve this evolutionary design problem.

Fit with the literature. The concept of loss-sensitivity we propose is based on hightened "attention" to losses relative to gains, which results in lower cognitive noise for losses in the framework of the NCM. This mechanism is consistent with

[2]The reason for this optimality is the same for which Bayesian estimators are considered preferable for small samples in a machine learning setup—see e.g. Bishop, 2006, chapter 3, for a formal discussion and proof. Intuitively, while the Bayesian estimator will result in systematic bias arising from regression to the mean of the prior, it will also drastically reduce the variance of the estimator. The trade-off, known as the *bias-variance trade-off*, is favourable to the Bayesian estimator in the precise sense of minimizing its mean squared error over many trials.

several studies in the literature directly investigating attention to losses, and linking the results to rejection probabilities of mixed gambles. Tom, Fox, Trepel and Poldrack (2007) show that neural activations react more strongly to changes in losses than to changes in gains, and that rejection of even-odds gain-loss wagers is closely linked to these neural activation functions. Sokol-Hessner, Hsu, Curley, Delgado, Camerer and Phelps (2009) document that arousal (measured by skin conductance) is generally greater for losses than for gains, and that arousal differences are predictive of loss aversion. Pachur, Schulte-Mecklenbeck, Murphy and Hertwig (2018) show that loss aversion estimated in a PT setting correlates with relative attention to losses over gains, and that manipulating such attention impacts estimated loss aversion. Using rich eye-tracking data, Hirmas, Engelmann and Weele (2024) show that subjects who pay more attention to gains (as measured using dwelling times) are more sensitive towards gains, and hence more likely to accept any given gain-loss wager over a range of stakes. These studies thus provide direct support to loss-sensitivity being driven by differential attention to gains and losses.

Our results further cast new light on the recent debate about the empirical magnitude and theoretical meaning of loss-aversion. Ert and Erev (2013) have argued on empirical grounds that loss-attitudes are strongly stake-dependent, and that such stake-dependence is incompatible with PT (though see Bleichrodt and L'Haridon, 2023, for a paper disputing that stake-dependence challenges PT, as well as its very existence). Yechiam and Hochman (2013) have argued that loss aversion is an artifact of questionable measurement techniques, and that more balanced measures will typically produce loss neutrality (i.e. $\lambda = 1$; see also Yechiam and Zeif, 2025, for a re-analysis of the studies in Brown et al., 2024). Using a large-scale sample representative of the US, Chapman et al. (2024) document wide-spread gain seeking (the opposite of loss aversion, i.e. $\lambda < 1$) in a sample representative of the US. Our results partially reconcile these seemingly opposing patterns, and reinterpret them in the light of adaptive models of choice detailing the origin of observed choice behaviour in cognitive noise.

The concept of loss-sensitivity also reconciles several findings in the wider literature that may otherwise seem at odds. For instance, it reconciles the observation of moderate-stake gain seeking in the general population, documented by Chapman et al. (2024), with the observation that people are reluctant to invest into the stock market (the *equity premium puzzle*; Mehra and Prescott, 1985), which has been explained by people's excessive focus on short-term losses (Benartzi and Thaler, 1995). At the same time, the origin of seemingly loss-averse behaviour in cognitive frictions during information processing could explain why some standard explanations based on loss aversion, such as the gap in selling versus buying prices, has been found to correlate poorly with empirical measures of loss aversion (Chapman, Dean, Ortoleva, Snowberg and Camerer, 2023). To the extent that measured loss aversion is an expression of cognitive noise, and may be highly variable across stake levels, it should perform poorly as a predictor of how people behave in different contexts (see also Bouchouicha, Oprea, Vieider and Wu, 2025, for a similar case about risk attitudes more generally).

Our findings further add to a series of recent papers ascribing observed behaviour to cognitive frictions instead of preferences. Several papers have modelled systematic probability distortions documented in prospect theory as arising from cognitive noise (Zhang, Ren and Maloney, 2020; Khaw, Li and Woodford, 2023; Frydman and Jin, 2023; Vieider, 2024b; Netzer, Robson, Steiner and Kocourek, 2024; Bouchouicha et al., 2025). Enke and Graeber (2023) show that low confidence ratings in the valuation of lotteries are predictive of probability distortions. Oprea (2024) replicates typical probability distortions under risk using an experimental setting in which risk has been removed entirely, showing that the observed patterns are an outgrowth of complexity, rather than capturing attitudes towards risk. Oprea and Vieider (2024) show that noisy cognition can account for the opposite probability distortions observed when choice primitives are described versus experienced (Barron and Erev, 2003; Hertwig, Barron, Weber and Erev, 2004), and leverage this insight to close the description-experience gap. To the best of our knowledge, we are the first to empirically test the noisy cognition account of choices over mixed-domain gain-loss wagers.

2 Risk taking in mixed wagers is stake-dependent

In our first experiment, we aim to test whether loss attitudes are indeed stake-dependent. This will allow us to revisit the debate on whether stake-dependence exists at all (Ert and Erev, 2013; Bleichrodt and L'Haridon, 2023), and whether general population samples are generally gain seeking (Chapman et al., 2024). We will further discuss the extent to which such stake dependence, if any, is compatible with prospect theory. This will serve to set the stage for the discussion of generative models of loss attitudes to follow below.

2.1 Experimental stimuli and methods

We recruited 1000 respondents representative of the adult population of the United Kingdom on Prolific Academic, using the "Representative Survey" option. The sample reflects the demographic distribution of the adult population of the United Kingdom, and is stratified across age, sex, and ethnicity using data from the UK Office of National Statistics.

We use a binary choice experiment in which subjects decide repeatedly between wagers yielding a gain G or else a loss L, each with probability $p = 0.5$, or else 0 (the status quo).[3] We vary G and L orthogonally and symmetrically between £2 and £30 in steps of £2, and combine all gains with all losses. Choices are presented one per screen, and the order of presentation is randomized at the individual level. The position of the risky option on the screen (up or down) is also randomized. Online Appendix A.1 provides a full list of the choice stimuli, and Online Appendix H shows the instructions and a screenshot of the choice environment. Subjects

[3] Some of the models we will examine below are explicitly targeted at binary choice. Models such as PT, on the other hand, are independent of the measurement method, since they (implicitly) incorporate the principle of *procedure invariance*. Binary choices are furthermore the preferred measurement tool emerging from the recent debate on the existence of loss aversion (see e.g. Yechiam and Zeif, 2025 and Bleichrodt and L'Haridon, 2023)

received a fixed payment for their time following Prolific regulations, and passed some initial comprehension questions.

We deliberately use *hypothetical* choices in all our experiments. Hypothetical choices are widely used in PT for choices involving losses, given the difficulties of enforcing real losses. In their meta-analysis of the universe of loss aversion estimates, Brown et al. (2024) found no difference between hypothetical and real monetary payments. Recent large-scale experiments have found no (or very minor) differences between hypothetical and real incentives in a number of tasks, and explicitly recommend hypothetical choices where real incentives may be difficult or problematic to implement (Brañas-Garza, Estepa-Mohedano, Jorrat, Orozco and Rascón-Ramírez, 2021; Brañas-Garza, Jorrat, Espín and Sánchez, 2023; Gneezy, Halevy, Hall, Offerman and van de Ven, 2024). While some experiments have deducted losses from an endowment, such an endowment could introduce a confound. While choices could reflect attitudes towards losses if all subjects frame the choice extremely narrowly, *de facto* forgetting the endowment, even just a few subjects integrating the endowment with the choice quantities (thus treating outcomes framed as losses as gains) could bias the results. We thus consider hypothetical choices to make for a cleaner test of our hypotheses.

2.2 Nonparametric results

Figure 1 shows the nonparametric choice proportions of mean-preserving (*fair*) spreads (where $G = L = x$) over 0 by stake size (exact choice proportions are reported in Online Appendix D.1). This allows us to discuss the results using a behavioural definition of risk aversion, which corresponds to a dislike of mean-preserving spreads around the status quo of 0 (and thus a choice proportion of the status quo larger than 50%). There is clear stake-dependence of risk taking for fair spreads around 0. For stakes $x \leq £10$ we find evidence for risk *seeking*—a majority of subjects prefer the risky fair spread over the sure 0. For values $x > £10$ this tendency inverts, and subjects show increasing risk aversion as stakes increase (i.e. the choice proportion of the fair spread declines steadily as stakes x increase).

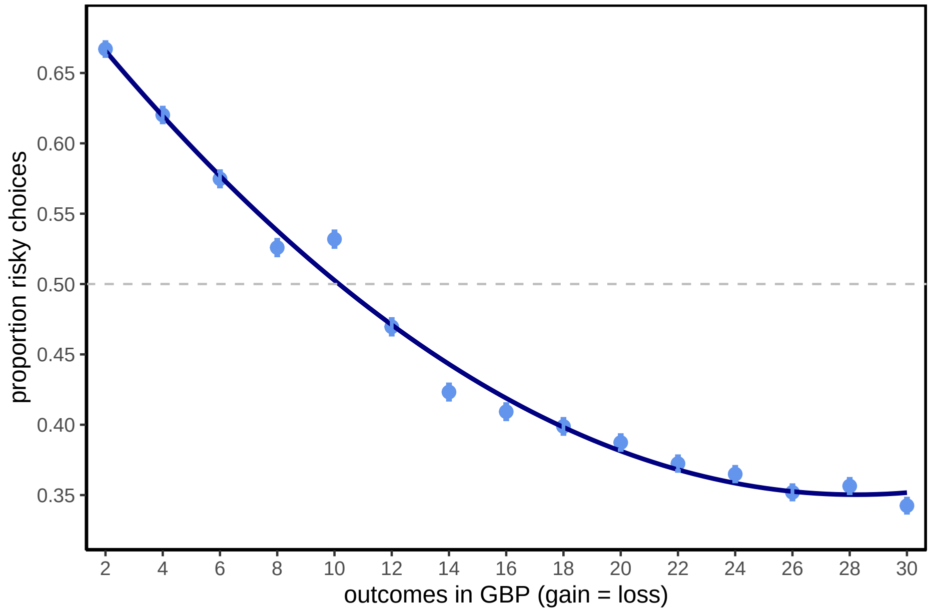

Figure 1: Risk taking by stake level, $G = L = x$

The figure shows choice proportions for a wager $(x, 0.5; -x)$ over 0 for values of x between £4 and £30 and their standard errors. The line is fitted by a polynomial regression of the second degree.

Our results thus replicate the widespread risk seeking documented by Chapman et al. (2024) for moderate-stake mixed wagers in a different subject population (UK vs US). At the same time we qualify that finding: while the great majority of our subjects are indeed risk seeking for small to moderate stake gain-loss wagers (including the stakes used by Chapman et al., 2024), most respondents turn to risk aversion as stakes get larger. We next examine to what extent this is consistent with loss aversion as modelled in PT.

Statistical tests. To test the loss-sensitivity hypothesis statistically using the full force of our data, we estimate a logistic regression including Bayesian hierarchical subject-level intercepts, which serve to cluster errors at the level of the individual respondent (see Online Appendix E.1 for details). The model includes a regression coefficient on the logarithm of the loss, $ln(L)$, as well as on the logarithm of the gain, $ln(G)$. The slope coefficient on the log-loss is significantly larger than the one on the log-gain (the difference between absolute value of loss regression coefficient and gain coefficient is 1.127, 95% credible interval [1.094 , 1.161]), indicating loss-sensitivity. Furthermore, the intercept is clearly positive, taking a value of 0.871 on the probability scale (which indicates risk taking proportions extrapolated to $G = L = 1$). Online Appendix E describes the regression model in detail, provides the code, and includes the full regression table.

2.3 Stake-dependence and prospect theory

PT models loss aversion through a constant $\lambda > 1$, capturing the added disutility of a loss over a monetarily equivalent gain. Risk aversion—in the behavioural sense of a choice of the sure 0 over a mean-preserving spread—will increase in stakes whenever the utility function for losses is steeper than the one for gains.[4] In the modern (cumulative) version of PT (Tversky and Kahneman, 1992), the definition of loss aversion is further impacted by the ratio of decision weights (which capture attitudes towards probabilities, and contribute to the explanation of risk taking) in addition to utility curvature (Schmidt and Zank, 2005). The coexistence of risk seeking for small stakes ($>$ 50% choices of the fair spread) and risk aversion for large stakes ($<$ 50% choices of the fair spread) we observe can be captured in PT by preference functionals exhibiting the following two characteristics:

1. Sensitivity towards gains decreases more quickly than sensitivity towards losses (i.e., utility is 'steeper' for losses). This ensures that the disutility attributed to a loss of $x > 0$ relative to the utility of an equally sized gain x increases in the stake size x, thus creating stake-dependence of risk aversion (i.e., the dislike of fair spreads increases in the stakes x).

2. To account for risk seeking ($>$ 50% choices of the fair spread) for small x, 'relative optimism' (the decision weight attributed to a 50% chance of winning x, over the decision weight attached to a 50% chance of losing x) must exceed the utility ratio.[5]

[4] As Kahneman and Tversky (1979) write on p. 279: "Indeed, most people find symmetric bets of the form $(x, .50; -x, .50)$ distinctly unattractive. Moreover, the aversiveness of symmetric fair bets generally increases with the size of the stake. That is, if $x > y > 0$, then $(y, .50; -y, .50)$ is preferred to $(x, .50; -x, .50)$. [...] Thus, the value function for losses is steeper than the value function for gains."

[5] Technically, a fair spread offering $\pm x$ with even odds will be preferred over the status quo if

Under commonly used behavioural definitions of loss aversion decision weights play no role.[6] Under such definitions, $\lambda < 1$ is thus the only way of accounting for small stake risk taking of the type we document under standard parametrizations.[7] More general definitions provide additional flexibility. As we show in some detail in Online Appendix F, none of the studies in the meta-analysis of Imai et al. (2025) predicts the patterns we document here precisely, and the estimates that come closest also have estimates of $\lambda < 1$. While PT can thus account for the type of stake-dependence we documented above, the required parameter combinations are rather unusual. This raises the question of where such parameter combinations may originate—a question on which PT is silent, since it treats the parameters governing behaviour as exogenous. Below, we will thus examine generative models providing cognitive micro-foundations for PT-like behaviour to determine whether they can shed light on the origins of these patterns.

2.4 Discussion of Results

Our data indicate stake-dependence in the acceptance proportion of fair spreads, combined with risk seeking (a majority of subjects accepting the fair spread) for small stakes. PT can account for these patterns with a combination of loss-

$w^+(0.5)u(x) \geq \lambda w^-(0.5)u(-x)$, where u captures decreasing sensitivity (to potentially different degrees for gains and losses), $w^+(0.5)$ is the decision weight for gains, and $w^-(0.5)$ the decision weight for losses. Risk seeking for a given stake size x in conjunction with its decrease in stakes will be observed whenever $\frac{w^+(0.5)}{w^-(0.5)} > \frac{\lambda u(-x)}{u(x)}$ and $\frac{u(-x)}{u(x)} > 1$, or equivalently, whenever $\frac{w^+(0.5)}{\lambda w^-(0.5)} > \frac{u(-x)}{u(x)} > 1$. The first inequality may obtain because of optimism for gains (i.e. the probability of winning is *over*weighted, $w^+(0.5) > 0.5$), because of optimism for losses (i.e. the probability of losing is *under*weighted, $w^-(0.5) < 0.5$), or because of "gain seeking" (the opposite of loss aversion, i.e. $\lambda < 1$)—or of course a combination of these motives.

[6]This is automatically the case for any model in which probability weighting is assumed to be the same for gains and for losses, as modelled in original PT and assumed in many empirical estimations. An alternative is that, even in cases where probability weighting is allowed to differ for gains and losses, it is assumed to be 'edited out' for mixed 50-50 gambles, as proposed e.g. by Schmidt and Zank (2005).

[7]By 'standard parametrizations' we mean the absence of multiple inflection points in the utility function, such as postulated e.g. by Friedman and Savage (1948) to save expected utility theory against its critics. As we show in Online Appendix F, this condition is indeed met by all existing PT estimations of loss aversion. While allowing for additional degrees of freedom in utility—such as arbitrary inflection points in utility—could 'save' loss aversion, our point here is that no studies investigating such functions has ever documented the conditions that would be required for this, making them ex ante implausible.

sensitivity (utility being 'steeper' for losses than for gains), and gain seeking (i.e. $\lambda < 1$). The unusual parameter combination needed to explain the results nevertheless raises a new set of questions: where might such parameter combinations originate? And can we identify some underlying reason for which the choice patterns we observed above may arise? Being a purely descriptive model, PT does not provide any answers to these questions. Below, we thus turn to *generative* models providing cognitive micro-foundations for choice behaviour over mixed gain-loss wagers. As we will see shortly, such models make additional testable predictions on how choices adapt to the distribution of gains and losses in the environment, which will allow us to pitch them directly against PT, as well as against each other.

3 Choice adaptation and stake dependence

If prospect theory does not tell us where loss-sensitivity may originate, which model does? Prime candidates are adaptive models of choice according to which the mind leverages information about the distribution of choice primitives in the environment to overcome cognitive frictions. To the best of our knowledge, there are two such models at a level of detail to provide precise predictions for choices over mixed gambles. The first is Decision-by-Sampling (*DbS*; Stewart, Chater and Brown, 2006). The second is the more recent noisy cognition model (*NCM*) of Khaw, Li and Woodford (2021). In this section, we describe the two models, and detail how they can account for the choice patterns encountered so far. In section 4, we will test the two models against each other, as well as against PT.

3.1 Decision-by-Sampling and loss attitudes

DbS is motivated by the psychological insight that assessing outcomes on an absolute, invariant scale, as assumed in expected utility theory and PT, would be cognitively challenging to the point of being empirically implausible (see also Tversky, 1969, for an early discussion of a similar point). While outcomes are perceived objectively according to this model, the environmental characteristics enter at the stage of determining their utility or value (see also Robson, 2001a, and Netzer, 2009, for a similar modelling intuition).

General model sketch. According to DbS, the utility of a given gain or loss depends on a small number of ordinal comparisons of the given outcome to be evaluated with outcomes sampled from memory. Utility is then constructed as the rank of the outcome in binary comparisons with sampled amounts. Let R be the rank of an outcome in comparisons with N samples. The outcome to be assessed, call it x, will have utility $u(x) = \frac{R-1}{N-1}$, i.e. the ordinal ranking of the outcome within the drawn samples directly gives its utility. The shape of utility will then be determined by the distribution of gains and losses in the environment (in memory). Stewart et al. (2006) examine bank transactions in the UK to illustrate typical monetary amounts people will experience in everyday life. They show the distribution of credits (monetary gains) to follow a power law, where moderately sized credits are frequent, and larger credits become increasingly less frequent. Ordinal comparisons to this distribution would then give rise to a canonical concave utility function over gains.[8]

Distributions of gains and losses, and loss-sensitivity. Losses are treated as separate from gains in the model. Examining *debits* on UK bank accounts, Stewart et al. (2006) show that the distribution of such debits once again follows a power law. The many small expenditures and few large expenditures will then give rise to a convex utility function. The crucial dimension for our purposes, however, is the comparison with gains. Losses tend, on average, to be much smaller than gains (lump-sum gains, such as salary payments, tend to be put towards many small expenses, such as everyday shopping and bills). This yields the prediction

[8]Note that similar distortions will also apply to the probability dimension. Stewart et al. (2006), however, do not discuss probability distortions as applying separately to gains and to losses, so that—at least implicitly—probability distortions are predicted to be identical for gains and losses and will thus cancel out when evaluating 50-50 gain-loss wagers as we do here.

that a given loss of x will typically have a higher relative rank than a given gain of the same monetary amount x. This immediately yields the prediction that losses receive a greater weight than monetarily equivalent gains.

The crucial element resides in the differential sensitivity towards gains and losses entailed by their ordinal ranking. This sets the model apart from PT, where loss aversion can obtain even for identical utility curvature for gains and losses (as is the case in the structural parameters estimated by Tversky and Kahneman, 1992). In particular, stake-dependence will be observed whenever the distribution of losses is narrower than that of gains, which ought to give rise to increased sensitivity towards losses as compared to gains. This indeed appears to be the case based on the distributions of debits and credits shown by Stewart et al. (2006). Aversion to mean-preserving spreads is thus purely a product of loss-sensitivity, and stake-dependence is an inbuilt feature of the model (given empirical distributions of gains and losses as typically experienced by people).

It is more difficult to determine whether risk seeking for small stakes is predicted by the model. In particular, what is required for small-stake risk seeking of the type we have observed is that the utility of a small gain of up to £10 be larger than the utility of a small loss. This could conceivably happen if the empirical distributions of gains and losses present different features for such small amounts, and in particular, if small gains are relatively more frequent than small losses.[9] The model can thus account for small-stake risk seeking, but its empirical plausibility will depend on the (unobserved) distribution of gains versus losses in memory.[10]

[9]This mechanism would then be akin to allowing for multiple inflection points in PT utility functions, as discussed in footnote 7. The main difference from PT is that in DbS, such multiple inflection points would need to be justified based on empirical outcome distributions of gains and losses. PT, on the other hand, has no mechanism on which such modelling choices would be based. It is in this sense that DbS provides micro-foundations for the function used in PT.

[10]An additional complication arises due to the error structure. Errors in DbS stem from the randomness of the handful of samples taken from memory to which outcomes are compared to obtain their ordinal ranking. This gives rise to a binomial variable attached to the utility or rank. If small gains and losses are both relatively infrequent, then such errors will be relatively influential, and will push choices towards 50-50 for very small stakes, given that they will overpower the relatively small differences in the ordinal ranks of gains and losses.

3.2 Noisy cognition and loss attitudes

The key insight underlying the noisy cognition model of Khaw et al. (2021) (henceforth: *NCM*) is that DMs cannot directly access the choice quantities presented to them. Instead, they will neurally 'code' such outcomes by means of neurons discharging action potentials (the electrical discharges which constitute the way in which information is coded and transmitted by the brain)—a process that will typically result in some noise in finite neuronal populations. Systematic biases will then arise from the optimal way the brain deals with such noise.

General model sketch. Take the decision of whether to accept or reject a binary wager offering a gain G with probability 0.5, or else a loss L.[11] The mind tries to implement a choice rule whereby the risky wager is chosen over the status quo whenever the expected gain exceeds the expected loss, i.e. $p\,G \geq (1-p\,)L$.[12] The mind will implement this choice rule based on the *inferences* it draws about the (a prior unknown) outcomes G and L from the noisy signals representing them (the action potentials encoding information mentioned above), call them r_g and r_ℓ, so that the DM will choose the wager whenever $\mathbb{E}[\,G\,|\,r_g\,] \geq \mathbb{E}[\,L\,|\,r_\ell\,]$.

Let the signals be single draws from two normal distributions:

$$r_g \sim \mathcal{N}(\,ln(G)\,,\,\nu_g^2\,)\;,\;r_\ell \sim \mathcal{N}(\,ln(L)\,,\,\nu_\ell^2\,), \tag{1}$$

where ν_g and ν_ℓ quantify the coding noise for gains and losses. The signals are correct *on average*, but any given draw may be affected by noise. To arrive at the mental choice rule, the signals will need to be decoded by combination with a prior. This will indeed be optimal (in the precise sense of minimizing the mean

[11] The assumption of even odds describes the great majority of the empirical evidence, including in our own experiments. While it has the merit of simplifying our exposition considerably, the predictions of the model are valid beyond this specific setting. Khaw et al. (2023) and Vieider (2024b) both present accounts allowing for nonlinear probability distortions, and which incorporate the account of loss aversion discussed here as a special case. Given that both models are defined over log-odds, 50-50 odds will drop out of the equation, so that the discussion presented here remains valid under these more general accounts.

[12] This choice rule is optimal in the sense that it maximizes lifetime wealth. Note that allowing for an additional transformation of outcomes using a utility function incorporating the decreasing marginal utility of wealth does not affect the conclusions of the model.

squared error of the estimator), inasmuch as information about the distribution of outcomes in the environment can be used to discipline the inferences drawn from the noisy signal. We start by parsimoniously assuming that gains and losses are drawn from a *common* prior. This will allow us to focus on the key insights of the model, and will not affect the key predictions (see below for details):

$$ln(G)\,,\,ln(L) \sim \mathcal{N}(\,\mu\,,\,\sigma^2\,). \tag{2}$$

Combining signals and prior to arrive at the posterior expectations of the two quantities, aggregating over repeated stimuli and substituting the resulting inferences into the optimal choice rule above (see Online Appendix B for the details of the derivation) yields the following stochastic choice equation:

$$Pr[(G,p;-L) \succ 0] = \Phi\left[\frac{\alpha\,ln(G) - \beta\,ln(L) + (\beta-\alpha)\,\widehat{\mu}}{\sqrt{\nu_g^2\alpha^2 + \nu_\ell^2\beta^2}}\right], \tag{3}$$

where $\alpha \triangleq \frac{\sigma^2}{\sigma^2+\nu_g^2}$ and $\beta \triangleq \frac{\sigma^2}{\sigma^2+\nu_\ell^2}$ are *discriminability* parameters capturing perceptual accuracy of gains and losses, and $\widehat{\mu}$ is a 'transformed prior mean'.[13] $Pr[(G,p;-L) \succ 0]$ is the probability of accepting the wager, and Φ is the standard normal cumulative distribution function. The model thus predicts that choice patterns concerning the acceptance or rejection of mixed wagers will be a stochastic function of the noisily inferred choice quantities.

The NCM and loss-sensitivity. In the NCM, loss-sensitivity takes the form of heightened "attention" towards losses compared to gains, here identified with lower coding noise for losses than for gains (i.e. $\nu_\ell < \nu_g$; see Online Appendix C for a stylized model of attention). This immediately implies that $\beta > \alpha$—as-if utility is 'steeper for losses than for gains'. For fair spreads offering $\pm x$ with

[13]The key quantity for the weight given to the signal can be shown to be $\frac{\nu_i}{\sigma}$, i.e. the relative coding noise attached to the outcome. This can loosely be interpreted as 'lack of attention', where increasing attention could decrease coding noise and hence increase discriminability. The 'transformed prior mean' is defined as $\widehat{\mu} \triangleq \mu + \frac{1}{2}\sigma^2$—see Online Appendix B for derivation and a discussion.

even odds, the numerator in equation (3) will simplify to $(\beta - \alpha)\ [\widehat{\mu} - ln(x)]$. The common outcome enters the equation with a negative sign, so that—if subjects are loss-sensitive in the sense of $\beta - \alpha > 0$—the choice probability of the wager, $Pr[(x, p; -x) \succ 0]$, will decline in the stake size x.

The degree of risk taking for small stakes, in turn, will depend on the transformed prior mean $\widehat{\mu}$. To see this, let us again focus on fair spreads, so that $G = L = x$. For $x = 1$, the choice probability of the fair spread depends purely on $(\beta - \alpha)\,\widehat{\mu}$. Assuming loss-sensitivity (in the sense of $\beta - \alpha > 0$), small-stake gain seeking will be observed as long as $\widehat{\mu} > 0$. Stochastic indifference (identified with a numerator equal to 0) between the status quo and a fair spread offering $\pm x$ with even odds will be observed at $exp(\,\widehat{\mu}\,)$, which thus creates a dividing line between 'small' and 'large' stakes according to the NCM.[14]

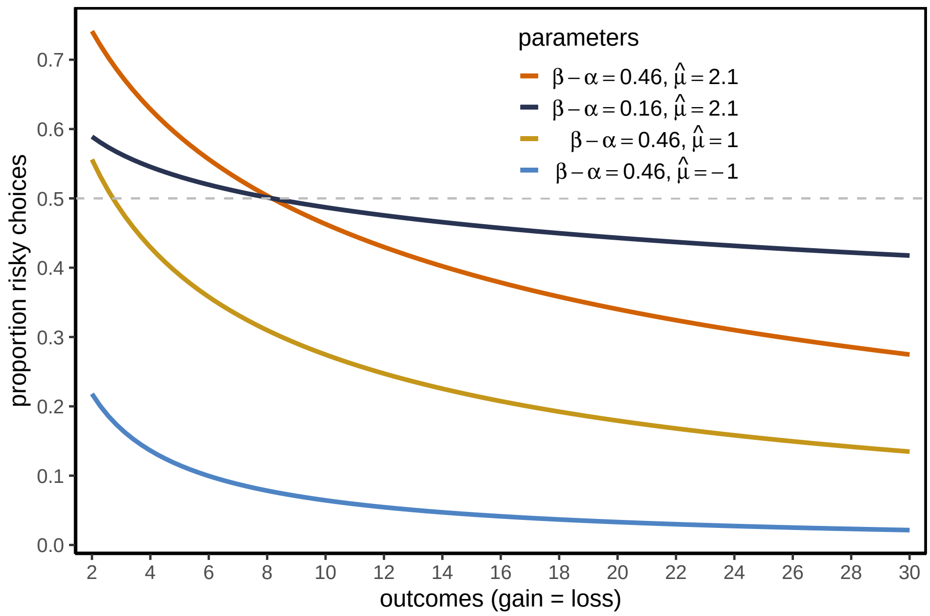

Figure 2: Risk taking by stake level, $G = L$

The figure illustrates choice proportions for a fair spread $(x, 0.5; -x)$ over 0 for different parametrizations of equation (3). Loss-sensitivity $\beta - \alpha$ determines the slope of the curve, and hence the degree of stake-dependence. The transformed common prior mean $\widehat{\mu}$ determines the 'fixed point', with the point of stochastic indifference between the choice options being described by $e^{\widehat{\mu}}$.

Figure 2 illustrates the choice patterns predicted for fair spreads. For $\beta - \alpha > 0$,

[14]To see this, define $\ln(\eta) \triangleq \widehat{\mu}$. For fair spreads the numerator in equation (3) thus becomes $(\beta - \alpha)\, ln(\eta/x)$. Values of $x < \eta$ will thus be uplifted towards 1 under loss sensitivity, i.e. as long as $0 < \beta - \alpha < 1$.

we expect the choice probability of the wager to decline in stake size. The larger this loss-sensitivity, the more pronounced the stake-dependence will be. The value of the 'transformed prior mean' $\widehat{\mu}$, in turn, determines the fixed point at which we will observe stochastic indifference between the wager and the sure outcome of 0. Small values of the mean then reflect "pessimistic expectations", which result in risk averse choices. For larger values of $\widehat{\mu}$, risk seeking for small stakes will coexist with risk aversion for large stakes.

Micro-foundations for loss aversion? The co-existence of stake-dependence and risk seeking for small stakes can also be organized by a more general model that allows for different priors for gains and losses. Let the prior for gains be $\mathcal{N}(\mu_g, \sigma_g^2)$, and the prior for losses $\mathcal{N}(\mu_\ell, \sigma_\ell^2)$. The numerator in equation (3) then becomes $\alpha\, ln(G) - \beta\, ln(L) - [(1-\beta)\widehat{\mu}_\ell - (1-\alpha)\widehat{\mu}_g]$. As shown by Khaw et al. (2021) in their Online Appendix D, we can now derive cognitive microfoundations for loss aversion as modelled in PT by defining $\lambda \triangleq exp[(1-\beta)\,\widehat{\mu}_\ell - (1-\alpha)\,\widehat{\mu}_g]$. Substituting $ln(\lambda)$ for $(1-\beta)\,\widehat{\mu}_\ell - (1-\alpha)\,\widehat{\mu}_g$ and exponentiating the numerator, one can see that the choice probability of the wager will be proportional to $G^\alpha - \lambda\, L^\beta$. This provides an account of the origin of loss aversion as modelled in PT, where loss aversion (in the sense of $\lambda > 1$) would obtain from a systematically larger (weighted) prior for losses than for gains.

Intuitively, loss aversion (in the sense of $\lambda > 1$) will thus be observed whenever $\widehat{\mu}_\ell > \widehat{\mu}_g$, i.e. whenever losses are expected to be larger than gains on average.[15] Having a prior for losses that is 'hard-wired' to be larger than the one for gains, however, may arguably be suboptimal for adaptation, since it would necessarily put limits on the possibility to correctly learn the distribution of stimuli in the environment. The small-stake risk seeking we find indeed suggests the opposite pattern of $(1-\beta)\widehat{\mu}_\ell < (1-\alpha)\,\widehat{\mu}_g$ (see Online Appendix B for a precise derivation and for a more extensive discussion). The latter seems empirically plausible, since

[15] Technically, this condition is neither necessary nor sufficient in the presence of loss-sensitivity. The necessary and sufficient condition is $(1-\beta)\,\widehat{\mu}_\ell > (1-\alpha)\,\widehat{\mu}_g$. This shows how loss-sensitivity and the prior means will interact in producing the observed choice patterns according to the model.

it correctly reflects the distributions of credits and debits people are exposed to on a daily basis according to the data of Stewart et al. (2006).[16]

Both a common prior specification with $\widehat{\mu} > 0$, and a specification based on separate priors (with $(1-\beta)\,\widehat{\mu}_\ell < (1-\alpha)\,\widehat{\mu}_g$) can thus naturally account for the patterns we have documented above. Both also seem plausible from an evolutionary and empirical point of view (i.e. taking into account the distribution of typical earnings and expenditures). Which specification makes more accurate predictions is empirically testable—a point to which we turn in the next section.

4 Coding noise versus sampling noise

DbS and the NCM both incorporate loss-sensitivity as a mechanism potentially driving loss attitudes. Both furthermore incorporate mechanisms leveraging knowledge about the statistical distribution of choice primitives (gains and losses) in the environment to overcome cognitive frictions arising in the choice process. To a casual observer, this may suggest that the two models will make similar predictions about behaviour. As it turns out, however, the two models make radically different—and even diametrically opposite—predictions in some specific cases. Here, we exploit this insight to pitch the two adaptive models against each other. We further exploit the combination of predictions emerging from the models for adaptation to gains and losses to pitch the adaptive account against prospect theory. Note that, while DbS has been tested in the past (see Walasek and Stewart, 2015, for a test of loss attitudes), the tests generally manipulated gains and losses *jointly*, which makes them ill-suited to tease apart the predictions of the two models.[17]

[16]One could make an argument that gains will be larger than losses on average across most situations people will encounter in the daily lives. For instance, losses in experiments tend to be smaller than gains, since subjects would not willingly partake in bets offering negative expected values. The last few centuries have further witnessed unprecedented growth in per capita income, which again points to gains being larger than losses on average.

[17]Some of these tests have further proven controversial. Walasek and Stewart (2015) showed that exposing subjects to large losses and small gains made loss aversion disappear, whereas exposing them to small losses and large gains increased loss aversion. The study results were

4.1 Experiment II: Adaptation and risk taking

In experiment II we aim to test a core mechanism underlying the predictions of both DbS and the NCM—that choice patterns observed over mixed gain-loss gambles should be systematically swayed by the distribution of gains and losses in the immediate decision environment. Specifically, we randomly vary whether subjects are exposed to small or large losses, or to small or large gains, in a between subject design with 4 treatment conditions. Jointly, these tests can distinguish between the predictions of DbS, the NCM, and PT.

Experimental stimuli and treatments. The treatment consists of presenting subjects with 100 pure-gain or pure-loss choices in the first part of the experiment, which constitutes the *adaptation phase*. The adaptation choices consist of forced binary choices between a 50-50 wager involving an amount x (which could be a gain or loss, depending on the treatment) obtaining with probability 0.5 or else 0, and a sure amount c (which has the same sign as x). The treatment manipulates whether subjects are exposed to gains or losses, and whether they are exposed to *small* amounts or *large* amounts. There are thus 4 treatment conditions: 1) small gains; 2) large gains; 3) small losses; and 4) large losses. Subjects assigned to small amounts make choices for x ranging between £3 and £15. Subjects randomly assigned to the large amount conditions make choices in which x ranges between £4 and £40. The sure amounts c are chosen such as to be distributed symmetrically around the expected value of the lottery, so that choices are meaningful (see Online Appendix A.2 for details of the stimuli). All choices are presented randomly, and the position of the options on the screen is counter-balanced in both the adaptation and test trials.

later challenged by André and de Langhe (2021b) based on the parametric analysis in the paper, which was shown to spuriously produce the reported results even for simulated data without treatment intervention, since different stimuli were used for the analysis of different treatments. See, however, the reply by Walasek, Mullett and Stewart (2021), which showed that the same patterns also appear in the nonparametric choice data when looking at stimuli that are common across the treatments, and the counter-reply by André and de Langhe (2021a). Results reported by Stewart, Reimers and Harris (2014) similarly did not hold up in a subsequent adversarial collaboration by Alempaki, Canic, Mullett, Skylark, Starmer, Stewart and Tufano (2019), which showed them to be an artifact of the parametric analysis techniques adopted.

After the adaptation phase, subjects are given a new set of instructions introducing the mixed gambles. These instructions are very short, simply introducing the new choice situation. The common test stimuli are constructed in a way similar to experiment I. To cast a wider net while keeping the experiment at a reasonable length, however, we now vary the outcomes from £4 to £40 in steps of £4 (while again keeping the probabilities fixed at 50-50). This also ensures that the adaptation stimuli do not exceed the range of the test stimuli. All gains are again crossed with all losses, resulting in 100 choices. The overall number of choices in the experiment is thus 200, which compares favourably with typical experiments in the literature.

We ran the experiment on Prolific UK with a total of $N = 400$ subjects (on average 100 subjects per treatment condition). The experiment took about 15 minutes and was thus shorter than experiment I (since we dropped the lengthy questionnaire). Subjects were paid £10 per hour according to Prolific regulations, and the choices themselves were hypothetical, just like in experiment I. Indeed, in this context providing an endowment could be even more problematic if even just some subjects integrate it with the payoffs from the choices, since the model predictions rely on cleanly separating gains from losses in the adaptation phase.

Predictions of Decision-by-Sampling. DbS predicts the distribution of outcomes in the environment to directly affect utility curvature. This happens because outcomes observed in previous choices ought to influence the distribution of the samples drawn from memory, which in turn affects the rank of a given outcome x. The predictions will differ for whether subjects are exposed to small versus large gains or small versus large losses:

> *Gain adaptation*: Exposing subjects to large gains should result in a lower rank for a given gain x than exposing subjects to small gains. Subjects exposed to large gains thus ought to be more risk averse in subsequent mixed outcome choices than subjects exposed to small gains. This simply happens because of the lower rank attributed to the given gain x by subjects exposed

to large gains, which makes any given gain appear as “less important”.

Loss adaptation: Exposing subjects to small losses ought to result in higher ranks being attributed to a given loss x than exposing subjects to large losses. It follows that subjects exposed to large losses should be less risk averse over subsequent mixed gambles than subjects exposed to small losses.

Predictions of the Noisy Cognition Model. Similarly to DbS, the NCM predicts that the distribution of outcomes in the environment should impact subsequent choices. This happens because these outcomes will contribute to the learning of the prior distribution. We derive the following predictions:

Gain adaptation: Subjects exposed to large gains ought to take more risk in subsequent mixed choices than subjects exposed to small gains. This happens because exposing subjects to larger gains ought to shift their prior mean upwards. It is important to note that this prediction holds regardless of which specification we use, i.e. regardless of whether there are different priors for gains and losses (in which case it produces an upward shift in $\widehat{\mu}_g$) or whether the prior is common to gains and losses. For gain adaptation we thus obtain *the exact opposite prediction* compared to DbS, thus allowing us to test the two models against each-other.

Loss adaptation: The predictions for the loss adaptation treatments are less clear-cut. Based on the different priors, exposing subjects to larger losses ought to increase the mean of the loss prior, thus increasing risk aversion. The more parsimonious common prior specification flips the prediction to an *increase* in risk taking (a *decrease* in risk aversion).[18]

Prospect Theory Predictions. PT makes no explicit predictions about how

[18] The reason for the flip in predictions is interesting. Exposing subjects to larger losses would now affect the common prior mean. The latter, in turn affects behaviour in interaction with loss-sensitivity, since it enters the equation predicting choice behaviour with a positive sign only as long as $(\beta - \alpha) > 0$. In other words, a larger prior mean predicts more risk tolerance because it is less important for losses than for gains, given that the signal is more precise for losses and will hence be given more weight.

the choice quantities in the environment might impact choice. The only conceivable effect could come from the adaptation of the reference point. Even allowing for shifts in the reference point the predictions are unclear. Supposedly, larger gains might shift the reference point further upwards, thus increasing the range of outcomes falling into the loss domain. Larger losses ought to have the opposite effect, increasing the range of outcomes falling into the gain domain.[19] The problem, however, is that any predictions would depend on the precise parameter values used to generate them. Since there is currently no agreement even on whether subjects are loss averse or gain seeking, this makes deriving predictions *almost* impossible.

There is, however, one restriction on plausible PT predictions in terms of risk aversion (i.e. choices for the status quo over fair spreads) we *can* derive in this case: if choice proportions react to "upwards shifts in the reference point" upon exposure towards large versus small gains, then "downward shifts in the reference point" ought to elicit reactions in the exact opposite direction.[20] That is, while changes in behaviour following the gain adaptation or loss adaptation treatments cannot be predicted in isolation, jointly the two treatments should elicit shifts in choices that—if any—go in opposite directions.

4.2 Results

We again start by examining choice behaviour for fair spreads, with $G = L = x$. This will allow us to revisit the stake-dependence found above, and to obtain nonparametric insights into the level of risk taking across treatments.

Nonparametric results. Panel A of Figure 3 shows the nonparametric choice

[19]Note that our adaptation stimuli are constructed in such a way as to always fall within the range of the stimuli encountered in the test set. Presumably, all choices will thus still apply to mixed domain choices even after shifts in the reference point.

[20]It is important to note that this prediction holds whatever the definition of loss aversion, and whatever shape the PT functionals take. The only thing that is required is that the shape of probability weighting and utility do not change jointly with the reference point. PT indeed has no mechanism that would allow for such a change. If one does want to allow for such a change *ex post factum*, however, then we are happy to stipulate that PT is meta-physically alive and can account for our evidence (and indeed any type of choice pattern one might observe).

proportions of fair spreads after exposing subjects to small versus large *gains* in the adaptation phase. Risk taking is fairly low across both treatments. This may be due to priors being generally pessimistic for gains—see Oprea and Vieider (2024) for an explicit discussion of pessimistic priors for gains—or to the use of a somewhat different subject pool (generic Prolific subjects rather than a representative sample). There is, however, a clear treatment effect: subjects exposed to large gains in the adaptation phase are much more risk seeking than subjects exposed to small gains. This is consistent with the prediction of the NCM, but goes in the opposite direction of the effects predicted by DbS.

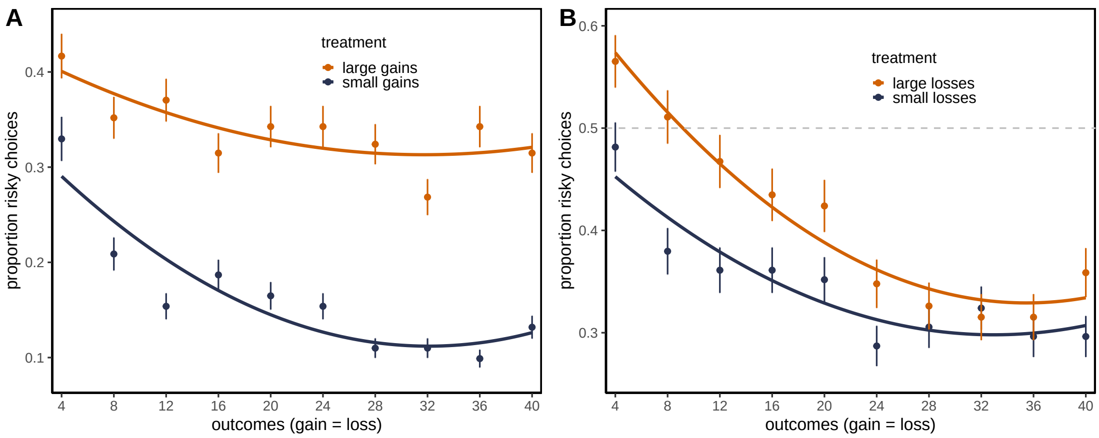

Figure 3: Risk taking proportions for fair spreads across treatments

Acceptance proportions of risky wager for fair spreads, where $G = L$. Panel A shows the results for the gain adaptation treatments, and panel B the results for the loss adaptation treatments.

Panel B in Figure 3 shows risk taking proportions for fair spreads after the loss adaptation treatments (for a table with choice proportions and standard errors, see Online Appendix D.2). We observe more risk taking in the large loss treatment than in the small loss treatment. In particular, we see risk seeking for $G = L = 4$ and $G = L = 8$ in the large loss treatment, after which we see risk aversion as stakes increase further. In the small loss treatment, however, we see consistent risk aversion for all stake levels. These results correspond to the DbS prediction, as well as to the prediction of the NCM based on a common prior. Jointly with the results for gains, they do however allow us to determine whether PT with shifting reference points can account for the results. It cannot: PT predicts shifts that go in opposite directions upon exposing subjects to large losses or to large gains, whereas we observe shifts in the *same* direction.

Statistical tests. To statistically examine the treatment effects in the whole dataset, we can run the same logistic regression used above, but with treatment-specific aggregate parameters (corresponding to treatment fixed effects—see Online Appendix E.1 for details). We statistically confirm loss-sensitivity across all treatments (see Online Appendix E.2 for complete regression tables). Importantly from the perspective of the current test, however, the aggregate intercept (capturing risk taking proportions at $ln(G) = ln(L) = 0$) for subjects exposed to large gains is significantly larger than the one for subjects exposed to small gains (difference of 1.494, 95% CrI [0.243 , 2.749]). Subjects exposed to large losses furthermore have a significantly larger intercept than subjects exposed to small losses (difference of 1.348, 95% CrI [0.281 , 2.412]). This confirms the results shown above from a statistical point of view using choices from the entire test set.

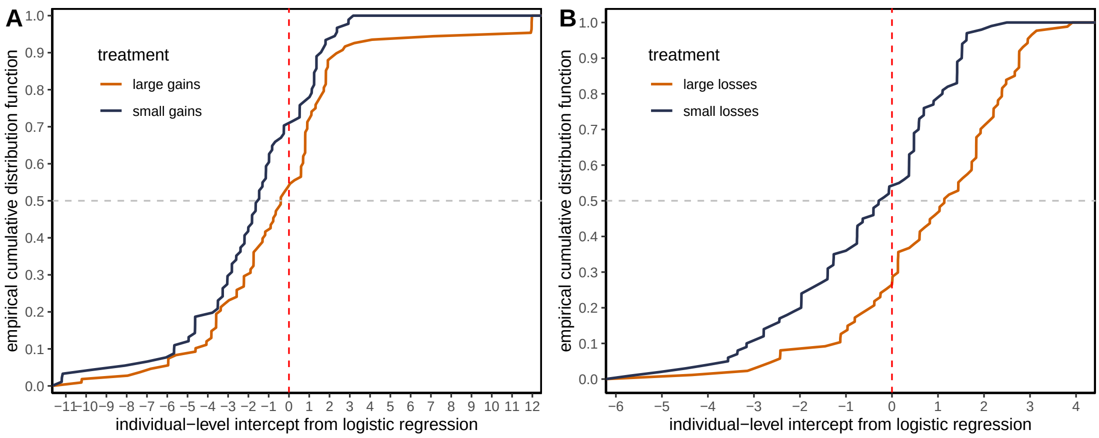

Figure 4: Individual intercepts of logistic regression by treatment

The figure plots the empirical cumulative distribution function of the individual-level intercepts of the logistic regression, separated by treatment. Panel A shows the distribution of individual-level intercepts in the gain adaptation treatments, and panel B in the loss adaptation treatments. The intercepts can be interpreted as risk taking propensity when $ln(G) = ln(L) = 0$, measured on a log-odds scale. The intercept can then be conceptualized as being proportional to $(\beta - \alpha) \times \widehat{\mu}$ in the common prior version of the NCM. Some outliers have been cut from the figure to improve the visualization.

We can also examine the results in a different way, by looking at the individual-level intercepts across treatments rather than at their hierarchical means. The intercepts, capturing risk taking when $ln(L) = ln(G) = 0$, have particular importance under the NCM, where they capture the risk taking tendency deriving from the prior. Figure 4 thus further plots the distributions by treatment of the individual-level hierarchical intercepts. Panel A plots the empirical cumulative distribution functions for the gain adaptation treatments. Respondents exposed to larger gains in the adaptation phase have clearly larger intercepts, which is confirmed by a Wicoxon ranksum test on the individual level mean intercepts ($p = 0.019$). Panel B shows the equivalent empirical cumulative distribution functions for the loss adaptation treatments. The distribution for large losses falls substantially to the right of the distribution for small losses. A Wilcoxon ranksum test on the individual-level means of the intercepts confirms the result of increased risk taking by subjects who have been exposed to larger losses in the adaptation phase ($p < 0.001$).

Discussion of Results. Our diagnostic test deployed to contrast the NCM predictions with DbS delivered a clear verdict. The predictions of the NCM, whereby exposing subjects to larger gains ought to increase the mean of the gain prior and thereby increase their risk taking, are clearly supported by the data. This also means that the DbS account is rejected, since it would predict the exact opposite patterns. The results can furthermore not be organized by changing reference points under PT, which—no matter what the parameter constellation may look like—predicts effects in opposite directions when the reference point shifts in opposite directions.

An interesting question concerns the deeper reasons behind the contrasting predictions of DbS and the NCM. While both models allow for different sensitivity towards gains and losses, in DbS the differential sensitivity is all there is: larger utility attributed to a loss x than to a gain x will necessarily trigger risk aversion. In the NCM, on the other hand, the as-if utility parameter coincides with the Bayesian evidence weight (i.e. the weight put on the signal relative to the prior expectation). Larger coding noise (and lower discriminability of as-if utility) thus trigger increased regression to the mean of the prior, and the behavioural predictions for fair spreads will thus depend on the position of the outcome x vis-a-vis that prior. The prediction of increased risk taking thus derives from the prior mean increasing upon exposure to larger gains.

Combination with the prior to decode noisy signals is a key feature of the NCM. Ultimately, the difference between the two models thus arises from the NCM being rooted in a constrained optimization framework: given noise in outcome assessments, the Bayesian estimator optimally deals with such noise by pooling the information in the signal and the information about the environment, in proportion to their respective informational contents (see Ma, Kording and Goldreich, 2023, chapter 4, for an intuitive discussion of the estimator's optimality in the context of Bayesian observer models; see Bishop, 2006, chapter 3, for a discussion of the optimality of the Bayesian estimator in a machine learning context and for a formal proof). Such optimal information aggregation would seem important from an evolutionary design perspective, and the foundation in a mechanism that is common to human and animal signal processing (including for sensori-motor tasks) promises to deliver unifying principles based on which one can derive predictions about behaviour.

5 Loss-sensitivity and cognitive ability

The results of Chapman et al. (2024) present a mystery: respondents with higher cognitive ability are found to be more loss averse. The same high-cognitive ability individuals, however, take more risk over pure gain wagers.

Here, we return to our data from experiment I to examine correlations between the parameters of the NCM model and observable characteristics of the subjects, and in particular, of cognitive ability. To this end, we structurally estimate equation (3), the only model that can explain the sum-total of the evidence we presented. We obtain individual-level parameters by estimating a Bayesian hierarchical model. We subsequently treat the estimated parameters as uncertain quantities, whereby each parameter is weighed by the inverse of its variance (as done e.g. in meta-analysis; see Vieider, 2024a for a discussion and tutorial). This allows to the the full uncertainty surrounding the individual-level parameters into account. Within this measurement-error framework, we then use an outlier-robust regression to regress the model parameter on observable characteristics of the respondents.

We use this model to regress the estimated parameters on observable subject characteristics, notably, the age, sex, and income of each respondent. Our main variable of interest is cognitive ability, which is measured as the number of correct answers to 11 questions testing the numeracy and general reasoning ability of the respondents administered in the final questionnaire. Online Appendix G contains the questions used to elicit these characteristics, explains the econometric code, and presents a robustness analysis to alternative specifications.

Table 1: Regression of NCM parameters

ind. var.	prior mean $\widehat{\mu}$	loss-sens. $\beta - \alpha$	gain discr. α
cognitive ability	**−0.107**	**−0.008**	**0.012**
	(0.044)	(0.004)	(0.004)
age	**0.084**	0.006	−0.005
	(0.042)	(0.004)	(0.004)
female	0.080	**0.038**	**−0.038**
	(0.094)	(0.009)	(0.008)
income	0.002	0.001	−0.001
	(0.029)	(0.003)	(0.002)
Constant	**1.878**	**0.038**	**0.956**
	(0.085)	(0.007)	(0.006)
Observations	994	994	994

The regression of individual-level coefficients on respondents' characteristics is based on an outlier-robust measurement error model (see Online Appendix G for details). Effects significant at the 5% level are highlighted in bold, and standard errors are reported in parentheses. The regressions are based on 994 respondents, since the sex identifier was missing in the Prolific data for the remaining 6 respondents. Continuous variables such as age, and categorical variables with large numbers of categories, such as income and cognitive ability, are entered as z-scores.

Table 1 reports our key results. The transformed prior mean, $\widehat{\mu}$, decreases in cognitive ability, indicating decreased risk taking by high cognitive ability respondents.

Cognitive ability further impacts loss-sensitivity. In particular, high cognitive ability respondents are less loss-sensitive. Jointly, these two effects show that individuals with high cognitive ability take less risk over mixed wagers, consistent with the results reported by Chapman et al. (2024).

Regression III further traces this effect back to high cognitive ability individuals having higher gain-discriminability (i.e., lower coding noise for gains). Together, these effects explain why high cognitive ability individuals would be less risk averse for gains, even though they are more risk averse for mixed wagers. Other than this, we find that older people take more risk on average. Women are more loss-sensitive, an effect that we trace back to reduced discriminability of gains, which implies that they may be more risk taking for small stakes, while taking less risk for larger stakes.

6 Discussion and Conclusion

Summary and interpretation of results. Loss aversion is a key concept in prospect theory, and in behavioural economics at large. Here, we have systematically re-examined the concept of loss aversion in light of recent adaptive models of choice providing cognitive microfoundations for choices over mixed-domain gain-loss wagers. Experiment I showed pronounced stake dependence, whereby risk seeking for small-stake mean-preserving spreads around 0 coexists with risk aversion for large stakes. We have shown that such behaviour can be rationalized by PT in principle, but that the required parameter combination clashes with the typical interpretation in the PT literature, according to which a dislike of mixed wagers is driven by loss aversion.

We thus argued that our results support an alternative concept, which we named *loss-sensitivity*. According to this concept, seemingly "loss averse" behaviour may emerge over some outcome range when decision-makers exhibit higher sensitivity towards losses than towards gains. While this feature can be organized by prospect theory functionals, it takes centre-stake in explicitly adaptive models of choice such as Decision-by-Sampling (Stewart et al., 2006) and the Noisy Cognition Model (Khaw et al., 2021). We thus devised 4 treatment conditions that could 1) separate the adaptive models from PT predictions, including when reference points can adapt to the choice environment; and 2) from each other. The results delivered a clear verdict in favour of the predictions emerging from the Noisy Cognition Model. The key mechanism underlying the difference in predictions between the adaptive models is interesting: the superior predictive performance of the Noisy Cognition Model emerges from its grounding in a constrained optimization problem, which is not a feature of the more intuitively-derived Decision-by-Sampling model.

The mechanism responsible for the patterns we document in the NCM consists of hightened sensitivity to losses compared to gains, which manifests in lower coding noise for losses. This is not an inbuilt feature of the model. It is, however, straightforward to endogenize coding noise in the model. Heng, Woodford and Polania (2023) propose a model in which the coding noise affecting the representation of a given choice stimulus is an endogenous product of exposure time to that stimulus. Coding noise will thereby be a decreasing function of exposure time. This creates an intriguing link with the literature showing that dwelling times on a particular attribute (the gain or the loss for even-odds mixed-domain wagers) are predictive of the acceptance probability of the wager (Pachur et al., 2018; Hirmas et al., 2024). Such an endogenous account of attention, linked to the loss-sensitivity account in the NCM, thus seems a promising topic for future research.

Implications for the literature. Our results help, first of all, to reconcile some apparently contradictory findings that have accumulated in the literature. Loss-sensitivity implies that apparent "gain seeking" can co-exist with apparent "loss aversion" over different stake ranges—the reason why we more neutrally described these attitudes as 'risk taking over mixed gain-loss wagers'. This may go some way towards explaining the high levels of between-study heterogeneity in loss aversion estimates documented by Brown et al. (2024). This also means that the widespread risk seeking observed by Chapman et al. (2024) in the US-American general population can be reconciled with large-spread risk aversion observed in financial markets (Mehra and Prescott, 1985), which has been explained with extreme risk aversion applied to narrowly framed investment windows (Benartzi and Thaler, 1995; Gneezy and Potters, 1997; Thaler, Tversky, Kahneman and Schwartz, 1997). Our findings can also help organize recent evidence that loss-aversion is empirically unrelated to the gap between selling and buying prices (Chapman et al., 2023). The role of cognitive noise in shaping choice behaviour over mixed wagers makes it indeed likely that this discrepancy may also originate in cognitive frictions—an issue that is clearly deserving of closer scrutiny.

At the face of it, Decision-by-Sampling is a model of noisy sampling, whereas the Noisy Cognition Model is a model of noisy coding. This distinction, however, is much less clear cut than one might at first think. In particular, the NCM is silent about the precise mechanism by which likelihood and prior are combined to arrive at a posterior. One possible mechanism consists of taking samples from the prior.[21] Even in such a setup, however, the predictions of the two models would remain distinct: the reason for this is, once more, the optimization basis of the NCM, which would persist regardless of the way in which the posterior obtains. Nevertheless, this also shows how the mechanisms discussed in both models may well need to be combined to obtain a realistic description of the neural processes underlying the choice predictions of the noisy coding model.[22]

[21] The other one is that the mind stores precise representations of the statistical quantities involved, which can then be recombined using a process akin to the one described in the model derivation. Note, however, that a combination with evidence sampled from the prior would only need a very slight modification to the formal setup of the NCM, and would not change any of its predictions: it would simply result in the introduction of additional stochasticity resulting from the sampling process.

[22] One reason discussed in the literature to exclude "Bayesian models" as plausible explanations of choice is variability in observed choice behaviour. Note, however, that the noisy cognition model clearly falls into a different category of Bayesian model from those targeted by this criticism, since choices will inherently be noisy because of the noise incorporated in the initial signal. Just like the sampling-based distinction discussed above, the Bayesian versus non-Bayesian distinction is thus not meaningful in such generic terms.

Neural and evolutionary mechanisms. Increased sensitivity to losses seems meaningful from an evolutionary point of view. Robson (2001b) and Netzer (2009) argue powerfully that, in order to maximize evolutionary fitness (identified with number of surviving offspring), evolution would have endowed humans with a proximate utility function over consumption geared towards this goal. This would be effective inasmuch as consumption is observable in each period, whereas fitness would only be observed over very long time periods (and given that a steady consumption profile is needed to maximize fitness). Losses in consumption then ought to be treated differently from gains in the mapping from consumption utility into fitness utility to guarantee inter-temporal stability of the consumption profile: large losses can do more damage to fitness than large gains can do good. Tom et al. (2007) indeed showed that neural activation signals changed with the size of losses to a greater extent than they did for equivalent gains, thus suggesting increased sensitivity towards losses.[23]

This brings us to the question about the deeper origins of the mechanisms described in the noisy cognition model. One possibility is that increased attention to losses is hard-wired into the brain by evolution. An alternative—and possibly complementary—explanation, however, may suggest that the degree to which attention is increased towards losses (and the ranges over which this is the case), ought to be itself adapted to the distribution of stimuli in the environment (see e.g. Heng, Woodford and Polania, 2020, for a formalization of this idea). Increased attention to losses may also arise because losses tend to be smaller than gains on average, as indeed discussed in DbS. Ultimately, the mechanisms described in the NCM and the DbS may thus be combined to increase our understanding of how adaptation to local circumstances can be leveraged to overcome cognitive frictions—a conclusion that is much in the spirit of our inspirational quote at the beginning of this paper.

[23]Such an account is also consistent with neural evidence that attributes a significant role in mixed-domain gain-loss decisions to brain regions such as the amygdala (Martino, Camerer and Adolphs, 2010). The amygdala is generally thought to be responsible to fear responses, which may then trigger increased attention to 'dangerous' situations, and large losses in general.

References

Abdellaoui, Mohammed, Han Bleichrodt, and Olivier l'Haridon (2008) 'A tractable method to measure utility and loss aversion under prospect theory.' *Journal of Risk and uncertainty* 36, 245–266

Alempaki, Despoina, Emina Canic, Timothy L Mullett, William J Skylark, Chris Starmer, Neil Stewart, and Fabio Tufano (2019) 'Reexamining how utility and weighting functions get their shapes: A quasi-adversarial collaboration providing a new interpretation.' *Management Science* 65(10), 4841–4862

André, Quentin, and Bart de Langhe (2021a) 'How (not) to test theory with data: Illustrations from walasek, mullett, and stewart (2020).' *Working Paper*

— (2021b) 'No evidence for loss aversion disappearance and reversal in walasek and stewart (2015).' *Journal of Experimental Psychology: General* 150(12), 2659

Barron, Greg, and Ido Erev (2003) 'Small feedback-based decisions and their limited correspondence to description-based decisions.' *Journal of behavioral decision making* 16(3), 215–233

Benartzi, Shlomo, and Richard H Thaler (1995) 'Myopic loss aversion and the equity premium puzzle.' *The quarterly journal of Economics* 110(1), 73–92

Bishop, Christopher M. (2006) *Pattern recognition and machine learning,* vol. 4 (Springer)

Bleichrodt, Han, and Olivier L'Haridon (2023) 'Prospect theoryâs loss aversion is robust to stake size.' *Judgment and Decision Making* 18, e14

Bouchouicha, Ranoua, Ryan Oprea, Ferdinand M. Vieider, and Jilong Wu (2025) 'Is prospect theory really a theory of choice?' Technical Report, Ghent University Discussion Papers

Brañas-Garza, Pablo, Diego Jorrat, Antonio M Espín, and Angel Sánchez (2023) 'Paid and hypothetical time preferences are the same: Lab, field and online evidence.' *Experimental Economics* 26(2), 412–434

Brañas-Garza, Pablo, Lorenzo Estepa-Mohedano, Diego Jorrat, Victor Orozco, and Ericka Rascón-Ramírez (2021) 'To pay or not to pay: Measuring risk preferences in lab and field.' *Judgment and Decision Making* 16(5), 1290–1313

Brown, Alexander L., Taisuke Imai, Ferdinand M. Vieider, and Colin F. Camerer (2024) 'Meta-analysis of empirical estimates of loss aversion.' *Journal of Economic Literature* 62(3), 485–616

Chapman, Jonathan, Erik Snowberg, Stephanie W Wang, and Colin Camerer (2024) 'Looming large or seeming small? attitudes towards losses in a representative sample.' *Review of Economic Studies, forthcoming*

Chapman, Jonathan, Mark Dean, Pietro Ortoleva, Erik Snowberg, and Colin Camerer (2023) 'Willingness to accept, willingness to pay, and loss aversion.' Technical Report, National Bureau of Economic Research

Enke, Benjamin, and Thomas Graeber (2023) 'Cognitive uncertainty.' *The Quarterly Journal of Economics* 138(4), 2021–2067

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

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

Frydman, Cary, and Lawrence J Jin (2023) 'On the source and instability of probability weighting.' Technical Report, National Bureau of Economic Research

Gelman, Andrew, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin (2014) *Bayesian data analysis,* vol. 3 (CRC press Boca Raton, FL)

Gneezy, Uri, and Jan Potters (1997) 'An Experiment on Risk Taking and Evaluation Periods.' *The Quarterly Journal of Economics* 112(2), 631 – 45

Gneezy, Uri, Yoram Halevy, Brian Hall, Theo Offerman, and Jeroen van de Ven (2024) 'How real is hypothetical? a high-stakes test of the allais paradox.' Technical Report

Gonzalez, Richard, and George Wu (1999) 'On the shape of the probability weighting function.' *Cognitive psychology* 38(1), 129–166

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

_ (2023) 'Efficient numerosity estimation under limited time.' *BioRxiv* pp. 2023–

07

Hertwig, Ralph, Greg Barron, Elke U Weber, and Ido Erev (2004) 'Decisions from experience and the effect of rare events in risky choice.' *Psychological science* 15(8), 534–539

Hirmas, Alejandro, Jan B Engelmann, and Joël van der Weele (2024) 'Individual and contextual effects of attention in risky choice.' *Experimental Economics* pp. 1–28

Imai, Taisuke, Salvatore Nunnari, Ferdinand M. Vieider, and Jilong Wu (2025) 'Meta-analysis of empirical estimates of loss aversion.' *Working Paper*

Kahneman, Daniel, and Amos Tversky (1979) 'Prospect Theory: An Analysis of Decision under Risk.' *Econometrica* 47(2), 263 – 291

Khaw, Mel Win, Ziang Li, and Michael Woodford (2021) 'Cognitive imprecision and small-stakes risk aversion.' *The Review of Economic Studies* 88(4), 1979–2013

_ (2023) 'Cognitive imprecision and stake-dependent risk attitudes.' Technical Report

Köbberling, Veronika, and Peter P. Wakker (2005) 'An index of loss aversion.' *Journal of Economic Theory* 122(1), 119 – 131

Kőszegi, Botond, and Matthew Rabin (2007) 'Reference-Dependent Risk Attitudes.' *The American Economic Review* 97(4), 1047–1073

Lobel, Robert Eugene, Marcelo Cabus Klotzle, Paulo Vitor Jordao Da Gama Silva, and Antonio Carlos Figueiredo Pinto (2017) 'Prospect theory: A parametric analysis of functional forms in brazil.' *Revista de Administração de Empresas* 57(5), 495–509

Ma, Wei Ji, Konrad Paul Kording, and Daniel Goldreich (2023) *Bayesian Models of Perception and Action: An Introduction* (MIT press)

Martino, Benedetto De, Colin F. Camerer, and Ralph Adolphs (2010) 'Amygdala damage eliminates monetary loss aversion.' *Proceedings of the National Academy of Sciences* 107(8), 3788–3792

Mehra, Rajnish, and Edward C Prescott (1985) 'The equity premium: A puzzle.'

Journal of monetary Economics 15(2), 145–161

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

Netzer, Nick, Arthur Robson, Jakub Steiner, and Pavel Kocourek (2024) 'Risk perception: measurement and aggregation.' *Journal of the European Economic Association* p. jvae053

Oprea, Ryan (2024) 'Decisions under risk are decisions under complexity.' *American Economic Review, forthcoming*

Oprea, Ryan, and Ferdinand M. Vieider (2024) 'Minding the gap: On the origins of probability weighting and the description-experience gap.' *Working Paper*

Pachur, Thorsten, Michael Schulte-Mecklenbeck, Ryan O Murphy, and Ralph Hertwig (2018) 'Prospect theory reflects selective allocation of attention.' *Journal of experimental psychology: general* 147(2), 147

Prelec, Drazen (1998) 'The probability weighting function.' *Econometrica* pp. 497–527

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

_ (2001b) 'Why would nature give individuals utility functions?' *Journal of Political Economy* 109(4), 900–914

Schmidt, Ulrich, and Horst Zank (2005) 'What is Loss Aversion?' *Journal of Risk and Uncertainty* 30(2), 157–167

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

Sokol-Hessner, Peter, Ming Hsu, Nina G Curley, Mauricio R Delgado, Colin F Camerer, and Elizabeth A Phelps (2009) 'Thinking like a trader selectively reduces individuals' loss aversion.' *Proceedings of the National Academy of Sciences* 106(13), 5035–5040

Stewart, Neil, Nick Chater, and Gordon DA Brown (2006) 'Decision by sampling.' *Cognitive psychology* 53(1), 1–26

Stewart, Neil, Stian Reimers, and Adam JL Harris (2014) 'On the origin of utility, weighting, and discounting functions: How they get their shapes and how to change their shapes.' *Management Science* 61(3), 687–705

Thaler, Richard H, Amos Tversky, Daniel Kahneman, and Alan Schwartz (1997) 'The effect of myopia and loss aversion on risk taking: An experimental test.' *The quarterly journal of economics* 112(2), 647–661

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

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

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

Vieider, Ferdinand M. (2024a) 'Bayesian estimation of decision models.' Technical Report

_ (2024b) 'Decisions under uncertainty as bayesian inference on choice options.' *Management Science, forthcoming* 70(12), 9014–9030

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

Walasek, Lukasz, and Neil Stewart (2015) 'How to make loss aversion disappear and reverse: tests of the decision by sampling origin of loss aversion.' *Journal of experimental psychology: general* 144(1), 7

Walasek, Lukasz, Timothy L Mullett, and Neil Stewart (2021) 'Acceptance of mixed gambles is sensitive to the range of gains and losses experienced, and estimates of lambda (λ) are not a reliable measure of loss aversion: Reply to andré and de langhe (2021).'

Yechiam, Eldad, and Dana Zeif (2025) 'A re-meta-analysis of loss aversion: Beyond the sum of all individual studies.' *Working Paper*

Yechiam, Eldad, and Guy Hochman (2013) 'Losses as modulators of attention: review and analysis of the unique effects of losses over gains.' *Psychological bulletin* 139(2), 497

Zhang, Hang, Xiangjuan Ren, and Laurence T Maloney (2020) 'The bounded rationality of probability distortion.' *Proceedings of the National Academy of Sciences* 117(36), 22024–22034

A Appendix: Stimuli

A.1 Experiment I

The first experiment consists of 128 trials in total, with the stimuli list presented in Table 2.

ID	Gain	Loss	ID	Gain	Loss	ID	Gain	Loss	ID	Gain	Loss
1	**2**	**-2**	33	**12**	**-14**	65	**18**	**-24**	97	**26**	**-20**
2	**2**	**-4**	34	**12**	**-16**	66	**18**	**-26**	98	**26**	**-24**
3	**4**	**-2**	35	**12**	**-18**	67	**18**	**-28**	99	**26**	**-26**
4	**4**	**-4**	36	**12**	**-22**	68	**20**	**-10**	100	**26**	**-28**
5	**4**	**-6**	37	**12**	**-24**	69	**20**	**-14**	101	**28**	**-14**
6	**4**	**-8**	38	**12**	**-26**	70	**20**	**-16**	102	**28**	**-18**
7	**4**	**-10**	39	**14**	**-8**	71	**20**	**-18**	103	**28**	**-20**
8	**6**	**-4**	40	**14**	**-10**	72	**20**	**-20**	104	**28**	**-22**
9	**6**	**-6**	41	**14**	**-12**	73	**20**	**-22**	105	**28**	**-26**
10	**6**	**-8**	42	**14**	**-14**	74	**20**	**-24**	106	**28**	**-28**
11	**6**	**-10**	43	**14**	**-16**	75	**20**	**-26**	107	**28**	**-30**
12	**6**	**-12**	44	**14**	**-18**	76	**20**	**-28**	108	**30**	**-16**
13	**8**	**-4**	45	**14**	**-20**	77	**20**	**-30**	109	**30**	**-20**
14	**8**	**-6**	46	**14**	**-22**	78	**22**	**-12**	110	**30**	**-22**
15	**8**	**-8**	47	**14**	**-28**	79	**22**	**-14**	111	**30**	**-24**
16	**8**	**-10**	48	**16**	**-8**	80	**22**	**-16**	112	**30**	**-28**
17	**8**	**-12**	49	**16**	**-10**	81	**22**	**-18**	113	**30**	**-30**
18	**8**	**-14**	50	**16**	**-12**	82	**22**	**-20**	114	**2**	**-2**
19	**8**	**-16**	51	**16**	**-14**	83	**22**	**-22**	115	**4**	**-4**
20	**8**	**-18**	52	**16**	**-16**	84	**22**	**-24**	116	**6**	**-6**
21	**10**	**-4**	53	**16**	**-18**	85	**22**	**-28**	117	**8**	**-8**
22	**10**	**-6**	54	**16**	**-20**	86	**22**	**-30**	118	**10**	**-10**
23	**10**	**-8**	55	**16**	**-22**	87	**24**	**-12**	119	**12**	**-12**
24	**10**	**-10**	56	**16**	**-24**	88	**24**	**-16**	120	**14**	**-14**
25	**10**	**-12**	57	**16**	**-30**	89	**24**	**-18**	121	**16**	**-16**
26	**10**	**-14**	58	**18**	**-8**	90	**24**	**-20**	122	**18**	**-18**
27	**10**	**-16**	59	**18**	**-12**	91	**24**	**-22**	123	**20**	**-20**
28	**10**	**-20**	60	**18**	**-14**	92	**24**	**-24**	124	**22**	**-22**
29	**12**	**-6**	61	**18**	**-16**	93	**24**	**-26**	125	**24**	**-24**
30	**12**	**-8**	62	**18**	**-18**	94	**24**	**-30**	126	**26**	**-26**
31	**12**	**-10**	63	**18**	**-20**	95	**26**	**-12**	127	**28**	**-28**
32	**12**	**-12**	64	**18**	**-22**	96	**26**	**-18**	128	**30**	**-30**

Table 2: The list of risky wagers. In each trial, subjects must choose to either accept or reject a binary wager. Each wager offers a potential gain G and an equal chance of a potential loss L, with the values of G and L symmetrically distributed across trials. An additional 15 trials are repeated under the condition $G = L$. All monetary outcomes are presented in GBP.

A.2 Experiment II

The second experiment consists of two phases: an adaptation phase followed by a test phase. In the loss conditions, the adaptation phase includes 100 pure-loss choices, where subjects have to choose between a risky wager $(-x, 0.5; 0)$ and a sure loss c, with subjects randomly assigned to either a small loss or a large loss treatment. In the gain conditions, the adaptation phase includes 100 pure-gain choices, where subjects have to choose between a risky wager $(x, 0.5; 0)$ and a sure gain c, with subjects randomly assigned to either a small gain or a large gain treatment. The stimuli list for the adaptation trials is presented in Table 3. The 100 test trials are identical across conditions. The stimuli list for the test trials is presented in Table 4.

Small treatment						Large treatment					
ID	x	c	ID	x	c	ID	x	c	ID	x	c
1	**2**	**1**	51	**11**	**9**	1	**4**	**2**	51	**32**	**4**
2	**3**	**1**	52	**11**	**10**	2	**8**	**2**	52	**32**	**6**
3	**3**	**2**	53	**12**	**1**	3	**8**	**4**	53	**32**	**8**
4	**4**	**2**	54	**12**	**2**	4	**8**	**6**	54	**32**	**10**
5	**4**	**3**	55	**12**	**3**	5	**12**	**2**	55	**32**	**12**
6	**5**	**1**	56	**12**	**4**	6	**12**	**4**	56	**32**	**14**
7	**5**	**2**	57	**12**	**5**	7	**12**	**6**	57	**32**	**16**
8	**5**	**4**	58	**12**	**6**	8	**12**	**8**	58	**32**	**18**
9	**6**	**1**	59	**12**	**7**	9	**12**	**10**	59	**32**	**20**
10	**6**	**3**	60	**12**	**8**	10	**16**	**2**	60	**32**	**22**
11	**6**	**4**	61	**12**	**9**	11	**16**	**4**	61	**32**	**24**
12	**6**	**5**	62	**12**	**10**	12	**16**	**6**	62	**32**	**26**
13	**7**	**1**	63	**12**	**11**	13	**16**	**8**	63	**32**	**28**
14	**7**	**2**	64	**13**	**1**	14	**16**	**10**	64	**32**	**30**
15	**7**	**3**	65	**13**	**2**	15	**16**	**12**	65	**36**	**2**
16	**7**	**4**	66	**13**	**3**	16	**16**	**14**	66	**36**	**4**
17	**7**	**5**	67	**13**	**4**	17	**20**	**2**	67	**36**	**6**
18	**7**	**6**	68	**13**	**5**	18	**20**	**4**	68	**36**	**8**
19	**8**	**1**	69	**13**	**6**	19	**20**	**6**	69	**36**	**10**
20	**8**	**2**	70	**13**	**7**	20	**20**	**8**	70	**36**	**12**
21	**8**	**3**	71	**13**	**8**	21	**20**	**10**	71	**36**	**14**
22	**8**	**4**	72	**13**	**9**	22	**20**	**12**	72	**36**	**16**
23	**8**	**5**	73	**13**	**10**	23	**20**	**14**	73	**36**	**18**
24	**8**	**6**	74	**13**	**11**	24	**20**	**16**	74	**36**	**20**
25	**8**	**7**	75	**13**	**12**	25	**20**	**18**	75	**36**	**22**
26	**9**	**1**	76	**14**	**1**	26	**24**	**2**	76	**36**	**24**
27	**9**	**2**	77	**14**	**2**	27	**24**	**4**	77	**36**	**26**
28	**9**	**3**	78	**14**	**3**	28	**24**	**6**	78	**36**	**28**
29	**9**	**4**	79	**14**	**4**	29	**24**	**8**	79	**36**	**30**
30	**9**	**5**	80	**14**	**5**	30	**24**	**10**	80	**36**	**32**
31	**9**	**6**	81	**14**	**6**	31	**24**	**12**	81	**36**	**34**

Continued from previous page

Small treatment						Large treatment					
ID	x	c	ID	x	c	ID	x	c	ID	x	c
32	**9**	**7**	82	**14**	**7**	32	**24**	**14**	82	**40**	**2**
33	**9**	**8**	83	**14**	**8**	33	**24**	**16**	83	**40**	**4**
34	**10**	**1**	84	**14**	**9**	34	**24**	**18**	84	**40**	**6**
35	**10**	**2**	85	**14**	**10**	35	**24**	**20**	85	**40**	**8**
36	**10**	**3**	86	**14**	**11**	36	**24**	**22**	86	**40**	**10**
37	**10**	**4**	87	**14**	**12**	37	**28**	**2**	87	**40**	**12**
38	**10**	**5**	88	**14**	**13**	38	**28**	**4**	88	**40**	**14**
39	**10**	**6**	89	**15**	**1**	39	**28**	**6**	89	**40**	**16**
40	**10**	**7**	90	**15**	**2**	40	**28**	**8**	90	**40**	**18**
41	**10**	**8**	91	**15**	**3**	41	**28**	**10**	91	**40**	**20**
42	**10**	**9**	92	**15**	**5**	42	**28**	**12**	92	**40**	**22**
43	**11**	**1**	93	**15**	**6**	43	**28**	**14**	93	**40**	**24**
44	**11**	**2**	94	**15**	**7**	44	**28**	**16**	94	**40**	**26**
45	**11**	**3**	95	**15**	**8**	45	**28**	**18**	95	**40**	**28**
46	**11**	**4**	96	**15**	**9**	46	**28**	**20**	96	**40**	**30**
47	**11**	**5**	97	**15**	**10**	47	**28**	**22**	97	**40**	**32**
48	**11**	**6**	98	**15**	**11**	48	**28**	**24**	98	**40**	**34**
49	**11**	**7**	99	**15**	**13**	49	**28**	**26**	99	**40**	**36**
50	**11**	**8**	100	**15**	**14**	50	**32**	**2**	100	**40**	**38**

Table 3: In Experiment II, subjects are randomly assigned to either a small or a large treatment condition during the adaptation phase. In the loss conditions, subjects have to choose between a risky wager $(-x, 0.5; 0)$ and a sure loss c, whereas in the gain conditions, subjects have to choose between a risky wager $(x, 0.5; 0)$ and a sure gain c. All monetary outcomes are presented in GBP.

ID	**Gain**	**Loss**	ID	**Gain**	**Loss**	ID	**Gain**	**Loss**	ID	**Gain**	**Loss**
1	**4**	**-4**	26	**12**	**-24**	51	**24**	**-4**	76	**32**	**-24**
2	**4**	**-8**	27	**12**	**-28**	52	**24**	**-8**	77	**32**	**-28**
3	**4**	**-12**	28	**12**	**-32**	53	**24**	**-12**	78	**32**	**-32**
4	**4**	**-16**	29	**12**	**-36**	54	**24**	**-16**	79	**32**	**-36**
5	**4**	**-20**	30	**12**	**-40**	55	**24**	**-20**	80	**32**	**-40**
6	**4**	**-24**	31	**16**	**-4**	56	**24**	**-24**	81	**36**	**-4**
7	**4**	**-28**	32	**16**	**-8**	57	**24**	**-28**	82	**36**	**-8**
8	**4**	**-32**	33	**16**	**-12**	58	**24**	**-32**	83	**36**	**-12**
9	**4**	**-36**	34	**16**	**-16**	59	**24**	**-36**	84	**36**	**-16**
10	**4**	**-40**	35	**16**	**-20**	60	**24**	**-40**	85	**36**	**-20**
11	**8**	**-4**	36	**16**	**-24**	61	**28**	**-4**	86	**36**	**-24**
12	**8**	**-8**	37	**16**	**-28**	62	**28**	**-8**	87	**36**	**-28**
13	**8**	**-12**	38	**16**	**-32**	63	**28**	**-12**	88	**36**	**-32**
14	**8**	**-16**	39	**16**	**-36**	64	**28**	**-16**	89	**36**	**-36**
15	**8**	**-20**	40	**16**	**-40**	65	**28**	**-20**	90	**36**	**-40**
16	**8**	**-24**	41	**20**	**-4**	66	**28**	**-24**	91	**40**	**-4**
17	**8**	**-28**	42	**20**	**-8**	67	**28**	**-28**	92	**40**	**-8**
18	**8**	**-32**	43	**20**	**-12**	68	**28**	**-32**	93	**40**	**-12**
19	**8**	**-36**	44	**20**	**-16**	69	**28**	**-36**	94	**40**	**-16**
20	**8**	**-40**	45	**20**	**-20**	70	**28**	**-40**	95	**40**	**-20**
21	**12**	**-4**	46	**20**	**-24**	71	**32**	**-4**	96	**40**	**-24**
22	**12**	**-8**	47	**20**	**-28**	72	**32**	**-8**	97	**40**	**-28**
23	**12**	**-12**	48	**20**	**-32**	73	**32**	**-12**	98	**40**	**-32**
24	**12**	**-16**	49	**20**	**-36**	74	**32**	**-16**	99	**40**	**-36**
25	**12**	**-20**	50	**20**	**- 40**	75	**32**	**-20**	100	**40**	**-40**

Table 4: The list of risky wagers. In each trial, subjects must choose to either accept or reject a binary wager. Each wager offers a potential gain G and an equal chance of a potential loss L, with the values of G and L symmetrically distributed across trials. The test trials are identical across conditions. All monetary outcomes are presented in GBP.

B Appendix: Derivation of the Noisy Coding Model

In this section, we provide further details regarding the derivations underlying the equations presented in the main text.

Assume subjects are presented with a choice between accepting a wager $(G, p; -L)$ or rejecting it, where $G, L > 0$, and $p = 1 - p = 0.5$ throughout. By applying optimal Bayesian updating, we combine the likelihoods in equation (1) with the prior in equation (2), yielding the following posterior distributions of logarithmic gains and losses, conditional on mental signals:

$$\ln(G) \mid r_g \sim \mathcal{N}\left(\frac{\sigma^2}{\sigma^2 + \nu_g^2} r_g + \frac{\nu_g^2}{\sigma^2 + \nu_g^2}\mu, \frac{\nu_g^2 \sigma^2}{\sigma^2 + \nu_g^2}\right)$$

$$\ln(L) \mid r_\ell \sim \mathcal{N}\left(\frac{\sigma^2}{\sigma^2 + \nu_\ell^2} r_\ell + \frac{\nu_\ell^2}{\sigma^2 + \nu_\ell^2}\mu, \frac{\nu_\ell^2 \sigma^2}{\sigma^2 + \nu_\ell^2}\right)$$

Defining the parameters $\alpha \triangleq \frac{\sigma^2}{\sigma^2 + \nu_g^2}$, $\beta \triangleq \frac{\sigma^2}{\sigma^2 + \nu_\ell^2}$, $\widehat{\mu} \triangleq \mu + \frac{\sigma^2}{2}$. Following the properties of the log-normal distribution, the posterior means of gains and losses are:

$$\begin{aligned}
\mathbb{E}[G \mid r_g] &= \exp\left\{\frac{\sigma^2}{\sigma^2 + \nu_g^2} r_g + \frac{\nu_g^2}{\sigma^2 + \nu_g^2}\mu + \frac{1}{2}\frac{\nu_g^2 \sigma^2}{\sigma^2 + \nu_g^2}\right\} \\
&= \exp\left\{\alpha r_g + (1 - \alpha)\widehat{\mu}\right\}
\end{aligned}$$

$$\begin{aligned}
\mathbb{E}[L \mid r_\ell] &= \exp\left\{\frac{\sigma^2}{\sigma^2 + \nu_\ell^2} r_\ell + \frac{\nu_\ell^2}{\sigma^2 + \nu^2}\mu + \frac{1}{2}\frac{\nu_\ell^2 \sigma^2}{\sigma^2 + \nu_\ell^2}\right\} \\
&= \exp\left\{\beta r_\ell + (1 - \beta)\widehat{\mu}\right\}
\end{aligned}$$

Substituting the posterior expectations into the mental choice rule, the decision maker will choose the wager whenever $\mathbb{E}[G \mid r_g] > \mathbb{E}[L \mid r_\ell]$, resulting in:

$$\exp\left\{\alpha r_g + (1 - \alpha)\widehat{\mu}\right\} > \exp\left\{\beta r_\ell + (1 - \beta)\widehat{\mu}\right\}$$

Since any monotonic transformation leaves this choice rule unaltered, we take the natural logarithm and rewrite it as:

$$\begin{aligned}
\alpha r_g - \beta r_\ell &> (1 - \beta)\widehat{\mu} - (1 - \alpha)\widehat{\mu} \\
&> (\alpha - \beta)\,\widehat{\mu}
\end{aligned}$$

This implies that the wager should be chosen if and only if $\alpha r_g - \beta r_\ell$ exceeds the threshold $(\alpha - \beta)\,\widehat{\mu}$.

The noisy coding model suggests that, conditional on the objective stimuli (G, L), $\alpha r_g - \beta r_\ell$ is a Gaussian random variable:

$$\alpha r_g - \beta r_\ell \sim \mathcal{N}\left(\alpha \ln(G) - \beta \ln(L), \nu_g^2 \alpha^2 + \nu_\ell^2 \beta^2\right)$$

This leads to the following z-score:

$$z = \frac{\alpha r_g - \beta r_\ell - [\alpha \ln(G) - \beta \ln(L)]}{\sqrt{\nu_g^2 \alpha^2 + \nu_\ell^2 \beta^2}}$$

which follows a standard normal distribution. Comparing this z-score to the condition for accepting the wager gives the stochastic choice rule in equation (3):

$$\Pr\left[(G, 0.5; -L) \succ 0\right] = \Phi\left(\frac{\alpha \ln(G) - \beta \ln(L) + (\beta - \alpha)\,\widehat{\mu}}{\sqrt{\nu_g^2 \alpha^2 + \nu_\ell^2 \beta^2}}\right)$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution.

For fair spreads offering $\pm x$ with even odds we have the following equation:

$$\Pr\left[(x, 0.5; -x) \succ 0\right] = \Phi\left(\frac{(\beta - \alpha)\,[\widehat{\mu} - \ln(x)]}{\sqrt{\nu_g^2 \alpha^2 + \nu_\ell^2 \beta^2}}\right)$$

With the assumption of asymmetric attention, i.e., $\nu_\ell < \nu_g$, we have $0 < \alpha < \beta < 1$. Under this condition, the model predicts that the choice probability of the wager declines with increasing stake size. The transformed common prior mean $\widehat{\mu}$ determines the 'fixed point', representing the threshold of stochastic indifference between the choice options, given by $e^{\widehat{\mu}}$. When $\widehat{\mu} > 0$, it predicts a coexistence of risk seeking for small stakes ($0 < x < e^{\widehat{\mu}}$) and risk aversion for large stakes ($x > e^{\widehat{\mu}}$).

The co-existence of stake-dependence and risk seeking for small stakes can, of course, also be organized by a more general model that allows for different priors for gains and losses. Let the prior for gains be $\mathcal{N}(\mu_g, \sigma_g^2)$, and the prior for losses $\mathcal{N}(\mu_\ell, \sigma_\ell^2)$, the posterior means of gains and losses can be re-written as:

$$\mathbb{E}[G \mid r_g] = \exp\left\{\alpha r_g + (1 - \alpha)\widehat{\mu}_g\right\}$$

$$\mathbb{E}[L \mid r_\ell] = \exp\left\{\beta r_\ell + (1 - \beta)\widehat{\mu}_\ell\right\}$$

where $\alpha \triangleq \frac{\sigma_g^2}{\sigma_g^2 + \nu_g^2}$, $\beta \triangleq \frac{\sigma_\ell^2}{\sigma_\ell^2 + \nu_\ell^2}$, $\widehat{\mu}_g \triangleq \mu_g + \frac{\sigma_g^2}{2}$, $\widehat{\mu}_\ell \triangleq \mu_\ell + \frac{\sigma_\ell^2}{2}$.

Following the similar derivation as above, the stochastic choice rule is given as

$$\Pr\left[(G, 0.5; -L) \succ 0\right] = \Phi\left(\frac{\alpha \ln(G) - \beta \ln(L) - [(1 - \beta)\widehat{\mu}_\ell - (1 - \alpha)\widehat{\mu}_g]}{\sqrt{\nu_g^2 \alpha^2 + \nu_\ell^2 \beta^2}}\right)$$

By exponentiating the numerator, one can see that the choice probability of the wager will be proportional to $G^\alpha - \lambda\, L^\beta$, where

$$\lambda \triangleq \exp\left[(1 - \beta)\,\widehat{\mu}_\ell - (1 - \alpha)\,\widehat{\mu}_g\right]$$

It provides a micro-foundations of the origin of loss aversion modeled in PT as a kink in the utility function. When $(1-\beta)\widehat{\mu}_\ell < (1-\alpha)\widehat{\mu}_g$, it implies $\lambda < 1$ (i.e. gain seeking), which can also account for the small-stake risk seeking we find.

Khaw et al. (2021) discuss a specific restriction under which the ratio $\frac{\nu_g}{\sigma_g}$ is the same as $\frac{\nu_\ell}{\sigma_\ell}$, implying $\alpha = \beta$, then the stochastic choice rule simplifies to:

$$\Pr\left[(G, p; -L) \succ 0\right] = \Phi\left(\frac{\alpha \ln\left(\frac{G}{L}\right) - \left[(1-\alpha)(\widehat{\mu}_\ell - \widehat{\mu}_g)\right]}{\alpha\sqrt{\nu_g^2 + \nu_\ell^2}}\right)$$

For fair spreads, where $G = L = x$, the expression $\ln\left(\frac{G}{L}\right)$ drops out, and choices depend only on $\ln(\lambda) \triangleq (1-\alpha)(\widehat{\mu}_\ell - \widehat{\mu}_g)$, regardless of stake size. If $\mu_\ell \geq \mu_g$, $\sigma_\ell \geq \sigma_g$, and at least one inequality is strict, the model predicts $\lambda > 1$, implying absolute loss aversion at all stakes.

C A Model of Attention in the NCM

Here, we discuss a stylized model of how attention will manifest within the Noisy Cognition Model. We thereby start from the empirical literature on attention to gains and losses. Pachur et al. (2018) (using Mouselab data) and Hirmas et al. (2024) (using eye-tracking data) both implemented attention as dwelling time on a given attribute, and showed that relative dwelling times on losses versus gains were predictive of risk taking in mixed gain-loss gambles. Here, we integrate this notion into the NCM in a highly stylized discrete time version. Heng et al. (2023) provide a more sophisticated continuous-time model that, applied to attention, would result in very similar predictions to the ones we derive here.

Let us assume without loss of generality that there is a fixed number of neurons producing the signals in equation (1). We also assume that there is a fixed number of action potentials per second a neuron can fire. These two assumptions immediately imply a fixed mapping from the time spent considering an attribute to the precision of the error. For simplicity's sake, we will here focus directly on the signal gains *relative* to losses for an equal outcome x. This allows us to avoid notational clutter, but otherwise happens without loss of generality. The coding noise affecting a specific observation will now be dependent on the number of draws—i.e. the time spent considering the given attribute. We will now use r_{ij} to indicate a single draw j from the likelihood for a given stimulus x_i:

$$r_{ij} \sim \mathcal{N}(\, ln(x_i)\, ,\, \nu^2\,),$$

We will use $\overline{r}_i \triangleq \frac{1}{m_i} \sum_{j=1}^{m_i} r_{ij}$ to indicate the mean of the draws for a given stimulus x_i. Since the draws are independent conditional on the stimulus x_i, the variance of the signal will be $\nu_i^2 \triangleq \frac{\nu^2}{m_i}$, where m_i indicates the number of signals drawn for a given stimulus x_i. From this, it follows directly that stimuli that receive increased attention—in the sense of DMs dwelling on them for an increased time period—will have lower coding noise. The model, albeit highly stylized, thus captures the key mechanism behind loss-sensitivity we refer to in the main text.

D Appendix: Choice Proportions for Fair Spreads

D.1 Experiment I

Stakes (£)	Proportions (%)
2	66.70 (0.50)
4	62.01 (0.53)
6	57.48 (0.55)
8	52.59 (0.56)
10	53.19 (0.56)
12	46.96 (0.56)
14	42.32 (0.55)
16	40.93 (0.54)
18	39.88 (0.54)
20	38.73 (0.53)
22	37.24 (0.52)
24	36.49 (0.52)
26	35.19 (0.51)
28	35.64 (0.51)
30	34.25 (0.50)

Table 5: The table shows the acceptance proportions of risky wagers for fair spreads $(x, 0.5; -x)$, standard errors are in parentheses.

D.2 Experiment II: Losses

Stakes (£)	Small losses (%)	Large losses (%)
4	48.15 (2.40)	56.52 (2.56)
8	37.96 (2.27)	51.09 (2.61)
12	36.11 (2.22)	46.74 (2.60)
16	36.11 (2.22)	43.48 (2.56)
20	35.19 (2.19)	42.39 (2.55)
24	28.70 (1.97)	34.78 (2.37)
28	30.56 (2.04)	32.61 (2.29)
32	32.41 (2.11)	31.52 (2.25)
36	29.63 (2.01)	31.52 (2.25)
40	29.63 (2.01)	35.87 (2.40)

Table 6: The table shows the acceptance proportions of risky wagers for fair spreads $(x, 0.5; -x)$, standard errors are in parentheses. In the "large losses" treatment, subjects have been exposed to large losses in the adaptation phase; in the "small losses" treatment, they have been exposed to small losses. The test data are identical across treatments.

D.3 Experiment II: Gains

Stakes (£)	Small gains (%)	Large gains (%)
4	32.97 (2.32)	41.67 (2.34)
8	20.88 (1.73)	35.19 (2.19)
12	15.38 (1.36)	37.04 (2.24)
16	18.68 (1.59)	31.48 (2.08)
20	16.48 (1.44)	34.26 (2.17)
24	15.38 (1.36)	34.26 (2.17)
28	10.99 (1.03)	32.41 (2.11)
32	10.99 (1.03)	26.85 (1.89)
36	9.89 (0.93)	34.26 (2.17)
40	13.19 (1.20)	31.48 (2.08)

Table 7: The table shows the acceptance proportions of risky wagers for fair spreads $(x, 0.5; -x)$, standard errors are in parentheses. In the "large gains" treatment, subjects have been exposed to large gains in the adaptation phase; in the "small gains" treatment, they have been exposed to small gains. The test data are identical across treatments.

E Appendix: Logistic Regression

E.1 Logistic regression model

The probability of selecting the risky option is modeled using logistic regression, as follows:

$$Choice_i \sim \text{Bernoulli}(Prob_i)$$

where

$$\text{Logistic}^{-1}(Prob_i) = \gamma_0^j + \gamma_g \ln(G_i) + \gamma_l \ln(L_i).$$

Here, G_i and L_i represent the gain and loss outcomes, respectively, for the i-th trial. A logarithmic transformation is applied to both G_i and L_i in the regression.

We adopt a Bayesian hierarchical model that incorporates both subject-level and aggregate-level parameters. The subject-specific intercept γ_0^j for the j-th individual is modeled as:

$$\gamma_0^j \sim \mathcal{N}(\bar{\gamma}_0, \sigma),$$

where $\bar{\gamma}_0$ denotes the aggregate-level intercept, and σ captures heterogeneity across subjects. The priors for the aggregate-level parameters are specified as:

$$\bar{\gamma}_0 \sim \mathcal{N}(0, 5),$$

$$\sigma \sim \text{Exponential}(1),$$
$$\gamma_g \sim \mathcal{N}(0, 5),$$
$$\gamma_l \sim \mathcal{N}(0, 5).$$

These distributions are considerably wider than the range into which we would expect the estimated parameters to fall based on the scale of the data. This means that they are 'mildly regularizing'—they aid convergence of the estimation algorithm, without however swaying the results. Within each treatment group, the aggregate-level intercept $\bar{\gamma}_0$, as well as the slope coefficients γ_g and γ_l, are treated as fixed effects.

E.2 Estimation results

The estimation is performed using Stan for each experiment. We employ 4 chains, each with 10,000 iterations, of which the first 5,000 are warm-up iterations. Thus, a total of 20,000 iterations are used to draw the posterior estimates for each parameter. For Experiment II, the parameters for the small and large treatment conditions are estimated simultaneously and treated as 'fixed effects'. Details of the programming are provided in Appendix E.3. In this model, the aggregate-level intercept $\bar{\gamma}_0$, as well as the slope coefficients γ_g and γ_l, are assumed to be treatment-specific, i.e., fixed within each treatment group but vary across treatment groups, while σ is assumed to remain constant across treatments. Convergence is assessed by examining the R-hat statistics and checking for divergent iterations. The estimation results are summarized in Table 8.

We find clear evidence of loss sensitivity across all treatments, as the difference between the slope coefficient on the log-loss and the slope coefficient on the log-gain, i.e., $|\gamma_l| - |\gamma_g|$, is significantly greater than 0. Additionally, the aggregate intercept $\bar{\gamma}_0$ in the large treatment condition is significantly higher than in the low treatment condition. This holds both for the loss adaptation conditions (difference in posteriors: 1.348, 95% CrI [0.281, 2.412]), and for the gain adaptation conditions (difference in posteriors: 1.494, 95% CrI [0.243, 2.749]). This finding is further supported by a Wilcoxon rank-sum test conducted on the individual-level mean intercepts, yielding respondents exposed to the large treatment condition during the adaptation phase exhibit significantly higher intercepts both following the loss adaptation ($p = 1.3 \times 10^{-6}$) and the gain adaptation ($p = 0.019$).

	Representative	**Small Losses**	**Large Losses**	**Small Gains**	**Large Gains**
$\bar{\gamma}_0$	1.91	-0.11	1.24	-1.79	-0.30
	(1.69, 2.14)	(-0.83, 0.62)	(0.48, 2.01)	(-2.75, -0.85)	(-1.13, 0.52)
γ_g	5.97	4.14	3.27	4.45	3.92
	(5.89, 6.05)	(3.94, 4.35)	(3.09, 3.45)	(4.18, 4.72)	(3.71, 4.12)
γ_l	-7.09	-4.47	-4.01	-5.04	-4.34
	(-7.18, -7.00)	(-4.67, -4.27)	(-4.20, -3.82)	(-5.31, -4.79)	(-4.54, -4.14)
$\|\gamma_l\| - \|\gamma_g\|$	1.13	0.32	0.74	0.59	0.43
	(1.09, 1.16)	(0.19, 0.47)	(0.60, 0.88)	(0.41, 0.77)	(0.28, 0.57)
$\frac{\bar{\gamma}_0}{\|\gamma_l\| - \|\gamma_g\|}$	1.70	-0.52	1.66	-3.17	-0.83
	(1.51, 1.88)	(-3.68, 1.62)	(0.71, 2.55)	(-6.09, -1.21)	(-3.51, 1.06)

Table 8: Aggregate-level estimates with 95% confidence intervals shown in parentheses. $\bar{\gamma}_0$ represents the aggregate-level intercept, while γ_g and γ_l denote the slope coefficients for the log-gain and log-loss, respectively. Experiment I employs a representative sample, while Experiment II incorporates an adaptation phase featuring small or large treatment trials in either the gain or loss domain.

E.3 Stan Code

Below, we provide the Stan code used for the estimation. Vieider (2024a) provides a tutorial on how to use Stan and on how to launch the code from RStudio.

```
data {
  int<lower=0> N;
  int<lower=0> Nid;
  array[N] int id;
  array[N] int choice_risky;
  vector[N] lg;
  vector[N] ll;
  array[Nid] int<lower=1, upper=2> treat;
  array[N] int<lower=1, upper=2> treatn;
}

parameters {
  vector[Nid] alpha;          // Individual-level intercepts

  vector[max(treat)] alpha_hat;
  real<lower=0> alpha_sd;
```

```

  vector[max(treat)] gamma;
  vector[max(treat)] lambda;

}

model {
  // Priors for treatment means and standard deviations
  alpha_hat ~ normal(0,5);
  alpha_sd ~ exponential(1);
  lambda ~ normal(0,5);
  gamma ~ normal(0,5);

  // Individual-level priors based on treatment
  for (i in 1:Nid)
    alpha[i] ~ normal(alpha_hat[treat[i]], alpha_sd);

  // Likelihood for choice
  choice_risky ~ bernoulli_logit(  alpha[id] + gamma[treatn] .* lg
      + lambda[treatn] .* ll );
}

generated quantities{
  vector[max(treat)] lmg;
  vector[max(treat)] mu;

 for (t in 1:max(treat)){
  lmg[t] = - lambda[t] - gamma[t];
  mu[t] = alpha_hat[t] / lmg[t];
 }
}
```

F Appendix: Meta-analysis of Prospect Theory

F.1 Choice prediction from existing PT estimations

In the main text, we claim that the type of stake-dependence we document requires gain-seeking in the sense $\lambda < 1$ To show this, let acceptance of the mixed gain-loss wagers we use to be captured by the condition $w^+(0.5)\,u(x) > \lambda\,w^-(0.5)$ u(-x). Since utility for losses needs to be steeper than for gains to capture any stake dependence, capturing the co-existence of risk seeking and risk aversion over fair spreads we document in our data requires the following condition to hold:

$$\frac{1}{\lambda} \times \frac{w^+(0.5)}{w^-(0.g)} > \frac{u(-x)}{u(x)} > 1.$$

Original prospect theory Kahneman and Tversky (1979), as well as many behavioural definitions of loss aversion, either impose equal probability weighting for gains and losses, so that $w^+(0.5) = w^-(0.5)$, or assume that probability weighting is 'edited out' in decision making for even odds (Schmidt and Zank, 2005). In those cases, the condition obtains trivially: once choice patterns are stake-dependent, capturing this requires $\frac{u(-x)}{u(x)} > 1 \; \forall \; x$. Since $\frac{w^+(0.5)}{w^-(0.g)} 1$ by definition in such models, a necessary condition for risk seeking over small stakes will be $\frac{1}{\lambda} > 1$ (the sufficient condition is $\frac{1}{\lambda}) > \frac{u(-x)}{u(x)}$.

The modern (cumulative) version requires a different approach, since the relation is now influenced by the ratio of decision weights, $\frac{w^+(0.5)}{w^-(0.g)}$. We thus take the PT parameters in the meta-analysis of Imai et al. (2025), which encodes the universe of PT estimates. The dataset contains 174 estimates of utility and probability weighting function parameters for gains and losses, as well as loss aversion.[24] Of these, 52 estimates (30%) assume utility curvature to be the same for gains and losses, and are thus not informative for stake-dependence.

Figure 5 presents a scatter plot of the remaining 122 estimates with the 'relative utility' for £10 (capturing the difference in utilities attached to a loss versus a gain of £10 purely due to differential sensitivity) on the x-axis, and the 'relative optimism' (the inverse of λ multiplied by the ratio of decisions weights for gains versus losses) on the y-axis.[25] Of the 122 estimates, 31 (25%) exhibit sensitivity towards losses that decreases more rapidly than sensitivity towards gains (grey dots to the left of the vertical line at 1), thus failing the necessary condition for

[24] An estimate is defined as a single reported measure of PT functionals. Different estimates could thus stem from different papers, from different experiments within the same paper, or even from different measures obtained from the same experimental data (e.g., from different functional form assumptions, or from reporting the mean versus the median of individual-level estimates). We choose the level of the estimate for our analysis to cast as wide a net as possible, and to avoid excluding estimates that could account for our data.

[25] We use £10 to illustrate the patterns since this is the largest amount for which we find significant risk seeking in our data. Note, however, that the results are not very sensitive to the amount used, and the implications discussed below also hold if we were to use a smaller amount instead.

stake-dependence (this includes the estimates in the Online Appendix of Chapman et al., 2024). Of the remaining estimates, the great majority (96%) falls below the dashed ascending 45-degree line, thus failing condition 2) under which 'relative optimism' must exceed the utility ratio for risk seeking to be observed for small stakes.

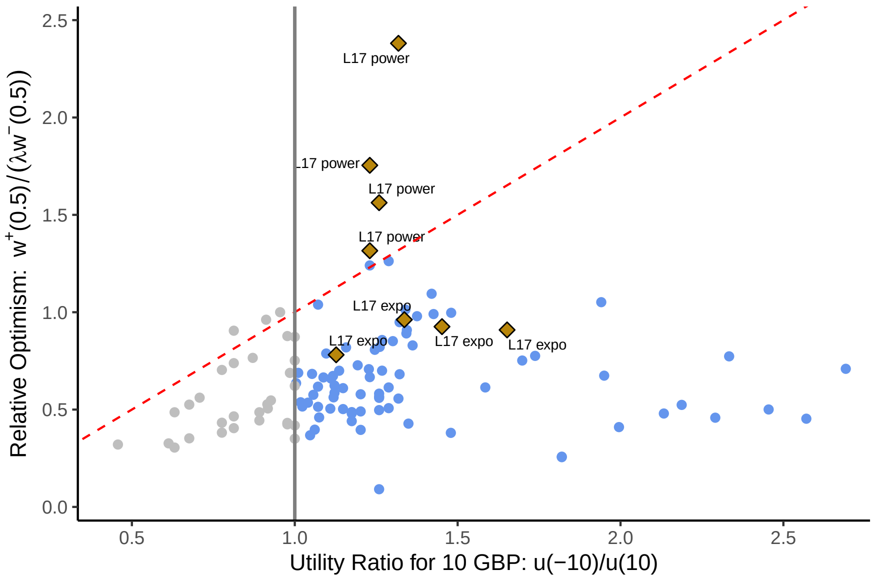

Figure 5: Scatter plot of PT parameter combinations

The figure shows a scatter plot of the PT parameter combinations under which stake dependence can coexist with gain seeking for small stakes. Estimates to the left of the vertical line fall foul of condition 1), which entails that $\frac{u(-10)}{u(10)} > 1$, and do thus not predict any stake dependence. Estimates that fall to the right of the vertical line but below the dashed ascending 45-degree line fall foul of the additional condition 2) according to which we need $\frac{w^+(0.5)}{\lambda w^-(0.5)} > \frac{u(-10)}{u(10)}$ to obtain risk seeking for small stakes. Only points above the red dashed line and to the right of the vertical solid line are qualitatively in line with the results we report. Note that a few outlying observations (which do not predict our results) have been cut from the figure to improve the visual display.

Only four estimates fulfill the condition under which the qualitative patterns we document ought to occur. All four estimates summarize the mean and median individual-level parameters from one and the same dataset by Lobel, Klotzle, Silva and Pinto (2017), combining power utility with different forms of probability weighting functions. Interestingly, the same paper includes four more estimates in which exponential utility is used instead (labelled as 'L17 expo' in the figure). These estimates reveal a rather different story, with all 4 of them falling below the dashed line. What distinguishes the 'L17 power' from the 'L17 expo' estimates is the loss aversion parameter, λ. Whereas the specifications using exponential utility all estimate mild loss aversion ($\lambda > 1$), the power utility specifications invariably imply substantial gain seeking ($\lambda \leq 0.76$). It is well-known in the PT literature that estimates of λ obtained with different power utility parameters for gains and losses are arbitrary, since they depend purely on the numerical scale of the outcomes (i.e. simply rescaling outcomes from pounds to pence in the estimation would result in substantially different estimates of λ; see Köbberling and Wakker, 2005; Wakker, 2010).

More to the point, the only estimates in the PT literature that can at least qualitatively account for our findings can do so because they estimate $\lambda < 1$, the same condition at which we point in the main text. Figure 6 further examines the fit of the meta-analytic evidence to the choice patterns we observe in fair spreads (see Online Appendix F.2 for details on the parameters). The function based on the parameters reported by Tversky and Kahneman (1992) falls very far from the choice proportions we observe, which are indicated by the large blue dots. It does not predict any stake dependence, since utility curvature for gains and losses is estimated to be equal. The estimated parameters instead predict consistent and strong risk aversion. The prediction arising from this parameter combination does not fall far from the predictions from the third quartile of our meta-analysis parameters. The median parameter values from the meta-analysis also predict universal risk aversion over fair spreads, as indeed does the first quartile of the parameters.

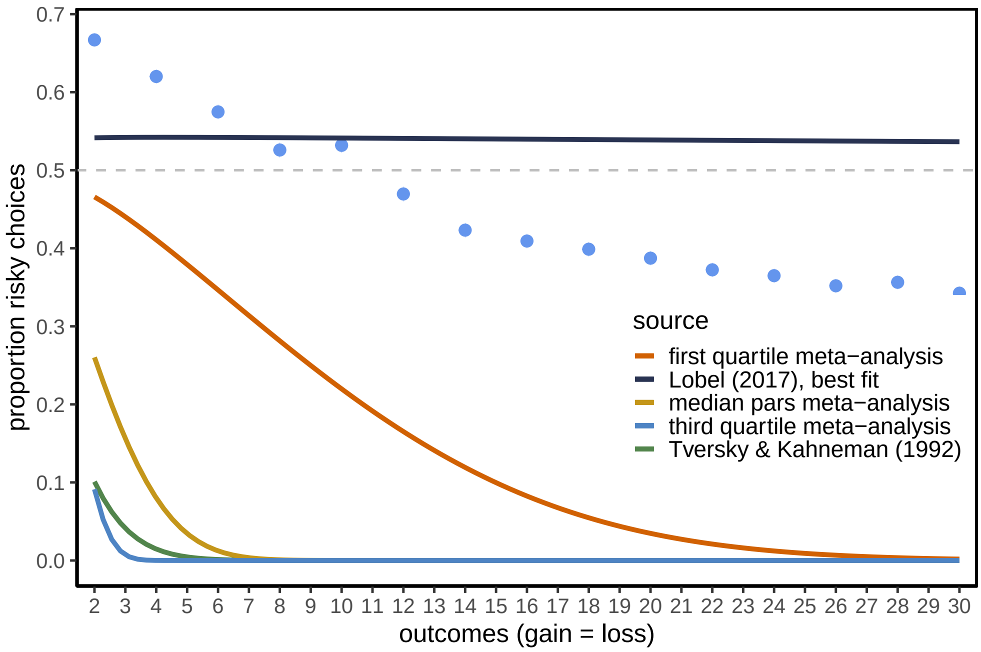

Figure 6: Fit of predictions from PT functionals to nonparametric choice data

The figure shows the nonparametric choice data (blue dots), together with functions predicting choice proportions based on a PT model (see Online Appendix F.2 for details). The Tversky and Kahneman (1992) relies on the famous PT parameters estimated in that paper. It produces substantial loss aversion throughout. The prediction from Lobel et al. (2017) is based on the topmost point from Figure 5, which is the closest to our data since it predicts the highest level of small-stake risk seeking.

The prediction from the study of Lobel et al. (2017) closest to our nonparametric choice proportions—the best predictor from the only experiment according to our meta-analysis that can at least in principle predict the qualitative patterns we observe—falls very far from the actual patterns. The reason is that stake-dependence is quite weak, resulting in a flat curve predicting universal mild risk seeking over the range of stakes we have in our study. While the patterns we observe can thus certainly be organized by *some* parameter combination under PT, we conclude that those combinations are highly unlikely based on the existing empirical evidence about PT functionals.

F.2 Dataset of PT literature estimates

Table 9 presents the database of Prospect Theory (PT) quantities displayed in Figure 5, which illustrates a scatter plot of meta-analytic PT parameter combinations under which stake dependence can coexist with gain-seeking behavior for small stakes. Two types of utility functions are employed to compute (dis)utilities for £10 gains and losses, following the utility specifications in the original papers from which we take our data:

1. In the CRRA case, the power utility function is used:

$$u(x) = \begin{cases} x^{\alpha} & \text{if } x \geq 0, \\ -\lambda(-x)^{\beta} & \text{if } x < 0. \end{cases}$$

2. In the CARA case, the exponential utility function is applied:

$$u(x) = \begin{cases} \frac{1}{\alpha}\left(1 - e^{-\alpha x}\right) & \text{if } x \geq 0, \\ -\frac{\lambda}{\beta}\left(1 - e^{\beta x}\right) & \text{if } x < 0. \end{cases}$$

Four types of probability weighting functions are applied to compute the decision weights for $p = 0.5$ gains and losses.[26] The functions $w^{+}(p)$ and $w^{-}(p)$ may share common parameters or have separate ones:

1. KT (Tversky and Kahneman, 1992):

$$w(p) = \frac{p^{\alpha}}{\left(p^{\alpha} + (1-p)^{\alpha}\right)^{\frac{1}{\alpha}}}$$

2. LLO (Log Odds Linear; Gonzalez and Wu, 1999):

$$w(p) = \frac{\beta p^{\alpha}}{\beta p^{\alpha} + (1-p)^{\alpha}}$$

3. Prelec I (Single-Parameter; Prelec, 1998):

$$w(p) = e^{-(-\ln(p))^{\alpha}}$$

4. Prelec II (Two-Parameter; Prelec, 1998):

$$w(p) = e^{-\frac{1}{\beta}(-\ln(p))^{\alpha}}$$

[26] Note that not all papers provide estimates of probability weighting functions. For instance, Abdellaoui, Bleichrodt and l'Haridon (2008) applied a semi-parametric method to directly elicit the decision weights at $p = 0.5$. Since these are all we need here, these studies are included in our analysis (even though they do not figure in our main source, the meta-analysis of Imai et al., 2025).

Paper	λ	utility	$u(10)$	$u(-10)$	pwf	$w^+(0.5)$	$w^-(0.5)$
Abdellaoui et al. (2008)	2.61	CRRA	7.24	11.48		0.46	0.45
Abdellaoui et al. (2011)	2.47	CRRA	6.17	9.12	LLO	0.41	0.44
Abdellaoui and Kemel (2014)	3.52	CRRA	6.03	10.96	LLO	0.43	0.47
Abdellaoui and Kemel (2014)	3.52	CRRA	6.03	10.96	Prelec II	0.41	0.46
Abdellaoui and Kemel (2014)	1.81	CRRA	9.12	11.48	LLO	0.50	0.49
Abdellaoui and Kemel (2014)	1.81	CRRA	9.12	11.48	Prelec II	0.49	0.49
Attema et al. (2018)	1.45	CRRA	8.13	5.75	Prelec II	0.42	0.51
Attema et al. (2019)	1.28	CRRA	9.12	9.12	Prelec I	0.46	0.48
Attema et al. (2019)	1.08	CRRA	7.94	6.46	Prelec I	0.45	0.47
Booij et al. (2009)	1.58	CRRA	7.23	6.70	LLO	0.44	0.51
Coricelli et al. (2018)	1.74	CRRA	2.24	2.75	Prelec II	0.46	0.38
Coricelli et al. (2018)	1.53	CRRA	1.61	2.16	Prelec II	0.51	0.33
Erner et al. (2013)	2.51	CRRA	14.02	8.59	LLO	0.43	0.52
Kemel and Paraschiv (2018)	1.96	CARA	9.52	9.90	Prelec II	0.45	0.42
Kemel and Paraschiv (2018)	3.15	CARA	9.47	10.05	Prelec II	0.43	0.35
L'Haridon and Vieider (2019)	1.17	CARA	8.21	11.02	Prelec II	0.52	0.50
L'Haridon and Vieider (2019)	1.22	CARA	7.94	11.32	Prelec II	0.57	0.47
L'Haridon and Vieider (2019)	1.28	CARA	7.94	10.92	Prelec II	0.60	0.48
L'Haridon and Vieider (2019)	2.17	CARA	8.81	9.44	Prelec II	0.53	0.39
L'Haridon and Vieider (2019)	1.27	CARA	8.13	10.25	Prelec II	0.53	0.51
L'Haridon and Vieider (2019)	1.47	CARA	8.82	10.02	Prelec II	0.52	0.51
L'Haridon and Vieider (2019)	1.72	CARA	8.28	10.50	Prelec II	0.55	0.46
L'Haridon and Vieider (2019)	1.56	CARA	8.12	8.15	Prelec II	0.55	0.55
L'Haridon and Vieider (2019)	1.18	CARA	8.12	10.73	Prelec II	0.56	0.50
L'Haridon and Vieider (2019)	2.22	CARA	9.14	9.67	Prelec II	0.58	0.46
L'Haridon and Vieider (2019)	1.27	CARA	7.81	11.56	Prelec II	0.57	0.45
L'Haridon and Vieider (2019)	1.05	CARA	7.47	10.61	Prelec II	0.56	0.48
L'Haridon and Vieider (2019)	1.57	CARA	9.22	9.31	Prelec II	0.52	0.48
L'Haridon and Vieider (2019)	1.88	CARA	7.46	6.83	Prelec II	0.55	0.56
L'Haridon and Vieider (2019)	1.34	CARA	7.95	10.35	Prelec II	0.57	0.50
L'Haridon and Vieider (2019)	1.59	CARA	8.29	8.16	Prelec II	0.58	0.53
L'Haridon and Vieider (2019)	1.93	CARA	9.04	10.11	Prelec II	0.54	0.50
L'Haridon and Vieider (2019)	2.36	CARA	8.52	8.67	Prelec II	0.63	0.50
L'Haridon and Vieider (2019)	2.79	CARA	10.10	10.86	Prelec II	0.51	0.40
L'Haridon and Vieider (2019)	1.83	CARA	8.75	9.52	Prelec II	0.59	0.49
L'Haridon and Vieider (2019)	1.41	CARA	8.35	10.41	Prelec II	0.53	0.47
L'Haridon and Vieider (2019)	1.62	CARA	8.26	9.85	Prelec II	0.55	0.46
L'Haridon and Vieider (2019)	1.80	CARA	8.63	9.64	Prelec II	0.62	0.51
L'Haridon and Vieider (2019)	1.49	CARA	8.18	8.62	Prelec II	0.54	0.53
L'Haridon and Vieider (2019)	1.27	CARA	8.13	10.92	Prelec II	0.55	0.48
L'Haridon and Vieider (2019)	1.54	CARA	8.07	10.98	Prelec II	0.57	0.45
L'Haridon and Vieider (2019)	1.90	CARA	9.09	8.34	Prelec II	0.51	0.53
L'Haridon and Vieider (2019)	2.24	CARA	8.12	10.74	Prelec II	0.63	0.41
L'Haridon and Vieider (2019)	1.26	CARA	7.89	10.00	Prelec II	0.54	0.51

Continued from previous page

Paper	λ	utility	$u(10)$	$u(-10)$	pwf	$w^+(0.5)$	$w^-(0.5)$
L'Haridon and Vieider (2019)	1.45	CARA	8.06	9.33	Prelec II	0.60	0.51
Lobel et al. (2017)	1.28	CARA	5.38	6.06	Prelec I	0.49	0.49
Lobel et al. (2017)	1.08	CARA	5.18	7.52	Prelec I	0.49	0.49
Lobel et al. (2017)	1.10	CARA	3.67	6.06	TK	0.49	0.49
Lobel et al. (2017)	1.04	CARA	5.38	7.19	TK	0.50	0.50
Lobel et al. (2017)	0.76	CRRA	1.91	2.34	Prelec I	0.41	0.41
Lobel et al. (2017)	0.57	CRRA	1.55	1.91	Prelec I	0.39	0.39
Lobel et al. (2017)	0.64	CRRA	1.70	2.14	TK	0.41	0.41
Lobel et al. (2017)	0.42	CRRA	1.32	1.74	TK	0.35	0.35
Nilsson et al. (2010)	0.75	CRRA	7.41	9.55	TK	0.46	0.49
Nilsson et al. (2010)	0.81	CRRA	8.13	10.00	TK	0.48	0.48
Rieger et al. (2016)	1.25	CRRA	3.98	3.98	TK	0.22	0.49
Rieger et al. (2016)	1.11	CRRA	3.02	4.79	TK	0.31	0.46
Rieger et al. (2016)	1.08	CRRA	2.57	2.82	TK	0.43	0.50
Rieger et al. (2016)	1.28	CRRA	2.34	2.34	TK	0.38	0.48
Rieger et al. (2016)	1.08	CRRA	3.63	3.89	TK	0.28	0.50
Rieger et al. (2016)	1.42	CRRA	2.75	3.55	TK	0.43	0.50
Rieger et al. (2016)	0.99	CRRA	3.47	2.34	TK	0.26	0.50
Rieger et al. (2016)	1.62	CRRA	2.63	6.76	TK	0.31	0.42
Rieger et al. (2016)	1.72	CRRA	3.47	7.94	TK	0.37	0.47
Rieger et al. (2016)	1.43	CRRA	3.47	3.55	TK	0.37	0.50
Rieger et al. (2016)	1.26	CRRA	2.63	2.34	TK	0.27	0.48
Rieger et al. (2016)	1.36	CRRA	3.39	2.14	TK	0.21	0.50
Rieger et al. (2016)	1.49	CRRA	3.63	2.82	TK	0.32	0.50
Rieger et al. (2016)	1.71	CRRA	3.24	7.94	TK	0.40	0.47
Rieger et al. (2016)	1.52	CRRA	2.24	2.34	TK	0.27	0.48
Rieger et al. (2016)	1.49	CRRA	3.16	4.07	TK	0.37	0.49
Rieger et al. (2016)	1.38	CRRA	2.57	3.09	TK	0.34	0.50
Rieger et al. (2016)	3.80	CRRA	6.31	7.94	TK	0.16	0.47
Rieger et al. (2016)	1.38	CRRA	2.63	3.09	TK	0.31	0.46
Rieger et al. (2016)	1.29	CRRA	3.16	2.00	TK	0.31	0.49
Rieger et al. (2016)	1.27	CRRA	2.24	2.19	TK	0.26	0.48
Rieger et al. (2016)	1.37	CRRA	2.45	3.09	TK	0.31	0.46
Rieger et al. (2016)	1.38	CRRA	2.57	3.09	TK	0.37	0.46
Rieger et al. (2016)	1.32	CRRA	2.63	3.09	TK	0.27	0.46
Rieger et al. (2016)	1.31	CRRA	2.63	2.34	TK	0.31	0.48
Rieger et al. (2016)	1.43	CRRA	2.63	3.55	TK	0.31	0.50
Rieger et al. (2016)	1.37	CRRA	1.82	3.55	TK	0.46	0.50
Rieger et al. (2016)	1.10	CRRA	2.00	1.62	TK	0.25	0.48
Rieger et al. (2016)	1.29	CRRA	3.02	2.04	TK	0.23	0.50
Rieger et al. (2016)	0.94	CRRA	2.88	2.88	TK	0.39	0.47
Rieger et al. (2016)	1.06	CRRA	2.09	2.04	TK	0.43	0.46
Rieger et al. (2016)	1.14	CRRA	2.04	2.34	TK	0.26	0.45
Rieger et al. (2016)	1.95	CRRA	5.13	3.98	TK	0.37	0.49

Continued from previous page

Paper	λ	utility	$u(10)$	$u(-10)$	pwf	$w^+(0.5)$	$w^-(0.5)$
Rieger et al. (2016)	1.47	CRRA	2.95	7.94	TK	0.49	0.47
Rieger et al. (2016)	0.99	CRRA	2.75	2.14	TK	0.33	0.47
Rieger et al. (2016)	1.05	CRRA	5.13	2.34	TK	0.16	0.48
Rieger et al. (2016)	1.27	CRRA	2.45	2.82	TK	0.39	0.50
Rieger et al. (2016)	1.59	CRRA	2.95	3.55	TK	0.31	0.50
Rieger et al. (2016)	2.31	CRRA	3.02	52.48	TK	0.35	0.21
Rieger et al. (2016)	3.19	CRRA	3.02	52.48	TK	0.35	0.21
Rieger et al. (2016)	1.41	CRRA	2.45	2.00	TK	0.28	0.49
Rieger et al. (2016)	1.31	CRRA	2.09	2.63	TK	0.38	0.50
Rieger et al. (2016)	1.28	CRRA	2.75	4.79	TK	0.46	0.46
Rieger et al. (2016)	1.63	CRRA	2.75	5.50	TK	0.33	0.49
Rieger et al. (2016)	1.60	CRRA	3.02	6.61	TK	0.40	0.48
Rieger et al. (2016)	1.37	CRRA	2.34	3.09	TK	0.38	0.50
Rieger et al. (2016)	1.33	CRRA	1.82	3.09	TK	0.46	0.46
Rieger et al. (2016)	1.01	CRRA	7.41	7.94	TK	0.49	0.47
Rieger et al. (2016)	1.61	CRRA	2.75	3.09	TK	0.43	0.46
Rieger et al. (2016)	1.51	CRRA	3.55	11.48	TK	0.39	0.50
Rieger et al. (2016)	1.06	CRRA	2.75	3.09	TK	0.33	0.50
Rieger et al. (2016)	1.36	CRRA	2.63	3.09	TK	0.31	0.46
Rieger et al. (2016)	1.29	CRRA	3.63	3.55	TK	0.28	0.50
Sagemüller and Mußhoff (2020)	0.95	CRRA	2.24	4.35	Prelec I	0.55	0.55
Stolk-Vos et al. (2022)	1.90	CRRA	10.28	21.93	Prelec I	0.40	0.44
Stolk-Vos et al. (2022)	1.36	CRRA	8.97	20.94	Prelec I	0.42	0.40
Vieider et al. (2014)	1.51	CARA	9.66	10.73	Prelec II	0.48	0.48
Vieider et al. (2014)	2.39	CARA	9.11	10.10	Prelec II	0.60	0.50
Vieider et al. (2016)	0.78	CRRA	9.53	33.50	Prelec II	0.46	0.37
Vieider et al. (2016)	0.58	CRRA	6.58	23.82	Prelec II	0.50	0.39
Vieider et al. (2016)	0.90	CRRA	14.89	54.45	Prelec II	0.42	0.31
Zeisberger et al. (2012)	1.37	CRRA	10.00	8.13	TK	0.49	0.49
Zeisberger et al. (2012)	1.37	CRRA	9.77	8.51	TK	0.49	0.47
Tversky and Kahneman (1992)	2.25	CRRA	7.59	7.59	TK	0.42	0.45
Enke and Shubatt (2023)	1.00	CRRA	5.75	5.50	LLO	0.47	0.47
Enke and Shubatt (2023)	1.04	CRRA	6.31	5.75	LLO	0.47	0.47
Erev et al. (2010)	1.50	CRRA	7.76	9.55	TK	0.46	0.46

Table 9: Functions predicting the coexistence of gain-seeking behavior for small stakes and stake dependence within a Prospect Theory (PT) framework. The parameter λ represents loss aversion (or gain seeking). $u(10)$ and $u(-10)$ denote the (dis)utilities for £10 gains and losses, respectively. $w^+(0.5)$ and $w^-(0.5)$ represent the decision weights at $p = 0.5$ for gains and losses, respectively.

Table 10 presents the parameters of Prospect Theory (PT) to predict the stake-dependent choice proportions displayed in Figure 6, which examines the fit of the meta-analytic evidence to the choice patterns we observe in fair spreads.

	ρ^+	ρ^-	π^+	π^-	λ
Tversky & Kahneman (1992)	0.88	0.88	0.42	0.45	2.25
Lobel et al. (2017)	0.12	0.24	0.35	0.35	0.42
First quartile meta-analysis	0.70	0.89	0.41	0.40	0.88
Median meta-analysis	0.90	1.06	0.424	0.443	1.444
Third quartile meta-analysis	0.97	1.34	0.47	0.47	1.81

Table 10: Functions predicting stake-dependent choice proportions based on a Prospect Theory (PT) model. The parameters ρ^+ and ρ^- represent the power utility coefficients for gains and losses, respectively. π^+ and π^- denote the decision weights at $p = 0.5$ for gains and losses, respectively. The parameter λ captures loss aversion (or gain seeking).

G Structural Estimation of Noisy Cognition Model

We estimate the Noisy Cognition Model structurally using a Bayesian hierarchical model coded in Stan. Vieider (2024a) provides a tutorial on how to code such model, and how to launch them from R using CmdStanR. The code below estimates the model using parallelization of the simulations in a single chain to speed up convergence:

```
functions {
  real partial_sum(array[] int choice_risky_slice, int start, int end, array[] real lg, array[] real ll,
                      vector alpha, vector gamma, vector lambda, vector sigma, array[] int id) {
    int size = end - start + 1;              // Size of the slice
    vector[size] pv;

    // Loop through the slice and compute the logit probabilities
    for (i in 1:size) {
      int idx = start + i - 1;               // Adjusted index within the original data
      pv[i] = (alpha[id[idx]] + gamma[id[idx]] * lg[idx] - lambda[id[idx]] * ll[idx])/sigma[id[idx]];
    }

    return bernoulli_lpmf(choice_risky_slice | Phi(pv) );  // Log-probabilities
  }
}
data {
  int<lower=0> N;                // Total number of observations
  int<lower=0> Nid;              // Number of individuals
  array[N] int id;               // Individual IDs
  array[N] int choice_risky;     // Binary outcome
  array[N] real lg;                  // Log gain
  array[N] real ll;                  // Log loss
  int<lower=0> GS;
}
parameters {
  vector<lower=0>[3] tau;                     // Population-level standard deviations
  cholesky_factor_corr[3] L;                  // Cholesky factor of correlation matrix
  array[Nid] vector[3] Z;                     // Uncorrelated individual-level parameters
  vector[3, K] means;                     // Regression coefficients for orthogonal predictors
}
transformed parameters {
  array[Nid] vector[3] theta;                 // Individual-level parameters
  vector[Nid] alpha;                          // Individual \(\alpha\)
  vector<lower=0,upper=1>[Nid] gamma;                          // Individual \(\gamma\)
  vector<lower=0,upper=1>[Nid] lambda;                         //
```

```
    Individual \(\lambda\)
  vector<lower=0>[Nid] sigma;
  vector<lower=0>[Nid] nug2;
  vector<lower=0>[Nid] nul2;

  // Compute individual-level parameters
  for (n in 1:Nid) {
    theta[n] = means + diag_pre_multiply(tau, L) * Z[n];
    alpha[n] = theta[n, 1];
    gamma[n] = inv_logit( theta[n, 2] );
    lambda[n] = inv_logit( theta[n, 3] );

    nug2[n] = 1/gamma[n] - 1;
    nul2[n] = 1/lambda[n] - 1;
    sigma[n] = sqrt( gamma[n]^2 * nug2[n] + lambda[n]^2 * nul2[n] );
  }
}
model {
  tau ~ exponential(1);                     // Weakly informative prior for standard deviations
  L ~ lkj_corr_cholesky(4);                 // LKJ prior for correlation matrix (eta=4)

  // Priors for individual-level parameters
  for (n in 1:Nid)
    Z[n] ~ std_normal();

  to_vector(beta) ~ normal( 0, 10 );

  // Parallelized likelihood computation using reduce_sum
  target += reduce_sum(partial_sum, choice_risky, GS, lg, ll, alpha, gamma, lambda, sigma, id);
}
generated quantities{
vector[Nid] ls;
vector[Nid] lsr;
vector[Nid] mu;
for (n in 1:Nid) {
  ls[n] = lambda[n] - gamma[n];
  lsr[n] = lambda[n]/gamma[n];
  mu[n] = alpha[n]/ls[n];
  }

}
```

We use the estimated parameters to regress them on observable characteristics using a measurement error model. Let a generic parameter estimated for individual i be denoted by θ_i. The model assumes that θ_i is a noise measure of the true underlying parameter $\widehat{\theta}_i$, which is unobserved. The measurement error in θ_i is thereby given by sd_i, the standard deviation of the Bayesian posterior distribution of the parameters. Examining parameter distributions, we found them to be well-fit by normal distributions, so that we model the measurement error as $\theta_i \sim$

$\mathcal{N}(\widehat{\theta}_i, sd_i^2)$. $\widehat{\theta}_i$ is subsequently regressed on observable characteristics of the decision makers using a student-t distribution, which is outlier-robust (Gelman, Carlin, Stern, Dunson, Vehtari and Rubin, 2014). The following code implements the mode in Stan:

```
data {
    int<lower=1> N;
    vector[N] theta;
    vector[N] se;

   // Design matrix used for regression
    int<lower=0> K; // dimension of the design matrix
    matrix[N,K] x; // design matrix to impute SEs
}
parameters {
    vector[N] theta_hat;
    real<lower=0> sigma;
    vector[K] beta;
}
model {
    // Priors
    sigma ~ normal(0, 10);
    beta ~ normal(0, 10);

    // Measurement error model
    theta ~ normal(theta_hat, se);

    // Regression model:
    theta_hat ~ student_t( 3 , x * beta  , sigma );
}
```

The effect of cognitive ability is robust to dropping the covariates reported in the main text, or indeed to introducing them one by one.

H Appendix: Experiment Screenshots

Instructions

Please read these instructions carefully. On the next screen, we will ask you to answer a few questions about these instructions as a comprehension check. If you correctly answer all of them, you will earn an extra £0.25 bonus.

We will ask you to make a series of choices between different options entailing monetary payoffs. Your choices are anonymous, and cannot be traced back to you personally. This study conforms to ethics guidelines, and all information you are given is true.

The choices we will show you involve both gains and losses. For example, suppose your earnings were determined by a question that asked you to decide between

1. *Win £24 or lose £20 with equal probability;*
2. *£0 for certain.*

If you chose the first option, and a losing balls is drawn, you would lose 20 pounds; if a winning ball is drawn, you would earn 24 pounds. **If you chose the second option**, then you would receive 0 pounds (nothing).

All choices we show you involve lotteries involving a gain or a loss, versus a sure amount of 0 pounds (nothing). All choices are hypothetical. We would nevertheless like to ask you to **carefully consider the amounts involved**, and to choose as though those amounts could be won or lost for real. There are no incorrect choices, we are purely interested in your preferences.

Which of the following options do you prefer?

Win **£26** or **lose £26 with equal probability**;

£0 for certain.

Figure 7: Screenshots of the experimental instructions and an example of the choice trial.

What level of education do you complete?

- No formal education
- GCSE
- A level/BTEC
- Undergraduate Degree
- Postgraduate Degree

Which of the following best describes your current occupation?

- Employed
- Not employed, looking for work
- Not employed, NOT looking for work
- Studying
- Retired
- Disabled, not able to work

Which of the following best describes your personal income last year?

- Below £10,000
- £10,001 to £20,000
- £20,001 to £30,000
- £30,001 to £40,000
- £40,001 to £50,000
- Above £50,001

Figure 8: Screenshots of questionnaire questions on education, employment and income.

1. Imagine that we roll a fair, six-sided die 1,000 times. Out of 1,000 rolls, how many times do you think the die would come up as an even number (2, 4, 6)?

2. In the HOTPICKS LOTTERY, the chances of winning a £10 prize are 1%. What is your best guess about how many people would win a £10 prize if 1,000 people each buy a single ticket from HOTPICKS?

3. In the ACME PUBLISHING SWEEPSTAKES, the chance of winning a car is 1 in 1,000. What proportion of tickets of ACME PUBLISHING SWEEPSTAKES win a car? Please indicate the probability in percent (%):

4. Out of 1,000 people in a small town 500 are members of a choir. Out of these 500 members in the choir 100 are men. Out of the 500 inhabitants that are not in the choir 300 are men. What is the probability that a randomly drawn man from the town is a member of the choir? Please indicate the probability in percent (%):

5. Imagine we are throwing a loaded die (6 sides). The probability that the die shows a 6 is twice as high as the probability of each of the other numbers. On average, out of 70 throws, how many times would the die show the number 6?

Figure 9: Screenshots of questionnaire questions on numeracy test.

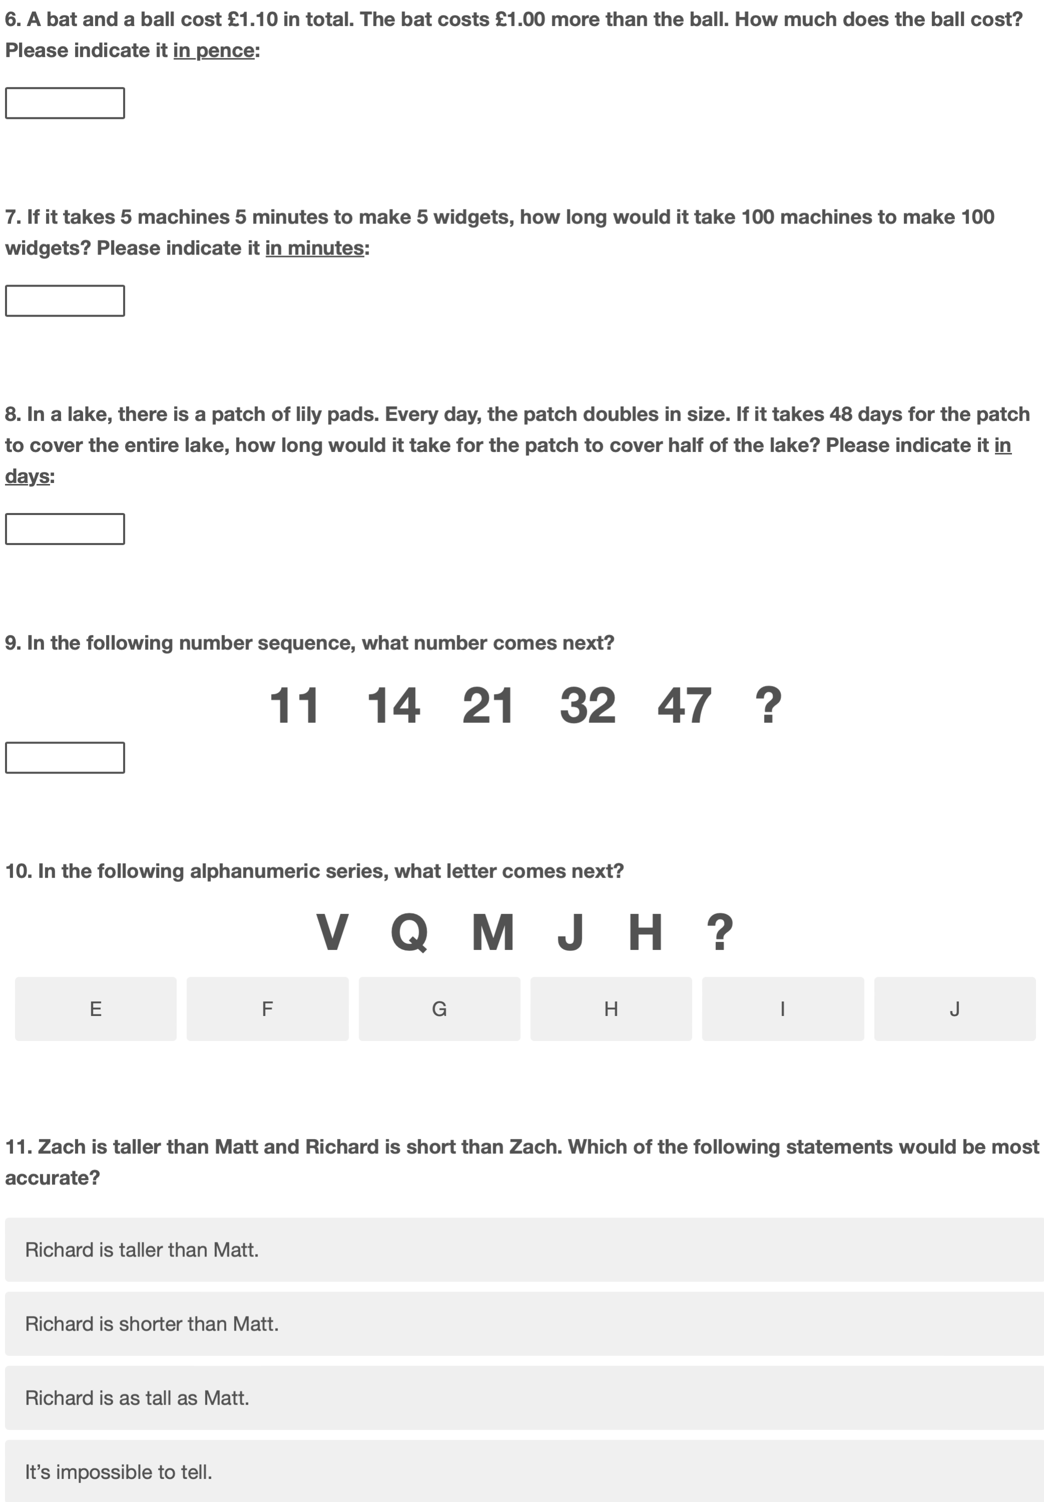

Figure 10: Screenshots of questionnaire questions on CRT and ICAR test.