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*

March 31, 2026

Abstract

Loss aversion—the tendency to weigh losses more heavily than comparable gains—is a key concept in behavioural economics. Yet empirical estimates of loss aversion vary widely across studies and often diverge from the parameter values required to explain important economic phenomena. To reconcile these seemingly contradictory findings, we re-examine behaviour over mixed gain–loss wagers through the lens of *generative* models of decision-making, which derive behavioural predictions from constrained optimization conditional on noisy internal representations of outcomes. We design diagnostic experimental manipulations that distinguish between two prominent generative models—Decision by Sampling and the Noisy Cognition Model—and prospect theory. The results support the Noisy Cognition Model over Decision-by-Sampling and prospect theory. Our findings reconcile seemingly contradictory patterns in the literature within a coherent account based on a single set of underlying parameters.

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 value equally sized gains—is one of the central concepts in behavioural economics. Prospect theory (PT) represents loss aversion as a kink in the utility function at the reference point (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992; Köbberling and Wakker, 2005). The concept has been used to explain a

*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.

range of important economic phenomena, including systematic differences between willingness to pay and willingness to accept (the *WTP–WTA gap*; Kahneman et al., 1990; 1991), the status quo bias (Samuelson and Zeckhauser, 1988; Kahneman et al., 1991), small-stakes risk aversion (Rabin, 2000; Rabin and Thaler, 2001), and financial market puzzles such as the equity premium puzzle (Mehra and Prescott, 1985; Benartzi and Thaler, 1995; Gneezy and Potters, 1997).

Hundreds of studies have estimated the PT parameter governing this kink, and Brown et al. (2024) place its meta-analytic average just below 2. At the same time, measured parameters vary widely, and this heterogeneity remains largely unexplained by observable moderators. Ert and Erev (2013) have argued that attitudes toward losses are stake-dependent in ways incompatible with PT—though Bleichrodt and L'Haridon (2023) question this interpretation. Using large-scale representative samples, Chapman et al. (2024) document widespread *gain-seeking* in mixed gambles—the opposite of loss aversion. Chapman et al. (2023) fail to find robust correlations between measured loss aversion and the WTP–WTA gap. Taken together, these findings raise questions about whether a single stable parameter can adequately summarize attitudes toward gains and losses.

One way to organize these findings is to move beyond reduced-form representations of preferences and consider *generative* models of decision-making. Rather than treating preferences as primitives, such models derive behavioural predictions from constrained optimization conditional on noisy internal representations of choice primitives. In this perspective, the reduced-form parameters of PT emerge as outcomes of optimal adaptation processes rather than stable tastes. This approach has two advantages: it can reconcile some of the seemingly contradictory findings discussed above within a coherent framework, while also predicting systematic changes in behaviour across environments that would be difficult to reconcile with interpretations based on stable tastes.

Risk-taking in mixed gain-loss gambles is stake-dependent. Figure 1 shows choice proportions for fair spreads offering $\pm x$ with 50-50 probabilities over the status quo, for different stake levels x. We collected the data from a sample representative of the UK population, with $N = 1{,}000$ subjects making hypothetical choices and $N = 506$ making incentivized choices (see caption for details).

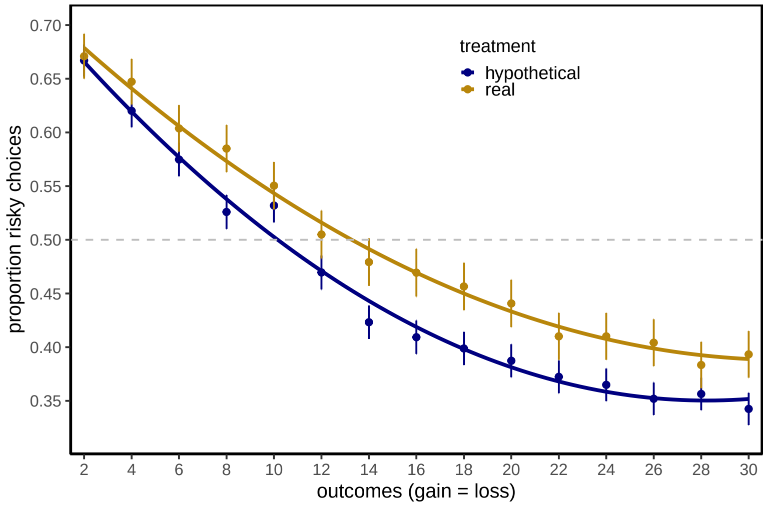

Figure 1: Risk taking by stake level for fair spreads around 0

The figure shows choice proportions for a wager offering x with probability 0.5 and $-x$ with probability 0.5 over the status quo of 0, for values of $x \in \{2, 4, \ldots, 30\}$ (in GBP), with 95% confidence intervals. The line is fitted by a second-degree polynomial regression. Gains G and losses L each ranged from £2 to £30 in steps of £2, with every gain crossed with every loss, yielding 128 choices in total (including some repeated choice tasks). The experiment was conducted using Prolific UK's representative sample service. Subjects in the hypothetical condition were paid for their time according to Prolific regulations at a rate of £10 per hour. Subjects in the real condition received the same fixed payment, with 1 in 10 subjects additionally paid for one randomly selected choice; gains and losses were integrated with an initial endowment equal to the maximum possible loss (£30). The experiment lasted approximately 20 minutes, including a questionnaire and tests of cognitive ability. See Online Appendix A for details.

The figure reveals two key results. First, choice proportions for the fair spread over zero are strongly stake-dependent, declining monotonically as the stake size x increases from £2 to £30. Second, for small stakes $x \leq \pounds 10$, we observe systematic risk-seeking. We will refer to this stake-dependence in choice proportions as *loss-sensitivity*—a precise mathematical definition is provided in Section 2. Our results replicate the finding of risk-seeking over mixed wagers with moderate stakes by Chapman et al. (2024) for a different representative population (UK versus US), while qualifying it: risk-seeking prevails only for small stakes, giving way to increasing risk aversion as stakes grow.[1]

[1]The stake-dependence we document cannot be explained by a stable distribution of loss-averse preferences combined with additive stochastic choice noise. Under such a model, choice proportions for fair spreads would remain below 0.5 for all stake levels, since the noise merely smooths a distribution centered on loss aversion. The observed risk-seeking at small stakes therefore requires a genuine change in the underlying choice tendency with stake size. We return

Providing real incentives from an endowment leads to a consistent rightward shift in choice proportions, indicating higher risk-taking throughout (statistical analysis in A.2). This is consistent with Li and Vieider (2025), whose meta-analysis finds that mixed gambles are the primary exception to the otherwise negligible effect of incentives on individual decision-making—with the increase in risk-taking under real incentives likely driven by house-money effects. Importantly, loss-sensitivity is robust across incentive conditions.

Theoretical implications of loss-sensitivity. Any model accounting for the patterns described above must incorporate heightened sensitivity to losses relative to gains. PT can accommodate loss-sensitivity through steeper utility for losses than gains—a possibility explicitly discussed by Kahneman and Tversky (1979). The specific pattern we document, however—risk-seeking at small stakes giving way to risk aversion at larger ones—runs counter to the received wisdom that people uniformly dislike mixed gambles regardless of stake size.[2] We are interested in a deeper question: where do these patterns originate? PT is silent on this, treating its parameters as exogenous taste primitives. The generative models we examine instead endogenize these parameters, thereby generating predictions about how they should vary across decision environments.

To this end, we turn to two adaptive models of decision-making that provide endogenous accounts of the origins of risk-taking over mixed gain–loss wagers. The first is Decision by Sampling (*DbS*; Stewart et al., 2006), the second the Noisy Cognition Model (*NCM*; Khaw et al., 2021). Both provide micro-foundations for PT-like choice patterns, but differ fundamentally in where cognitive frictions enter the decision process and how they are resolved. These differences lead the two models to make diametrically opposite predictions in specific choice situations—an insight we exploit to test them against each other.

to this point in the Conclusion.

[2] Kahneman and Tversky (1979) write on p. 279: "the value function for losses is steeper than the value function for gains"—a characterization consistent with loss-sensitivity as we define it. However, the same passage also states that "most people find symmetric bets of the form $(x, .50; -x, .50)$ distinctly unattractive", suggesting uniform risk aversion regardless of stake size. It is precisely this second claim that our data challenge: we find systematic risk-seeking for small stakes, with risk aversion emerging only as stakes increase.

In DbS, noise arises at the stage of utility assignment: outcomes are evaluated not on an absolute scale but through ordinal comparison with a small sample of outcomes drawn from memory, reflecting the gains and losses typically encountered in the environment. The utility of an outcome is determined by its rank among these sampled outcomes, so that the distribution of outcomes in memory directly shapes subsequent choices. Loss-sensitivity arises when the distribution of losses is more concentrated than the distribution of gains, so that ranks increase more rapidly for losses than for gains as stakes grow.

In the NCM, by contrast, noise enters at the perceptual stage: the decision maker forms noisy internal representations of the gain and loss and optimally infers the true magnitudes by combining the internal representations with a Bayesian prior summarizing expectations about the distribution of outcomes in the choice environment. Loss-sensitivity arises when losses are perceived more accurately than gains. Prior expectations furthermore determine whether a given outcome is over- or underestimated, naturally allowing for the coexistence of risk-seeking at small stakes and risk aversion at larger ones—a feature we exploit in the experimental tests below.

Empirical tests of choice adaptation. DbS and the NCM both postulate adaptive mechanisms that exploit information about the environment to overcome cognitive frictions. Despite this commonality, the two models make sharply divergent predictions about how adaptation affects observed choices in specific situations.

To test these predictions, in Experiment II subjects are exposed to choices involving small versus large gains in an initial adaptation phase, followed by a common test set of mixed gain–loss wagers. In DbS, exposure to larger gains should *decrease* the utility attributed to a given gain x, since it will have a lower rank among the sampled comparisons, leading to increased risk aversion over mixed gambles. The NCM makes the opposite prediction: larger gains in the environment shift prior expectations upward, increasing risk-taking over mixed gain–loss gambles.

We find that subjects exposed to large gains in the adaptation phase display systematically higher levels of risk-taking in subsequent mixed choices than subjects exposed to small gains—consistent with the NCM prediction and contrary to DbS. This result holds regardless of whether choices are incentivized or not.

Visual displays as a test of perceptual noise. The adaptation experiment above cannot conclusively separate the NCM from PT: even though PT makes no clear ex ante predictions about adaptation effects, reference-point adjustments could in principle generate the observed patterns under some parameter combinations. We therefore introduce an additional experiment that provides a clean test of the NCM against both PT and DbS simultaneously.

The manipulation consists of providing a visual aid conveying the magnitudes of potential gains and losses alongside their numerical description. From the perspective of PT, this display is redundant: behaviour is driven by stable preferences applied to objectively perceived outcomes, and presentational format leaves the value function unchanged. DbS makes the same prediction—no effect of the visual aid—but for a different reason: outcomes are perceived objectively and the gain G and loss L are each ranked independently against memory samples, so the display format cannot affect utility assignment. Under the NCM, by contrast, the presentation mode matters directly: loss-sensitivity arises from asymmetric noise in the internal representation of gains and losses, and a visual display that makes their magnitudes more salient is expected to reduce this asymmetry, attenuating loss-sensitivity and flattening the stake-dependence of risk-taking.

We find precisely this: the visual display significantly reduces stake-dependence, consistent with the NCM and inconsistent with both PT and DbS. Even under a perceptual interpretation of PT parameters—arguably present in the original formulation of Kahneman and Tversky (1979)—PT does not *predict* this pattern. This highlights our interpretation of the NCM as endogenizing PT-style parameters: accepting the psychophysical account proposed by the NCM restores predictive content to the descriptive framework precisely in the contexts where it has been lacking.

A unified account of behaviour. The NCM provides a unified account of behaviour across different domains. As we discuss in more detail in Section 5, the same set of NCM parameters can account both for small-stake risk aversion over pure gains as documented by Khaw et al. (2021) and for the findings in this paper. The framework also sheds light on the cognitive ability puzzle highlighted by Chapman et al. (2024), who find risk aversion over pure gains to be negatively associated with cognitive ability, while loss aversion is positively associated with cognitive ability. Under the NCM this is not a contradiction: identical parameters can explain both findings. Extensions of the NCM further predict systematic changes not only in risk-taking, but also in patience, probability sensitivity, and time-delay sensitivity (Oprea and Vieider, 2026). The same mechanism also helps reconcile several findings that might otherwise appear contradictory, such as moderate-stake gain-seeking in representative samples (Chapman et al., 2024) alongside reluctance to invest in the stock market—the equity premium puzzle (Mehra and Prescott, 1985; Benartzi and Thaler, 1995).

Fit with the literature. The concept of loss-sensitivity we propose is based on heightened attention to losses relative to gains, resulting in lower cognitive noise for losses than for gains in the NCM framework. This mechanism is consistent with several studies investigating differential attention to gains and losses. At the neural and physiological level, Tom et al. (2007) show that neural responses react more strongly to changes in losses than gains and that rejection of mixed wagers is closely linked to these activation patterns, while Sokol-Hessner et al. (2009) document stronger arousal responses to losses than gains that predict loss aversion. At the behavioural level, Pachur et al. (2018) show that loss aversion correlates with relative attention to losses and that manipulating attention affects estimated loss aversion, and Hirmas et al. (2024) show using eye-tracking that subjects paying more attention to gains are more likely to accept mixed wagers.

Our findings also contribute to a growing literature that attributes observed behaviour to cognitive frictions rather than stable preferences. Several papers explain the systematic probability distortions documented in PT as consequences of cognitive noise (Zhang et al., 2020; Khaw et al., 2023; Frydman and Jin, 2023; Vieider, 2024; Netzer et al., 2025; Bouchouicha et al., 2025). Related work shows that confidence in lottery valuations predicts probability distortions (Enke and Graeber, 2023), that typical distortions can arise in experimental environments without risk (Oprea, 2024), and that noisy cognition helps explain the description–experience gap (Oprea and Vieider, 2024; Barron and Erev, 2003; Hertwig et al., 2004). We contribute to this literature by providing the first empirical tests of the noisy cognition account for choices over mixed gain–loss wagers and by discriminating between competing generative models of such choices.

2 Theoretical foundations of loss-sensitivity

In this section we examine the theoretical foundations of loss-sensitivity. We first define and characterize the concept formally, then discuss how it can be accommodated within models of utility maximization before turning to generative models that derive the pattern from principles of constrained optimization.

2.1 Loss-sensitivity: Definition and Characterization

Consider a decision maker (DM) choosing whether to accept or reject a binary wager that yields a gain $G > 0$ with probability $\frac{1}{2}$ and a loss $-L$ (where $L > 0$) otherwise, relative to a status quo of 0. We use this choice environment throughout to derive and discuss model predictions.[3]

Loss-sensitivity. For any loss $L > 0$, let $G^*(L)$ denote the minimum gain required for the DM to prefer the gamble to 0. We normalize $G^*(0) = 0$, and assume that $G^*(L)$ is continuously differentiable and strictly increasing in L.

Definition. *A decision maker exhibits* loss-sensitivity *if the function* $G^*(L)$ *has*

[3]The assumption of even odds describes the great majority of the empirical evidence, including our own experiments, and simplifies the exposition considerably. Both DbS and the NCM can generate nonlinear probability distortions, and the qualitative predictions regarding loss-sensitivity extend to more general specifications under either model. We abstract from probability distortions throughout to maintain focus on the gain–loss dimension.

elasticity greater than 1 *for all* $L > 0$*, that is,*

$$\varepsilon(L) \equiv \frac{d\ln G^*(L)}{d\ln L} = \frac{L}{G^*(L)} \cdot \frac{dG^*(L)}{dL} > 1 \quad \textit{for all } L > 0. \tag{1}$$

Loss-sensitivity requires that a given percentage increase in L be compensated by a larger percentage increase in G to preserve indifference. An equivalent characterization is that the ratio $G^*(L)/L$ is strictly increasing in L: larger losses require more than proportionally larger compensation to restore indifference.

An additional feature of the empirical pattern in figure 1 is the existence of a threshold $L^* > 0$ defined by $G^*(L^*) = L^*$. For smaller losses ($L < L^*$), the DM exhibits risk-seeking, whereas for larger losses ($L > L^*$) she exhibits risk aversion. The threshold L^* is the fixed point of $G^*(L)$ on the 45-degree line: below it, the compensating gain required falls short of the loss, so fair spreads are accepted; above it, the required compensation exceeds the loss, so fair spreads are rejected.

Relation to prospect theory. PT can accommodate loss-sensitivity through different combinations of its parameters, including utility curvature, probability weighting, and the loss aversion coefficient. Empirical applications have largely emphasized the kink in the value function, though heightened sensitivity to losses was already present in the original formulation of Kahneman and Tversky (1979). The stake-dependence component of loss-sensitivity is broadly consistent with the bulk of PT parameter estimates in the literature (Imai et al., 2025): steeper utility for losses than gains naturally generates increasing risk aversion as stakes grow. Risk-seeking at small stakes is a different matter. Regardless of how these patterns are organized within PT, however, the model does not *predict* it: PT is silent on why its parameters should take the values required, and on why they should vary with the informational environment. This motivates our interest in generative models.

Expected utility accounts. Alaoui and Penta (2026) provide a framework capturing rich reversals in risk attitudes around an exogenous reference point x_0 within an otherwise standard expected utility framework. We provide the technical details and examine the implied utility curvature empirically in Online Appendix B. The continuous case of *Success attachment*—as well as the discontinuous cases of both Success attachment and Failure avoidance—are compatible with small stake risk-seeking and loss-sensitivity, as well as the utility curvature patterns we document. This provides a benchmark case against which to examine the empirical tests of the generative models in experiments II and III.

2.2 Decision by Sampling and loss attitudes

DbS is motivated by the important psychological insight that assessing outcomes on an absolute, invariant scale—as assumed in expected utility theory and PT—would be cognitively demanding. Outcomes are perceived objectively in this model; environmental characteristics enter instead at the stage of utility assignment (see also Robson, 2001a; Netzer, 2009, for related modelling approaches).

Model sketch. According to DbS, the utility of a given outcome depends on ordinal comparisons with outcomes sampled from memory. Let F_g and F_ℓ denote the environmental distributions of gains and losses stored in memory, respectively. When evaluating a gain G, the DM compares it to N samples $\widehat{G}_1, \ldots, \widehat{G}_N$ drawn i.i.d. from F_g; the internal representation of G is its empirical rank,

$$r_g = \frac{1}{N} \sum_{i=1}^{N} \mathbf{1}\{G > \widehat{G}_i\}.$$

Conditional on G, the indicators $\mathbf{1}\{G > \widehat{G}_i\}$ are Bernoulli random variables with success probability $F_g(G)$, so that by the central limit theorem[4],

$$r_g = F_g(G) + \varepsilon_g, \qquad \varepsilon_g \sim \mathcal{N}\left(0,\ \frac{F_g(G)\left[1 - F_g(G)\right]}{N}\right),$$

[4] The normal approximation is accurate for moderate values of N except when $F_g(G)$ is extremely close to 0 or 1. In such low-discriminability regions, stochasticity in choice may come to dominate the difference in expected ranks appearing in the numerator—a point we return to in the Conclusion.

and analogously for r_ℓ, with $\widehat{L}_1, \ldots, \widehat{L}_N$ drawn i.i.d. from F_ℓ.[5] The DM accepts the wager whenever r_g exceeds r_ℓ, yielding a stochastic choice rule:

$$\Pr[\text{accept}] = \Phi\left(\frac{F_g(G) - F_\ell(L)}{\sqrt{\sigma_g^2 + \sigma_\ell^2}}\right),$$

where Φ is the standard normal CDF, $\sigma_g^2 \triangleq \frac{F_g(G)[1-F_g(G)]}{N}$ and $\sigma_\ell^2 \triangleq \frac{F_\ell(L)[1-F_\ell(L)]}{N}$. The resulting stochastic choice rule is a Probit. The numerator reflects the expected ranks of G and L in their respective environmental distributions, while the denominator reflects the variability in those ranks arising from the finite number of samples drawn from memory.

Loss-sensitivity in DbS. Stewart et al. (2006) further argue that the environmental distributions of gains and losses follow power laws. In our notation, the probability that a memory draw $\widehat{G}_i$ exceeds a given gain G, and the probability that a memory draw $\widehat{L}_i$ exceeds a given loss L, are given by:

$$\Pr[\widehat{G}_i > G] = \left(\frac{G}{G_0}\right)^{-\rho_g}, \quad \Pr[\widehat{L}_i > L] = \left(\frac{L}{L_0}\right)^{-\rho_\ell}, \tag{2}$$

where $G_0, L_0 > 0$ are scale parameters governing the location of the distributions. These expressions define the upper tails of the distributions; the corresponding CDFs are given by $F_g(G) = 1 - (G/G_0)^{-\rho_g}$ and $F_\ell(L) = 1 - (L/L_0)^{-\rho_\ell}$. The ordering $\rho_\ell > \rho_g > 0$ is based on the bank account data reported by Stewart et al. (2006): the loss distribution is more concentrated than the gain distribution.

The function $G^*(L)$ is determined by the condition that the gain G and the loss L receive equal rank in their respective environmental distributions—the point at which the DM is indifferent between accepting and rejecting the wager. In the large-sample limit (where sampling noise becomes negligible), this indifference condition reduces to $F_g(G^*(L)) = F_\ell(L)$: the gain $G^*(L)$ occupies the same percentile of the gain distribution as L occupies in the loss distribution. Substituting

[5] Stewart et al. (2006) use a slightly different normalization, $r_x = \frac{R-1}{N-1}$, where R is the count of samples strictly below x. For large N the difference is negligible and absorbed by the CLT approximation.

the power-law distributions and solving for $G^*(L)$ yields

$$G^*(L) = G_0 \left(\frac{L}{L_0} \right)^{\rho_\ell / \rho_g}. \tag{3}$$

The elasticity of $G^*(L)$ is therefore constant and equal to $\varepsilon(L) = \rho_\ell / \rho_g$. Loss-sensitivity in the sense of Definition 1 thus arises whenever $\rho_\ell > \rho_g$. Given the empirical distributions of debits and credits documented by Stewart et al. (2006), this implies an elasticity greater than one.

Small-stake risk-seeking. The fixed point L^*—the stake level at which the DM is indifferent between accepting and rejecting a fair spread—is determined by the condition $G^*(L^*) = L^*$, i.e. the point at which the indifference curve crosses the 45-degree line. Substituting equation (3) and solving yields

$$L^* = \left(\frac{L_0^{\rho_\ell}}{G_0^{\rho_g}} \right)^{\frac{1}{\rho_\ell - \rho_g}}. \tag{4}$$

A positive fixed point exists whenever $L_0 > G_0$. For stakes above their respective scale parameters, the power-law exponents govern the rate at which rank increases with monetary size: since $\rho_\ell > \rho_g$, ranks increase more rapidly for losses than for gains, yielding loss-sensitivity. Below the scale parameters, this ordering can reverse. When $L_0 > G_0$, gains rank higher than losses of the same size over a range of small stakes, $F_g(x) > F_\ell(x)$, generating risk-seeking. The condition $L_0 > G_0$ therefore implies that the two distributions cross at L^*, with fair spreads accepted below this point and rejected above it.

2.3 Noisy cognition and loss attitudes

The key insight underlying the noisy cognition model (*NCM*) of Khaw et al. (2021) is that decision makers cannot directly access the choice quantities presented to them. Instead, outcomes are internally encoded through noisy neural representations. Systematic choice regularities then emerge from the optimal way the brain processes these signals, exploiting prior information about the distribution of outcomes in the environment.

Model sketch. The DM follows a choice rule under which the wager is accepted whenever the expected gain exceeds the expected loss. For 50-50 odds, this reduces to accepting whenever $G \geq L$.[6] Since the true values of G and L are not directly accessible, the DM infers them from noisy signals r_g and r_ℓ. The wager is accepted whenever $\mathbb{E}[\, G \mid r_g \,] \geq \mathbb{E}[\, L \mid r_\ell \,]$.[7]

Let the signals be draws from two normal distributions:

$$r_g \sim \mathcal{N}(\ln G, \nu_g^2), \quad r_\ell \sim \mathcal{N}(\ln L, \nu_\ell^2), \tag{5}$$

where ν_g and ν_ℓ quantify the coding noise for gains and losses respectively. The signals are unbiased *on average*, but any given draw may be affected by noise. The DM could in principle compare the signals directly, accepting the wager whenever $r_g \geq r_\ell$—yielding an equation formally resembling the DbS choice rule. Inference about the true quantities G and L can however be improved by leveraging expectations about the decision environment. Take a Bayesian prior:

$$\ln G,\ \ln L \sim \mathcal{N}(\mu, \sigma^2) \tag{6}$$

Combining the signals with the prior yields posterior expectations for G and L. Averaging over the distribution of signals for a given stimulus and substituting these expectations into the optimal choice rule (see Online Appendix C) yields:

$$\Pr[accept] = \Phi\left[\frac{\alpha \ln G - \beta \ln L + (\beta - \alpha)\widehat{\mu}}{\sqrt{\alpha^2 \nu_g^2 + \beta^2 \nu_\ell^2}}\right], \tag{7}$$

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 the perceptual accuracy of gains and losses respectively, and $\widehat{\mu} \triangleq \mu + \frac{1}{2}\sigma^2$ is the transformed prior mean.

[6] This choice rule is optimal in the sense that it maximizes lifetime wealth; allowing for an additional utility transformation incorporating decreasing marginal utility of wealth does not affect the conclusions of the model.

[7] This decision rule extends naturally to non-even odds when combined with a model of probability distortions defined over log-odds (Khaw et al., 2023; Vieider, 2024).

The NCM and loss-sensitivity. In the NCM, loss-sensitivity arises from heightened attention to losses relative to gains, captured by lower coding noise for losses than for gains ($\nu_\ell < \nu_g$; see Online Appendix D for a stylized model of attention). This immediately implies $\beta > \alpha$—as-if utility is steeper for losses than for gains. The ratio β/α maps directly onto the elasticity of $G^*(L)$ defined in Definition 1: it can be shown that $\varepsilon(L) = \beta/\alpha$ for all L, so that $\varepsilon(L) > 1$ is equivalent to $\beta > \alpha$.[8] For fair spreads ($G = L = x$), the numerator in equation (7) simplifies to $(\beta - \alpha)[\widehat{\mu} - \ln x]$, and the acceptance probability declines monotonically in x whenever $\beta > \alpha$, yielding stake-dependence consistent with Definition 1.

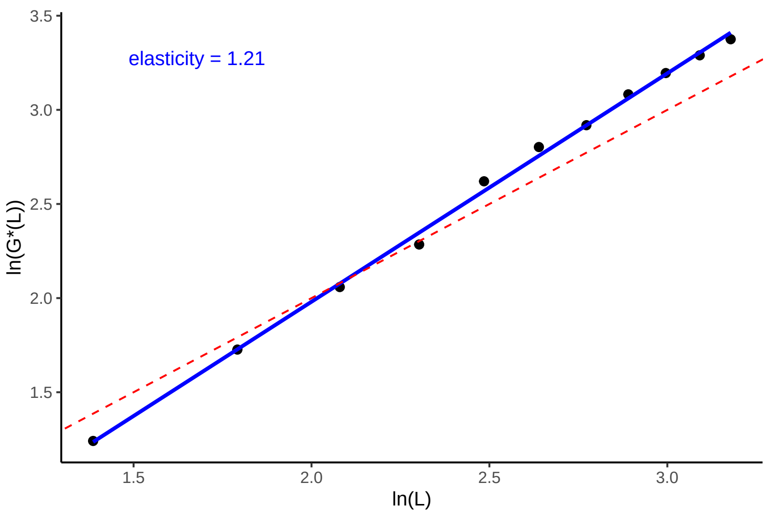

Figure 2: Elasticity in the representative data

The figure plots the log of $G^*(L)$—the interpolated gain value that makes a decision maker indifferent between taking the wager and rejecting it—against $\log(L)$.

Interestingly, both DbS under power-law environmental distributions and the NCM predict a *constant* elasticity of $G^*(L)$—equal to ρ_ℓ/ρ_g under DbS and β/α under the NCM. A regression of $G^*(L)$ on L in log space should therefore yield a constant slope if either model is correct. The $G^*(L)$ curve we estimate from our data, shown in Figure 2, is indeed approximately log-linear, consistent with this prediction.

The transformed prior mean $\widehat{\mu}$ determines the fixed point $L^* = \exp(\widehat{\mu})$ introduced in Definition 1: stochastic indifference between the wager and the status quo

[8] To see this, note that the indifference condition $\alpha \ln G^*(L) = \beta \ln L + (\alpha - \beta)\widehat{\mu}$ implies $\ln G^*(L) = (\beta/\alpha) \ln L + ((1-\beta)/\alpha)\, \widehat{\mu}$, so that $\frac{d \ln G^*(L)}{d \ln L} = \beta/\alpha$.

occurs at $x = \exp(\widehat{\mu})$, with risk-seeking for $x < \exp(\widehat{\mu})$ and risk aversion for $x > \exp(\widehat{\mu})$. Pessimistic prior expectations (low $\widehat{\mu}$) thus produce risk aversion throughout, while sufficiently optimistic expectations generate the coexistence of risk-seeking at small stakes and risk aversion at large stakes that we document empirically.

Micro-foundations for loss aversion. The common prior specification can be extended to allow separate priors for gains and losses. Let the prior for gains be $\mathcal{N}(\mu_g, \sigma_g^2)$ and for losses $\mathcal{N}(\mu_\ell, \sigma_\ell^2)$. The numerator in equation (7) then becomes $\alpha \ln G - \beta \ln L - [(1-\beta)\widehat{\mu}_\ell - (1-\alpha)\widehat{\mu}_g]$. Define $\lambda \triangleq \exp[(1-\beta)\widehat{\mu}_\ell - (1-\alpha)\widehat{\mu}_g]$. Substituting yields a choice rule proportional to $G^\alpha - \lambda L^\beta$, thereby providing cognitive micro-foundations for the PT loss-aversion parameter λ.

Loss aversion in the sense of $\lambda > 1$ obtains whenever $(1-\beta)\widehat{\mu}_\ell > (1-\alpha)\widehat{\mu}_g$, or equivalently $\frac{\widehat{\mu}_\ell}{\widehat{\mu}_g} > \frac{1-\alpha}{1-\beta}$—showing how loss-sensitivity and prior expectations jointly determine whether loss aversion arises. The small-stake risk-seeking we document implies the opposite inequality, $(1-\beta)\widehat{\mu}_\ell < (1-\alpha)\widehat{\mu}_g$. Intuitively, this is consistent with gains being larger than losses on average in typical environments—as suggested by the distributions of credits and debits documented by Stewart et al. (2006)—though this is neither a necessary nor sufficient condition.

Both a common prior with $\widehat{\mu} > 0$ and a specification with separate priors satisfying $(1-\beta)\widehat{\mu}_\ell < (1-\alpha)\widehat{\mu}_g$ can therefore account for the patterns we document. Which specification generates more accurate predictions is testable—a point which we examine in some detail in Online Appendix E.

3 Experiment II: Environmental Adaptation

DbS and the NCM both incorporate loss-sensitivity and exploit statistical regularities in the environment to mitigate cognitive frictions in the choice process. Despite these commonalities, the two models make sharply divergent predictions in specific choice situations. Here, we exploit these differences to test the models against each other.

In Experiment II we test a core mechanism underlying both DbS and the NCM: choice patterns over mixed gain–loss gambles should be systematically influenced by the distribution of gains in the immediate decision environment.[9] We therefore randomly vary whether subjects are exposed to small or large gains, under either hypothetical or real incentives, in a 2×2 between-subject design. Focusing on gains ensures that the test is diagnostic regardless of the prior specification in the NCM, since both the common-prior and separate-prior versions make identical predictions for gain adaptation—an issue we examine in more detail in Online Appendix E.

Experimental stimuli and treatments. The main treatment consists of presenting subjects with 126 pure-gain 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 a gain x obtaining with probability 0.5 or else 0, and a sure amount c. The treatment manipulates whether subjects are exposed to *small* (mean = £19.78, range £2–£32) or *large* (mean = £79.11, range £8–£128) gains, with the large-gains distribution being a fourfold scaling of the small-gains distribution. The stimuli were designed so that the variance of the logged outcomes, $\ln(x)$ and $\ln(c)$, is identical across conditions, ensuring that the two distributions differ only in their mean and not in their spread on the log scale. The distribution of stimuli is approximately log-normal in both conditions.[10] See Online Appendix F.1 for full details.

Orthogonal to the main adaptation treatment, we varied whether choices were

[9] Adaptation to different stimulus distributions provides a diagnostic test between the two models. Previous tests of DbS have manipulated gains and losses *jointly* (Walasek and Stewart, 2015), which is insufficient to separate their predictions; our design instead manipulates the gain environment while holding the loss environment fixed. Earlier tests of DbS have also proven methodologically contentious. Walasek and Stewart (2015) report that jointly manipulating gains and losses affected risk-taking over mixed wagers in the predicted direction, but André and de Langhe (2021) show that the parametric results were an artifact of using different stimuli across treatments (but see also the reply by Walasek et al., 2021). A subsequent adversarial collaboration by Alempaki et al. (2019) reached similar conclusions for related designs by Stewart et al. (2014).

[10] The use of approximately log-normal distributions serves to align the experimental environment with the scale on which the NCM operates. Under DbS, this implies that the shape of the rank function remains unchanged, with only the effective scale of the distribution shifting.

hypothetical or incentivized. In the incentivized condition, one in ten subjects was randomly selected to play one of their choices for real and received an endowment of £30—equal to the maximum possible loss in the test phase—from which gains and losses were added or deducted accordingly. The use of a random incentive scheme rather than paying all subjects follows standard practice and does not affect choice behaviour (Li and Vieider, 2025).

After the adaptation phase, subjects received brief instructions introducing the mixed gain–loss choices. The test stimuli follow the same structure as in Experiment I, with outcomes ranging from £2 to £30 in steps of £2 and probabilities fixed at 50–50. To keep the test phase shorter than the adaptation phase, we retained only choices for which $|G - L| \leq 4$, yielding 84 test choices. These near-diagonal choices are the most informative about loss-sensitivity, since gains and losses are closely matched and asymmetries in their treatment largely determine acceptance rates. We recruited $N = 400$ subjects on Prolific UK (approximately 100 per treatment cell) at a rate of £10 per hour following Prolific regulations.

Predictions of Decision-by-Sampling. DbS predicts that the distribution of outcomes encountered in the environment directly affects the ranking of outcomes, because previously observed outcomes determine the samples drawn from memory against which a given outcome x is evaluated. Exposing subjects to larger gains shifts the comparison set upward, lowering the rank attributed to any given gain x and making it appear less favorable relative to losses. This effect operates through a shift in the effective scale of the gain distribution—captured by an increase in the parameter G_0—which reduces the range of stakes over which gains dominate losses in rank terms, thereby increasing risk aversion over subsequent mixed gain–loss choices relative to subjects exposed to smaller gains.[11]

Predictions of the Noisy Cognition Model. Like DbS, the NCM predicts that the distribution of previously encountered outcomes affects subsequent choices—

[11] This prediction further assumes that adaptation can occur quickly within the course of an experiment, following an equivalent assumption by Walasek and Stewart (2015). If such rapid adaptation does not occur, DbS instead predicts no change.

but the direction of the prediction is exactly opposite. Exposing subjects to larger gains shifts the prior mean upward, increasing the probability of accepting any given mixed wager.[12] For gain adaptation, the NCM and DbS therefore make opposite predictions, providing a clean test between the two models.[13]

Results. Panel A of Figure 3 shows nonparametric choice proportions for fair spreads after exposure to small versus large gains in the adaptation phase. The choice proportions are aggregated across incentive conditions. There is a clear treatment effect: subjects exposed to large gains are distinctly more risk-seeking than those exposed to small gains. This pattern is consistent with the prediction of the NCM but contradicts the prediction of DbS.[14]

To statistically examine treatment effects across the full dataset, we estimate a reduced-form Probit regression of the form $\Phi(\gamma_0 + \gamma_g \ln(G) + \gamma_\ell \ln(L))$, with hierarchical intercepts and gain slopes, and treatment-specific loss slopes. The model is estimated using Bayesian methods in `Stan` (see Online Appendix H for details). This specification allows us to recover individual-level slope ratios $|\gamma_\ell/\gamma_g|$, which capture loss-sensitivity, alongside individual-level intercepts γ_0 within a single model applied consistently across all experiments.

Loss-sensitivity is confirmed across all treatments (see Online Appendix H.2 for

[12] Unlike DbS, where outcomes are perceived objectively and noise enters only at the stage of rank-based utility assignment, the NCM requires that environmental distributions be learned from noisy posterior inferences—the only information available to the decision maker. As a consequence, the learned prior need not faithfully reproduce the true statistics of the environment. In particular, the accuracy of learning depends on the precision of the inference itself: because coding noise is higher for gains than for losses ($\nu_g^2 > \nu_\ell^2$), learning from a gain environment will be slower and less accurate than learning from a loss environment.

[13] Prospect Theory makes no explicit predictions about how prior exposure to different gain distributions should affect subsequent mixed choices. The only conceivable channel would be adaptation of the reference point, but even this yields no clear prediction: larger gains might shift the reference point upward, increasing the range of outcomes falling into the loss domain, with ambiguous net effects on risk-taking. This makes it difficult to derive refutable ex ante predictions from PT in this context.

[14] One feature of the risk-taking proportions is that we no longer observe the large levels of small-stake risk-taking seen in the representative data shown in the introduction. A plausible explanation is that subjects enter the experiment with pessimistic expectations, which are then updated over the course of the experiment through noisy posterior inferences. Learning speed in such a system depends on signal precision, so learning in gain environments will be slower than in loss or mixed environments. This hypothesis is further supported by results on loss adaptation—shown in Online Appendix E.1—after which mixed choices exhibit high levels of small-stake risk-taking consistent with those observed in the representative subject pool.

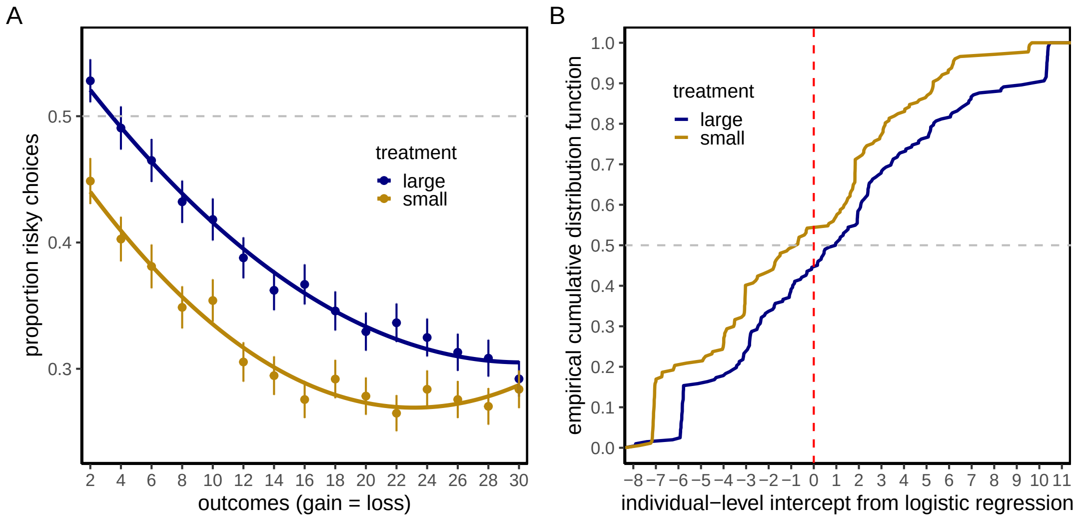

Figure 3: Risk taking proportions after small vs large gain adaptation

Panel A shows acceptance proportions of the risky wager for fair spreads around 0, where $G = L$. Panel B shows empirical cumulative distribution functions of mean individual-level intercepts estimated in a probit regression with hierarchical intercepts qnd slopes and treatment-specific aggregate parameters. Some outlying observations may be cut from the graph for a better visual experience.

full regression tables). Importantly, the aggregate intercept—capturing risk-taking when $\ln(G) = \ln(L) = 0$ (i.e. $G = L = 1$)—is credibly larger for subjects exposed to large gains than for those exposed to small gains (difference = 0.984, 95% CrI $[0.375, 1.598]$). This pattern is consistent with the NCM prediction that exposure to larger gains shifts the prior mean upward. Panel B of Figure 3 shows the empirical distribution of individual-level intercepts separately by treatment. The distribution for the large-gains treatment is shifted monotonically to the right of the small-gains distribution, confirmed by a Wilcoxon rank-sum test on individual posterior mean intercepts ($p < 0.001$).

Real incentives have similar effect as in Experiment I (Li and Vieider, 2025): when losses are administered from an endowment, subjects are more risk-seeking, but the effect falls short of conventional significance thresholds (Wilcoxon rank-sum test, $p = 0.123$). Importantly, the adaptation treatment effect holds regardless of incentive condition, both for hypothetical choices ($p = 0.008$) and for real incentives ($p = 0.018$). Online Appendix H.2 presents separate results by incentive condition and additional statistical tests.

Discussion of Results. Our diagnostic test yields a clear result. The NCM predicts that exposure to larger gains shifts the prior mean upward, increasing risk-taking in subsequent mixed choices—a prediction supported by the data. By contrast, the opposite prediction of DbS is not supported. In Online Appendix E, we further examine whether the prior in the NCM is common to gains and losses, or whether separate priors are required.

An interesting question concerns the deeper reason behind the contrasting predictions of DbS and the NCM. While both models allow for differential sensitivity to gains and losses, in DbS behaviour is entirely determined by the relative ranking of gains and losses in the environment. As a result, assigning a higher rank to a loss x than to a gain of equal magnitude implies risk aversion at that stake level. In the NCM, by contrast, the as-if utility weight coincides with the Bayesian evidence weight—that is, the weight placed on the signal relative to the prior. Lower coding precision therefore induces stronger regression toward the prior mean, so that behavioural predictions for fair spreads depend critically on the position of x relative to that prior. The prediction of increased risk-taking after exposure to larger gains thus arises from the upward shift in the prior mean μ, rather than from any change in sensitivity to gains or losses per se.

The integration of prior information to decode noisy signals is a central feature of the NCM. Ultimately, the difference between the two models reflects the NCM's foundation in a constrained optimisation framework: given noise in outcome assessments, the Bayesian estimator pools signal and prior information in proportion to their respective informational contents (Bishop, 2006; Ma et al., 2023). Such optimal information aggregation may be important from an evolutionary perspective. Moreover, grounding the model in a mechanism shared across human and animal signal processing—including sensorimotor tasks—suggests the possibility of unifying principles from which behavioural predictions can be derived.[15]

[15] This does not imply that the rank-based comparison mechanism underlying DbS lacks merit. As a description of the cognitive process by which outcomes are evaluated relative to a reference distribution, DbS captures a psychologically important feature of behaviour. The issue is therefore not with the mechanism itself, but with the literal predictions that follow from treating it as a fully specified model of choice. We return to this point in the Conclusion.

4 Experiment III: Visual Outcome Displays

Experiment II discriminated between the NCM and DbS, showing that adaptation to different gain environments supports the NCM account. It could not, however, conclusively separate these adaptive accounts from PT. We therefore introduce a new experiment that provides a clean separation between the NCM and PT. It also offers an additional diagnostic test of the NCM against DbS.

The experiment manipulates whether subjects receive a visual aid conveying the size of gains and losses in addition to their numerical description. From the perspective of PT, this manipulation is irrelevant: behaviour is driven by stable preferences applied to numerically described outcomes, and presentational format leaves the value function unchanged. The same prediction—no effect of the visual aid—follows from DbS, but for a different reason: outcomes are perceived objectively and G and L are ranked *independently* against their respective memory samples, so the visual display does not affect rank-based utility assignment.

Under the NCM, by contrast, the presentational mode matters directly. Loss-sensitivity arises from an asymmetry in the internal representation of G and L, and a visual aid that makes gains and losses more directly comparable is expected to reduce this asymmetry. The predicted consequence is an attenuation of loss-sensitivity and a flattening of the stake-dependence of risk-taking. We exploit this contrast to test the NCM against both PT and DbS.

Experimental stimuli and treatments. The main manipulation varied whether gains and losses were described numerically—as in Experiments I and II—or whether a visual aid was additionally provided. Figure 4 shows screenshots of the decision environment in the two conditions. In the visual condition, coloured horizontal bars proportional in length to the amounts at stake accompanied the numerical descriptions of G and L. Orthogonal to the main treatment, we once again varied whether choices were hypothetical or played for real, following the same procedures as in Experiments I and II.

The experimental procedures were otherwise identical to those in Experiments I

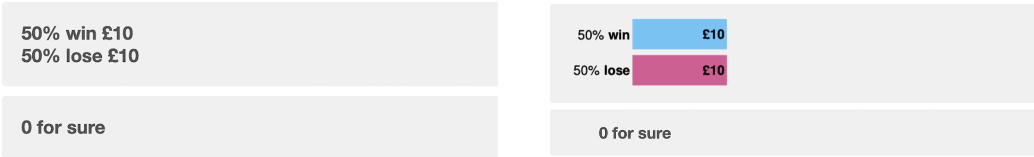

Figure 4: Textual versus visual displays of choices

Decision environment in the textual condition (left hand side) and in the visual condition (right hand side) for a 50-50 chance at winning or losing £10.

and II. We targeted $N = 400$ subjects on Prolific UK, using the same test stimuli as in Experiment I. As no lengthy questionnaire was included, the experiment was shorter, with a median completion time of approximately 12 minutes.

Predictions of the Noisy Cognition Model. Loss-sensitivity in the NCM emerges from an asymmetry in the internal representation of G and L, reflected in the ratio $\beta/\alpha > 1$. We hypothesize that the visual aid reduces this asymmetry, yielding a compression of β/α toward 1 and a corresponding flattening of the stake-dependence of risk-taking.

Predictions of PT and DbS. Under the standard interpretation of PT, attitudes toward gains and losses are governed by stable taste functionals applied to objectively perceived outcomes. Presentational format is therefore irrelevant, and the visual aid should have no effect on choice. DbS yields the same prediction for a different reason: outcomes are perceived objectively and gains and losses are ranked independently against their respective memory samples, so the visual display does not affect rank-based utility assignment.

Results. Figure 5 shows the effect of the presentation-format manipulation. Panel A reports choice proportions for fair spreads across stake levels. In the textual condition, risk-seeking for small stakes gives way to risk aversion as stakes increase, replicating the pattern observed in Experiments I and II. The visual condition exhibits the same qualitative pattern but with a markedly flatter slope, consistent with the NCM prediction that the visual aid reduces the asymmetry in the internal representation of G and L and thereby attenuates loss-sensitivity.

Panel B shows the empirical cumulative distribution function of $\ln(|\gamma_\ell/\gamma_g|)$—the

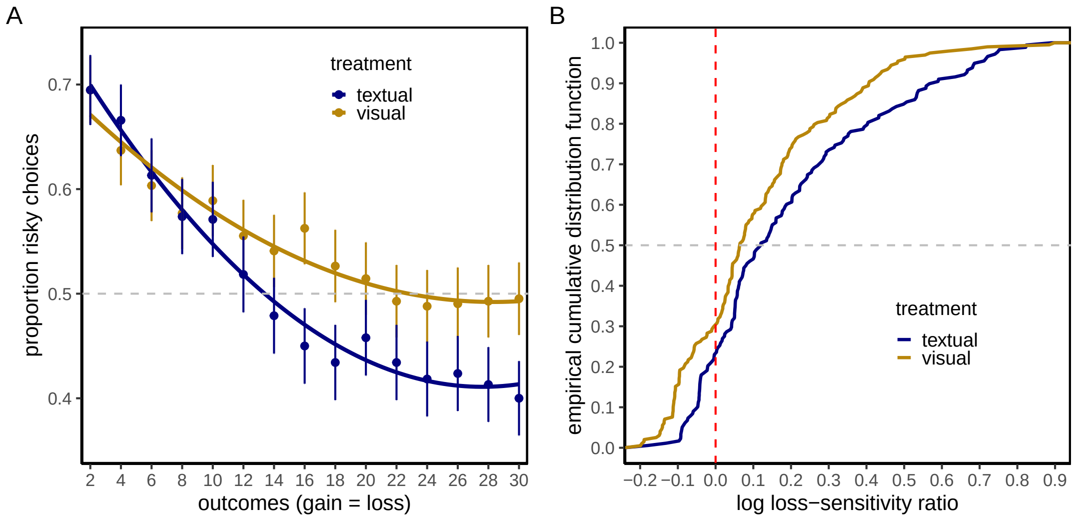

Figure 5: Risk taking proportions for described vs visually displayed outcomes

Panel A shows acceptance proportions of the risky wager for fair spreads around 0, where $G = L$. Panel B shows empirical cumulative distribution functions of logged mean individual-level loss-sensitivity ratios (i.e. $ln\,(\gamma_\ell/\gamma_g)$) estimated in a probit regression with hierarchical intercepts and slopes and treatment-specific population parameters.

log of the reduced-form counterpart of β/α in the NCM—estimated from the Probit regression described above. The population-level slope ratio decreases in the visual condition (-0.124, 95% CrI $[-0.190, -0.059]$), consistent with the NCM prediction. The individual-level eCDFs confirm this result: while the great majority of subjects exhibit loss-sensitivity in both conditions, the distribution in the visual condition is clearly shifted to the left, indicating reduced loss-sensitivity ($p < 0.001$, Wilcoxon rank-sum test on individual posterior mean slope ratios).

Providing real incentives increases risk-taking, but the effect falls short of conventional significance thresholds ($p = 0.128$, Wilcoxon rank-sum test). Importantly, the effect of the visual manipulation is robust across incentive conditions: the slope ratio is credibly reduced both in the hypothetical condition ($p = 0.028$) and in the real-incentive condition ($p < 0.001$). Online Appendix H.2 presents separate results by incentive condition and additional analyses.

Discussion of Results. The results highlight our interpretation of the NCM as endogenizing PT-style parameters. Even under a perceptual interpretation of PT—where parameters are allowed to originate in perceptual processes rather than reflecting stable tastes—PT lacks the explicit machinery to predict how a visual aid should affect choices. The NCM provides precisely this machinery: by tracing loss-sensitivity to asymmetric perceptual noise, it generates a clear prediction about the direction and nature of the visual-manipulation effect, which is borne out in the data.

5 Loss-sensitivity and cognitive ability

We have so far focused primarily on aggregate patterns. Here, we analyze the representative-sample experiment in greater detail, examining between-subject heterogeneity. This allows us to revisit a puzzle documented by Chapman et al. (2024): respondents with higher cognitive ability appear both more loss averse and more willing to take risks over pure gain wagers. As we show, the NCM organizes this apparent contradiction within a unified set of model parameters.

Data analysis. We analyze individual-level heterogeneity in the representative-sample experiment using the reduced-form hierarchical Probit model employed throughout the paper. The linear index $\gamma_0 + \gamma_g \ln(G) + \gamma_\ell \ln(L)$ maps directly into the NCM: loss-sensitivity corresponds to $\beta/\alpha = |\gamma_\ell/\gamma_g|$, while the weighted contribution of the prior mean satisfies $(\beta - \alpha)\widehat{\mu} = \gamma_0$. These two quantities are our main variables of interest, capturing loss-sensitivity and small-stake risk-taking at $G = L = 1$, respectively. We estimate the regressions using a Bayesian measurement-error model applied to the individual-level posterior parameter estimates, thereby accounting for posterior uncertainty in γ_0 and γ_ℓ/γ_g. To ensure robustness to outliers, we employ a Student-t likelihood with 3 degrees of freedom (Gelman et al., 2014). Details are provided in Online Appendix I.

Loss-sensitivity and cognitive ability. We observe considerable heterogeneity across subjects. Our main variable of interest is cognitive ability, measured as the number of correct answers to 11 questions on numeracy and general reasoning ability administered in the final questionnaire.[16] Online Appendix I reports the questions, details the econometric specification, and provides additional analysis

[16] The battery consists of 5 numeracy questions, 3 questions from the Cognitive Reflection Test (Frederick, 2005), and 3 questions from the International Cognitive Ability Resource (Condon and Revelle, 2014).

and figures of parameter heterogeneity.

Table 1: Regression of NCM parameters

	Loss-sensitivity ($\|\gamma_\ell/\gamma_g\|$)				**Small-stake risk-taking** (γ_0)			
	I	II	III	IV	V	VI	VII	VIII
cognitive ability	**−0.015**	**−0.020**	**−0.018**	−0.009	**−0.405**	**−0.450**	**−0.443**	**−0.353**
	(0.005)	(0.006)	(0.006)	(0.006)	(0.061)	(0.067)	(0.066)	(0.070)
cog. ability^2		**−0.013**	**−0.013**	**−0.013**		−0.120	−0.122	−0.113
		(0.005)	(0.006)	(0.006)		(0.065)	(0.062)	(0.064)
income			**−0.013**	−0.006			−0.073	−0.009
			(0.006)	(0.006)			(0.066)	(0.066)
age				**0.013**				**0.131**
				(0.006)				(0.064)
female				**0.079**				**0.833**
				(0.012)				(0.131)
real incentives	−0.020	−0.019	−0.020	−0.020	0.059	0.064	0.056	0.105
	(0.011)	(0.012)	(0.012)	(0.012)	(0.132)	(0.130)	(0.132)	(0.137)
Constant	**1.128**	**1.141**	**1.143**	**1.105**	**1.167**	**1.287**	**1.294**	**0.844**
	(0.007)	(0.009)	(0.009)	(0.010)	(0.076)	(0.100)	(0.098)	(0.123)
Observations	1506	1506	1506	1460	1506	1506	1506	1460

The regression of individual-level posterior means on respondents' characteristics is based on an outlier-robust Bayesian measurement error model (see Online Appendix I for details). Effects significant at the 5% level are highlighted in bold, and standard errors are reported in parentheses. Continuous variables are entered as z-scores. The reduced sample in reg IV and reg VIII reflects missing age and sex identifiers for 46 respondents.

Table 1 reports the results. Regressions I–IV relate loss-sensitivity to cognitive ability and a set of controls. Cognitive ability enters negatively across all specifications, with the effect remaining stable as controls are added. The quadratic term is also negative, indicating that the effect strengthens at higher ability levels: subjects who perform well on our tests pay more equal attention to gains and losses. Women and older subjects exhibit higher loss-sensitivity. Turning to small-stake risk-taking (reg. V–VIII), cognitive ability again enters negatively throughout: higher-ability subjects are less inclined to accept the lottery at small stakes. The effect remains robust to the inclusion of demographic controls. Women and older subjects exhibit higher small-stake risk-taking.

Taken together, cognitive ability is associated with both lower loss-sensitivity and lower small-stake risk-taking, consistent with higher-ability subjects paying more equal attention to gains and losses.

Discussion of results. The results above paint a coherent picture. Subjects who perform well on our cognitive ability tests are both less loss-sensitive and less willing to accept small-stake mixed gambles. This pattern helps organize the seemingly paradoxical correlations documented by Chapman et al. (2024). Within the NCM there is no contradiction. If higher cognitive ability reduces internal noise in the representation of gains ν_g—and hence increases α—risk-taking for gains increases, since α governs as-if utility curvature over gains.[17] At the same time, the ratio β/α shrinks toward 1, reducing small-stake risk-taking in mixed gambles. The two effects therefore follow coherently from a single mechanism.

A consistent set of NCM parameters with $\beta > \alpha$ and $\widehat{\mu} > 0$ can also organize a broader set of findings. First, it explains the loss-sensitivity and small-stake risk-taking documented in Experiment I: $\beta > \alpha$ generates loss-sensitivity directly, while $\widehat{\mu} > 0$ shifts the decision threshold in favour of the lottery for small stakes. Second, $\alpha < 1$ also explains the small-stake risk aversion documented by Khaw et al. (2021) for pure gain lotteries, since it governs as-if utility curvature.[18] Third, these parameter values are also consistent with the treatment effects documented in Experiments II and III.[19]

[17]In reality, both ν_g and ν_ℓ may be lower for higher-ability individuals. The result nevertheless holds as long as the proportional reduction in ν_g exceeds that in ν_ℓ, which seems plausible given the higher baseline level of ν_g relative to ν_ℓ.

[18]In that setting, the prior mean $\widehat{\mu}$ drops out of the decision equation, since both the gain and the status quo of zero are evaluated relative to the same prior.

[19]Experiment II shifts the prior mean, modulating risk-taking propensity without affecting loss-sensitivity β/α. Experiment III directly modulates relative sensitivity to gains and losses by reducing perceptual noise through a visual aid manipulation. This interpretation is consistent with Oprea and Vieider (2026), who show that reducing outcome noise for pure gain lotteries increases risk-taking; Experiment III extends this logic to mixed gambles, where the same mechanism generates a distinct prediction about loss-sensitivity.

6 Discussion and Conclusion

Loss aversion is a cornerstone of behavioural economics. Here, we have re-examined risk-taking in mixed gambles through the lens of adaptive models of choice. Across a representative sample and a series of targeted treatment conditions, our results support an alternative concept—*loss-sensitivity*—whereby seemingly loss-averse behaviour over some stake range emerges from heightened perceptual sensitivity to losses relative to gains. This perspective helps reconcile several apparently contradictory findings in the literature: the coexistence of gain-seeking for small stakes (Chapman et al., 2024), highly variable loss-aversion estimates in prospect-theory estimations (Brown et al., 2024), and small-stake risk aversion for pure gains (Khaw et al., 2021). The noisy cognition model provides a single, coherent set of parameters that organizes these findings simultaneously.

Interpretation and implications. Decision-by-Sampling and the Noisy Cognition Model are both adaptive accounts of how decision makers cope with cognitive frictions in information processing. The difference in their predictions ultimately stems from the inference mechanism underlying the NCM. In DbS, differences in the experienced magnitudes of gains and losses translate directly into differences in utility through the ordinal ranking of outcomes relative to typical gains or losses. In the NCM, by contrast, noisy signals are decoded using a prior that summarizes the learned structure of the environment. This Bayesian decoding is what generates the distinct behavioural predictions of the NCM.

There is a second distinction between the models arising from the Bayesian inference mechanism, concerning variability in choices across trials. Prospect theory makes no predictions about stochastic choice and is typically augmented with a random-utility term; when options become difficult to discriminate, as can occur for small stakes, this creates well-known difficulties in mapping stochastic choice behaviour onto taste parameters (Ballinger and Wilcox, 1997; Wilcox, 2011; Apesteguia and Ballester, 2018; McGranaghan et al., 2024). Decision-by-Sampling predicts a similar qualitative pattern as a behavioural feature: noise will monotonically increase as discriminability declines in the tails of the environmental distributions, with choice reverting to 50-50 indifference in the limit. This, however, does not occur in the Noisy Cognition Model. Under Bayesian inference, choices should stabilize when signals become very noisy, since decision makers will rely increasingly on relatively stable priors in drawing their inferences. Choice variability should instead be maximized at intermediate levels of signal noise, generating an inverse-U relationship between outcome discriminability and stochastic choice. Vieider (2026) examines this prediction and finds support for the Bayesian inverse-U pattern. Together with the present findings, this provides convergent evidence that Bayesian inference serves as an organizing principle of cognition under uncertainty more broadly.

A natural question concerns the role of Decision by Sampling in light of the results reported here. While our tests do not support the behavioural predictions of DbS as a model of choice, the mechanism of noisy value coding it proposes remains highly complementary to the noisy number coding underlying the Noisy Cognition Model. Indeed, DbS has been shown to implement efficient adaptation to environmental distributions under certain conditions (Bhui and Gershman, 2018). One way of viewing DbS is therefore as a mechanism that achieves efficient value attribution by sampling from values stored in memory (Shadlen and Shohamy, 2016). From this perspective, it could be combined with the Bayesian inference mechanism of the Noisy Cognition Model, unifying the power of noisy inference with the neurobiological realism of adaptive value coding.

Deeper mechanisms and open questions. Heightened sensitivity to losses carries a natural evolutionary interpretation. Robson (2001b) and Netzer (2009) argue that evolution would have endowed humans with a proximate utility function over consumption geared towards fitness maximization. Since fitness is only observable over very long time horizons, whereas consumption is observable in each period, inter-temporal stability of the consumption profile becomes a proximate goal—and large losses in consumption can do more damage to fitness than equivalent gains can do good. Loss-sensitivity may thus reflect an evolutionarily calibrated asymmetry in the mapping from consumption to fitness utility.

This evolutionary perspective also raises a broader question about the nature of the parameters commonly used to describe choice behaviour. If loss-sensitivity arises from perceptual and inferential mechanisms rather than stable tastes, then many parameters traditionally interpreted as preferences may instead reflect properties of cognitive processing. Understanding how such parameters emerge from the interaction of perception, inference, and learning may therefore be key to developing a more unified account of decision making under uncertainty.

References

Alaoui, Larbi, and Antonio Penta (2026) 'Attitudes towards success and failure.' *American Economic Journal: Microeconomics* 18, 105–145

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

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

Ballinger, T Parker, and Nathaniel T Wilcox (1997) 'Decisions, error and heterogeneity.' *The Economic Journal* 107(443), 1090–1105

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

Bhui, Rahul, and Samuel J Gershman (2018) 'Decision by sampling implements efficient coding of psychoeconomic functions.' *Psychological Review* 125(6), 985

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

Bleichrodt, Han, and Olivier L'Haridon (2023) 'Prospect theory's 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

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* 92, 2893–2922

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

Condon, David M, and William Revelle (2014) 'The international cognitive ability resource: Development and initial validation of a public-domain measure.' *Intelligence* 43, 52–64

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

Frederick, Shane (2005) 'Cognitive reflection and decision making.' *Journal of Economic perspectives* 19(4), 25–42

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

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

Kahneman, Daniel, Jack L Knetsch, and Richard H Thaler (1990) 'Experimental tests of the endowment effect and the coase theorem.' *Journal of Political Economy* 98(6), 1325–1348

_ (1991) 'Anomalies: The endowment effect, loss aversion, and status quo bias.' *Journal of Economic Perspectives* 5(1), 193–206

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

Li, Yuchi, and Ferdinand M. Vieider (2025) 'Incentive effects on decisions under risk and over time: A meta-analysis of experimental studies.' *Working Paper*

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

McGranaghan, Christina, Kirby Nielsen, Ted O'Donoghue, Jason Somerville, and Charles D Sprenger (2024) 'Distinguishing common ratio preferences from common ratio effects using paired valuation tasks.' *American Economic Review* 114(2), 307–347

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 (2025) 'Risk perception: measurement and aggregation.' *Journal of the European Economic Association* 23(4), 1309–1349

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 the description-experience gap.' *Working Paper*

_ (2026) 'The new psychophysics of risk and time.' *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

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

Rabin, Matthew, and Richard H. Thaler (2001) 'Anomalies: Risk Aversion.' *The Journal of Economic Perspectives* 15(1), 219–232

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

Samuelson, William, and Richard Zeckhauser (1988) 'Status quo bias in decision making.' *Journal of Risk and Uncertainty* 1(1), 7–59

Shadlen, Michael N, and Daphna Shohamy (2016) 'Decision making and sequential sampling from memory.' *Neuron* 90(5), 927–939

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

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

Vieider, Ferdinand M. (2024) 'Decisions under uncertainty as bayesian inference on choice options.' *Management Science* 70(12), 9014–9030

_ (2026) 'The Structure of Stochastic Choice.' Technical Report, Ghent University Discussion Papers

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).' *Journal of Experimental Psychology: General* 150, 2666–

2670

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

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 Additional Details for Experiment I

A.1 Stimuli

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 Choice proportions for fair spreads

Stakes (£)	Real(%)	Hypo(%)	χ^2 test
2	67.1	66.7	0.828
4	64.7	62.0	0.146
6	60.4	57.5	0.127
8	58.5	52.6	0.002
10	55.0	53.2	0.336
12	50.5	47.0	0.067
14	47.9	42.3	0.003
16	46.9	40.9	0.002
18	45.7	39.9	0.002
20	44.1	38.7	0.005
22	41.0	37.2	0.045
24	41.0	36.5	0.016
26	40.4	35.2	0.005
28	38.3	35.6	0.146
30	39.3	34.2	0.006

Table 3: The table shows the acceptance proportions of risky wagers for fair spreads $(x, 0.5; -x)$ under both hypothetical and real conditions. The last column reports p-values from Pearson's χ^2 tests comparing the difference between two choices proportions at each stake level.

A.3 Questionnaire and tests of cognitive ability

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 6: 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 7: Screenshots of questionnaire questions on numeracy test.

6. A bat and a ball cost £1.10 in total. The bat costs £1.00 more than the ball. How much does the ball cost? Please indicate it in pence:

[ ]

7. If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets? Please indicate it in minutes:

[ ]

8. In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake? Please indicate it in days:

[ ]

9. In the following number sequence, what number comes next?

11 14 21 32 47 ?

[ ]

10. In the following alphanumeric series, what letter comes next?

V Q M J H ?

E	F	G	H	I	J

11. Zach is taller than Matt and Richard is short than Zach. Which of the following statements would be most accurate?

- Richard is taller than Matt.
- Richard is shorter than Matt.
- Richard is as tall as Matt.
- It’s impossible to tell.

Figure 8: Screenshots of questionnaire questions on CRT and ICAR test.

B Reference-dependent Expected Utility Theory

B.1 Representation characterizations

Alaoui and Penta (2026) study four attitudes that capture reversals in risk attitudes around an exogenous reference point x_0 within an otherwise standard expected utility framework. They characterize local reversals of risk attitude for mixed lotteries as the probability of success changes. To account for our findings, the reference point needs to coincide with the status quo, i.e. $x_0 = 0$, given that neither $x_0 < 0$ nor $x_0 > 0$ can simultaneously account for small stake risk seeking and monotonic stake-dependence of the type we document.

Amongst the four attitudes, *Failure avoidance* in its continuous instantiation — the case closest to a loss averse value function in PT — and *Success seeking* are immediately ruled out because they imply local concavity around the reference point of 0, thus implying small stake risk aversion.[20] *Success attachment* and *Failure resignation*—as well as "*Aspiration*", which results from their intersection—are locally convex at least under some specifications and can thus account for small stake risk-seeking of the type we document. So can the *discontinuous* case of *Failure avoidance*, for $K \equiv u(x_0^+) - u(x_0^-)$ if $u(x_0^+) > |u(x_0^-)|$, that is, if the upward jump on the gain side dominates the relative downward jump in the loss side—a case that is consistent with the model, even though Alaoui and Penta (2026) do not explicitly discuss it.[21]

The model can also account for stake effects of the type we document. Particularly, the "*Aspiration*" attitude combines concave utility for losses with concave utility for gains, thereby directly fulfills the condition for loss-sensitivity. *Failure resig-*

[20] Continuous failure avoidance—the case equivalent to a loss averse value function under PT—is ruled out because the model does not allow for probability weighting and imposes discipline on utility curvature under the form of consistent curvature to either side of the reference point. PT itself does not impose these restrictions, so that it indeed has a variety of ways of fitting the patterns we document ex post.

[21] The condition $u(x_0^+) > |u(x_0^-)|$ we impose is stricter than the condition $K > 0$ imposed by Alaoui and Penta (2026), which implies $u(x_0^+) > u(x_0^-)$. This ensures that the discontinuity K straddles the references point, thus facilitating our discussion of small stake risk seeking, but otherwise happens without loss of generality.

nation offers more flexibility on utility for gains, and thus additionally requires utility for losses to be steeper than for gains to capture loss-sensitivity. *Success attachment* in continuous instantiation can also capture loss-sensitivity as long as utility for losses is steeper than utility for gains—and the same is the case for both *Success attachment* and *Failure avoidance* in the discontinuous case.

B.2 Utility shape

We next examine whether the empirical patterns in this paper are compatible with the predictions of Alaoui and Penta (2026). To this end, we use the pure gain and pure loss choices from the hypothetical adaptation treatments to examine utility curvature. We focus on the small stake adaptation, since that gives us choices for prizes in the range from £2 to £32, which coincide with the range over which we measure choices over mixed gambles. Choices are invariably between a prize x, payed with a 50% probability, and a sure outcome c.

We estimate a simple aggregate expected utility model, with $u(x) = x^\rho$. We normalize the utility difference between the lottery and the sure amount, $px^\rho - c^\rho$ (and the opposite difference for losses) by σx^ρ, where σ is the error SD of the random utility error, and the multiplication with the utility of the largest outcome serves to impose monotonicity on the mapping from binary choices to the expected utility representation by contextualizing it (Wilcox, 2011).

We estimate the model in Stan using the following code (where for losses we simply replace choices of the risky option with choices of the sure option):

```
data{
    int<lower=1> N;
    array[N] real x;
    array[N] real c;
    array[N] int choice_risky;
}
parameters{
    real rho;
    real<lower=0> sigma;
```

```
}
model{
    rho ~ normal( 1 , 1 );
    sigma ~ normal( 0.5 , 1 );

    for ( i in 1:N )
        choice_risky[i] ~ bernoulli_logit( (0.5 * pow(x[i], rho) -
            pow(c[i], rho)) / (sigma * pow(x[i], rho) ) );
}
```

Results. For gains, we estimate a concave utility function with $\rho_g = 0.663$ and a credible interval of $[0.651, 0.675]$. For losses, we estimate $\rho_l = 0.817$, with a credible interval of $[0.803, 0.832]$. Utility is thus concave for gains, convex for losses, and steeper for losses than for gains, i.e. $\rho_g < \rho_l < 1$. These parameter estimates are consistent with the notion of loss sensitivity and provide independent evidence from the pure outcome domains.

Conclusions. Taken together, the results imply:

1. The special case of *Aspiration* is not supported by the data, as it requires concave utility for losses, whereas the utility we find is convex. The same holds for *Failure resignation*, which also requires concave utility for losses.

2. The continuous case of *Success attachment* is fully compatible with our data patterns, can account for loss-sensitivity, small stake risk-taking, and the utility curvature we document for gains and losses.

3. Both *Success attachment* and *Failure avoidance* in their discontinuous cases can account for our empirical patterns if i) utility for losses is steeper than utility for gains; and ii) $u(x_0^+) > |u(x_0^-)|$, that is, if the upward jump on the gain side at the discontinuity dominates the downward drop on the loss side.

C 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 (5) with the prior in equation (6), 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 (7):

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

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:

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

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:

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

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.

D 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 (5). 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.

E Common Prior versus Separate Priors

According to equation (7), under the common-prior specification, the acceptance of risky choices is proportional to

$$\alpha \ln G - \beta \ln L + (\beta - \alpha)\widehat{\mu},$$

whereas under the separate-priors specification for gains and losses, it is proportional to

$$\alpha \ln G - \beta \ln L - \left[(1-\beta)\widehat{\mu}_\ell - (1-\alpha)\widehat{\mu}_g\right].$$

These expressions imply that, in the gain-adaptation treatments of Experiment II, the common-prior and separate-priors specifications generate the same qualitative prediction: subjects exposed to larger gains should subsequently make more risky choices in mixed lotteries than subjects exposed to smaller gains. The intuition is straightforward as exposure to larger gains should shift the relevant prior mean upward. Under the separate-priors specification, this operates through an increase in the gain prior mean, $\widehat{\mu}_g$; under the common-prior specification, it operates through an increase in the common prior mean, $\widehat{\mu}$. Hence, the prediction is identical across the two specifications.

By contrast, the loss-adaptation treatments yield sharply divergent predictions, which allows the two specifications to be empirically distinguished. In particular:

- Under the separate-priors specification, exposure to larger losses should increase the mean of the loss prior, $\widehat{\mu}_\ell$, thereby *increasing* risk aversion.
- Under the common-prior specification, exposure to larger losses should increase the common prior mean, $\widehat{\mu}$. In turn, this implies an *increase* in risk taking—that is, a *decrease* in risk aversion—just as in the gain-adaptation treatments.

Therefore, whereas the two specifications are observationally equivalent in the gain-adaptation treatments, they generate opposite predictions in the loss-adaptation treatments, providing a clean test between them.

E.1 Loss adaptation experiment

Experimental stimuli and treatments. The treatment manipulation consisted of presenting subjects with pure-loss choices in the first part of the experiment, which served as the *adaptation phase*. This phase varied whether subjects were exposed to *small* or *large* losses. The stimuli were otherwise identical to those used in the adaptation phase of Experiment II (see Appendix F.1), with the only modification being that gains were replaced by losses. The mixed-choice trials in the subsequent test phase were likewise identical to those in Experiment II. The experiment was conducted on Prolific UK with a total sample of $N = 201$ subjects, corresponding to approximately 100 subjects per treatment condition.

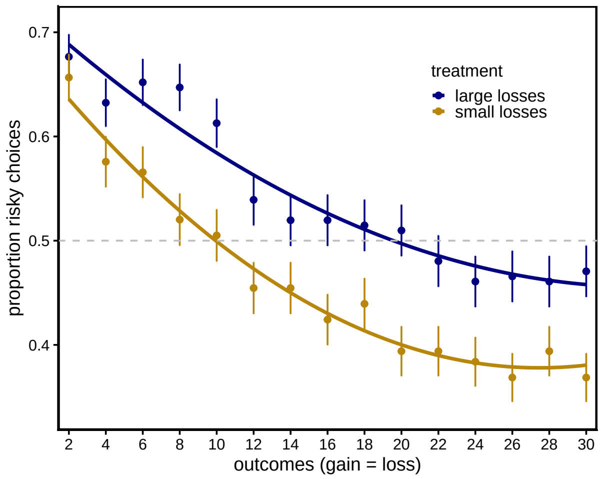

Figure 9: Acceptance proportions of risky wager for fair spreads where $G = L$, after small vs large loss adaptation.

Results. Figure 9 reports the proportion of risky choices for fair spreads in test phase following the loss-adaptation treatments. Subjects in the large-loss treatment exhibited a higher proportion of risky choices than subjects in the small-loss treatment. This pattern is consistent with the prediction of the NCM under the common-prior specification and therefore inconsistent with the separate-priors specification, which predicts the opposite pattern.

An interesting feature is that risk seeking is again pronounced—particularly in the large loss condition. The contrast with the gain adaptation treatment—and indeed with experiments I and III, which contain only mixed gain-loss choices—is noteworthy. A plausible explanation is that subjects enter the experiment with a pessimistic prior, expecting very small payoffs, and that learning subsequently proceeds at different rates based on losses vs gains. This is fully consistent with the logic of the NCM, where learning has to proceed based on noisy inferences about the choice quantities. If gains are inferred less precisely, the learning rate based on gains will be lower than the learning rate from losses, thus explaining the different levels of risk-taking in the gain-adaptation versus loss-adaptation treatments.

E.2 Joint manipulation of gains and losses

A potential concern with our pure gain and pure loss adaptation treatments is that subjects could infer something about the dimension they do not see from the stimuli they experience. While this does not affect our diagnostic gain adaptation treatment in the main text, it could affect our conclusions on whether a common prior can account for our results in the NCM. To address this question, we ran additional experiments in which subjects are either exposed to i) small gains and large losses; or ii) large gains and small losses. These ought to elicit large behavioural differences if there are separate priors for gains and losses, but should result in approximately equal choice proportions across conditions if the prior is common to gains and losses.

We started with a hypothetical experiment, using the same stimuli and procedures as for our pure gain and pure loss adaptation experiments. Other than in those experiments, however, we exposed subjects to either i) small gains and large losses; or ii) large gains and small losses. We ran this experiment twice: once for hypothetical incentives, as in our preferred specification; and once in an incentivized condition. In the incentivized condition we incentivized the mixed choices only, given the difficulties of incentivizing the loss adaptation treatment, where the largest loss amounts to £128. We administered the experiment to 400 subjects on Prolific UK (200 subjects hypothetical, 200 incentivized), following procedures analogous to those used in the experiments described in the paper. Gains and losses were presented in separate blocks and introduced by short instructions. The order of presentation of these blocks was randomized.

Results. Figure 10 shows the results for the mean-preserving spreads around 0. We observe the usual increase in risk-taking when moving from hypothetical incentives (panel A) to real incentives (panel B). In neither case, however, is there a level effect as observed when manipulating gains or losses separately.

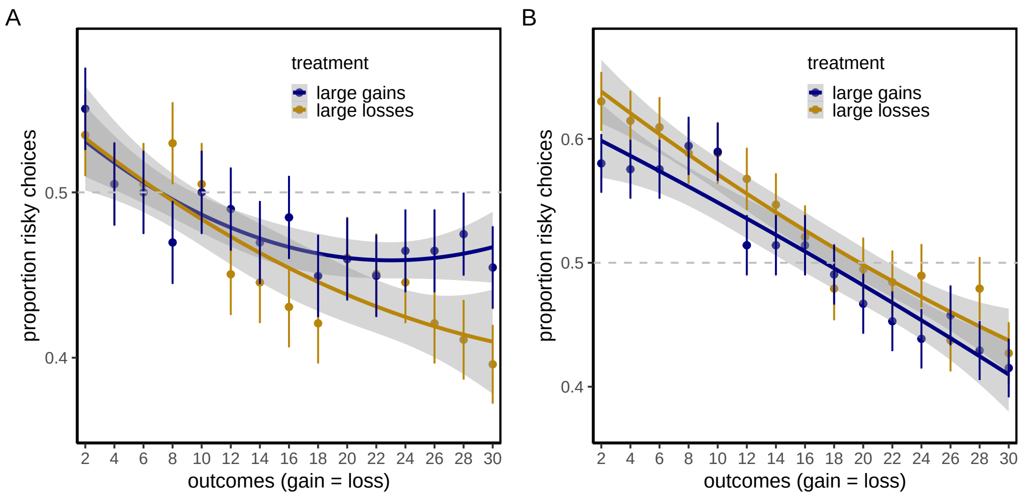

Figure 10: Adaptation to gains and losses. Hypothetical (A) and real (B)

These results thus add to the pure loss adaptation results in supporting a common prior interpretation for the NCM. Such an interpretation is indeed consistent with the modelling assumptions for gains, where the lottery prize x and the sure amount c are also assumed to have a common prior—even though the distribution of x dominates that of c. We discuss potential rationales for the common prior in the main text.

F Additional Details for Experiment II

F.1 Stimuli

The second experiment consists of two phases: an adaptation phase followed by a test phase. The adaptation phase includes 126 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 (range £2–£32) or a large gain treatment (range £8–£128). The stimuli list for the adaptation trials is presented in Table 4. The test trials are identical across conditions, following the same structure as in Experiment I. To keep the test phase shorter than the adaptation phase, we retained only choices for which $|G - L| \leq 4$ presented in Table 2, yielding 84 test choices.

Small treatment						Large treatment					
ID	x	c	ID	x	c	ID	x	c	ID	x	c
1	**2**	**1**	64	**20**	**11**	1	**8**	**4**	64	**80**	**44**
2	**3**	**1**	65	**20**	**12**	2	**12**	**4**	65	**80**	**48**
3	**3**	**2**	66	**20**	**13**	3	**12**	**8**	66	**80**	**52**
4	**4**	**1**	67	**20**	**14**	4	**16**	**4**	67	**80**	**56**
5	**4**	**2**	68	**20**	**15**	5	**16**	**8**	68	**80**	**60**
6	**4**	**3**	69	**20**	**16**	6	**16**	**12**	69	**80**	**64**
7	**6**	**1**	70	**20**	**17**	7	**24**	**4**	70	**80**	**68**
8	**6**	**2**	71	**20**	**18**	8	**24**	**8**	71	**80**	**72**
9	**6**	**3**	72	**20**	**19**	9	**24**	**12**	72	**80**	**76**
10	**6**	**4**	73	**24**	**1**	10	**24**	**16**	73	**96**	**4**
11	**6**	**5**	74	**24**	**2**	11	**24**	**20**	74	**96**	**8**
12	**8**	**1**	75	**24**	**3**	12	**32**	**4**	75	**96**	**12**
13	**8**	**2**	76	**24**	**4**	13	**32**	**8**	76	**96**	**16**
14	**8**	**3**	77	**24**	**5**	14	**32**	**12**	77	**96**	**20**
15	**8**	**4**	78	**24**	**6**	15	**32**	**16**	78	**96**	**24**
16	**8**	**5**	79	**24**	**7**	16	**32**	**20**	79	**96**	**28**
17	**8**	**6**	80	**24**	**8**	17	**32**	**24**	80	**96**	**32**
18	**8**	**7**	81	**24**	**9**	18	**32**	**28**	81	**96**	**36**
19	**10**	**1**	82	**24**	**10**	19	**40**	**4**	82	**96**	**40**
20	**10**	**2**	83	**24**	**11**	20	**40**	**8**	83	**96**	**44**
21	**10**	**3**	84	**24**	**12**	21	**40**	**12**	84	**96**	**48**
22	**10**	**4**	85	**24**	**13**	22	**40**	**16**	85	**96**	**52**
23	**10**	**5**	86	**24**	**14**	23	**40**	**20**	86	**96**	**56**
24	**10**	**6**	87	**24**	**15**	24	**40**	**24**	87	**96**	**60**
25	**10**	**7**	88	**24**	**16**	25	**40**	**28**	88	**96**	**64**
26	**10**	**8**	89	**24**	**17**	26	**40**	**32**	89	**96**	**68**
27	**10**	**9**	90	**24**	**18**	27	**40**	**36**	90	**96**	**72**
28	**12**	**1**	91	**24**	**19**	28	**48**	**4**	91	**96**	**76**
29	**12**	**2**	92	**24**	**20**	29	**48**	**8**	92	**96**	**80**
30	**12**	**3**	93	**24**	**21**	30	**48**	**12**	93	**96**	**84**

Continued from previous page

Small treatment						Large treatment					
ID	x	c	ID	x	c	ID	x	c	ID	x	c
31	**12**	**4**	94	**24**	**22**	31	**48**	**16**	94	**96**	**88**
32	**12**	**5**	95	**24**	**23**	32	**48**	**20**	95	**96**	**92**
33	**12**	**6**	96	**32**	**1**	33	**48**	**24**	96	**128**	**4**
34	**12**	**7**	97	**32**	**2**	34	**48**	**28**	97	**128**	**8**
35	**12**	**8**	98	**32**	**3**	35	**48**	**32**	98	**128**	**12**
36	**12**	**9**	99	**32**	**4**	36	**48**	**36**	99	**128**	**16**
37	**12**	**10**	100	**32**	**5**	37	**48**	**40**	100	**128**	**20**
38	**12**	**11**	101	**32**	**6**	38	**48**	**44**	101	**128**	**24**
39	**16**	**1**	102	**32**	**7**	39	**64**	**4**	102	**128**	**28**
40	**16**	**2**	103	**32**	**8**	40	**64**	**8**	103	**128**	**32**
41	**16**	**3**	104	**32**	**9**	41	**64**	**12**	104	**128**	**36**
42	**16**	**4**	105	**32**	**10**	42	**64**	**16**	105	**128**	**40**
43	**16**	**5**	106	**32**	**11**	43	**64**	**20**	106	**128**	**44**
44	**16**	**6**	107	**32**	**12**	44	**64**	**24**	107	**128**	**48**
45	**16**	**7**	108	**32**	**13**	45	**64**	**28**	108	**128**	**52**
46	**16**	**8**	109	**32**	**14**	46	**64**	**32**	109	**128**	**56**
47	**16**	**9**	110	**32**	**15**	47	**64**	**36**	110	**128**	**60**
48	**16**	**10**	111	**32**	**16**	48	**64**	**40**	111	**128**	**64**
49	**16**	**11**	112	**32**	**17**	49	**64**	**44**	112	**128**	**68**
50	**16**	**12**	113	**32**	**18**	50	**64**	**48**	113	**128**	**72**
51	**16**	**13**	114	**32**	**19**	51	**64**	**52**	114	**128**	**76**
52	**16**	**14**	115	**32**	**20**	52	**64**	**56**	115	**128**	**80**
53	**16**	**15**	116	**32**	**21**	53	**64**	**60**	116	**128**	**84**
54	**20**	**1**	117	**32**	**22**	54	**80**	**4**	117	**128**	**88**
55	**20**	**2**	118	**32**	**23**	55	**80**	**8**	118	**128**	**92**
56	**20**	**3**	119	**32**	**24**	56	**80**	**12**	119	**128**	**96**
57	**20**	**4**	120	**32**	**25**	57	**80**	**16**	120	**128**	**100**
58	**20**	**5**	121	**32**	**26**	58	**80**	**20**	121	**128**	**104**
59	**20**	**6**	122	**32**	**27**	59	**80**	**24**	122	**128**	**108**
60	**20**	**7**	123	**32**	**28**	60	**80**	**28**	123	**128**	**112**
61	**20**	**8**	124	**32**	**29**	61	**80**	**32**	124	**128**	**116**
62	**20**	**9**	125	**32**	**30**	62	**80**	**36**	125	**128**	**120**
63	**20**	**10**	126	**32**	**31**	63	**80**	**40**	126	**128**	**124**

Table 4: In Experiment II, subjects are randomly assigned to either a small or a large gain treatment during the adaptation phase, to choose between a risky wager $(x, 0.5; 0)$ and a sure gain c. All monetary outcomes are presented in GBP.

F.2 Choice proportions for fair spreads

Stakes (£)	Small(%)	Large(%)	χ^2 test
2	44.9	52.8	0.025
4	40.3	49.1	0.013
6	38.1	46.5	0.017
8	34.9	43.2	0.016
10	35.4	41.8	0.064
12	30.5	38.8	0.015
14	29.5	36.2	0.043
16	27.6	36.7	0.006
18	29.2	34.6	0.104
20	27.8	32.9	0.118
22	26.5	33.6	0.028
24	28.4	32.5	0.210
26	27.6	31.3	0.248
28	27.0	30.8	0.237
30	28.4	29.2	0.797

Table 5: The table shows the acceptance proportions of risky wagers for fair spreads $(x, 0.5; -x)$ under the small- and large-gains adaptation treatments. The last column reports p-values from Pearson's χ^2 tests comparing the difference between the two choice proportions at each stake level.

F.3 Adaptation effects by incentive condition

In the main text, we have shown the aggregate effect of the adaptation treatment across incentive conditions. Figure 11 instead shows the effect of adaptation to large versus small gains separately by incentive condition—for hypothetical incentives (panel A), and for real incentives (panel B).

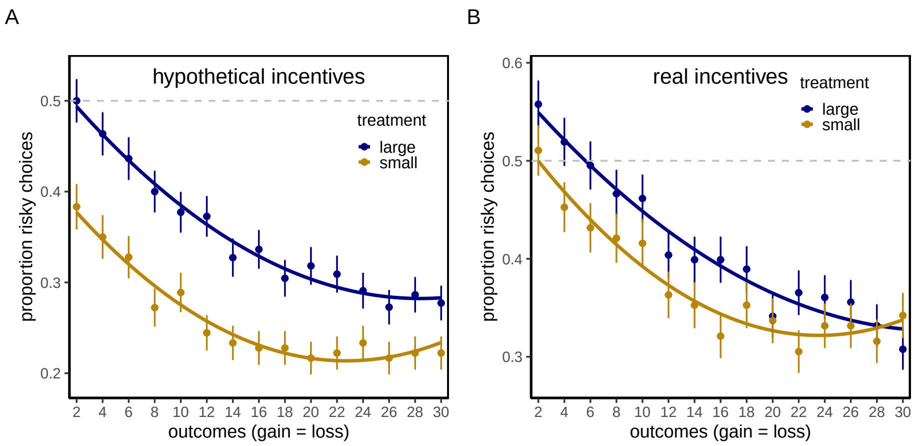

Figure 11: Risk-taking after gain adaptation by incetive condition.

G Additional details for experiment III

Just like in experiment II, in the main text we have shown the aggregate treatment effect of the visual manipulation while aggregating across incentive conditions. Figure 12 instead shows the effect of the visual display separately by incentive condition—for hypothetical incentives (panel A), and for real incentives (panel B).

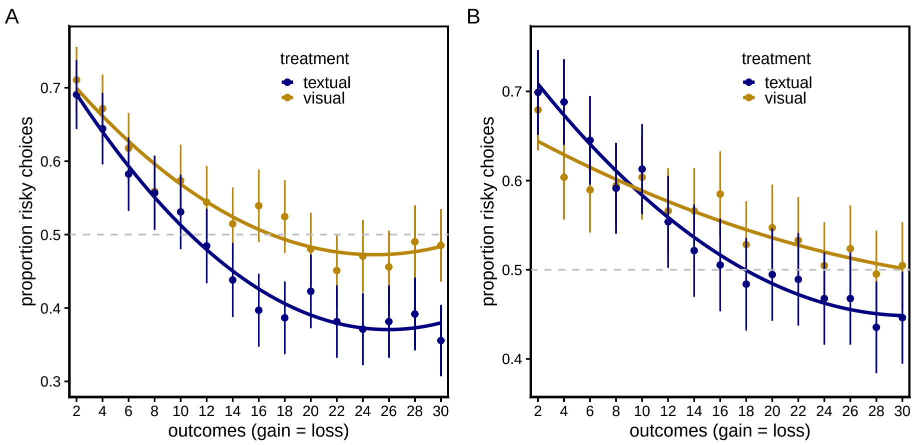

Figure 12: Risk-taking after gain adaptation by incetive condition.

H Probit Regression

H.1 Probit regression model

The probability of selecting the risky option is modeled using a reduced-form Probit regression, as follows:

$$Choice_i \sim Bernoulli(\text{Prob}_i)$$

where

$$\text{Prob}_i = \Phi\left\{\gamma_0^j + \gamma_g^j \ln(G_i) + \gamma_l \ln(L_i)\right\}$$

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 intercepts γ_0^j and gain slopes γ_g^j for the j-th individual is modeled as:

$$\gamma_0^j \sim \mathcal{N}(\bar{\gamma}_0, \sigma_0), \qquad \gamma_g^j \sim \mathcal{N}(\bar{\gamma}_g, \sigma_g)$$

where $\bar{\gamma}_0$ and $\bar{\gamma}_g$ denotes the aggregate-level intercept and gain slope, σ_0 and $\sigma_g$captures heterogeneity across subjects. The priors for the aggregate-level parameters are specified as:

$$\bar{\gamma}_0 \sim \mathcal{N}(0,2), \qquad \bar{\gamma}_g \sim \mathcal{N}(1,1), \qquad \gamma_l \sim \mathcal{N}(-1,1)$$

$$\sigma_0 \sim \text{Exponential}(1), \qquad \sigma_g \sim \text{Exponential}(1)$$

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 $\bar{\gamma}_g$ and γ_l, are treated as fixed effects.

H.2 Estimation results

The estimation is performed using Stan for each experiment. We employ 4 chains, each with 2,000 iterations, of which the first 1,000 are warm-up iterations. Thus, a total of 4,000 iterations are used to draw the posterior estimates for each parameter. In this model, the aggregate-level intercept $\bar{\gamma}_0$, as well as the slope coefficients $\bar{\gamma}_g$ and γ_l, are assumed to be treatment-specific, i.e., fixed within each treatment group but vary across treatment groups, while σ_0 and σ_g are assumed to remain constant across treatments. The parameters across different treatment conditions are estimated simultaneously, details of the programming are provided in Appendix H.3. Convergence is assessed by checking R-hat, ESS and divergent iterations. The estimation results are summarized in Table 6 for Experiment II and Table 7 for Experiment III.

	$\bar{\gamma}_0$	$\bar{\gamma}_g$	γ_l	$\lvert\gamma_l/\bar{\gamma}_0\rvert$
Small	-0.373	2.583	-3.074	1.191
	[-0.825, 0.075]	[2.380, 2.780]	[-3.208, -2.941]	[1.120, 1.270]
Large	0.611	2.457	-3.151	1.284
	[0.186, 1.026]	[2.261, 2.645]	[-3.287, -3.013]	[1.207, 1.369]
Δ	0.984			0.092
	[0.375, 1.598]			[-0.016, 0.202]
Hypo-Small	-0.506	2.283	-2.808	1.233
	[-1.101, 0.081]	[2.008, 2.566]	[-2.998, -2.619]	[1.122, 1.365]
Hypo-Large	0.399	2.423	-3.155	1.305
	[-0.178, 0.974]	[2.163, 2.684]	[-3.345, -2.968]	[1.203, 1.430]
Δ	0.905			0.072
	[0.097, 1.714]			[-0.096, 0.241]
Real-Small	-0.211	2.822	-3.297	1.170
	[-0.847, 0.434]	[2.534, 3.100]	[-3.483, -3.105]	[1.082, 1.268]
Real-Large	0.844	2.449	-3.120	1.277
	[0.211, 1.490]	[2.166, 2.734]	[-3.309, -2.924]	[1.168, 1.405]
Δ	1.055			0.107
	[0.147, 1.975]			[-0.042, 0.259]
Hypo	-0.062	2.406	-3.033	1.262
	[-0.510, 0.381]	[2.211, 2.606]	[-3.170, -2.892]	[1.182, 1.350]
Real	0.350	2.632	-3.197	1.216
	[-0.092, 0.791]	[2.436, 2.835]	[-3.329, -3.060]	[1.148, 1.291]
Δ	0.412			-0.046
	[-0.209, 1.009]			[-0.159, 0.062]

Table 6: Aggregate-level estimates with 95% confidence intervals shown in brackets. $\bar{\gamma}_0$ represents the aggregate-level intercept, while $\bar{\gamma}_g$ and γ_l denote the slope coefficients for the log-gain and log-loss, $|\gamma_l/\bar{\gamma}_0|$ captures loss-sensitivity. Experiment II incorporates an adaptation phase featuring small or large gain treatment, under either hypothetical or real condition.

	$\bar{\gamma}_0$	$\bar{\gamma}_g$	γ_l	$\lvert\gamma_l/\bar{\gamma}_0\rvert$
Textual	1.364	3.186	-3.859	1.212
	[0.989, 1.745]	[3.024, 3.346]	[-3.955, -3.760]	[1.159, 1.268]
Visual	0.822	3.810	-4.142	1.087
	[0.467, 1.184]	[3.651, 3.969]	[-4.237, -4.046]	[1.051, 1.125]
Δ	-0.543			-0.124
	[-1.059, -0.022]			[-0.190, -0.059]
Hypo-Textual	1.519	2.958	-3.745	1.268
	[0.973, 2.077]	[2.733, 3.193]	[-3.882, -3.604]	[1.186, 1.358]
Hypo-Visual	0.822	3.566	-4.032	1.131
	[0.276, 1.371]	[3.335, 3.792]	[-4.170, -3.899]	[1.072, 1.197]
Δ	-0.697			-0.136
	[-1.458, 0.071]			[-0.248, -0.030]
Real-Textual	1.191	3.396	-3.961	1.167
	[0.677, 1.718]	[3.176, 3.621]	[-4.103, -3.813]	[1.104, 1.233]
Real-Visual	0.798	3.996	-4.216	1.055
	[0.329, 1.255]	[3.793, 4.207]	[-4.345, -4.086]	[1.011, 1.102]
Δ	-0.393			-0.112
	[-1.086, 0.318]			[-0.190, -0.035]
Hypo	1.172	3.280	-3.889	1.186
	[0.803, 1.561]	[3.118, 3.436]	[-3.987, -3.794]	[1.139, 1.239]
Real	0.988	3.759	-4.132	1.100
	[0.617, 1.364]	[3.601, 3.924]	[-4.232, -4.035]	[1.061, 1.140]
Δ	-0.184			-0.087
	[-0.708, 0.344]			[-0.150, -0.024]

Table 7: Aggregate-level estimates with 95% confidence intervals shown in brackets. $\bar{\gamma}_0$ represents the aggregate-level intercept, while $\bar{\gamma}_g$ and γ_l denote the slope coefficients for the log-gain and log-loss, $\lvert\gamma_l/\bar{\gamma}_0\rvert$ captures loss-sensitivity. Experiment III manipulates whether gains and losses were described numerically in text or a visual aid was additionally provided, under either hypothetical or real condition.

H.3 Stan Code

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

```
functions {
  real partial_sum(array[] int choice_slice,
                   int start, int end,
                   vector alpha, vector gamma, vector lambda,
                   vector lg, vector ll,
                   array[] int id, array[] int treatn) {
    return bernoulli_lpmf(choice_slice |
      Phi_approx(alpha[id[start:end]] +
                 gamma[id[start:end]] .* lg[start:end] +
                 lambda[treatn[start:end]] .* ll[start:end]));
  }
}

```

```
data {
  int<lower=0> N;
  int<lower=0> Nid;
  array[N] int<lower=1, upper=Nid> id;
  array[N] int<lower=0, upper=1> 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;
  vector[max(treat)] alpha_hat;
  real<lower=0> alpha_sd;

  vector[Nid] gamma;
  vector[max(treat)] gamma_hat;
  real<lower=0> gamma_sd;

  vector[max(treat)] lambda;
}

model {
  alpha_hat ~ normal(0, 2);
  alpha_sd ~ exponential(1);
  gamma_hat ~ normal(1, 1);
  gamma_sd ~ exponential(1);
  lambda ~ normal(-1, 1);

  for (i in 1:Nid) {
    alpha[i] ~ normal(alpha_hat[treat[i]], alpha_sd);
    gamma[i] ~ normal(gamma_hat[treat[i]], gamma_sd);
  }

  target += reduce_sum(partial_sum, choice_risky, 128,
                       alpha, gamma, lambda,
                       lg, ll, id, treatn);
}

generated quantities {
  vector[Nid] lgr;
  vector[max(treat)] lgr_hat;
  real lgr_diff;
  real alpha_diff;

  for (i in 1:Nid)
    lgr[i]   = abs(lambda[treat[i]] / gamma[i]);

  for (t in 1:2)
    lgr_hat[t] = abs(lambda[t] / gamma_hat[t]);

  lgr_diff = lgr_hat[2] - lgr_hat[1];
  alpha_diff = alpha_hat[2] - alpha_hat[1];
}
```

I Heterogeneity representative sample

We estimate the Noisy Cognition Model structurally using the same probit model used throughout the paper and described above. Vieider (2024) provides a tutorial on how to code such model, and how to launch them from R using CmdStanR.

Individual heterogeneity. Figure 13 shows the empirical cumulative distribution functions of loss-sensitivity γ_ℓ/γ_g (panel A) and of small-stake risk seeking γ_0. Loss-sensitivity is slightly more pronounced in the hypothetical treatment—likely due to ceiling effects in risk-taking at small stakes in the incentivized treatment—but is the clear majority pattern regardless of incentivization. A clear majority of subjects has a positive intercept, capturing small-stake risk-seeking. Beyond this, heterogeneity across subjects for both measures is substantial.

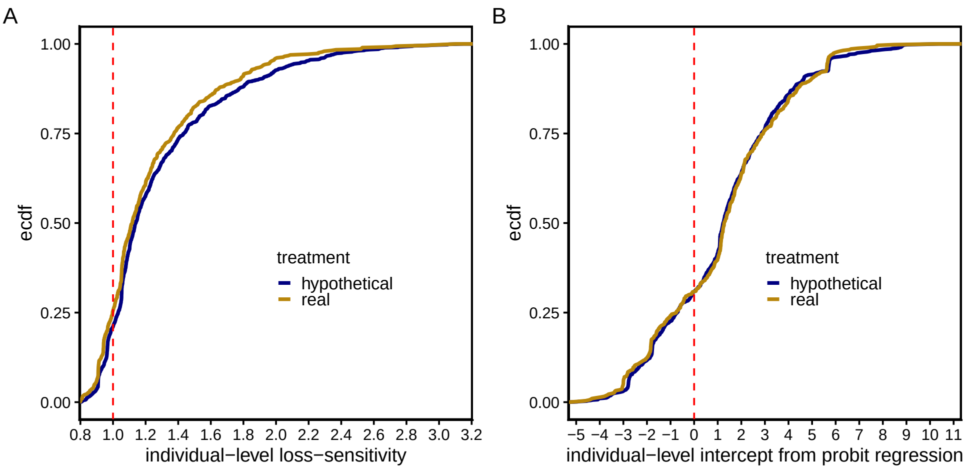

Figure 13: Population heterogeneity of loss-sensitivity and risk-taking

Empirical cumulative distribution functions of loss-sensitivity γ_ℓ/γ_g (panel A) and of small-stake risk seeking γ_0.

Regression analysis. 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 et al., 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.

References

Alaoui, Larbi, and Antonio Penta (2026) 'Attitudes towards success and failure.' *American Economic Journal: Microeconomics* 18, 105–145

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)

Heng, Joseph A, Michael Woodford, and Rafael Polania (2023) 'Efficient numerosity estimation under limited time.' *BioRxiv* pp. 2023–07

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

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

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

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

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