Ambiguity by Degree*

Ferdinand M. Vieider[1]

[1]*RISLαβ, Department of Economics, Ghent University ,*
ferdinand.vieider@ugent.be

November 4, 2026

Abstract

I here propose to define ambiguity as the imprecision of mental probability representations, rather than as a categorical feature of the outcome-generating process. An objectively described risk may thereby be more ambiguous than an Ellsberg urn with unknown composition. The model gives this notion empirical content through a *revealed-ambiguity* signature that identifies the relative ambiguity of two probability representations from observed choice patterns. Across two pre-registered experiments with $N \approx 1800$ subjects, I show that objectively described simple risks can be revealed as more ambiguous than Ellsberg-urns; and that conventional Ellsberg-ambiguity falls between two descriptions of simple risks. Ambiguity is thus not special—and neither is risk. This suggests that a single parsimonious model of representational ambiguity can organize behaviour under uncertainty.

J'ay dit plus d'une fois qu'il faudroit une nouvelle espece de Logique qui traiteroit des degrés de probabilité [...] une balance necessaire pour peser les apparances.[1]

Gottfried Wilhelm Leibniz (1704/1765), Book IV, Ch. XVI, §9

1 Introduction

In this paper I propose to define ambiguity as a property of internal probability representations, rather than as a categorical feature of the outcome-generating process as in Ellsberg (1961). Building on a rapidly growing literature attributing choice behaviour to cognitive frictions in information processing, I start from the premise that probabilities need to be internally represented by the brain before they can inform decisions. The precision of these representations is inherently limited, and can further vary with seemingly superficial features of presentation mode or description, not only with whether the

*I gratefully acknowledge funding by the Research Foundation–Flanders under grant nr. G064826N: "Noisy Cognition as a Unified Driver of Economic Decisions". This paper supersedes an earlier note titled "Statistical Foundations of Ambiguity Attitudes".

[1]"I have said more than once that we would need a new kind of logic dealing with degrees of probability [...] a balance necessary for weighing appearances." A shorter version of this epigraph appears at the beginning of John Maynard Keynes' (1921) *Treatise on Probability.* Keynes further writes: "If two probabilities are equal in degree, ought we, in choosing our course of action, to prefer that one which is based on a greater body of knowledge?" (Ch. XXVI, §7, cited by Chew, Miao and Zhong, 2023).

underlying probability is objectively described or not. This dissociates representational precision from the conventional distinction between risk and ambiguity. Both instead lie on a continuum of imprecision—or *representational ambiguity*—in internal probability representations. A fully described simple risk may therefore be either less or more ambiguous than an Ellsberg-type unknown outcome-generating process. Ambiguity is not special—and neither is risk: ambiguity differs *by degree*.

The definition proposed here goes to the heart of the historical debate on ambiguity. The Ellsberg paradox—whereby a draw to win a fixed prize from a known-composition urn is preferred to a draw from an unknown-composition urn regardless of the winning colour—challenged the idea that choices under uncertainty could be organized by a single additive probability measure, thus violating subjective expected utility theory (*SEU*; Savage, 1954). This has motivated a proliferation of models explaining the SEU anomaly through preferences over whether the probabilities governing outcomes are precisely known or not. The present account suggests a more parsimonious alternative. The canonical Ellsberg patterns thus emerge from the same cognitive frictions that generate familiar departures from expected utility under fully described risk, without requiring separate primitive preferences over knowledge of the outcome-generating process. This creates scope for considerable theoretical simplification: a single model can organize seemingly distinct departures from SEU, recover the classical benchmark as representational noise vanishes, and predict when familiar behavioural patterns emerge, disappear, or reverse.

Risk and ambiguity are not categorically distinct. Two preregistered experiments illustrate the key conceptual rupture. Experiment I holds fixed an Ellsberg-type urn whose composition is hidden, while varying only the presentation mode of a fully described comparison probability. The latter is presented either visually, using ten consecutively numbered balls, or numerically as fractions describing the number of winning balls in larger urns—either with denominators between 60 and 90, or between 230 and 310 (see Section 3.1 for screenshots and details). Panel A of Figure 1 illustrates the key finding. The baseline comparison between the Ellsberg urn and the visual urn produces two familiar patterns of categorical ambiguity[2]: ambiguity aversion at 50–50, and increasing

[2]I use the terms *Ellsberg-type ambiguity* and *categorical ambiguity* throughout to distinguish the concept of an unknown outcome-generating process from the continuous concept of ambiguity used here, unless the distinction is clear from context. These terms can also include vagueness in naturally occurring

choice of the described risk as the matched probability of winning rises.[3] Describing the same risks through fractions, however, *reverses both patterns*. The unknown-composition urn is now chosen over the fully described risk, and choice proportions of the unknown urn increase with the probability of winning. Merely changing the presentation mode of simple risk can therefore reverse the conventional Ellsberg pattern.

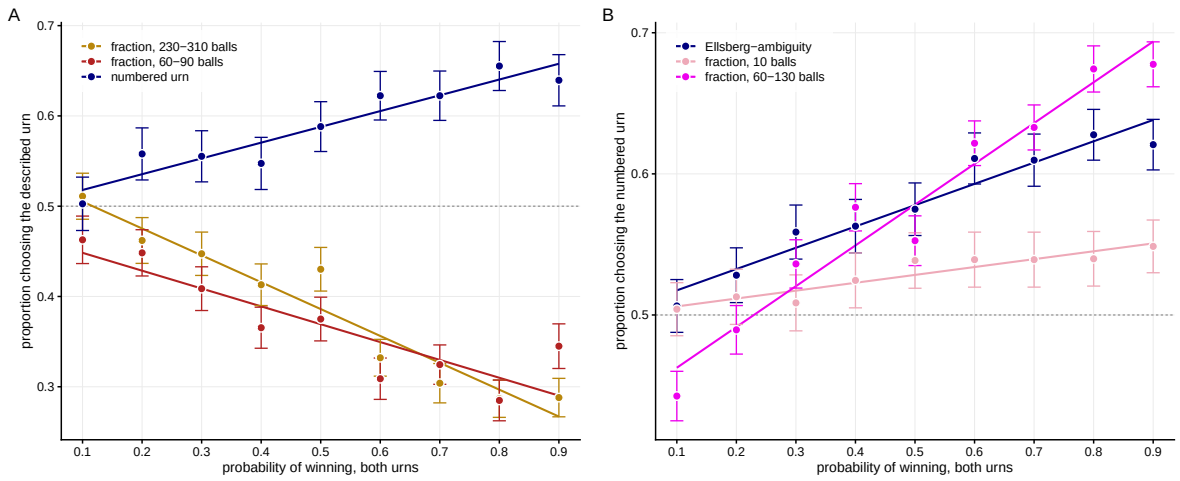


Figure 1: Choice proportion of risky option for $p_r = p_a$

Panel A shows choice proportions of the risky option by probability and treatment arm, jointly with 95% confidence intervals. The lines are linear fits. Panel B shows choice proportions of the visually displayed risk over either the Ellsberg-tye urn, or risks expressed using fractions. All choice concerns matched probabilities, in the sense of Raiffa (1961), i.e. have the same number of winning 'colours' in the options.

Experiment II inverts the comparison. I now hold fixed a visually displayed urn with ten consecutively numbered balls and compare it to three alternatives: the Ellsberg-type urn from Experiment I; an objectively equivalent risk expressed by simple fractions such as 1/10, 2/10, ...; and equivalent risks expressed by fractions describing larger urns. Panel B of Figure 1 illustrates the results. Remarkably, the characteristic patterns associated with ambiguity emerge already when comparing two objectively equivalent, fully described urns, each containing ten balls: aversion to the ambiguity arising from the fraction format at 50–50, and increasing aversion to that ambiguity as the matched probability of winning rises. The latter pattern becomes stronger still for fractions describing larger urns, surpassing even the ambiguity-insensitivity observed in the Ellsberg benchmark.[4] In terms of observed behaviour, conventional Ellsberg ambiguity thus falls

probabilities, which again are not special according to the present account.

[3]The probabilities underlying these choices are always matched in the sense of Raiffa (1961): the risky urn and the Ellsberg urn have identical numbers of winning "colours".

[4]The use of fractions is motivated by research on numerical cognition documenting greater difficulty in processing fractions (DeWolf, Grounds, Bassok and Holyoak, 2014). In the present design, the visual display facilitates probability assessment through the grouping of consecutively numbered balls (Ciccione

between two descriptions of simple risk. Ambiguity differs by degree, not by category: *unknown outcome-generating processes are not behaviourally special.*

Ambiguity as imprecision in probability representations. The experiments previewed above are motivated by a model of cognitive frictions in probability processing. The key premise is that probabilities are internally coded by the brain with finite precision. This builds on a foundational insight from neuroscience: information about the outside world is not directly available to the brain, but must first be encoded in internal representations before it can inform decisions. The fact that a probability is objectively described therefore does not imply that it is internally represented with perfect precision. Rather, representational precision may depend both on the outcome-generating process and on seemingly irrelevant features of how probabilities are presented—for instance, whether an objectively identical probability is displayed visually or described by a fraction.[5] I refer to these influences jointly as the *presentation mode*.[6]

Imprecise probability representations are disciplined through Bayesian inference: noisy internal representations are combined with prior beliefs to infer the underlying probabilities. The decision maker then chooses the lottery with the higher posterior probability of winning. Greater representational ambiguity implies stronger regression toward the prior mean. Consequently, when two lotteries with matched winning probabilities but different levels of representational ambiguity are compared, on average the more ambiguous representation yields a higher inferred probability below the prior mean, and a lower inferred probability above it. Under the disciplining assumption of a common prior, relative choice proportions rotate around the prior mean: the less ambiguous representation becomes increasingly attractive as the matched probability rises. This yields a *revealed-ambiguity* characterization: the representation whose choice share rises with the matched winning probability is revealed to be the less ambiguous one. Under a pessimistic prior with a mean below 0.5—here interpreted as a learned cautious default in the absence of reliable

and Dehaene, 2020), whereas fractions require integrating information from numerator and denominator. Larger fractions with varying denominators may further increase processing demands by preventing reliance on a fixed numerical reference.

[5] Average representational ambiguity, as well as heterogeneity in sensitivity to sources and presentation modes, may also vary across decision makers. I leave such heterogeneity largely implicit in the theory development and focus instead on variation across probability representations.

[6] Here and throughout, *representation* refers to internal information processing by the brain, and thus to a latent feature of the decision process; *presentation* refers to external features of how an outcome-generating process or probability is presented or described.

information—the less ambiguous representation is also preferred when objective winning probabilities are matched at 50–50, delivering canonical ambiguity aversion.

Relation to standard models. The model and experimental evidence challenge key assumptions in standard models of ambiguity attitudes. Despite the prima facie resemblance between representational ambiguity and sets of priors, the similarity is primarily formal. In canonical multiple-prior models, multiplicity of beliefs reflects uncertainty about the probabilities governing the outcome-generating process (Gilboa and Schmeidler, 1989). Here, dispersion instead reflects imprecision in the internal representation of a fully described probability. This distinction is directly relevant for the experiments: the Ellsberg process is held fixed while only the presentation mode of a simple, fully described risk changes. Canonical multiple-prior models reduce to expected utility for such risks and therefore cannot account for the resulting reversals.[7]

The present model connects more closely to approaches that relax the categorical distinction between risk and Ellsberg-type ambiguity. Segal (1987) models ambiguity as a two-stage lottery and explains ambiguity attitudes through failures to reduce compound lotteries. Klibanoff, Marinacci and Mukerji (2005) model ambiguity through a distribution over possible probability measures, with ambiguity attitudes arising from a transformation of the expected utilities associated with those measures. These approaches have the advantage of accommodating the strong empirical association between behaviour under categorical ambiguity and compound risk (e.g. Halevy, 2007; Dean and Ortoleva, 2019; Chew et al., 2023). Yet both retain simple risk as a presentation-invariant benchmark. They therefore cannot account for the experimental reversals documented here, which are generated solely by changing how a fully described simple risk is presented. Accommodating these results would require allowing presentation-dependent uncertainty even about objectively described probabilities, raising the question of what determines this uncertainty and its systematic effects on choice.

Source-dependent preferences provide a more flexible approach, permitting choice to vary across different sources of uncertainty even when point beliefs are held fixed (Tversky and

[7]Another canonical approach relaxes additivity of beliefs by representing them through non-additive capacities (Schmeidler, 1989). Here, I maintain additivity of mean beliefs and locate ambiguity instead in the precision of their internal representation. For fully described risks, the relevant benchmark in Schmeidler (1989) remains expected utility, so the same conclusion applies.

Fox, 1995; Chew and Sagi, 2008). Their flexibility, however, raises a fundamental question: what distinguishes one source of uncertainty from another? Recent work highlights this conceptual difficulty (Baillon, Bleichrodt, Li and Wakker, 2025).[8] The issue is particularly stark in Experiment II. Here, choices between draws from one and the same urn containing ten balls generate canonical ambiguity signatures solely because the winning probability is displayed visually for one option and expressed as a fraction for the other. Accounting for this finding through source-dependent preferences would therefore require treating two formats for describing the winning probability in one and the same urn as distinct sources of uncertainty. At that point, however, the notion of a source ceases to provide an independent explanation of behaviour: it becomes an ex-post label for precisely the presentation effects I endeavour to explain here. The present model instead provides a cognitive microfoundation for these effects through differences in the precision of internal probability representations, generating systematic predictions about their direction and strength that source-dependent preferences alone cannot supply.

Implications for ambiguity preferences. Identifying ambiguity with a continuum of representational imprecision rather than as a categorical property of the outcome-generating process has the potential to reconcile several seemingly disparate findings in the literature. For instance, the model naturally accommodates the close association between Ellsberg-type ambiguity and compound risk (Halevy, 2007; Dean and Ortoleva, 2019; Chew et al., 2023), while allowing different forms of ambiguity and compound lotteries to leave distinct behavioural footprints (Chew, Miao and Zhong, 2017). It also suggests why training in compound-lottery reduction may separate responses to compound risk from ambiguity (Wu, Fehr, Hofland and Schonger, 2024): such interventions may remove one source of representational ambiguity without eliminating the other. Closely related evidence is provided by de Clippel, Moscariello, Ortoleva and Rozen (2025), who find that the association between compound-lottery aversion and conventional ambiguity aversion is particularly pronounced among subjects experiencing greater cognitive uncertainty about compound lotteries. This finding fits naturally into the present account. Compound lotteries need not be represented with greater ambiguity merely because they

[8] Baillon et al. (2025) treat all risk as a single source in their formal analysis, but subsequently qualify this in the discussion as a tractability assumption that holds only approximately. I discuss the implications of this qualification in Section 2.4.

are compound. A relationship with Ellsberg-type ambiguity aversion should emerge primarily when the probabilities underlying compound lotteries are themselves represented imprecisely, activating greater reliance on cautious prior beliefs in both settings.

The added value of procedural models. The approach developed here belongs to a growing literature that derives systematic choice patterns from cognitive frictions in information processing (see Enke, 2024; Vieider, 2025a; Oprea and Woodford, 2026, for recent reviews). Such models provide cognitive microfoundations for descriptive behavioural approaches, but also go meaningfully beyond them: By modelling the processes underlying decisions, they can generate hypotheses that do not emerge from descriptive models (and that arguably could be interpreted as contradicting them). In the present paper, the notion of imprecise probability representations motivates the hypothesis that Ellsberg-type ambiguity may not constitute a distinct category of uncertainty in the first place. Such attempts to ground behavioural regularities in underlying mechanisms are not entirely foreign to behavioural economics: for instance, Kahneman and Tversky (1979) drew heavily on psychophysical insights when developing prospect theory, before its functionals came to be widely re-interpreted as capturing stable tastes towards risk.

Procedural models also promise considerable theoretical simplification while simultaneously increasing explanatory power. Frydman and Jin (2022; 2025) show how efficient adaptation of cognitive precision can generate systematic contextual changes in risk-taking and probability-dependence that challenge conventional descriptive models. Oprea and Vieider (2024) use a Bayesian model that is closely related to the one developed here to explain the description–experience gap (Hertwig, Barron, Weber and Erev, 2004), a longstanding paradox under prospect theory, and leverage its predictions to eliminate that gap experimentally. Oprea and Vieider (2026) extend the same framework used here to risky and intertemporal choice, showing how manipulations of information presentation can attenuate, eliminate, or even reverse canonical patterns from the descriptive literature, including probability- and delay-insensitivity, small-stakes risk aversion and impatience, present bias, and subadditivity. The same underlying information-processing mechanisms can thus account for regularities traditionally described through distinct preference functionals, and resolve paradoxes arising from them, while predicting how these regularities change with the decision environment.

These findings suggest a version of the Lucas critique for descriptive models of behaviour. When the regularities described by such models truly arise from information-processing frictions, their parameters need not remain stable across environments, even when underlying preferences remain unchanged. The problem is reminiscent of the distinction between *as-if* and *substantive* rationality. In an experiment with induced preferences, de Clippel, Oprea and Rozen (2026) find that most subjects behave consistently with maximizing some utility function, while only a minority maximize the utility function assigned to them. Behaviour that can be rationalized by a preference functional need not therefore reveal true preferences. The present paper extends this argument to ambiguity: the same cognitive mechanism that generates canonical patterns under simple risk can also produce the behavioural signatures traditionally attributed to primitive ambiguity preferences. Explicitly modelling information processing can thus reveal common foundations underlying seemingly distinct behavioural phenomena, reducing the need for separate preference functionals while at the same time improving predictions across contexts. More fundamentally, isolating cognitive distortions in choices may help redirect attention toward the welfare-relevant preferences over consequences that motivated the early economic literature (Samuelson, 1937; Vickrey, 1945; Alaoui and Penta, 2025).

2 Modelling Framework

2.1 Ambiguity as imprecision in probability representations

Consider a decision maker (DM) assessing an event with objective probability $p \in (0,1)$.[9] The central premise is that an objective probability need not be represented internally with perfect precision. Instead, the DM has access only to a noisy internal representation of p, which is then used to inform inference and choice. Representational noise may arise from perceptual and information-processing constraints in approximate number assessment (Dehaene, 2003; Khaw, Li and Woodford, 2021; 2026); more generally, it may reflect the difficulty of extracting the decision-relevant implications of a described probability. Different description formats can thereby facilitate or hinder internal representation and Bayesian inference (Gigerenzer and Hoffrage, 1995; Cosmides and Tooby, 1996).

I here model the internal representation as a binomial success count. For presenta-

[9] For Ellsberg-type ambiguity, objectivity is established through exchangeability in the spirit of Raiffa (1961). I refer to such equal objective probabilities under exchangeability as *matched probabilities*.

tion mode i, let the signal (or success count) be $s_i \sim \text{Bin}(\kappa_i, p)$, and define the represented probability as the relative success count, $\tilde{p}_i = s_i/\kappa_i$. The representation is thus unbiased, $\mathbb{E}[\tilde{p}_i \mid p] = p$. This constitutes a disciplining assumption: objectively matched probabilities are represented identically on average, and the expected representations of complementary probabilities sum to one. Differences across presentation modes are thereby restricted to the precision of the representation, rather than arbitrary systematic biases in probability encoding. The representation is, however, noisy, with $\mathbb{V}[\tilde{p}_i \mid p, \kappa_i] = p(1-p)\kappa_i^{-1}$. The parameter κ_i therefore indexes *Bernoulli-equivalent precision*: it is the number of independent Bernoulli observations that would generate the same sampling variation.[10] Larger κ_i implies a more faithful representation of p, and as $\kappa_i \to \infty$ the distribution of $\tilde{p}_i$ collapses to a Dirac–*delta* at the true probability. This representation is not meant to imply that the brain literally draws κ_i independent binary observations; κ_i instead quantifies the information content of the internal representation.

Definition: I define the ambiguity of probability representation i as $A_i \triangleq \kappa_i^{-1}$.[11] Ambiguity is therefore a latent property of the internal representation of a probability. Representation j is *more ambiguous* than representation i iff $A_j > A_i$.

The index i may refer to different outcome-generating processes, but also to different presentation modes of the same fully described probability. A fully described risk can therefore have positive representational ambiguity, and two objectively identical risks can have different values of A_i when their probabilities are described differently. Ambiguity is thus continuous rather than categorical: Ellsberg-type uncertainty, natural uncertainty, and fully described risks may all differ in A_i, and described risks need not sit at an extreme of this continuum. Conventional ambiguity is one source of imprecision in probability representations, but only one among many. In this sense, ambiguity is not special.

[10] I here reserve the term *precision* for κ_i, rather than for inverse variance. The distinction matters because the variance of a Bernoulli representation also depends mechanically on p through $p(1-p)$. Thus two representations with the same κ_i have the same effective precision by the present definition, even though their raw variances differ across probabilities.

[11] One can equivalently define ambiguity as $A_i = \mathbb{V}[\tilde{p}_i \mid p, \kappa_i]\,[p(1-p)]^{-1}$. This highlights that representational ambiguity measures the imprecision of the internal probability representation after scaling out the sampling variation associated with the probability level itself.

Bayesian inference and posterior probabilities. Let us further assume that the DM is *Bayesian*. Bayesian is here used in the sense of statistical decision theory (Berger, 2013), rather than in the Savagian sense of updating probability point estimates. Assume the DM has a prior distribution

$$p \sim \text{Beta}(\kappa_0 p_0 \,,\, \kappa_0(1-p_0))\,, \tag{1}$$

with p_0 capturing the prior mean and κ_0 the prior precision. The prior need not reflect the statistical distribution of probabilities in the experimental environment, as sometimes assumed in the literature. Instead, it may be interpreted as a cautious default assessment brought to unfamiliar decision environments, potentially shaped by prior experience. A prior mean $p_0 < 0.5$ can then generate familiar patterns of risk aversion, as shown in related accounts of risky choice (Oprea and Vieider, 2024; 2026; Vieider, 2026).

To the extent that probability representations are ambiguous, i.e. $A_i > 0$, it is optimal to combine the information conveyed by the representation with prior information.[12,13] For a particular realization of the internal representation, the posterior expectation is

$$\widehat{p}_i = \omega_i \tilde{p}_i + (1-\omega_i)p_0, \tag{2}$$

where $\omega_i \triangleq \kappa_i/(\kappa_i+\kappa_0) = 1/(1+A_i\kappa_0)$ is the Bayesian evidence weight. The weight increases with representational precision κ_i, and decreases with ambiguity $A_i \triangleq \kappa_i^{-1}$.

The immediate consequence is regression toward the prior mean: $\widehat{p}_i = p_0 + \omega_i(\tilde{p}_i - p_0)$. Since $0 < \omega_i < 1$ for $A_i > 0$, represented probabilities below p_0 are over-inferred, whereas those above p_0 are under-inferred. Greater representational ambiguity A_i lowers ω_i and thereby strengthens both effects. Ambiguity therefore generates overestimation of small probabilities and underestimation of large probabilities whenever the prior lies in the interior of the probability range.

[12]This is optimal in the precise sense that it minimizes mean squared prediction error over repeated representations; see Bishop (2006), chapter 4.

[13]Bayesian updating is assumed to operate at a subconscious level. This is consistent with the idea that Bayesian inference is a fundamental principle of brain functioning (Knill and Pouget, 2004; Ma, Kording and Goldreich, 2023; Oprea and Woodford, 2026). See also Hohwy (2013) for a popular science treatment.

Stochastic choice. The posterior assessment $\widehat{p}_i$ is latent because it depends on the noisy representation $\tilde{p}_i$. To give the model empirical content, I characterize its distribution over repeated internal representations of the same probability. Its first two moments—the moments of the *response distribution*; cf. Ma et al., 2023, chapters 3 & 4—are

$$\mathbb{E}[\widehat{p}_i \mid p, A_i] = \omega_i p + (1-\omega_i)p_0 = p_0 + \omega_i(p - p_0), \tag{3}$$

$$\mathbb{V}[\widehat{p}_i \mid p, A_i] = \omega_i^2 \, \mathbb{V}[\tilde{p}_i \mid p, \kappa_i] = p(1-p)\frac{\kappa_i}{(\kappa_i + \kappa_0)^2}. \tag{4}$$

This allows to characterize choices between options with different degrees of ambiguity, A_i and A_j. Consider two processes with the same objective winning probability p, and assume that their internal signals are independently generated. For sufficiently large precisions κ_i and κ_j, the Binomial errors are approximately normal, so the difference between posterior assessments $\widehat{p}_i$ and $\widehat{p}_j$ is approximately normal as well. Suppose further that both processes yield the same fixed outcome x upon success and zero otherwise. The common prize then drops out of the comparison[14], yielding:

$$P[i \succ j] = \Phi\left[\frac{(\omega_i - \omega_j)(p - p_0)}{\sqrt{p(1-p)\left[\frac{\kappa_i}{(\kappa_i+\kappa_0)^2} + \frac{\kappa_j}{(\kappa_j+\kappa_0)^2}\right]}}\right], \tag{5}$$

where $P[i \succ j]$ denotes the probability that process i is chosen over process j. The resulting Probit index measures discriminability: choice becomes more systematic as the expected posterior assessments move farther apart relative to the dispersion of their response distributions. Stochastic choice therefore depends on the cardinal difference between posterior assessments, not merely their ordering.

Normality is not essential to the stochastic-choice mechanism. Let F_{ij} denote the distribution of the latent response difference $d_{ij} = \widehat{p}_i - \widehat{p}_j$ implied by the underlying Beta–Binomial updating process. Choice probabilities are then determined by the probability mass above the comparison threshold, with ties broken randomly. For sufficiently precise internal representations, however, the Binomial signals become approximately Gaussian, and so does their difference. This approximation is generally reasonably accurate already

[14] This assumes that noise in outcome encoding is independent of noise in probability representation. I maintain this simplifying assumption throughout because it yields a particularly transparent characterization. It is, however, not required for the qualitative results presented here to obtain.

at moderate levels of representational precision, particularly away from the probability boundaries. I henceforth maintain normality as a convenient approximation that yields tractable choice probabilities and sharp revealed-ambiguity characterizations.

Illustration. Figure 2 illustrates the mapping between latent ambiguity and observable choice proportions implied by equation (5). Two features emerge clearly. First, the prior mean p_0 determines the probability at which the two processes are chosen equally often: whenever $A_i \neq A_j$, $P[i \succ j] = 0.5 \;\leftrightarrow\; p = p_0$. A larger p_0 therefore shifts the 50–50 crossing to the right. Second, holding the prior fixed, relative ambiguity determines both the direction and the strength of the choice pattern. When $A_i < A_j$, the choice proportion of process i over j increases monotonically with p. Increasing the ambiguity contrast $A_j - A_i$ makes the pattern more pronounced, generating a steeper relationship with p.

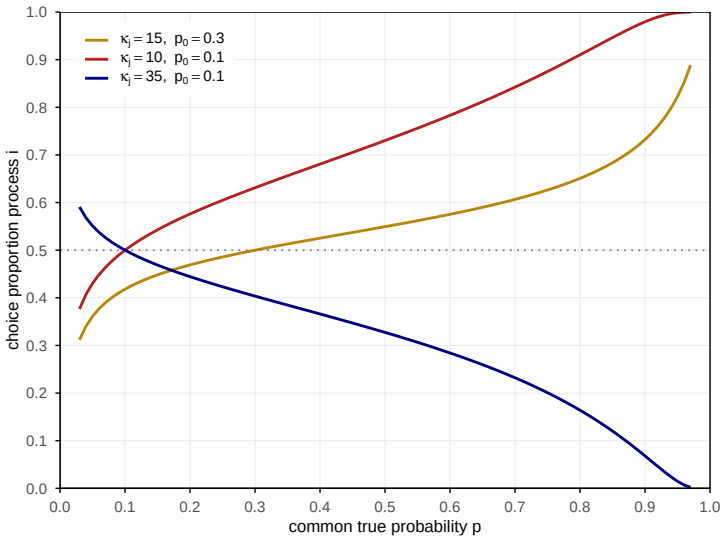


Figure 2: Choice proportions under different ambiguity levels

The figure shows the choice proportion of process i over process j as a function of the matched probability p, based on equation (5). The illustrations keep $\kappa_i = 20$ and $\kappa_0 = 10$ fixed while varying the prior mean p_0 and the precision κ_j of process j.

The choice-proportion curve, however, need not slope upward. In the blue curve, $\kappa_j = 35 > \kappa_i = 20$, so that $A_j < A_i$. The direction of the pattern therefore reverses: process i is chosen less than half of the time at $p = 0.5$, and its choice proportion decreases as p rises. The underlying logic is unchanged. Whichever representation is less ambiguous is relatively favoured as the matched winning probability moves farther above the prior mean, and relatively disfavoured as it moves farther below it.

2.2 Revealed ambiguity

The stochastic choice equation provides a behavioural characterization of relative representational ambiguity without requiring its latent levels to be observed. Relative ambiguity is revealed by how this choice probability changes with p.

Proposition 1 (Revealed ambiguity)**.** *Consider two representations i and j with ambiguity levels A_i and A_j, evaluated by a decision maker under a common prior with mean $p_0 \in (0,1)$ and precision $\kappa_0 > 0$. Let $q_{ij}(p) \triangleq P[i \succ j \mid p]$ denote the probability of choosing i over j when both alternatives offer the same winning probability p. Then:*

1. Ordinal signature. *The choice probability $q_{ij}(p)$ is strictly increasing in p if $A_j > A_i$, strictly decreasing if $A_j < A_i$, and constant at $1/2$ if $A_j = A_i$.*

2. Strength of revealed ambiguity. *Holding the benchmark representation i and the common prior fixed, the slope of $q_{ij}(p)$ at $p = p_0$ is strictly increasing in A_j.*

The first part provides an ordinal revealed-ambiguity signature: a choice probability increasing in p reveals j as more ambiguous than i, whereas a decreasing choice probability reveals the reverse. The second part goes beyond this sign criterion. Against a fixed benchmark i, the slopes at the common crossing point rank comparison representations by ambiguity, even when their slopes have the same sign. Differences in these slopes also quantify the strength of the revealed-ambiguity effect on choice. Differences in slopes thus provide a cardinal comparison of relative ambiguity against a common benchmark.

The common-prior assumption is natural when comparing objectively matched probabilities, and serves as a parsimonious disciplining benchmark. In particular, it implies that choice proportions remain constant at one half when the two representations have equal ambiguity, and that changes in relative ambiguity rotate the choice function around a common crossing point. The central relationship between ambiguity and the slope of choice probabilities does, however, not require this assumption. Even with different prior means, the sign of probability-sensitivity at $p = 0.5$ continues to reveal the relative ambiguity ordering, provided that the prior precision κ_0 is the same. The strength ordering likewise survives for slopes of the latent Probit index, although differences in elevation may affect the corresponding slopes of observed choice proportions.[15]

[15] With different prior means, the ordinal signature need not retain its sign throughout the probability range. Distinct priors also allow systematic choice patterns between equally ambiguous representations,

Familiarity and learning. Beyond this characterization, representational ambiguity need not be determined by presentation mode alone. If familiarity or experience improves the precision with which a given representation is encoded, its ambiguity level should decline. This gives a sharp comparative static:

Corollary 1 (Learning and ambiguity reversal)**.** *Hold the ambiguity of representation i fixed at $A_i > 0$. Any reduction in A_j monotonically reduces its revealed ambiguity relative to i. If learning reduces A_j through A_i, the positive slope signature first vanishes and then reverses. Under a pessimistic prior $p_0 < 0.5$, ambiguity aversion at matched 50–50 probabilities likewise vanishes and reverses. (Proof in Online Appendix A.2)*

This implication connects naturally to evidence on familiarity and learning.[16] Decision makers familiar with conventionally ambiguous sources often display less ambiguity aversion, and sometimes ambiguity seeking, together with greater sensitivity to likelihoods (Heath and Tversky, 1991). In the present framework, such patterns arise naturally if familiarity makes the conventionally ambiguous source more precisely represented than the risk benchmark. Ert and Trautmann (2014) find that sampling can shift choices toward the ambiguous option while increasing sensitivity to the winning probability, a pattern consistent with an ambiguity-ordering reversal. In Oprea and Vieider (2024), we show that sampling also changes behaviour toward fully described risks, consistent with the broader premise that even described probabilities may be represented with finite precision.

2.3 Noisy probabilities and categorical ambiguity

The framework developed above applies directly to standard findings in the ambiguity literature. The main results follow once Ellsberg-type ambiguity—objective uncertainty about the underlying outcome-generating process—is associated with a more ambiguous probability representation in the sense defined above.

The Ellsberg 2-colour problem and ambiguity aversion. In the canonical 2-colour Ellsberg problem, subjects choose between an urn containing 50 red and 50 black

which are ruled out under the common-prior benchmark.

[16] As Keynes (1921, Ch. VI, §1) put it, "As the relevant evidence at our disposal increases [...] we have a more substantial basis upon which to rest our conclusion [...] an accession of new evidence increases the weight of an argument."

balls and an urn with unknown colour composition. Choice patterns typically favour the known-proportion urn regardless of the winning colour, directly violating additivity of subjective probabilities under expected utility. The model naturally accommodates this pattern: with free colour choice, both options have matched probability $p = .5$ (Raiffa, 1961), while the conventionally ambiguous urn is represented with greater ambiguity, $A_a > A_r$. Ambiguity aversion, $P[r \succ a] > .5$, therefore arises whenever $p_0 < .5$.[17] A prior mean below .5 has a natural interpretation as learned *caution* in highly uncertain situations: when probability information is imprecise, posterior beliefs are pulled toward a conservative default. This echoes the classical idea that the weight placed on a probability assessment should depend on the amount of information supporting it (Keynes, 1921).

The same condition, $p_0 < .5$, also generates risk aversion when probabilities are fully described (Vieider, 2024; Oprea and Vieider, 2024; 2026; Vieider, 2026). The resulting link between risk aversion and the Ellsberg paradox is reminiscent of accounts that do not require ambiguity-specific preferences, such as Segal (1987). A related intuition is developed by Halevy and Feltkamp (2005). There, however, the prediction arises because the one-shot task is perceived as a *bundle* of choice problems across which bets can be "hedged." Here, no such misperception of the decision problem is required: relative representational ambiguity and regression toward a cautious prior are sufficient.

N colours and ambiguity-insensitivity. The same logic extends immediately to N-colour settings in which the number of winning colours—and hence the matched probability p—varies across choices (e.g., Abdellaoui, Baillon, Placido and Wakker, 2011). With a fixed ambiguity ordering across probabilities, $A_a > A_r$, the revealed-ambiguity characterization implies that $P[r \succ a]$ increases monotonically in p: *ambiguity-insensitivity* means that the ambiguous option is chosen more often for $p < p_0$, whereas the risky option is chosen predominantly for $p > p_0$.

The magnitude of these effects should naturally depend on the objective uncertainty range, as in Becker and Brownson (1964), and on familiarity with the source of uncertainty, as in Heath and Tversky (1991). As familiarity increases, one should therefore

[17] Ambiguity seeking can still emerge at the individual level for subjects with $p_0 > .5$. This does not require the unknown urn to be represented as less ambiguous, since the relative ambiguity ordering is identified by the slope criterion rather than by choice at $p = .5$.

observe greater sensitivity to probabilities and—under a pessimistic prior, $p_0 < 0.5$—less ambiguity aversion. Both follow naturally if familiarity and learning reduce representational ambiguity. With sufficient learning, both patterns may reverse, revealing the conventionally ambiguous process as the less ambiguous representation.

Gains versus losses. The same mechanism immediately extends to losses. The Bayesian assessment of the probability is unchanged, but its effect on choice reverses because a higher probability is desirable for a gain and undesirable for a loss. Thus, whenever representational ambiguity lowers the posterior assessment of a matched probability, it makes the more ambiguous option less attractive for gains but more attractive for losses. The model therefore predicts the reflection pattern documented in the literature: ambiguity aversion for gains becomes ambiguity seeking for losses, and vice versa (Baillon and Bleichrodt, 2015; Trautmann and van de Kuilen, 2015).[18]

The Ellsberg 3-colour problem. The 3-colour problem requires one additional step. Consider an urn containing 90 balls, 30 of which are red, while the remaining 60 are green or blue in unknown proportions. Let q denote the unknown share of green balls among these 60. By exchangeability, its mean is 1/2, so that $\mathbb{E}[p_R] = \mathbb{E}[p_G] = \mathbb{E}[p_B] = 1/3$. In the first comparison, between betting on red and blue, both events therefore have matched probability $p = 1/3$. The probability of red is fully described, whereas the probability of blue additionally depends on the unknown green–blue split. Its representation consequently carries greater ambiguity, $A_B > A_R$. The revealed-ambiguity characterization then predicts a preference for red whenever $p_0 < 1/3$.

In the second comparison, between red-or-green and blue-or-green, the relative ambiguity ordering reverses. The probability of blue-or-green is fully determined at 2/3, whereas $p_{R\cup G} = 1/3 + 2/3q = 1 - p_B$ continues to depend on the unknown green–blue split. Hence $A_{R\cup G} > A_{B\cup G}$. For $p_0 < 2/3$, the model therefore predicts a preference for blue-or-green. The familiar 3-colour Ellsberg pattern thus follows because adding the common event reverses which side of the choice carries greater representational ambiguity. Crucially, the argument requires only a reversal in *relative* ambiguity: fully described probabilities

[18] This argument assumes that gains and losses are perfectly discriminated, as in related noisy-cognition models (Khaw et al., 2021; Vieider, 2024). Differences in attention or representational precision across domains could attenuate the symmetry without changing the basic mechanism.

may themselves retain positive representational ambiguity throughout.

2.4 Relation to the ambiguity literature

The definition of ambiguity proposed here bears a formal resemblance to approaches that replace a unique subjective probability by richer belief distributions, including sets of priors (Gilboa and Schmeidler, 1989), revealed ambiguity perception (Ghirardato, Maccheroni and Marinacci, 2004), and second-order probability distributions (Segal, 1987; Klibanoff et al., 2005). The key difference concerns the interpretation of this uncertainty and its role in generating choice. In these models, ambiguity is tied to unknown outcome-generating processes or compound lotteries, while simple risk either collapses to an expected-utility benchmark (Gilboa, 1987; Ghirardato et al., 2004; Klibanoff et al., 2005) or is treated as presentation-invariant (Segal, 1987). Here, by contrast, even a fully described risky probability is represented with finite precision, and two described risks may differ in ambiguity simply because their probabilities are presented in more or less accessible formats. Expected utility for risk emerges only in the limiting case as $A_r \to 0$.

This difference yields predictions suitable for empirical testing. For $A_r > 0$, the mean posterior assessment of a risky probability is $\mathbb{E}[\widehat{p}_r \mid p] = p_0 + \omega_r(p - p_0)$, so probability assessment and risk taking depend on the precision with which risk itself is represented. Likewise, the choice share of a conventionally ambiguous option relative to a risky benchmark depends on $(\omega_r - \omega_a)(p - p_0)$ through equation (5). Changing only the presentation of the risky option can therefore shift—and potentially reverse—the revealed-ambiguity pattern. This is the variation I exploit in the experiments below.

Source-dependence. A different approach retains point beliefs about uncertain events but allows their transformation into decision weights to depend on the source of uncertainty (Tversky and Fox, 1995; Chew and Sagi, 2008). Let p_m denote the risky probability that generates stochastic indifference relative to an Ellsberg-type ambiguous probability

p.[19] Equality of average posterior assessments implies

$$p_m = p_0 + \frac{\omega_a}{\omega_r}(p - p_0).$$

This mapping resembles a neoadditive source function (Chateauneuf, Eichberger and Grant, 2007), with fixed point p_0 and slope ω_a/ω_r.[20] The resemblance is again primarily formal. In source-dependent models, source-specific weighting transforms point beliefs into decision weights. Here, the same shape arises from Bayesian inference, with ω_a/ω_r determined by the relative precision of the two probability representations. The model therefore provides a cognitive foundation for source-like probability weighting while shifting the central question to what determines relative representational precision.

The status of fully described risk makes this distinction particularly stark. Source-dependent models typically treat risk as a distinct source. Baillon et al. (2025) treat all risk as a single source in their main analysis, but subsequently qualify this as a tractability assumption that holds only approximately in selected settings.[21] This illustrates the difficulty of determining where one source ends and another begins. No such categorical source partition is required here. A fully described probability is simply another internal representation, characterized by its own ambiguity A_i. Even the winning probability associated with one and the same urn may be represented with different precision depending on its presentation format. Ellsberg-type uncertainty may consequently be less ambiguous than a fully described risky urn if the latter is presented in a less immediately accessible format, reversing the conventional ordering of revealed ambiguity.

It is in this sense that the present model provides cognitive microfoundations for source-specific behaviour. The revealed-ambiguity criterion disciplines what counts as a distinct source, extending naturally from different chance devices to seemingly superficial dif-

[19] Source theory typically defines the related object through a matching probability (Baillon et al., 2025). The present setting instead uses binary choice. The distinction matters because matching and choice can trigger different cognitive processes (Bouchouicha, Oprea, Vieider and Wu, 2024; Khaw et al., 2026); Li and Wu (2026) show that binary choice, probability matching, and event matching produce distinct empirical patterns under ambiguity.

[20] To see this more directly, write $w(p) = a + bp$. Then $b = \omega_a/\omega_r$ and $a = (1 - \omega_a/\omega_r)p_0$.

[21] Baillon et al. (2025, Section 10.2, p. 8777) state: "The assumption holds approximately for emotion-neutral risky events and outcomes, and we focus on those." Draws from urns containing numbered balls would seem paradigmatic examples of such 'emotion-neutral' events. Although the authors acknowledge possible exceptions involving complex arithmetic (endnote 9), this provides little justification for treating visual and simple fractional descriptions of the winning probability in one and the same ten-ball urn as distinct sources of uncertainty.

ferences in presentation mode or description. It does so parsimoniously: rather than assigning a separate two-parameter weighting function to each source, the model generates source-like choice patterns from differences in representational precision and permits directional predictions about presentation effects, which I exploit in the preregistered experiments below. The resulting mapping from latent representation to stochastic choice is uniquely identified from binary-choice data (Vieider, 2025b), in contrast to probability weighting under prospect theory (Wilcox, 2011; Apesteguia and Ballester, 2018). Finally, the cognitive mechanism clarifies the status of source functions themselves. Whereas the literature variously interprets them as preference functionals or allows cognitive factors to shape them, here source dependence emerges directly from Bayesian inference over representations of differing precision. Learning and familiarity then affect source-dependent behaviour through the same mechanism rather than requiring a separate interpretation.

2.5 Log-odds mapping

Previous noisy-cognition models of probability distortions under risk have often used probabilities in log-odds space (Zhang, Ren and Maloney, 2020; Vieider, 2024; Khaw et al., 2026). This format has a particularly close formal connection to functional forms commonly used in the prospect-theory literature. To see this, let $\tilde{\ell}_i \sim \mathcal{N}(\ell_p, \nu_i^2)$, where $\ell_p \triangleq \ln \frac{p}{1-p}$, and define the precision of representation i as $\tau_i \triangleq \nu_i^{-2}$. Further let prior beliefs take the form $\ell_p \sim \mathcal{N}(\ell_0, \sigma^2)$, where $\ell_0 \triangleq \ln \frac{p_0}{1-p_0}$, with prior precision $\tau_0 \triangleq \sigma^{-2}$.[22] Bayesian updating yields the posterior mean $\widehat{\ell}_i = \gamma_i \tilde{\ell}_i + (1-\gamma_i)\ell_0$, where $\gamma_i \triangleq \frac{\tau_i}{\tau_i+\tau_0}$. The posterior weight placed on the internal representation is thus determined by its precision relative to prior precision, exactly paralleling $\omega_i = \frac{\kappa_i}{\kappa_i+\kappa_0}$ in the probability-space model. Comparing two representations i and j of the same probability p, and averaging over repeated representations, gives

$$P[i \succ j] = \Phi\left[\frac{(\gamma_i - \gamma_j)\left(\ln \frac{p}{1-p} - \ln \frac{p_0}{1-p_0}\right)}{\sqrt{\gamma_i^2\nu_i^2 + \gamma_j^2\nu_j^2}}\right].$$

[22] Here precision has its conventional Gaussian meaning of inverse variance. This is not a substantive departure from the Bernoulli-equivalent definition above. In log-odds space, the signal variance ν_i^2 is constant, so inverse variance itself provides a probability-independent measure of representational precision. In probability space, by contrast, raw variance also changes mechanically with $p(1-p)$, which is why precision there was defined in Bernoulli-equivalent units.

Unlike in the probability-space representation above, response variance is now constant along the log-odds line; probability dependence of stochastic choice enters through the systematic component in the numerator alone.[23]

The corresponding stochastic-indifference mapping has the same functional form as linear-in-log-odds probability-weighting functions used in prospect theory (Goldstein and Einhorn, 1987; Gonzalez and Wu, 1999; Zhang and Maloney, 2012). Let p_m denote the probability presented through mode i that generates 50–50 choice against probability p presented through mode j. Equality of mean posterior assessments yields

$$\ln\left(\frac{p_m}{1-p_m}\right) = \frac{\gamma_j}{\gamma_i}\ln\left(\frac{p}{1-p}\right) + \ln(\delta_{ij}), \tag{6}$$

where $\delta_{ij} \triangleq [p_0/(1-p_0)]^{1-(\gamma_j/\gamma_i)}$. When representation j is less precise than representation i, $\gamma_j/\gamma_i < 1$, producing the familiar inverse-S mapping. The Bayesian model therefore generates a linear-in-log-odds transformation between probability representations. If risk is used as the benchmark, this mapping is formally analogous to the relative source function $w_r^{-1} \circ w_a$, where w_r and w_a denote the probability-weighting functions associated with risky and ambiguous sources, respectively. Here, however, the transformation arises from differences in the precision of internal probability representations, rather than from independently specified source-dependent weighting functions.

Mapping between formats. The comparison with the probability-space model is revealing. In the model developed above, differences in mean posterior assessments are linear in $p - p_0$; the nonlinear probability dependence of choice arises instead from heteroscedastic response variance. Under homoscedastic log-odds noise, this allocation of curvature is reversed: the denominator of the choice index is constant across probabilities, while the systematic component in the numerator is nonlinear in p. The same Bayesian logic can therefore generate closely related choice patterns across probabilities through two different channels—heteroscedastic noise in probability space or nonlinear mean assessments under homoscedastic log-odds noise. This raises the question of how

[23] Using $\tau_i = \nu_i^{-2}$, the contribution of representation i to response variance can equivalently be written as $\tau_i/(\tau_i + \tau_0)^2$, directly paralleling $\kappa_i/(\kappa_i + \kappa_0)^2$ in the probability-space formulation. The latter is additionally multiplied by $p(1-p)$, reflecting the intrinsic Bernoulli variance of probability representations. Thus the two formulations share the same Bayesian precision-weighting structure but differ in how representational variance changes across the probability domain.

the precision assumptions in the two formats are related.

"Constant precision" on one representational scale therefore does not imply constant precision on another. Fisher information $\mathcal{I}_i(p)$ provides a common local measure of precision and allows the noise profile imposed on one scale to be expressed in the units of the other. In probability space, $\mathcal{I}_i^p(p) = \kappa_i[p(1-p)]^{-1}$, while under Gaussian log-odds noise, $\mathcal{I}_i^\ell(p) = [\nu_i^2 p^2(1-p)^2]^{-1}$. Equating the two yields $\nu_i^2(p) = [\kappa_i p(1-p)]^{-1}$. Thus, constant Bernoulli-equivalent precision corresponds to log-odds noise that increases as probabilities move away from even odds ($p = 0.5$). Conversely, $\kappa_i(p) = [\nu_i^2 p(1-p)]^{-1}$: constant log-odds noise requires greater Bernoulli-equivalent precision as probabilities move away from even odds. Constant precision in probability space and in log-odds space therefore imply different allocations of representational resources across probabilities.

The preceding comparison also clarifies the role of efficient coding in determining representational precision. A fixed log-odds code already implies greater Bernoulli-equivalent precision as one moves away from even odds, without requiring any adaptation to the decision environment. Efficient coding provides an additional mechanism through which precision may vary across probabilities. For instance, Frydman and Jin (2025) show how allocating greater precision to probabilities expected to occur more frequently generates systematic changes in the probability sensitivity of lottery valuations as those expectations change. Efficient coding is therefore complementary to the present account: while the choice of representation itself determines a baseline precision profile, adaptation to experience may further alter that profile.[24]

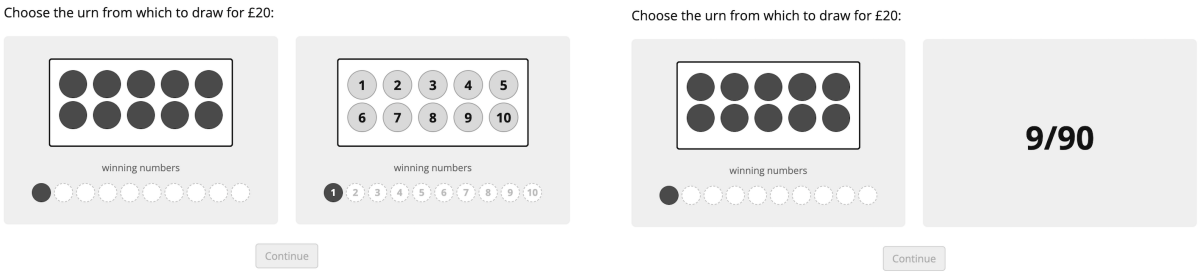
3 Experiment I: Which option is more ambiguous?

Experiment I probes a simple hypothesis: neither ambiguity-insensitivity nor ambiguity aversion are unique to Ellsberg-type ambiguity. In particular, I aim to test whether changing only how described risks are presented can attenuate or even reverse both regularities, thereby revealing the described risk as the more ambiguous option.

[24] See Polania and Hare (2026) for an efficient-coding account exploiting Fisher information to model cross-attribute allocation of a common representational budget.

3.1 Experimental procedures

Participants made a series of binary choices between two urns, each offering a prize of £20 if a ball carrying a winning number was drawn, and nothing otherwise. One urn was always *Ellsberg-ambiguous*: it contained ten balls bearing unknown numbers between 1 and 10, so that some numbers could occur multiple times and others not at all. To rule out concerns about urn stacking, participants were told that the urns had been precomposed and that at the end of the study they would rank the numbers 1 to 10 from most to least preferred. Their top-ranked numbers would then be designated as the winning numbers in every choice they had made. Accordingly, the row below the ambiguous urn displayed ten markers, of which k were filled, indicating only *how many* numbers would win rather than which numbers they were.



(a) visual condition **(b)** FRAC2 condition

Figure 3: Screenshots of experimental choices

Screenshots of the choice environment in the experimental conditions. Panel A shows a screenshot from the visual condition; Panel B from the fraction conditon FRAC2.

Treatments (known urn). The treatment was administered between subjects and varied only how the described probability was presented. Figure 3 shows screenshots from two conditions. In the **visual condition**, the urn contained ten balls numbered 1 through 10, each appearing exactly once, with the winning numbers indicated on a row from 1 to 10 below the urn. In the two **fraction conditions**, the same probability was instead displayed as a fraction $mk/10m$, with the multiplier m drawn afresh on every trial from $\{6, 7, 8, 9\}$ in condition FRAC2 and from $\{23, 26, 29, 31\}$ in condition FRAC3. Thus, a probability of 0.3 appeared as one of 18/60, 21/70, 24/80, or 27/90 in FRAC2, and as one of 69/230, 78/260, 87/290, or 93/310 in FRAC3. Participants were told that the fraction gave the number of winning balls relative to the total number of balls in the urn, and that this total varied from choice to choice. Conditions were assigned at random and were identical in all respects except the presentation mode of the risky option.

Hypotheses. The visual condition serves as the baseline, in which I expect to reproduce the canonical ambiguity patterns: ambiguity aversion on average and in the 50–50 Ellsberg choice, together with ambiguity insensitivity, so that the choice proportion of the risky option increases with the matched winning probability.[25] The visual display should provide particularly accessible probability information. Frequency representations have been shown to facilitate Bayesian reasoning in a variety of settings (Gigerenzer and Hoffrage, 1995; Cosmides and Tooby, 1996), while perceptual grouping can facilitate rapid numerical assessment (Ciccione and Dehaene, 2020). In the present display, highlighting the winning numbers creates a visually distinct group, whose size can be determined simply by reading the rightmost highlighted number. Since the urn always contains ten balls, assessing the winning probability thus requires reading only a single number. Oprea and Vieider (2026) show that a closely related visual display increases sensitivity to probabilities in tradeoffs between lotteries and sure amounts of money. I therefore expect this display to generate relatively precise probability representations.

The central prediction is that describing the *risky option* using fractions rotates this relationship: ambiguity insensitivity should attenuate and may reverse into oversensitivity, thereby revealing the described risk as more ambiguous. Fractions have been documented to be substantially harder to process than decimals when comparing numerical magnitudes, partly because their two-component structure requires integrating numerator and denominator (DeWolf et al., 2014). Varying denominators may further increase processing demands by preventing reliance on a fixed numerical reference across choices. I therefore expect the rotation to be stronger in FRAC3 than in FRAC2, on the premise that numerically larger but objectively equivalent fractions are more difficult to process. Under the maintained common-prior assumption, treatment effects should rotate individual choice-proportion curves around the prior mean p_0, where the two options are chosen equally often.[26] For a cautious prior, $p_0 < 0.5$, a sufficiently strong rotation to reverse the slope should also reverse choice proportions at the matched 50–50 comparison, turning ambiguity aversion into ambiguity seeking.

[25] The design, hypotheses, and analysis were pre-registered at AsPredicted prior to data collection (https://aspredicted.org/gz4ax4.pdf). The experimental instructions are shown in Online Appendix C.1.

[26] A similar qualitative prediction obtains under separate priors, although the location of the crossing point then depends more strongly on the latent prior.

Stimuli. The ambiguous urn's mean winning probability p_a took the values $0.1, 0.2, \ldots, 0.9$. On each trial, the described urn offered $p_r = p_a + \delta$, with $\delta \in \{0, \pm 0.1, \pm 0.2\}$ and probabilities truncated to $[0.1, 0.9]$, yielding 39 distinct cells. Cells with $\delta = 0$—for which the two urns had identical winning probabilities, and on which the main analysis below focuses—were presented four times; all other cells were presented once, for a total of 66 choices. Four additional trials pitted two described urns, using the condition's own display format, against one another at winning probabilities of 0.9 and 0.1. These dominance trials served as attention checks.[27] All 70 trials were presented in random order, and the side of the screen on which each urn appeared was randomized independently on every trial. Selecting an option required clicking on the option and then on a CONTINUE button, allowing choices to be revised before submission.

Incentives. One participant in ten was selected at random at the end of the study to play one of their choices for real. Li and Vieider (2026) meta-analyze the literature on incentive effects in risk-taking, delay discounting, and decisions under ambiguity, and find that this probabilistic payment scheme produces behaviour indistinguishable from paying every participant—and, more generally, that incentive effects in these domains are usually negligible. For a selected participant, one choice was drawn at random and resolved. A ball was drawn from the urn they had chosen, and if it carried one of the winning numbers—their own top-ranked numbers in the case of the ambiguous urn—they received a bonus of £20 in addition to the fixed participation payment.

Participants. The experiment was conducted online with 604 participants recruited through Prolific and resident in the United Kingdom. Participants were randomly assigned to VISUAL ($N = 190$), FRAC2 ($N = 209$), or FRAC3 ($N = 205$). After providing informed consent, participants read the instructions and completed three comprehension questions, which had to be answered correctly before proceeding. The median session lasted 12.4 minutes.

Statistical testing. The main text presents nonparametric choice patterns with confidence intervals and distributions of individual-level sensitivity measures using matched

[27] Below, I include all subjects and choice in the main data analysis. Overall, 95% of subjects made one mistake or fewer. Excluding subjects who made more than one mistake in the catch trials does not alter the conclusions.

probability data. Additional patterns for offset trials and reduced-form Probit regressions are reported in Online Appendix B, and do not affect any of the conclusions presented here unless explicitly stated otherwise.

3.2 Results: Ellsberg-Ambiguity is not special

I start with the matched-probability choices, for which $p_r = p_a$. Panel A of Figure 4 plots the proportion choosing the risky option by winning probability and treatment. The visual condition reproduces the standard ambiguity patterns. Subjects are ambiguity averse on average and in the canonical 50–50 comparison. Moreover, ambiguity aversion increases steadily with the matched probability of winning: choices are approximately ambiguity neutral at $p_r = p_a = 0.1$, while the risky urn is chosen more than 60% of the time once the winning probability exceeds 0.5. The visual condition thus provides a clean replication of the canonical regularities in the literature: ambiguity aversion and ambiguity insensitivity.[28]

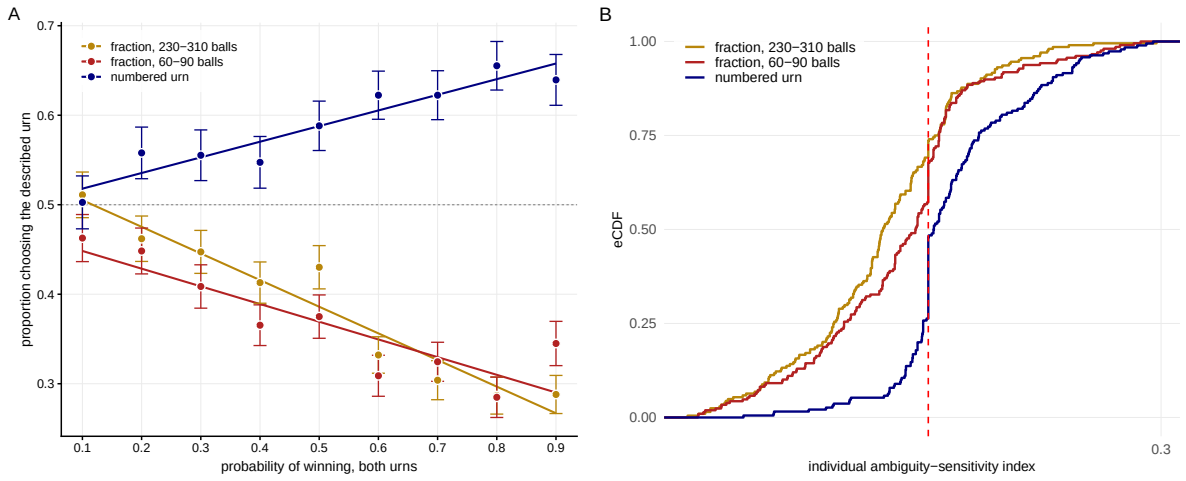


Figure 4: Choice proportion of risky option for $p_r = p_a$

Panel A shows choice proportions of the risky option by probability and treatment arm, jointly with 95% confidence intervals. The lines are linear fits. Panel B shows empirical cumulative distribution functions of an individual-level index of ambiguity-sensitivity, with 0 benchmarking perfect sensitivity.

When risk is displayed using fractions, however, the pattern changes sharply. Choice proportions of the risky option start at similar levels for $p_a = p_r = 0.1$, but then decline systematically as the probability of winning increases. Both canonical patterns observed in the visual baseline therefore *reverse*: The reflection around the 50–50 line indicates

[28] Some of the literature finds ambiguity *seeking* at small probabilities. This pattern arises primarily in choice list designs eliciting matching probabilities; see Li and Wu (2026) for a direct comparison. The discrepancy parallels differences between binary-choice and choice-list elicitation under risk (Bouchouicha et al., 2024).

that the same underlying behavioural patterns persist, but with the roles of the two options reversed. Once risk is expressed in fractions, *the described risky option is revealed as the more ambiguous option.*

Panel B of Figure 4 plots empirical cumulative distribution functions (*eCDFs*) of an individual-level index of ambiguity sensitivity, with zero corresponding to perfect sensitivity.[29] The figure serves two purposes. First, it shows that the aggregate patterns in Panel A are not an artifact of aggregation. Second, and more importantly for the revealed-ambiguity interpretation, the sign of probability sensitivity provides an ordinal measure of relative ambiguity at the individual level without requiring parametric estimation. A positive index reveals the Ellsberg urn as the more ambiguous representation, whereas a negative index reveals the described risk as more ambiguous.

In the baseline numbered-urn condition, 52% of subjects exhibit ambiguity insensitivity, compared with 26% exhibiting oversensitivity (the remaining 22% are classified as perfectly sensitive). The distribution lies significantly above zero according to a one-sample Wilcoxon signed-rank test ($p < 0.001$). The pattern reverses in the fraction conditions. Only 32% of subjects in FRAC2 and 26% in FRAC3 exhibit insensitivity, whereas 57% and 69%, respectively, exhibit oversensitivity; both distributions lie significantly below zero ($p < 0.001$). Thus the treatment changes not merely the aggregate choice function but the distribution of individual revealed-ambiguity orderings. Relative to the visual condition, both fraction treatments produce a pronounced shift toward oversensitivity ($p < 0.001$ in both cases). The shift is somewhat larger in FRAC3 than in FRAC2 ($p = 0.042$), consistent with the predicted dose effect, although this difference is small relative to the shift induced by moving from the numbered urn to fractions.

Implications. The first implication is straightforward: *Ellsberg-type ambiguity is not special.* Both canonical signatures of ambiguity reverse when only the presentation mode of the *risky probability* is changed. Once simple risk is displayed using fractions, the risky

[29] For each subject, I pair complementary probabilities (0.9 with 0.1, 0.8 with 0.2, and so forth) among the matched-probability choices and calculate the difference in described-urn choice proportions within each pair. The index averages these differences per unit of log-odds distance, weighting each pair by the square of that distance. It is zero when a common shift in both probabilities leaves relative choice unchanged, positive under insensitivity to Ellsberg-ambiguity, and negative under oversensitivity to Ellsberg-ambiguity. The construction does not require fitting a parametric choice model. An unweighted version yields the same substantive conclusions.

option acquires the behavioural signature conventionally associated with ambiguity: it is chosen less often on average, and its choice share falls as the matched winning probability rises. This is exactly what the model predicts. Which option behaves as "more ambiguous" depends on the relative noise with which the two probabilities are represented, not on whether one probability is objectively described and the other is not.

The second implication is more fundamental. In standard ambiguity models, the preference component applies specifically to lack of knowledge about the outcome-generating process. In the present experiment, that object is held fixed across treatments: the Ellsberg urn does not change. Only the presentation mode of the simple risky probability changes. Any preference term that applies specifically to the ambiguous outcome-generating process is therefore constant across treatments and cannot account for the observed reversal. The canonical Ellsberg patterns therefore provide no basis, by themselves, for identifying ambiguity aversion with a primitive preference toward a distinct category of uncertainty. This has direct implications for models in which simple risk provides an expected-utility benchmark, or more generally, a presentation-invariant benchmark. Such models predict no systematic change when only the presentation mode of that risk is altered. Nor is the reversal naturally explained as merely "more noise." The treatment effect is systematic and directional: it rotates the sensitivity pattern and shifts relative choice strongly enough to reverse both canonical signatures of ambiguity.

Similar considerations apply to approaches based on source-dependent probability distortions. One possible response is to allow risk itself to comprise multiple sources. This, however, immediately raises the question of what constitutes a "source of risk". The present experiment does not move between roulette wheels, card games, or other qualitatively distinct random devices: risk is always explicitly described by an urn with a precisely described proportion of winning balls. To accommodate the results, the visual and textual representations of the same risky urn would therefore already have to count as distinct sources. Behaviour moreover changes systematically across different urns describing that same risk. Source-dependent models can accommodate such patterns by assigning each presentation mode a different source function, but this places little restriction on how behaviour should change across modes. The present model instead links these changes through a single latent object: relative representational precision. A change in

that object has directional implications for probability sensitivity and, conditional on the prior, for ambiguity aversion.

4 Experiment II: Probability noise is continuous

Experiment I established that changing how a probability is described can reverse ambiguity attitudes, turning the canonical patterns in the literature into their opposites. Experiment II tests the continuum claim more directly by placing the standard risk–ambiguity comparison alongside two risk–risk comparisons.

4.1 Experimental procedures

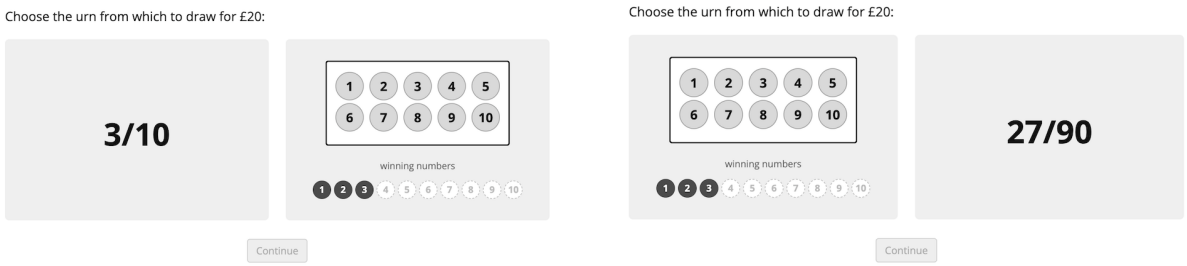
As in Experiment I, participants made a series of binary choices between two urns, each paying £20 if a ball carrying a winning number was drawn. One option was identical in every condition: a visually displayed urn containing ten balls numbered 1 through 10, each occurring exactly once. The winning numbers were shown as filled markers on a row from 1 to 10 below the urn. The between-subject treatment varied the object against which this numbered urn was compared:

- ELLSBERG: an urn containing ten balls carrying unknown numbers between 1 and 10, so that some numbers could occur repeatedly and others not at all. This condition is identical to the visual baseline condition in Experiment I and again serves as the benchmark. As before, the winning numbers were determined by the participant's own ranking, elicited after all choices had been completed.

- FRACS: an urn described by the fraction $k/10$, so that probabilities appeared as $1/10, 2/10, \ldots, 9/10$. This urn was likewise juxtaposed with the visual risky urn containing ten sequentially numbered balls (see Panel A in Figure 5).

- FRACM: an urn described by the fraction $mk/10m$, with the multiplier m drawn afresh on every trial from $\{6, 7, 8, 9, 11, 12, 13\}$, so that denominators ranged from 60 to 130. This condition therefore extends the FRAC2 treatment from Experiment I to a somewhat wider multiplier range. Here, however, the fraction urn is compared directly against the visual risky urn (see Panel B in Figure 5).

Hypotheses. Experiment I varied the presentation mode of the described option while holding the ambiguous urn fixed. Experiment II reverses this logic: the numbered urn



(a) visual condition **(b)** FRACM condition

Figure 5: Screenshots of experimental choices

Screenshots of the choice environment in the experimental conditions. Panel A shows a screenshot from the FRACS condition; Panel B from the FRACM conditon.

is held fixed while the option against which it is compared varies. Because the same numbered urn appears in all three conditions, treatment effects are measured relative to a common benchmark, allowing the relative ambiguity of the different representations to be compared directly. The cognitive rationale developed for Experiment I suggests that the visual display should generate particularly precise probability representations. Simple fractions with a fixed denominator of ten should generate higher representational ambiguity but remain comparatively accessible, whereas larger fractions with varying denominators result in heightened ambiguity. Crucially, neither need necessarily be less ambiguous than an Ellsberg-urn of unknown composition.

The key hypothesis is that conventional Ellsberg ambiguity lies on a continuum of representational ambiguity, between the two fraction formats. I therefore expect all three treatment urns to be revealed as more ambiguous than the numbered urn, but to different degrees. The simple fractions in FRACS should generate the weakest ambiguity-insensitivity and the weakest preference for the numbered urn. Based on Experiment I, the larger fractions in FRACM are expected to generate the strongest effects on both dimensions. The Ellsberg urn is predicted to fall between these two cases. Experiment II thus provides a direct test of the continuum hypothesis: conventional ambiguity should behave not as a separate category of uncertainty, but as an intermediate degree of representational imprecision.[30]

Stimuli, incentives, and participants. Incentives were identical to those in Experiment I. The stimuli were also the same, except that matched-probability choices at the

[30]These hypotheses were pre-registered before data collection (https://aspredicted.org/cc946k.pdf). The instructions are included in Online Appendix C.2.

endpoints 0.1 and 0.9 were repeated six rather than four times. I targeted 400 participants per treatment arm. The larger sample was motivated by the expectation that Experiment II would reveal more gradual shifts in slopes rather than the sharp reversals observed in Experiment I, and would therefore require greater power to detect treatment differences reliably. In total, 1,206 participants completed the study: 417 in AMB, 409 in FRACS, and 380 in FRACM.[31]

Statistical testing. As in Experiment I, the main text reports the key nonparametric choice patterns for matching probabilities and corresponding statistical tests. Additional results on offset trials and regression analyses are reported in Online Appendix B.

4.2 Results

In Experiment II, the visual urn containing balls numbered 1 through 10 is held fixed across conditions. To avoid confusion, I refer to this option as the *numbered urn*. I refer to the comparison option as the *ambiguous urn*, in line with my hypotheses.

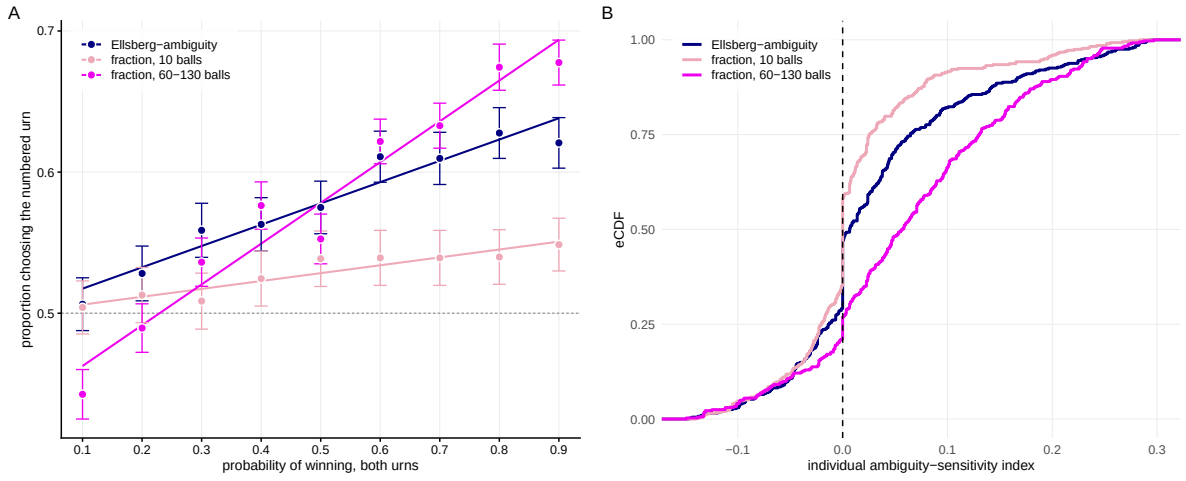


Figure 6: Choice proportion of $p_n = p_a$

Figure 6 shows the results. Panel A displays the aggregate choice patterns. The slopes are ordered monotonically across conditions. The shallowest slope occurs in FRACS. Remarkably, even this comparison already displays the canonical hallmarks of ambiguity: relative insensitivity to changes in the matched probability and aversion to the representational ambiguity induced by the fraction format, both on average and at 50–50—*even*

[31] I again include all observations in the analysis. Excluding subjects who made mistakes in the catch trials on more than one occasion does not affect the results.

though the two options are objectively identical in both probability and urn size and differ only in how that probability is described. At the other extreme, FRACM produces the steepest slope. Subjects prefer the fraction urn at small probabilities but increasingly switch toward the numbered urn as the common winning probability rises, producing the strongest probability dependence of the three conditions.[32] The slope in the Ellsberg condition is virtually indistinguishable from that in experiment I, and importantly, lies almost exactly between the slopes of the two fraction conditions. Ambiguity therefore does not occupy a distinct behavioural category. Instead, *conventional ambiguity falls between two different presentation modes of objectively known probabilities*, exactly as predicted if uncertainty about probabilities varies continuously with the precision of their internal representation.

Panel B complements the aggregate results by plotting the same individual-level sensitivity index used in Experiment I. The ordering is again monotonic. Insensitivity is weakest in FRACS, but remains significantly above zero according to a one-sample Wilcoxon signed-rank test ($p < 0.001$). Ellsberg-ambiguity produces significantly greater insensitivity than FRACS ($p = 0.002$), while FRACM produces greater insensitivity still ($p < 0.001$). Only a minority of subjects exhibit oversensitivity, with similar proportions across conditions. The main difference lies instead in the proportion of subjects classified as perfectly sensitive: this is highest in FRACS, lower under Ellsberg-ambiguity, and lowest in FRACM, mirroring the overall ordering of the sensitivity distributions.

Implications. The key implication is again simple: *conventional ambiguity does not stand out categorically, but is instead revealed to have representational ambiguity intermediate between two simple described risk presentation modes.* This reinforces the central conclusion from Experiment I. If Ellsberg-ambiguity is not a distinct representational category, then the observed choice patterns cannot naturally be interpreted as revealing a primitive preference toward knowledge about the outcome-generating process as such. Instead, the results support the Bayesian-inference account developed above, in which

[32] At 50–50, choice levels are similar for the Ellsberg condition and FRACM. The relatively low elevation of FRACM may reflect a ratio-bias effect (Denes-Raj, Epstein and Cole, 1995), whereby multiplying numerator and denominator by a common factor affects the perceived probability of winning. This hypothesis was not pre-registered, but is related to a mechanism discussed previously in Vieider (2026) for a related manipulation under risk, and shown by Filiz-Ozbay, Gulen, Masatlioglu and Ozbay (2022) for choices between risky and ambiguous urns of different sizes. Under the logic of the model, one would conceive of such an effect as a bias in the likelihood, rather than a shift in the prior.

behaviour reflects inference from imperfect internal representations of probabilities. The crucial point is that the representational ambiguity of a simple, fully described risk can be altered by the presentation mode of that probability. Simply increasing the urn size and describing it by a fraction is sufficient to reveal a described risk as more ambiguous than an Ellsberg-type urn.

Beyond the implications for standard ambiguity models discussed above, the results also bear on the literature comparing Ellsberg-ambiguity with compound risk. A large body of work studies situations in which the composition of an urn is itself generated by an upstream lottery over known compositions (e.g. Halevy, 2007; Dean and Ortoleva, 2019; Chew et al., 2023). The present results suggest that compound probability, too, need not constitute a distinct category. Instead, it can be viewed as another point on a continuum of ambiguity in probability representation. This also cautions against interpreting behavioural differences between compound risk and categorical ambiguity as evidence for distinct preferences toward such ambiguity. Reducing a compound lottery and understanding a randomization argument for ambiguity—such as that proposed by Raiffa (1961)—are cognitively different tasks, even when both ultimately imply the same objective probabilities. Training subjects to reduce compound lotteries may affect one type of representational ambiguity without touching the other, and could thereby even increase the contrast with unknown outcome-generating processes (Wu et al., 2024).

5 Conclusion

I have proposed a model in which ambiguity is a property of internal probability representations. Different outcome-generating processes—and different presentation modes of substantively identical simple risks—may thereby give rise to different degrees of representational precision. The model generates a *revealed-ambiguity signature*: ambiguity-insensitivity, defined as increasing aversion to the option with the more ambiguous representation as the winning probability rises, reveals which amongst two presentation modes is the more ambiguous. This signature provides a monotonic mapping from observed choice patterns into relative representational ambiguity.

Two experiments test this characterization. Experiment I compares an Ellsberg-type urn with objective risks described either visually or through fractions describing urns of

varying sizes. The visual-risk condition reproduces the canonical patterns of ambiguity aversion and ambiguity-insensitivity. Describing the same objective risks through fractions, however, reverses both patterns: the objectively described simple risks are now revealed as more ambiguous than the Ellsberg urn. Experiment II fixes the visual representation of risk and varies the comparison urn. Even simple fractions such as 1/10 and 2/10 generate the canonical ambiguity signatures relative to the visual risk presentation mode; conventional Ellsberg-ambiguity generates stronger signatures, while more fractions described by larger urn sizes generate stronger signatures still. Ellsberg-ambiguity therefore falls between different representations of objectively known risk. Ambiguity differs by degree, rather than by category.

An important theoretical tradition decomposes behaviour under ambiguity into *ambiguity perception* and *ambiguity attitudes* (Ghirardato et al., 2004). The evidence here suggests a more parsimonious possibility: ambiguity perception, combined with inference, may suffice to generate many of the canonical behavioural patterns. Bayesian decoding of imprecise probability representations produces systematic aversion and insensitivity without requiring a primitive preference toward ambiguity. The same logic applies under risk. As representational ambiguity vanishes, differences across presentation modes disappear and behaviour approaches the benchmark associated with precise probability representation. Learning and familiarity should therefore reduce revealed ambiguity, and may eventually reverse its ordering when one representation becomes more precise than another.

The broader implication concerns what experimental choices reveal about preferences. In the model, canonical patterns under both risk and Ellsberg-type ambiguity arise from Bayesian inference subject to noisy internal representations. Primitive attitudes toward either domain are therefore not required to explain these regularities. This conclusion complements a growing literature showing that systematic behaviour under risk and over time can arise from closely related information-processing frictions (e.g. Khaw et al., 2021; Frydman and Jin, 2022; Oprea and Vieider, 2026). The findings do not imply that stable preferences toward risk do not exist. Rather, they raise the question of how such preferences can be separated from the cognitive machinery through which small-stakes choices are evaluated. Determining when observed choices identify welfare-relevant preferences, and how those preferences should be measured, remains an important open

problem.

The broader implication is therefore constructive. A noisy-cognition account does not merely add another layer of behavioural flexibility; to the contrary, it can simplify the modelling landscape by replacing multiple domain-specific expected utility generalizations with a single process model. This, in turn, refocuses attention on the object of primary economic interest: the underlying preferences themselves. In that sense, the present framework offers a more unified positive theory of behaviour under uncertainty.

References

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

Alaoui, Larbi, and Antonio Penta (2025) 'What's in a u?' Technical Report, Universitat Pompeu Fabra

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

Baillon, Aurélien, and Han Bleichrodt (2015) 'Testing Ambiguity Models through the Measurement of Probabilities for Gains and Losses.' *American Economic Journal: Microeconomics* 7(2), 77–100

Baillon, Aurélien, Han Bleichrodt, Chen Li, and Peter P Wakker (2025) 'Source theory: A tractable and positive ambiguity theory.' *Management Science* 71(10), 8767â8782

Becker, Selwyn W, and Fred O Brownson (1964) 'What price ambiguity? or the role of ambiguity in decision-making.' *Journal of Political Economy* 72(1), 62–73

Berger, James (2013) *Statistical decision theory: foundations, concepts, and methods* (Springer Science & Business Media)

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

Bouchouicha, Ranoua, Ryan Oprea, Ferdinand M. Vieider, and Jilong Wu (2024) 'Is prospect theory really a theory of choice?' Technical Report, Ghent University Discussion Papers

Chateauneuf, Alain, Jürgen Eichberger, and Simon Grant (2007) 'Choice under Uncertainty with the Best and Worst in Mind: Neo-additive Capacities.' *Journal of Economic Theory* 137(1), 538–567

Chew, Soo Hong, and Jacob S Sagi (2008) 'Small worlds: Modeling attitudes toward sources of uncertainty.' *Journal of Economic Theory* 139(1), 1–24

Chew, Soo Hong, Bin Miao, and Songfa Zhong (2017) 'Partial ambiguity.' *Econometrica* 85(4), 1239–1260

_ (2023) 'Ellsberg meets Keynes at an urn.' *Quantitative Economics* 14(3), 1133–1162

Ciccione, Lorenzo, and Stanislas Dehaene (2020) 'Grouping mechanisms in numerosity perception.' *Open Mind* 4, 102–118

Cosmides, Leda, and John Tooby (1996) 'Are humans good intuitive statisticians after

all? rethinking some conclusions from the literature on judgment under uncertainty.' *cognition* 58(1), 1–73

de Clippel, Geoffroy, Paola Moscariello, Pietro Ortoleva, and Kareen Rozen (2025) 'Caution in the face of complexity.' Working Paper

de Clippel, Geoffroy, Ryan Oprea, and Kareen Rozen (2026) 'As if'

Dean, Mark, and Pietro Ortoleva (2019) 'The empirical relationship between nonstandard economic behaviors.' *Proceedings of the National Academy of Sciences* 116(33), 16262–16267

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

Denes-Raj, Veronika, Seymour Epstein, and Jonathan Cole (1995) 'The generality of the ratio-bias phenomenon.' *Personality and Social Psychology Bulletin* 21(10), 1083–1092

DeWolf, Melissa, Margaret A Grounds, Miriam Bassok, and Keith J Holyoak (2014) 'Magnitude comparison with different types of rational numbers.' *Journal of Experimental Psychology: Human Perception and Performance* 40(1), 71

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

Enke, Benjamin (2024) 'The cognitive turn in behavioral economics.' *Mimeo*

Ert, Eyal, and Stefan T Trautmann (2014) 'Sampling experience reverses preferences for ambiguity.' *Journal of Risk and Uncertainty* 49, 31–42

Filiz-Ozbay, Emel, Huseyin Gulen, Yusufcan Masatlioglu, and Erkut Y Ozbay (2022) 'Comparing ambiguous urns with different sizes.' *Journal of Economic Theory* 199, 105224

Frydman, Cary, and Lawrence J Jin (2022) 'Efficient coding and risky choice.' *Quarterly Journal of Economics* 137(1), 161–213

Frydman, Cary, and Lawrence J. Jin (2025) 'On the source and instability of probability weighting.' Working Paper

Ghirardato, Paolo, Fabio Maccheroni, and Massimo Marinacci (2004) 'Differentiating ambiguity and ambiguity attitude.' *Journal of Economic Theory* 118(2), 133–173

Gigerenzer, Gerd, and Ulrich Hoffrage (1995) 'How to improve bayesian reasoning without instruction: Frequency formats.' *Psychological Review* 102(4), 684–704

Gilboa, Itzhak (1987) 'Expected utility with purely subjective non-additive probabilities.' *Journal of Mathematical Economics* 16(1), 65–88

Gilboa, Itzhak, and David Schmeidler (1989) 'Maxmin expected utility with non-unique prior.' *Journal of Mathematical Economics* 18(2), 141–153

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

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

Halevy, Yoram (2007) 'Ellsberg Revisited: An Experimental Study.' *Econometrica* 75(2), 503–536

Halevy, Yoram, and Vincent Feltkamp (2005) 'A bayesian approach to uncertainty aversion.' *The Review of Economic Studies* 72(2), 449–466

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

Hertwig, Ralph, Greg Barron, Elke U Weber, and Ido Erev (2004) 'Decisions from experi-

ence and the effect of rare events in risky choice.' *Psychological science* 15(8), 534–539

Hohwy, Jakob (2013) *The predictive mind* (Oxford University Press, Oxford)

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

Keynes, John Maynard (1921) *A treatise on probability* (Courier Corporation)

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

_ (2026) 'Cognitive imprecision and stake-dependent risk attitudes.' Technical Report, NBER Working Paper 30417

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005) 'A Smooth Model of Decision Making under Ambiguity.' *Econometrica* 73, 1849–1892

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

Leibniz, Gottfried Wilhelm (1704/1765) *Nouveaux essais sur l'entendement humain* (Hachette)

Li, Yuchi, and Ferdinand M. Vieider (2026) 'Meta-analysis of incentive effects on risk-taking and delay-discounting.' *Journal of Political Economy Microeconomics, forthcoming*

Li, Yuchi, and Jilong Wu (2026) 'Ambiguity attitudes under matching and choice.' Technical Report, Ghent University Discussion Papers

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

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

Oprea, Ryan, and Ferdinand Vieider (2026) 'Information processing and the anomalies of risk and time.' Working Paper

Oprea, Ryan, and Michael Woodford (2026) 'Behavioral signatures of cognitive noise.' Technical Report, National Bureau of Economic Research

Polania, Rafael, and Todd Hare (2026) 'Efficient 'risk attiudes'.' *Preprint*

Raiffa, Howard (1961) 'Risk, Ambiguity and the Savage Axioms: Comment.' *Quarterly Journal of Economics* 75(4), 690–694

Samuelson, Paul A. (1937) 'A Note on Measurement of Utility.' *The Review of Economic Studies* 4(2), 155–161

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

Schmeidler, David (1989) 'Subjective probability and expected utility without additivity.' *Econometrica* pp. 571–587

Segal, Uzi (1987) 'The ellsberg paradox and risk aversion: An anticipated utility approach.' *International Economic Review* 28(1), 175–202

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

Tversky, Amos, and Craig R. Fox (1995) 'Weighing Risk and Uncertainty.' *Psychological Review* 102(2), 269–283

Vickrey, William (1945) 'Measuring Marginal Utility by Reactions to Risk.' *Econometrica* 13(4), 319

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

_ (2025a) 'Neuro-biological origins of prospect theory functionals.' In *Handbook of Prospect Theory,* ed. Mohammed Abdellaoui and Han Bleichrodt (Cheltenham, UK: Edward Elgar Publishing)

Vieider, Ferdinand M (2025b) 'The structure of stochastic choice.' *Working Paper*

Vieider, Ferdinand M. (2026) 'Sequential sampling of noisy information: Attention, stopping, and risk-taking.' Technical Report, Ghent University Discussion Papers

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

Wu, Keyu, Ernst Fehr, Sean Hofland, and Martin Schonger (2024) 'On the psychological foundations of ambiguity and compound risk aversion'

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

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

ONLINE APPENDIX

A Proofs

A.1 Proof of Proposition 1

Proof. From the stochastic choice equation,

$$q_{ij}(p) = \Phi\left[\Lambda_{ij}\frac{p-p_0}{\sqrt{p(1-p)}}\right], \qquad \Lambda_{ij} = \frac{\kappa_0(A_j - A_i)}{\sqrt{A_i(1+\kappa_0 A_j)^2 + A_j(1+\kappa_0 A_i)^2}}.$$

For the first part, let

$$h(p) = \frac{p-p_0}{\sqrt{p(1-p)}}.$$

Then

$$h'(p) = \frac{p(1-p_0) + p_0(1-p)}{2[p(1-p)]^{3/2}} > 0.$$

Since the standard normal density is strictly positive,

$$\operatorname{sgn}\left(\frac{\partial q_{ij}(p)}{\partial p}\right) = \operatorname{sgn}(\Lambda_{ij}) = \operatorname{sgn}(A_j - A_i).$$

If $A_j = A_i$, then $\Lambda_{ij} = 0$ and $q_{ij}(p) = 1/2$ for every $p \in (0,1)$.

For the second part, evaluating the derivative at $p = p_0$ yields

$$\left.\frac{\partial q_{ij}(p)}{\partial p}\right|_{p=p_0} = \frac{\phi(0)}{\sqrt{p_0(1-p_0)}}\Lambda_{ij}.$$

Write

$$D_{ij} = A_i(1+\kappa_0 A_j)^2 + A_j(1+\kappa_0 A_i)^2.$$

Holding A_i and κ_0 fixed,

$$\frac{\partial \Lambda_{ij}}{\partial A_j} = \frac{\kappa_0(1+\kappa_0 A_i)[\kappa_0 A_i^2 + 3\kappa_0 A_i A_j + 3A_i + A_j]}{2D_{ij}^{3/2}} > 0.$$

Thus the slope at p_0 is strictly increasing in A_j. □

A.2 Proof of Corollary (familiarity and learning)

Proof. The result follows directly from Proposition 1. Holding A_i fixed, Λ_{ij} is strictly increasing in A_j. Hence any reduction in A_j reduces its revealed ambiguity relative to i.

At $A_j = A_i$, $\Lambda_{ij} = 0$, so both revealed-ambiguity signatures vanish. Further reductions imply $A_j < A_i$ and hence $\Lambda_{ij} < 0$, reversing their direction. Under $p_0 < 0.5$, the corresponding 50–50 choice pattern reverses as well. □

B Additional results

B.1 Additional results Experiment I

Figure 7 shows the choice proportions for different risk offsets in experiment I (other effects are similar). Choice proportions react monotonically to changes in the risk probability while keeping the probability of the Ellsberg-urn, p_a, fixed. The elevation differences increase with the probability of winning, but are most pronounced when comparing The visual urn to both fraction conditions, while being absent between the fraction conditions.

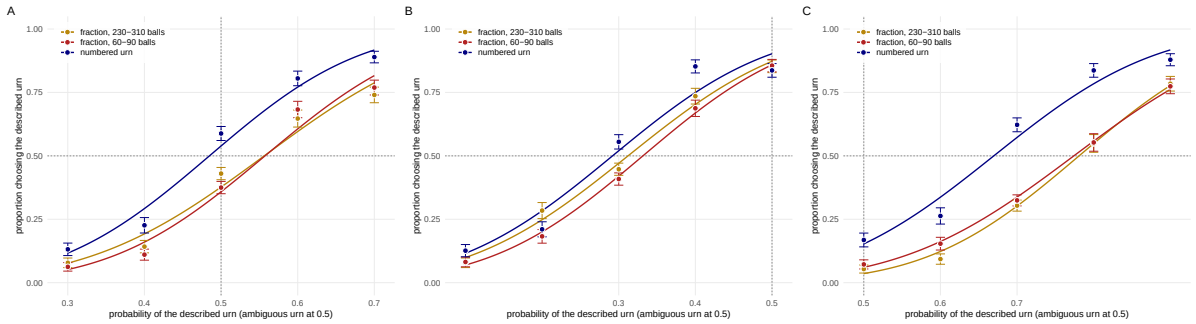


Figure 7: Risk offset effects in experiment I. Panel A: $p_a = 0.5$. Panel B: $p_a = 0.3$. Panel C: $p_a = 0.7$.

Hierarchical Probit specification. As a parametric complement to the nonparametric analyses, I estimate a hierarchical Probit model separately for each experiment, using both matched-probability and probability-offset trials. Let $y_{ns} = 1$ indicate that participant n chooses the reference urn on trial s, and let p^r_{ns} and p^a_{ns} denote the objective winning probabilities of the reference and comparison urns, respectively. The reference urn is the described urn in Experiment I and the numbered (visual) urn in Experiment II. Writing $g(n)$ for the treatment assigned to participant n, the choice probability is

$$\Pr(y_{ns} = 1) = \Phi\left(\frac{\alpha_n + \delta_n(p^r_{ns} - 1/2) + \gamma_{g(n)}(p^r_{ns} - p^a_{ns})}{\sqrt{\frac{1}{2}\left[p^r_{ns}(1-p^r_{ns}) + p^a_{ns}(1-p^a_{ns})\right]}}\right). \tag{7}$$

Here α_n governs choice elevation at matched 50–50 probabilities, δ_n governs local probability sensitivity at matched 50–50 probabilities, and $\gamma_{g(n)} > 0$ captures sensitivity to probability differences between the urns on offset trials.

On matched-probability trials, $p^r_{ns} = p^a_{ns} = p$, the denominator reduces to $\sqrt{p(1-p)}$, as in the Gaussian approximation to the theoretical choice equation. For offset trials, the denominator averages the Bernoulli variance terms of the two alternatives, rather than imposing the representation-specific variance weights of the structural model. The specification is therefore a theory-guided reduced form, not an exact structural likelihood for the offset trials.

The elevation and sensitivity coefficients vary across participants, with treatment-specific means:

$$\alpha_n = \alpha_{g(n)} + \sigma_\alpha z_{\alpha,n}, \qquad \delta_n = \delta_{g(n)} + \sigma_\delta z_{\delta,n}, \qquad z_{\alpha,n}, z_{\delta,n} \overset{\text{iid}}{\sim} \mathcal{N}(0,1).$$

I assign independent priors $\alpha_g \sim \mathcal{N}(0, 0.5^2)$, $\delta_g \sim \mathcal{N}(0,1)$, and $\gamma_g \sim \mathcal{N}^+(0, 5^2)$, where $\mathcal{N}^+$ denotes a normal distribution truncated below at zero. Both hierarchical standard deviations have Exponential(2) priors.

Hierarchical Probit estimates. As a parametric complement to the nonparametric analyses in the main text, Table 1 reports estimates from a hierarchical Probit model using both matched-probability and offset trials. The treatment-level coefficients α_t and δ_t govern choice elevation and probability sensitivity, respectively. At matched probabilities of 0.5, the latent Probit index for a participant with zero random effects is $2\alpha_t$, and its derivative with respect to the matched winning probability is $2\delta_t$.

The estimates reproduce both central treatment effects. In the visual condition, the elevation and probability-sensitivity coefficients are positive, corresponding to ambiguity aversion and ambiguity insensitivity. Both signs reverse in the fraction conditions. Relative to the visual condition, FRAC2 reduces the elevation coefficient by 0.274 and the sensitivity coefficient by 0.437, while FRAC3 reduces them by 0.256 and 0.584, respectively; all four posterior credible intervals exclude zero. The predicted dose effect appears primarily in probability sensitivity: the coefficient is 0.147 lower in FRAC3 than in FRAC2 (95% credible interval $[-0.276, -0.021]$), whereas the corresponding difference in eleva-

Table 1: Hierarchical Probit estimates, Experiment I

	Elevation coefficient α_t	Probability-sensitivity coefficient δ_t
Visual	0.101 [0.051, 0.152]	0.228 [0.132, 0.325]
FRAC2	−0.173 [−0.218, −0.128]	−0.209 [−0.300, −0.118]
FRAC3	−0.155 [−0.201, −0.110]	−0.356 [−0.446, −0.266]
FRAC2 − Visual	−0.274 [−0.341, −0.208]	−0.437 [−0.570, −0.304]
FRAC3 − Visual	−0.256 [−0.328, −0.190]	−0.584 [−0.717, −0.452]
FRAC3 − FRAC2	0.018 [−0.045, 0.080]	−0.147 [−0.276, −0.021]

Notes: Entries report posterior means and 95% credible intervals. The model is estimated using both matched-probability and offset trials. The coefficients α_t and δ_t govern choice elevation and probability sensitivity at matched probabilities of 0.5. The offset coefficients are included in estimation but omitted from the table.

tion is small and imprecisely estimated (0.018, 95% credible interval $[-0.045, 0.080]$). Thus, the hierarchical Probit estimates confirm the reversal of both canonical ambiguity signatures and the stronger probability-sensitivity effect under the more complex fraction format.

B.2 Additional results Experiment II

Figure 8 shows the choice proportions for different offsets in experiment II (other effects are similar). Note that here the visual urn is held fixed, and I thus plot offsets for the treatment urn. Choice proportions react monotonically to changes in the probability while keeping the probability of the numbered urn, p_n, fixed. The elevation differences differ relatively little, and not systematically. The Ellsberg-urn exhibits a slightly flatter slope than the two fraction conditions throughout. Notice that this is fully consistent with the model: choice stochasticity is nonmonotonic in representational ambiguity levels.

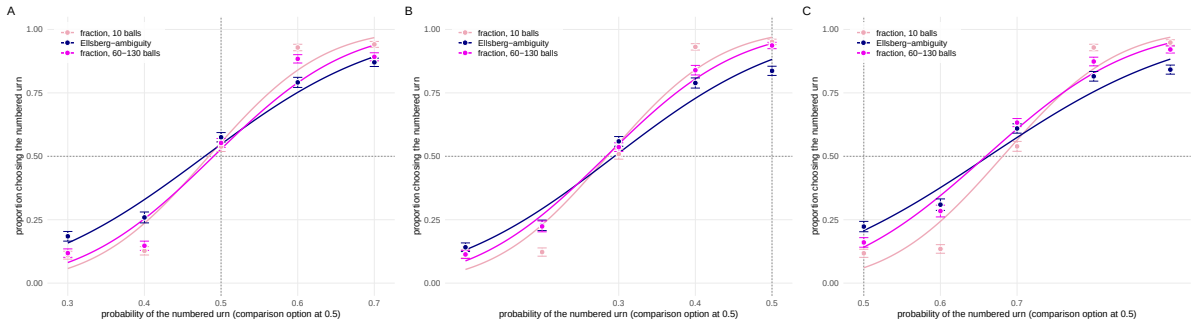


Figure 8: Risk offset effects in experiment II. Panel A: $p_n = 0.5$. Panel B: $p_n = 0.3$. Panel C: $p_n = 0.7$.

Hierarchical Probit estimates. Table 2 reports the corresponding hierarchical Probit estimates for Experiment II. The estimates reproduce the ordering observed in the

nonparametric analysis. Probability sensitivity increases monotonically from the simple fraction format ($\delta = 0.109$), through conventional Ellsberg ambiguity ($\delta = 0.191$), to the more complex fraction format ($\delta = 0.318$). The difference between FRACM and FRACS is estimated at 0.209 (95% credible interval $[0.122, 0.297]$), while FRACM also produces greater probability sensitivity than conventional ambiguity (difference 0.127, $[0.036, 0.214]$). The difference between conventional ambiguity and FRACS is positive but estimated less precisely (0.082, $[-0.002, 0.167]$).

The elevation coefficients exhibit the same ordinal pattern: 0.056 in FRACS, 0.095 under conventional ambiguity, and 0.102 in FRACM. Pairwise differences in these coefficients are estimated less precisely, however, with all 95% credible intervals including zero. The parametric analysis therefore most clearly reinforces the central sensitivity result: conventional ambiguity occupies an intermediate position between two different formats for describing objectively known risk.

Table 2: Hierarchical Probit estimates, Experiment II

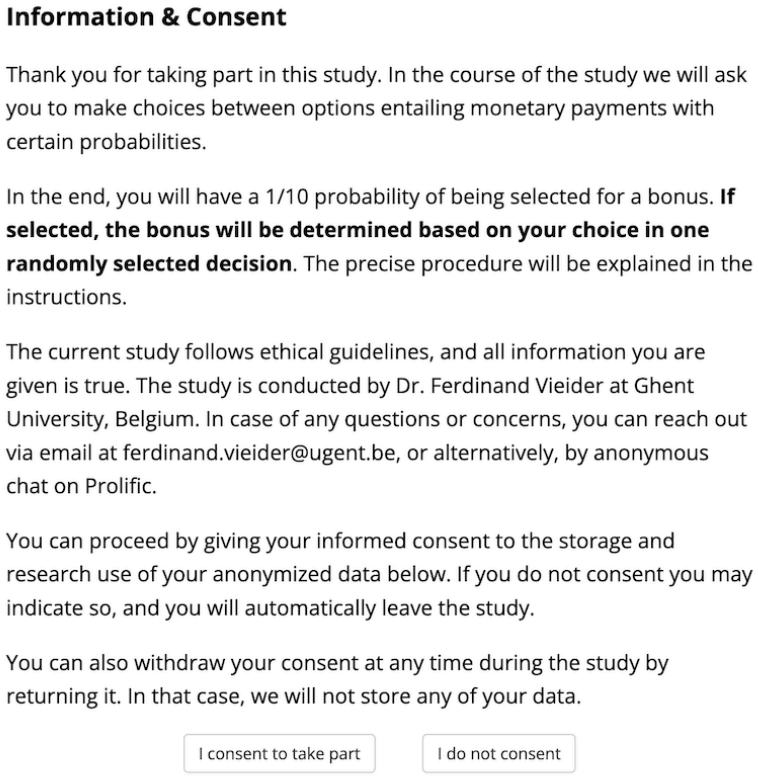
	Elevation coefficient α_t	Probability-sensitivity coefficient δ_t
FRACS	0.056 [0.023, 0.091]	0.109 [0.049, 0.170]
FRACM	0.102 [0.065, 0.139]	0.318 [0.255, 0.381]
Ellsberg	0.095 [0.059, 0.131]	0.191 [0.131, 0.252]
FRACM − FRACS	0.046 [−0.005, 0.099]	0.209 [0.122, 0.297]
Ellsberg − FRACM	−0.007 [−0.058, 0.045]	−0.127 [−0.214, −0.036]
Ellsberg − FRACS	0.039 [−0.011, 0.091]	0.082 [−0.002, 0.167]

Notes: Posterior means and 95% credible intervals in brackets. Estimates are from a hierarchical Probit including matched-probability and offset trials. The reported coefficients describe treatment-level choice elevation and probability sensitivity at matched probabilities; the offset coefficient is included in estimation but omitted from the table.

C Instructions

C.1 Instructions Experiment I (frac2)

Informed Consent



Information & Consent

Thank you for taking part in this study. In the course of the study we will ask you to make choices between options entailing monetary payments with certain probabilities.

In the end, you will have a 1/10 probability of being selected for a bonus. **If selected, the bonus will be determined based on your choice in one randomly selected decision**. The precise procedure will be explained in the instructions.

The current study follows ethical guidelines, and all information you are given is true. The study is conducted by Dr. Ferdinand Vieider at Ghent University, Belgium. In case of any questions or concerns, you can reach out via email at ferdinand.vieider@ugent.be, or alternatively, by anonymous chat on Prolific.

You can proceed by giving your informed consent to the storage and research use of your anonymized data below. If you do not consent you may indicate so, and you will automatically leave the study.

You can also withdraw your consent at any time during the study by returning it. In that case, we will not store any of your data.

I consent to take part I do not consent

Instructions frac2

Instructions

In this study you will repeatedly choose between two urns. A ball is drawn from the urn you choose. **If a ball carrying a winning number is drawn, you win £20**, else you win nothing. The prize is the same in every choice and for both urns.

One of the two urns always contains **10 balls carrying unknown numbers between 1 and 10. Some of these numbers may be present repeatedly, and some numbers may be absent**. The composition of these urns was fixed before the study began.

For this urn, **the winning numbers are numbers you choose yourself — but you will choose them at the END of the study, after all your decisions.** At the end you will rank the numbers 1 to 10 from your most preferred to your least preferred, and your top-ranked numbers will then be the winning numbers in every choice you made. **Nobody, including us, knows which numbers those will be while you are choosing: you decide them afterwards.**

Below the urn, a row of ten markers shows **how many** numbers win in that choice: if 3 markers are filled, your 3 top-ranked numbers will win.

The other urn is described by a fraction shown on screen. **The fraction gives the number of winning balls out of the total number of balls in that urn**: if it shows 69/230, then 69 of its 230 balls are winning balls. The total number of balls in this urn varies from choice to choice.

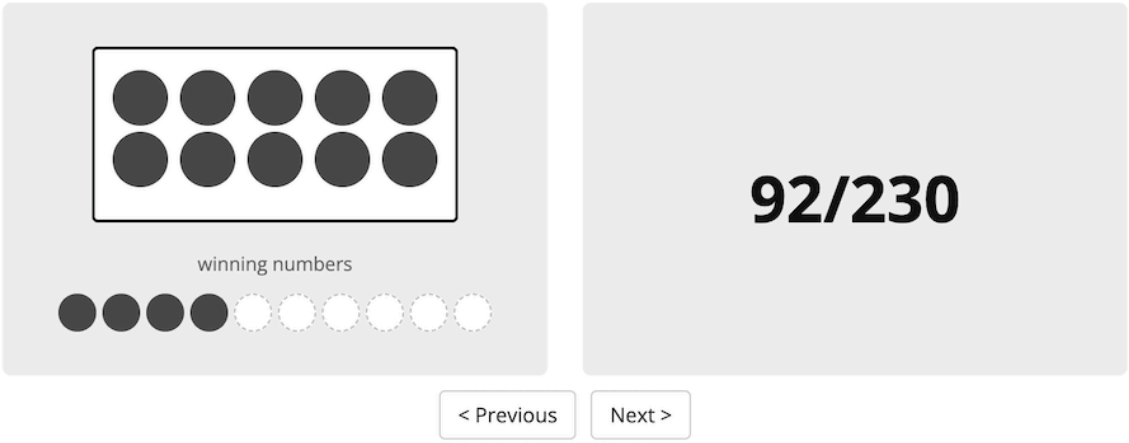
< Previous Next >

Example

Below is an example of a choice screen.

On the left is the urn with unknown numbers: 4 markers are filled, which means that **the 4 numbers you will rank highest at the end** are its winning numbers.

On the right, the urn is described by the fraction 92/230: 92 of its 230 balls are winning balls.



Incentive reminder

To repeat: the winning numbers of the unknown-numbers urn are your own — you will choose them at the very end of the study, once all your decisions are made.

One in ten participants will be selected at the end of the study for a bonus payment based on one randomly selected choice, which will be played out for real.

Before you can move on to the study, we ask you to reply to some comprehension questions. If you are unsure about any aspect of the task, please read these instructions again before proceeding.

< Previous Next >

Control questions

An urn shows 10 dark balls with unknown numbers. Which of the following is true?*

- ○ Each number from 1 to 10 appears exactly once.
- ○ The balls carry unknown numbers between 1 and 10; any combination is possible.
- ○ The urn is empty.

Which numbers count as winning numbers in the urn with unknown numbers?*

- ○ Numbers chosen at random by the computer in each choice.
- ○ Numbers I choose myself at the end of the study: my top-ranked numbers, as many as there are filled markers.
- ○ Always the numbers 1, 2 and 3.

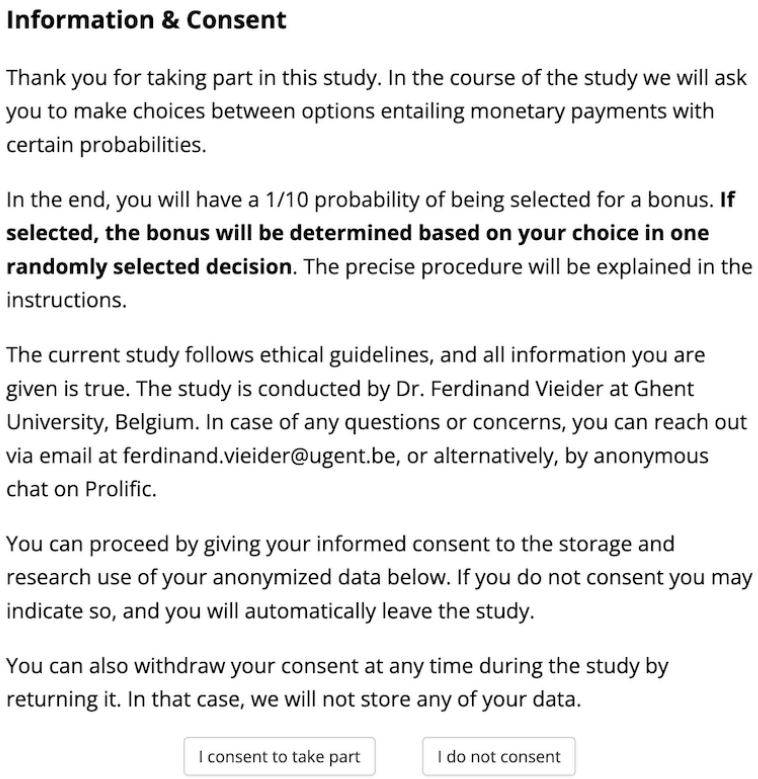
What prize do you stand to win from the two urns?*

- ○ The prize is £10 for the unknown urn, and £20 for the other urn.
- ○ The prize is £20 for the unknown urn, and £18 for the other urn.
- ○ The prize is £20 regardless of which urn I choose.

Continue

C.2 Instructions experiment II (fracS)

Informed Consent



Information & Consent

Thank you for taking part in this study. In the course of the study we will ask you to make choices between options entailing monetary payments with certain probabilities.

In the end, you will have a 1/10 probability of being selected for a bonus. **If selected, the bonus will be determined based on your choice in one randomly selected decision**. The precise procedure will be explained in the instructions.

The current study follows ethical guidelines, and all information you are given is true. The study is conducted by Dr. Ferdinand Vieider at Ghent University, Belgium. In case of any questions or concerns, you can reach out via email at ferdinand.vieider@ugent.be, or alternatively, by anonymous chat on Prolific.

You can proceed by giving your informed consent to the storage and research use of your anonymized data below. If you do not consent you may indicate so, and you will automatically leave the study.

You can also withdraw your consent at any time during the study by returning it. In that case, we will not store any of your data.

I consent to take part I do not consent

Instructions fracS

Instructions

In this study you will repeatedly choose between two urns. A ball is drawn from the urn you choose. **If a ball carrying a winning number is drawn, you win £20**, else you win nothing. The prize is the same in every choice and for both urns.

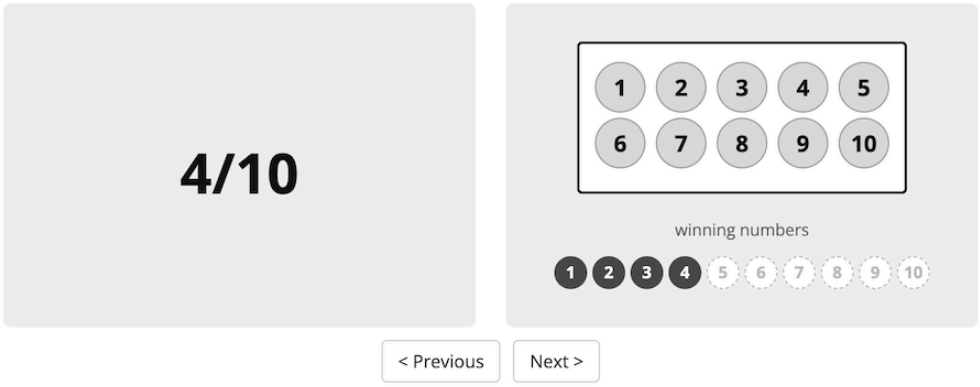
One of the two urns contains **10 balls numbered from 1 through 10, with each number occurring exactly once**. For this urn, **the winning numbers are always the first numbers in order**: if 3 numbers win, they are 1, 2 and 3. Below the urn, a row of the numbers 1 to 10 shows them as filled.

The other urn is described by a fraction shown on screen. **The fraction gives the number of winning balls out of the total number of balls in that urn**: if it shows 3/10, then 3 of its 10 balls are winning balls.

< Previous Next >

Example

Below is an example of a choice screen.



Incentive reminder

One in ten participants will be selected at the end of the study for a bonus payment based on one randomly selected choice, which will be played out for real.

Before you can move on to the study, we ask you to reply to some comprehension questions. If you are unsure about any aspect of the task, please read these instructions again before proceeding.

< Previous Next >

Control questions

An urn is described by the fraction 7/10. What does this mean?*

- ○ The urn contains 10 balls, of which 7 are winning balls.
- ○ The urn contains 7 balls, of which 10 are winning balls.
- ○ The urn contains 17 balls in total.

In the urn with the balls numbered 1 to 10, which numbers win?*

- ○ Numbers I choose myself at the end of the study.
- ○ Always the first numbers in order: if 3 numbers win, they are 1, 2 and 3.
- ○ Different numbers in every choice, drawn at random.

What prize do you stand to win?*

- ○ £10 from one option and £20 from the other.
- ○ £20 from one option and £18 from the other.
- ○ £20 regardless of which option I choose.

Continue