

Endogenous Risk Attitudes*

Nick Netzer

Arthur Robson

University of Zurich

Simon Fraser University

Jakub Steiner

Pavel Kocourek

University of Zurich and Cerge-Ei

Cerge-Ei

July 22, 2021

Abstract

In a model inspired by neuroscience, we show that constrained optimal perception encodes lottery rewards using an S-shaped encoding function and over-samples low-probability events. The implications of this perception strategy for behavior depend on the decision-maker's understanding of the risk. The strategy does not distort choice in the limit as perception frictions vanish when the decision-maker fully understands the decision problem. If, however, the decision-maker underrates the complexity of the decision problem, then risk attitudes reflect properties of the perception strategy even for vanishing perception frictions. The model explains adaptive risk attitudes and probability weighting as in prospect theory and, additionally, predicts that risk attitudes are strengthened by time pressure and attenuated by anticipation of large risks.

*We have benefited from comments of Cary Frydman, Ryan Webb and various seminar audiences. Giorgi Chavchanidze and Pavel Ilinov provided excellent research assistantship. This work was supported by the Alfred P. Sloan Foundation and the NOMIS Foundation (Netzer), by SSHRC grant 435-2013-0426 (Robson) and by ERC grant 770652 (Steiner).

1 Introduction

Although economists usually take preferences as exogenous and fixed, there is compelling evidence that these change with the context. For choices over gambles, we know at least since Kahneman and Tversky (1979) that risk attitudes are not fixed: the steep part of the S-shaped utility function in prospect theory adapts to the status quo. Rabin’s (2000) paradox provides another challenge for stable risk attitudes: choices over small and large risks are best represented by distinct Bernoulli utility functions. Risk attitudes are further modulated by external factors such as time pressure or framing (e.g., Kahneman, 2011). An additional well-known anomaly involves the overweighting of small objective probability events relative to more likely events (Kahneman and Tversky, 1979).

In this paper, we explain these anomalies in a unifying way. In our model, endogenous risk attitudes and probability weighting are the consequence of constrained optimal perception of lotteries, combined with a possible misspecification of the structure of the risk.

Our decision-maker (DM) employs a noisy non-linear encoding function that maps rewards to their mental representations and samples many such representations of rewards on all lottery arms. She optimizes the perception strategy – the encoding function and the sampling frequencies of all arms – for a given distribution of decision problems. The model explains adaptive S-shaped encoding of rewards and over-sampling of small-probability events as jointly optimal. The implications of the perception strategy for behavior are subtle. As the perception data become rich and approximate full information, behavior becomes risk-neutral whenever the DM understands the structure of the risk she faces, and hence learns about it in a correctly specified model. However, the perception strategy induces non-trivial risk attitudes when the DM applies a simplifying model to the encountered risk. In that case, the encoding function takes on the role of a Bernoulli utility function, and over-sampling of small probability events translates into overweighting of small probabilities, generating risk attitudes akin to those from prospect theory. The model also makes comparative-statics predictions about the impact of the lottery stakes, time pressure and framing for behavior.

Our procedural-choice model is inspired by the literature on optimal coding from neuroscience. A risk-neutral DM chooses between a lottery and a safe option. For illustration, envision the vivid example from Savage (1954) in which choosing the lottery represents purchase of a convertible car, the enjoyment of which depends on random weather. The DM knows the probabilities of the lottery arms (the different weather conditions), observes the value of the safe option (the price of the car), but faces a friction in processing the rewards (the weather-dependent enjoyments). She learns about the reward vector by sampling signals, from her own memory, experience of others, provided by the car dealer, etc. Each

signal is a reward of a respective lottery arm encoded via a non-linear encoding function with a finite range, then perturbed by additive Gaussian noise. After she has observed the perception data, the DM forms a maximum-likelihood or Bayesian estimate of the value of the lottery given her own perception strategy, and then takes the a posteriori optimal decision.

Choice of the perception strategy is a specific form of an attention-allocation problem. Our DM is akin to an engineer who measures a physical input by reading off the position of a needle on a meter. She can choose the measurement function that maps the physical input to the needle position. If the needle position has a stochastic component, then she can increase the precision of her measurement for a specific range of inputs by making the measurement function steep in this range. Further, our DM can allocate attention to a specific lottery arm by sampling it frequently.

We analyze the limit of rich perception data, motivated by two considerations. First, while human perception is inherently noisy, imprecisions can be partially mitigated by collecting more data, in particular when the stakes are large. Our limiting results are a useful approximation when the noise is finite and small relative to the stakes of the decision problem. Second, the limit is tractable.

We first prove that the expected loss from misperception, relative to choice under complete information, is approximately the mean squared error in perception of the lottery value, integrated over all decision problems in which the lottery value ties with the safe option. The conditioning on ties arises endogenously. Accuracy of perception has instrumental value for choice, and choice is trivial except where the values of two options are nearly equal, given that information is nearly complete.

We then derive the perception strategy for which the mean squared error over ties goes to zero most quickly as perception data become rich. We show, for unimodal symmetric reward densities, that an S-shaped encoding function and over-sampling of small-probability lottery arms are jointly optimal.

The DM chooses the encoding function to be steep near the modal rewards and flatter towards the tails of the reward distribution. She thus perceives the reward values typical for her environment relatively precisely, at the expense of precision at the tail rewards. Conditioning on ties induces a statistical association between tail rewards and small-probability arms, because tail rewards from large probability arms typically result in decision problems with distinct values, rather than in approximate ties. Thus, among problems with approximately equal value pairs, the S-shaped encoding function fails to reliably measure the rewards from the low-probability arms often relative to the high-probability rewards. The DM with an S-shaped encoding function therefore struggles to estimate the rewards

from low-probability arms in relatively many decision problems that are close to a tie, since these are tail rewards relatively often. It is optimal to compensate for this by over-sampling unlikely arms. To illustrate, consider the decision whether to take a flight. The DM may (optimally) struggle to comprehend the consequences of a low-probability aviation accident, and hence her attention allocated to this extreme contingency is optimally large relative to its probability.

We then turn to the behavioral consequences of the perception strategy. If the DM learns from the perception data in a correctly specified model, then the perception strategy affects precision of perception rather than risk attitudes. For instance, a locally steep encoding function translates into locally precise decoding of the values rather than into high marginal utility. Similarly, oversampling of an arm translates for a well-specified DM into increased precision of the arm's perceived reward rather than in the increased subjective probability of the arm. We point out, though, that if the DM interprets perception data in a naturally misspecified model, then a tight connection between the perception strategy and risk attitudes arises.

To illustrate the main idea, consider again the engineer who observes the needle position on her meter and knows that the position is a non-linear function of the measured input. But now assume that the needle trembles due to stochasticity of the input. If the engineer correctly understands that the input is stochastic, then she inverts each observed needle position to obtain the corresponding input value, thus eventually learning the true input distribution. But what if the engineer incorrectly anticipates a deterministic input and attributes the tremble of the needle to zero-mean measurement noise? Such an engineer must conclude that the deterministic input corresponds to the average needle position. Her input estimate is the certainty equivalent of the input distribution under a Bernoulli utility function equal to the meter's non-linear measurement function.

Our results on the behavioral implications of the perception strategy are analogous to the plight of the misspecified engineer. For simplicity, consider a DM who incorrectly anticipates a riskless lottery that pays the same reward on all its arms. Like the engineer who incorrectly anticipates a deterministic input, this DM estimates a single reward value, the perturbed encoding of which supposedly generates her perception data. When the noise of the encoding is additive Gaussian, then the maximum-likelihood estimate of the encoded value of this single reward is the average of all the observed signals. As the sample size diverges, the estimate converges to a convex combination of the encoded values of the arms' true rewards, where the weight on each arm is its sampling frequency. Hence, the DM's estimate of the lottery value converges to the certainty equivalent of the lottery evaluated with a Bernoulli utility function equal to the encoding function and subjective probabilities equal to the

sampling frequencies.

We emphasize that these risk attitudes predict behavior but do not reflect preferences in a welfare sense. The DM displays non-degenerate risk attitudes as a consequence of her misspecification bias. Had she anticipated a risky lottery and employed the correctly specified model, she would asymptotically learn the true lottery and make risk-neutral choices. In the presence of misspecification, however, choice remains distorted even under rich perception data.

We provide two extensions that bridge the gap between the extreme cases of a correctly specified DM who anticipates all possible risk and a misspecified DM who anticipates no risk at all. In our first approach, the DM is aware that she may face risk but uses a coarse partitional model of the true state space, much like Savage's (1954) decision-maker employing a small-world model of the grand world. The finest partition corresponds to the correctly specified DM, while the coarsest partition corresponds to the DM who anticipates no risk. There are various reasons why a DM might employ a coarse model. She might have evolved in a simple environment and the complexity of the environment might have increased, making previously payoff-irrelevant contingencies relevant, without the DM adapting to the change. For instance, financial experts may have access to rich data on returns but may estimate expected performance of an asset in a coarse model that fails to include all economic contingencies. Alternatively, the DM might have been framed to believe that the decision problem involves less risk than it does (by a car dealer, for example).

We find that, in the limit of nearly complete information, the coarse DM makes risk-neutral choices whenever she faces risk that is measurable with respect to her partition. But, whenever she faces a lottery that is not measurable with respect to her partition, she makes a biased choice even as her perception data become rich. She treats the lottery as if she had risk-attitudes implied by her perception strategy towards those elements of the risk that she does not comprehend, and is risk-neutral with respect to those elements of the risk that she does comprehend.

In his discussion of small-world models, Savage (1954) makes normative arguments for why the coarse representation of the complex grand world should assign subjective values to the elements of the state space partition that are correct averages of the true rewards within each element. Our approach departs from Savage in that we explicitly model the process of learning about rewards. We argue that the DM is unlikely to learn the correct average rewards for each element of her partition. If she learns within the small-world model, then, instead of the average reward, her estimate converges to the certainty equivalent under her encoding function and subjective probabilities equal to her sampling frequencies.

In our second approach, the DM anticipates some risk but finds large risks unlikely.

We formalize this by taking a joint limit in which perception data become rich and the prior reward distribution gradually concentrates on the set of riskless lotteries. Perception distortions therefore remain large relative to the level of perceptual discrimination required in typical decision problems. We find risk attitudes akin to those of the DM who does not anticipate risk at all. We then study comparative statics of these risk attitudes by varying the relative speed at which the two limits are taken. Within the parametrization we examine, choice becomes risk-neutral when the DM anticipates large risk a priori. In the context of Rabin’s (2000) paradox, this implies that framing a decision problem as one which features high risk attenuates the DM’s risk preferences. On the contrary, the DM becomes risk-neutral when she collects enough data. Thus, the model predicts that risk attitudes are induced under time pressure, mirroring the observation of Kahneman (2011) that prospect theory applies to fast instinctive decisions rather than to slow deliberative choices.

We build on a rich literature in neuroscience and economics. To our knowledge, our paper is the first to jointly explain adaptive S-shaped utility and probability weighting in a unifying framework.

Our work derives ultimately from psychophysics, a field that originated in Fechner’s (1860) study of stochastic perceptual comparisons based on Weber’s data. We rely on the modeling framework of Thurstone (1927) who hypothesized that perception is a Gaussian perturbation of an encoded stimulus. A large literature in brain sciences and psychology views perception as information processing via a limited channel and studies the optimal encoding of stimuli for a given channel capacity. Laughlin (1981) derives and tests the hypothesis that optimal neural encoding under an information-theoretic objective encodes random stimuli with neural activities proportional to their cumulative distribution values. This implies S-shaped encoding for unimodal stimulus densities.¹

In economics, S-shaped perception of rewards has also been derived as the constrained optimal encoding of rewards that are perceived with noise (see, among others, Friedman (1989), Robson (2001), Rayo and Becker (2007), and Netzer (2009)). These models mostly study choices over riskless prizes and thus, unlike the S-shaped value function from prospect theory, the derived encoding functions are not directly relevant to choices over gambles. Indeed, encoding functions are often interpreted as hedonic anticipatory utilities rather than as Bernoulli utilities in this literature.

Neuroscience studies encoding adaptations under various optimization objectives such as maximization of mutual information between the stimulus and its perception, maximization of Fisher information, or minimization of the mean squared error of perception (see

¹See Attneave (1954) and Barlow (1961) for early contributions and Heng et al. (2020) for more recent work. Woodford (2020) provides a review of psychophysics from an economic perspective.

e.g. Bethge, Rotermund, and Pawelzik (2002) and Wang, Stocker, and Lee (2016)). Economics can help here by providing microfoundations for the most appropriate optimization objective for perceptions related to choice. Netzer (2009) studies maximization of the expected chosen reward, an objective rooted in the instrumental approach of economics to information. Schaffner et al. (2021) report that the optimal encoding function as in Netzer provides a better fit to neural data than do encodings derived under competing objectives. In a model that differs in details concerning the perception friction, we extend Netzer’s instrumental approach to choices over gambles, finding a connection to one of the above reduced-form objectives. That is, in the limit with rich perception data, maximization of the expected chosen reward is equivalent to the minimization of the expected mean squared error in the perceived lottery value, where the expectation is over all marginal comparisons in the statistical environment. We show that this conditioning on marginal comparisons implies optimal oversampling of low-probability contingencies; an effect that would not arise under reduced-form objectives that maximize unconditional measures of precision.²

Three recent papers study risk attitudes stemming from reward encoding in the presence of noise. Khaw et al. (2018) show theoretically and verify experimentally that exogenous logarithmic stochastic encoding and Bayesian decoding generates risk attitudes in an effect akin to reversion to the mean. Frydman and Jin (2019) and Juechems et al. (2021) allow for optimal encoding of the lottery reward and show both theoretically and experimentally that this encoding adapts to the distribution of the decision problems and that the adaptation affects choice. Relative to these papers, we analyze endogenous perception biases of probabilities alongside with perception of rewards. We also differ in the proposed source of distortions. The discussed models assume well-specified learning, and thus they approximate frictionless benchmark when noise is small relative to stakes.

We focus on the limit of small encoding noise. Besides being a useful approximation for large-stake problems, the limit facilitates tractability and allows us to jointly optimize encoding and sampling for general lotteries. Our focus on vanishing noise also uncovers a novel connection between coding and behavior. While the impact of coding on behavior must necessarily vanish when the decoding model is well-specified, as in the previous literature, the implications for behavior remain substantial if the cognitive model used for decoding oversimplifies the structure of the risk. We connect perception distortions to classical repre-

²Herold and Netzer (2015) derive probability weighting as the optimal correction for an exogenous distortive S-shaped value function, and Steiner and Stewart (2016) find probability weighting to be an optimal correction for naive noisy information processing. The present paper derives both S-shaped encoding and small-probability over-sampling in a joint optimization. Robson et al. (2021) is a dynamic version of Robson (2001) and Netzer (2009), that captures low-rationality, real-time adaptation of an hedonic utility function used to make ultimately deterministic choices.

sentations using a Bernoulli utility function and subjective probability weights.

Salant and Rubinstein (2008) and Bernheim and Rangel (2009) provide a revealed-preference theory of the behavioral and welfare implications of frames – payoff-irrelevant aspects of decision problems. We provide an account of how a specific frame – anticipation of the risk structure – affects choice and welfare. As in Kahneman, Wakker, and Sarin (1997), our model implies a distinction between decision and welfare utilities. In the case of the misspecified DM, the gap between the decision utility that she anticipates the lottery to pay and welfare utility – the true expected lottery reward – may be large. Our model facilitates an analysis of systematic mistakes in decision making as outlined in Koszegi and Rabin (2008) and, for the case of framing effects, Benkert and Netzer (2018).

We apply the statistical results of Berk (1966) and White (1982) on asymptotic outcomes of misspecified Bayesian and maximum-likelihood decoding of perception data, respectively. The recent concept of Berk-Nash equilibrium in Esponda and Pouzo (2016) is defined as a fixed point of misspecified learning. This has motivated a renewed interest in misspecification across economics. Heidhues et al. (2018) characterize a vicious circle of overconfident learning, Molavi (2019) studies the macroeconomic consequences of misspecification, Frick et al. (2021) rank the short- and long-run costs of various forms of misspecification, and Eliaz and Spiegel (2020) focus on political-economy consequences of misspecification. We study the interplay of encoding and misspecified decoding of rewards.

2 Decision Process

The DM chooses between a safe option of value s and a lottery with arms $i = 1, \dots, I$, $I \geq 1$, where each arm i has a positive probability p_i and pays reward $r_i \in [\underline{r}, \bar{r}]$, where $\underline{r} < \bar{r}$ are arbitrary bounds. For the sake of simplicity, we fix the vector of arm probabilities and let the DM observe it frictionlessly. We relax this assumption below. The lottery rewards and the safe option value are generated randomly. The DM observes the exact value of the safe option but faces frictions in the perception of the lottery rewards. We let $\mathbf{r} = (r_i)_i \in [\underline{r}, \bar{r}]^I$ denote the vector of the rewards, and since the vector of probabilities is fixed, we identify \mathbf{r} with $(p_i; r_i)_i$ and refer to it as to a *lottery*. The pair (\mathbf{r}, s) is the *decision problem*.

The goal of the DM is to choose the lottery if and only if its expected value $r = \sum_i p_i r_i$ exceeds s . This risk-neutrality with respect to rewards is an implicit assumption on the units of measurement in which the rewards are expressed. For instance, the rewards might be an appropriate concave function of monetary prizes if the DM chooses among monetary lotteries and money has diminishing returns.

The DM estimates the unknown lottery \mathbf{r} from a sequence of n signals, where each signal

is a monotone transformation of one of the arm rewards perturbed with additive noise: she observes signals $x_k = (\hat{m}_k, i_k)$, $k = 1, \dots, n$. We refer to the first component, \hat{m}_k , as to the *perturbed message*. The second component, i_k , indicates the arm the message \hat{m}_k pertains to. Each perturbed message is generated by encoding the reward r_{i_k} of arm i_k into *unperturbed message* $m(r_{i_k})$ and by perturbing it to $\hat{m}_k = m(r_{i_k}) + \hat{\varepsilon}_k$, where the noise term $\hat{\varepsilon}_k$ is independently and identically distributed (iid) standard normal. Each sampled arm i_k is one of the lottery arms $i = 1, \dots, I$, iid across k with positive probabilities π_i . The function $m : [\underline{r}, \bar{r}] \rightarrow [\underline{m}, \bar{m}]$ is strictly increasing and continuously differentiable; we refer to it as to the *encoding function*. We dub the π_i as *sampling frequencies* and refer to $(m(\cdot), (\pi_i)_i)$ as the *perception strategy*. The size of the sample, n , is exogenous.

After she has observed the n signals, the DM forms an estimate q_n of the lottery's value and chooses the lottery if and only if $q_n > s$. We consider both maximum-likelihood (ML) and Bayesian estimators, $q_n = q_n^{ML}$ or $q_n = q_n^B$. In the first case, the DM is endowed with a compact set $\mathcal{A} \subseteq [\underline{r}, \bar{r}]^I$ of lotteries she anticipates and concludes that she has encountered the lottery

$$\mathbf{q}_n^{ML} \in \arg \max_{\mathbf{r}' \in \mathcal{A}} \prod_{k=1}^n \varphi(\hat{m}_k - m(r'_{i_k}))$$

that maximizes the likelihood of the observed signals, where φ is the standard normal density. Finally, she sets $q_n^{ML} = \sum_i p_i q_{in}^{ML}$.³ In the second case, the DM is endowed with a prior belief over \mathcal{A} and sets $q_n^B = \mathbb{E}[\sum_i p_i r_i \mid (x_k)_{k=1}^n]$ as the posterior expected lottery value. Both these specifications will lead to same conclusions as n diverges since the impact of the DM's prior becomes negligible in this limit.

We study decision-makers who may employ simplifying models of risk in the spirit of the *small world* of Savage (1954). The DM anticipates, rightly or wrongly, distinctions among some of the lottery arms to be payoff-irrelevant. Let \mathcal{P} be a partition of the set of all the lottery arms $\{1, \dots, I\}$. The DM anticipates that $r_i = r_j$ for all pairs of arms $i, j \in J$ that belong to a same element J of the partition \mathcal{P} . That is, she anticipates lotteries from a set

$$\mathcal{A}_{\mathcal{P}} = \left\{ \mathbf{r} \in [\underline{r}, \bar{r}]^I : r_i = r_{i'} \text{ for all } i, i', J \text{ such that } i, i' \in J, J \in \mathcal{P} \right\}. \quad (1)$$

For instance, if $\mathcal{P} = \{\{1, \dots, I\}\}$ is the coarsest partition, then the DM anticipates only degenerate lotteries that pay a same reward at all their arms. We refer to such lotteries as *riskless* and call other lotteries *risky*. If, on the other extreme, $\mathcal{P} = \{\{1\}, \dots, \{I\}\}$ is the finest partition, then the DM anticipates that any reward vector is possible and $\mathcal{A}_{\mathcal{P}} = [\underline{r}, \bar{r}]^I$.

³The maximum-likelihood estimate exists since \mathcal{A} is compact. It is unique for the specifications below.

3 Optimal Perception

The perception strategy needs to adapt to the prevailing statistical circumstances if it is to allocate attention efficiently. An increase of the sampling frequency of an arm increases the DM’s attention to its reward, but reduces attention to the rewards on other arms. Similarly, making the encoding function steep in a neighborhood of a reward value reduces noise in this neighborhood but entails increased noise elsewhere.

We denote the partition \mathcal{P} the DM employs during the adaptation stage by \mathcal{J} . That is, the DM anticipates lotteries from $\mathcal{A}_{\mathcal{J}}$ where each element of the partition \mathcal{J} specifies a set of lottery arms that the DM deems as payoff-equivalent. Since the distinction between arms in each $J \in \mathcal{J}$ is redundant, we treat J as an index of an arm, refer to the rewards at arms $i \in J$ simply as r_J , and model the whole lottery $\mathbf{r} = (r_J)_{J \in \mathcal{J}}$ as having $|\mathcal{J}|$ arms, each with probability $p_J = \sum_{i \in J} p_i$. A perception strategy consists of the increasing encoding function $m(\cdot)$ and interior sampling frequencies $(\pi_J)_J \in \Delta(\mathcal{J})$. In Section 4, we will use the notation to study the DM whose model is a misspecified small-world model of the grand world; for instance because the world became more complex after adaptation but before choice.

The DM optimizes her perception strategy ex ante for a given distribution of the decision problems. Specifically, the rewards r_J are iid with a continuous density h , and the safe option s is drawn from a continuous density h_s independently of the lottery rewards; both densities have supports $[\underline{r}, \bar{r}]$.⁴ We characterize the expected loss for general perception strategies for diverging n in the next subsection and then solve for the loss minimizing strategy in Subsection 3.2.

3.1 Objective

We take the number n of signals to be large and abstract in this section from uncertainty over the number of perturbed messages sampled for each arm and from divisibility issues. That is, we suppose the number of messages sampled for each arm $J \in \mathcal{J}$ is precisely $\pi_J n$. We let $m_{J,n}$ be the average of the perturbed messages sampled for arm J . Then, $m_{J,n} - m(r_J)$ is normally distributed with mean 0 and variance $1/(n\pi_J)$, for each given value of r_J . Since the signal errors are Gaussian, the vector of average perturbed messages, $\mathbf{m}_n = (m_{J,n})_J$, is a sufficient statistic for the lottery rewards.

For $z \in \{B, ML\}$, let q_n^z be the Bayesian and ML estimator of the lottery value and let

$$L^z(n) = \text{E} [\max\{r, s\} - \mathbb{1}_{q_n^z > s} r - \mathbb{1}_{q_n^z \leq s} s]$$

⁴Since s may have a distinct density from that of r_J , the safe option may be, for instance, the value of an alternative lottery with each of its rewards drawn from h .

be its ex ante expected loss relative to choice under complete information; the expectation is over \mathbf{r} , s and q_n^z .

Proposition 1. *Assume the encoding function m is continuously differentiable, the reward density h is continuous, and the density of the safe option h_s is continuously differentiable. Then, the Bayesian and ML estimators generate the same asymptotic loss*

$$\lim_{n \rightarrow \infty} nL^z(n) = \frac{1}{2} \mathbb{E} \left[h_s(r) \sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J)} \right] \text{ for } z \in \{B, ML\}. \quad (2)$$

See proof in Appendix A. The difference between the Bayesian and ML estimators asymptotically vanishes because the prior information has negligible impact on the Bayesian DM who receives many signals. The limit loss characterization in (2) has an intuitive interpretation. It is the mean squared error in the perception of the lottery value (multiplied by n and divided by two) integrated over all decision problems in which the true lottery value r ties with s . The conditioning on the tie arises because the likelihood of large perception errors vanishes with increasing n , and small perception errors distort choice only in decision problems in which an approximate tie arises. In the limit, the set of decision problems in which perception errors have nontrivial behavioral consequences approaches the set of problems with exact ties.

To understand the relevance of the squared error for loss, fix the true and perceived lottery values to be r and q_n^z , respectively. The perception error distorts choice and causes loss if and only if the safe option s attains a value between r and q_n^z . When n is large, and hence the error is small, this occurs with approximate probability $h_s(r)|q_n^z - r|$. Conditional on the choice being distorted like this, the expected loss is approximately $|q_n^z - r|/2$ since s is approximately uniformly distributed between r and q_n^z . Hence the overall loss over all s is approximately $h_s(r)(q_n^z - r)^2/2$. Taking the expectation of $(q_n^z - r)^2$ with respect to q_n^z yields the mean squared error of the value estimate for a given lottery, and taking the expectation with respect to the lottery gives an average over all decision problems with a tie, $s = r$.

To understand the expression in (2) in detail, consider the maximum-likelihood estimator (MLE); the Bayesian estimator differs only by a negligible term. The MLE of the reward r_J is

$$q_{J,n}^{ML} = m^{-1}(m_{J,n}), \quad (3)$$

and, since $m_{J,n} \sim \mathcal{N}(m(r_J), 1/(n\pi_J))$, the mean squared error of $q_{J,n}^{ML}$ is approximately

$1/(n\pi_J m'^2(r_J))$.⁵ The mean squared error of the value estimate is then approximately

$$\frac{1}{n} \sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J)}.$$

The mean squared error and hence the loss go to zero as $n \rightarrow \infty$, for any perception strategy. However, the perception strategy influences how fast the loss vanishes. Motivated by the characterization from (2), we define the *information-processing problem* as follows.

$$\min_{m'(\cdot) > 0, (\pi_J)_J > 0} \mathbb{E} \left[\sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J)} \mid r = s \right] \quad (4)$$

$$\text{s.t.:} \quad \int_{\underline{r}}^{\bar{r}} m'(\tilde{r}) d\tilde{r} \leq \bar{m} - \underline{m} \quad (5)$$

$$\sum_{J \in \mathcal{J}} \pi_J = 1. \quad (6)$$

The objective in (4) equals the asymptotic loss characterized in (2), up to a factor that is independent of the perception strategy.⁶ We let the DM control the derivative $m'(\cdot)$ and restrict it to be positive – this restricts the encoding function to be increasing and differentiable. Constraint (5) is implied by the finite range of the encoding function – the encoding function cannot be steep everywhere. Constraint (6) together with the restriction to positive sampling frequencies requires $(\pi_J)_J$ to be a probability distribution over \mathcal{J} ; the DM must also treat sampling frequencies as a scarce resource.

3.2 Optimization

Reward density $f(x)$ is *unimodal* and *symmetric* around the mode $r_m = (\underline{r} + \bar{r})/2$ if it is strictly decreasing on $(r_m, \bar{r}]$ and $f(r_m + x) = f(r_m - x)$ for all x . (We assume symmetry of the reward density because unimodality is preserved by summation for symmetric densities but not in general.) We say that the perception strategy $(m(\cdot), (\pi_J)_J)$ is optimal if $(m'(\cdot), (\pi_J)_J)$ solves the information-processing problem.

Proposition 2. *There is a unique optimal perception strategy. If the densities h and h_s are continuous, unimodal and symmetric around r_m , then*

⁵Equation (3) holds if $m_{J,n} \in [\underline{m}, \bar{m}]$. If $m_{J,n} < \underline{m}$ or $m_{J,n} > \bar{m}$, then the MLE of r_J is \underline{r} or \bar{r} , respectively. Note, that $P(m_{J,n} \in [\underline{m}, \bar{m}]) \rightarrow 1$ as n diverges for all $r_J \in (\underline{r}, \bar{r})$.

⁶This factor is two divided by the ex ante likelihood of tie.

1. *the optimal encoding function is S-shaped: It is convex below and concave above r_m , and*
2. *the DM over-samples the low-probability arms: For any two arms J, J' such that $p_J < p_{J'}$, $\frac{\pi_J}{p_J} > \frac{\pi_{J'}}{p_{J'}}$. In particular, when the lottery has two arms, then $\pi_J > p_J$ for the arm with probability $p_J < 1/2$ and vice versa for the high-probability arm.*

Since we have allowed the DM to condition her perception strategy on the partition \mathcal{J} and the arm probabilities $(p_J)_J$, the optimal encoding function depends on these. However, we show in Appendix B.3 that both claims of the proposition extend to a setting in which the DM has incomplete information about the partition and the probabilities when she chooses the encoding function, and she optimizes the sampling frequencies at the interim stage after she observes the partition and the probabilities.

The proof of Proposition 2 in Appendix B follows from the first-order conditions. The outline is as follows. Let

$$h_J(\tilde{r}) = h(\tilde{r}) \frac{\mathbb{E}[h_s(r)|r_J = \tilde{r}]}{\mathbb{E}[h_s(r)]}$$

be the density of the reward r_J at the arm J conditional on a tie between the lottery value and the safe option (the expectations are over \mathbf{r}). The first-order condition for the slope $m'(\tilde{r})$ of the encoding function is

$$2 \sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J m'^3(\tilde{r})} h_J(\tilde{r}) = \lambda \tag{7}$$

for each reward value \tilde{r} , where λ is the shadow price of the constraint (5). The left-hand side of (7) is the marginal benefit of the increase in the slope $m'(\tilde{r})$ at the reward value \tilde{r} . Such an increase reduces the DM's mean squared error in her perception of the lottery value if the reward r_J attains the value \tilde{r} at one of the arms $J \in \mathcal{J}$. This marginal reduction affects her choice if the value of the lottery r ties with s . Each summand on the left-hand side is proportional to the marginal reduction of the MSE in the perception of the lottery value multiplied by the likelihood that $r_J = \tilde{r}$ and that $r = s$. The constraint (5) implies that, at the optimum, the marginal benefit of a slope increase is equal across all reward values \tilde{r} .

We show in the Appendix B.1 that the density of the reward conditional on a tie is, for each arm, unimodal with the same mode as the unconditional reward density. The first-order condition (7) then implies that the optimal slope is proportional to a monotone transformation of a sum of unimodal functions that all have their maxima at the unconditional reward

mode,

$$m'(\tilde{r}) \propto \left(\sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J} h_J(\tilde{r}) \right)^{\frac{1}{3}}, \quad (8)$$

establishing Claim 1 of the Proposition.

Let us now turn to Claim 2. We show in Appendix B.2 that the DM wishes to over-sample arms about which she expects to be poorly informed. When optimizing at the ex ante stage, the DM conditions on the event of a tie because a marginal change of the sampling frequency affects choice only at ties. Conditional on a tie, the rewards at different arms are no longer identically distributed. The condition $\sum_J p_J r_J = s$ is relatively uninformative about the low-probability rewards, and hence the posterior distributions at ties of these are more spread-out compared to the high-probability rewards (see Appendix B.1). Recall that the DM measures reward r_J relatively poorly if the slope $m'(r_J)$ is low. Because m is relatively flat at tail rewards, the DM at the ex ante stage expects to be poorly informed about the rewards of low-probability arms and compensates by over-sampling these arms.⁷

Optimal over-sampling arises from our microfoundation of the optimization objective (4). Had the DM minimized the unconditional mean squared error, the effect would not arise because, unconditionally, all rewards are identically distributed. By taking the instrumental perspective that focuses on the payoff consequences of the perception errors in choice problems, we obtain an objective that conditions on ties and induces over-sampling as the optimal adaptation.

Proposition 2 generalizes Netzer (2009). When $|\mathcal{J}| = 1$, then our DM chooses between two riskless rewards, r and s . Both Netzer and we find that when r and s are independently drawn from a same density h , then the slope of the optimal encoding function is proportional to $h^{2/3}(r)$. To see this in our framework, note that the reward density conditional on a tie is proportional to $h^2(r)$ for $|\mathcal{J}| = 1$ and the result then follows from (8).⁸

In the absence of further frictions, choice approaches that under complete information as the number of signals diverges. In the next section, we allow for the possibility that the DM's model is misspecified: some of the lottery arms that she deems to be payoff-equivalent may differ in their rewards. We find that the DM who applies a simplified model of risk exhibits risk attitudes dictated by properties of the perception strategy.

⁷This argument relies on rewards at all arms being encoded with the same encoding function.

⁸The perception friction assumed in Netzer differs from that assumed here in technical details. Netzer studies the limit of increasingly fine partitions of the reward space, whereas we take the limit of vanishing additive encoding noise.

4 Behavior

The implications of the perception strategy for behavior depend on the DM’s degree of understanding of the risk. Consider the example from Savage (1954) mentioned earlier. The DM is contemplating the purchase of a convertible car for price s . The payoff from the purchase depends on the random weather; it is r_1 if the car is driven in sunny conditions and r_2 for rainy conditions. The upcoming weather is unknown, making the purchase a binary lottery. Let the probabilities of either weather type be one half.

The DM learns the values of r_1 and r_2 by sampling n signals. For each $k = 1, \dots, n$, she observes the weather $i_k \in \{1, 2\}$ and a message $\hat{m}_k = m(r_{i_k}) + \hat{\varepsilon}_k$ where m is the encoding function and the $\hat{\varepsilon}_k$ are iid standard normal. The sampling frequency of each weather type is one half, thus matching the actual probabilities. Each signal might derive from the DM’s own experience with a convertible, the experience of her peers, information provided by the car dealer, etc.

Consider two varieties of DM – *fine* and *coarse* – who differ in their anticipation of the risk structure. The fine DM knows that the weather is payoff-relevant and hence anticipates that the purchase will lead to one of two possibly distinct reward values (r_1, r_2) . The coarse DM employs a small-world model: she anticipates, as in Savage’s example, that the convertible will lead to “definite and sure enjoyments”, so she anticipates a degenerate lottery (r, r) .

Their distinct models of risk lead the two DMs to distinct conclusions even when they employ the same perception strategy and observe identical data. The fine DM asymptotically learns $m(r_i)$ for $i = 1, 2$ from the empirical distribution of the perturbed messages, inverts the encoding function and learns the true reward pair. Her estimate of the expected reward thus converges to the true expected reward and she makes the risk-neutral choice. See the left-hand graph in Figure 1.

The coarse DM observes the same empirical signal distribution but, since she omits the weather from her model of risk, she seeks a single message which best accounts for all the observed signals. For Gaussian additive errors, the single message that maximizes the likelihood of the observed data is the empirical average message, which almost surely converges to $(m(r_1) + m(r_2))/2$. Hence, the DM’s asymptotic estimate of the reward from driving the convertible is the certainty equivalent of the risky reward under the Bernoulli utility $u(\cdot) = m(\cdot)$ and equal probabilities. See the right-hand graph in Figure 1.

There are various paths that could have led the fine and the coarse DMs to their respective decision procedures. They could have evolved in a simple environment in which all the lotteries were measurable with respect to the coarsest partition $\mathcal{J} = \{\{1, 2\}\}$ of the set of arms $\{1, 2\}$. As outlined in the previous section, they both then optimized their encoding




Figure 1: Asymptotic estimated lottery value q of the fine (left) and the coarse (right) DMs.

functions in this environment. Afterwards, their environments became more complex so they currently encounter risky lotteries with $r_1 \neq r_2$. The fine DM has refined her anticipation and understands that she may now encounter a risky lottery. In contrast, the coarse DM has not made such an adjustment and continues to anticipate riskless lotteries only. It is plausible that real-world decision-makers are sometimes not aware of all contingencies that affect their payoffs. We will provide an example of a financial investor who omits a relevant variable from her econometric model of the return below. Alternatively, both DMs may have evolved in a risky environment with partition $\mathcal{J} = \{\{1\}, \{2\}\}$ and optimized their encoding functions for risky binary lotteries. Afterwards, the coarse DM was (incorrectly) assured that her next lottery will be riskless (possibly by a strategically interested party), while the fine DM was not told this. Finally, both DM's may know that they encounter a risky lottery but the coarse DM has chosen the coarse estimation procedure due to its simplicity. The coarse procedure consists of applying the inverse encoding function to the average of all perturbed messages, whereas the procedure of the fine DM requires applying the coarse procedure to each arm separately and then computing the lottery value.

This section takes the DM's perception strategy as given; it could have been optimized as in Section 3 or established by any different process. Subsection 4.1 extends the present binary example to arbitrary lotteries and sampling frequencies. A further generalization in Subsection 4.2 considers a DM who employs an arbitrary partitional model of risk; such a DM has some but only partial awareness of the risk she faces. Subsection 4.3 then focuses on a DM who anticipates risk but believes that large differences between rewards across the lottery arms are a priori unlikely. As in the case of the DM who anticipates no risk, this generates non-trivial risk attitudes.

4.1 Surprising Risk

We characterize here the behavior of a DM who has not anticipated any risk. She anticipates a lottery from the set

$$\mathcal{A} = \{\mathbf{r} \in [\underline{r}, \bar{r}]^I : r_i = r_j \text{ for all arms } i, j\}.$$

After she encounters a lottery, she observes data generated by her perception strategy, forms the ML or Bayesian estimate of the encountered lottery from \mathcal{A} , and chooses the lottery if and only if its estimated value exceeds s . The DM learns in a misspecified model – she may encounter an unanticipated risky lottery.

To describe her behavior, we say that the DM’s choice is represented by a Bernoulli utility $u(\cdot)$ and probabilities $(\rho_i)_{i=1}^I$ if in each decision problem (\mathbf{r}, s) such that

$$\sum_{i=1}^I \rho_i u(r_i) > [<] u(s),$$

the probability that the DM chooses the lottery \mathbf{r} converges to 1 [0] as $n \rightarrow \infty$.⁹

Proposition 3. *Let the DM form the ML or Bayesian estimate of the lottery. When she anticipates a riskless lottery, the DM’s choice is represented by a Bernoulli utility equal to the encoding function, $u(\cdot) = m(\cdot)$, and probabilities given by the sampling frequencies, $\rho_i = \pi_i$ for $i = 1, \dots, I$.*

The proposition follows from the result on misspecified ML estimation by White (1982) and from Berk (1966) for the Bayesian DM. These authors let an agent observe n iid signals from a signal density and form the estimate from a set of hypothesized signal densities that may fail to include the true density. They prove that the estimate almost surely converges to the minimizer of the Kullback-Leibler divergence from the true signal density as n diverges (if the minimizer is unique).

To apply White’s and Berk’s results in our setting, consider a DM who encounters a lottery \mathbf{r} . She observes the empirical distribution of approximately $\pi_i n$ signals drawn iid from $\mathcal{N}(m(r_i), 1)$ for each arm i . Since the DM has anticipated a riskless lottery, she forms an estimate of a single unperturbed message m_n , a perturbation of which has generated the observed data. White’s and Berk’s results imply that $m^* = \lim_{n \rightarrow \infty} m_n$ almost surely minimizes the Kullback-Leibler divergence from the true signal density. For Gaussian errors, this implies $m^* = \sum_i \pi_i m(r_i)$ almost surely. Thus, the DM’s estimate of the lottery value

⁹The probability is evaluated with respect to the stochastic signal sequence $(\hat{m}_k, i_k)_{k=1}^n$.

almost surely converges to the “certainty equivalent” $m^{-1}\left(\sum_{i=1}^I \pi_i m(r_i)\right)$. See Appendix C where we prove Proposition 4; Proposition 3 is then a special case.

The behavior of the DM who anticipates a riskless lottery is governed by the sampling frequencies rather than by the true arm probabilities. Indeed, this DM believes that the true probabilities are payoff-irrelevant. In contrast, the sampling frequencies govern the proportions of her data generated by each arm and hence her estimate of the encoded riskless reward she thinks she has encountered.

Example (estimating financial returns): An investor chooses between a safe asset with return s and a risky asset with return $\rho(\mathbf{x}, \mathbf{y})$ that depends on vectors of variables \mathbf{x} and \mathbf{y} . She employs a misspecified model: she neglects the role of variables \mathbf{y} , believing that the return is $\tilde{\rho}(\mathbf{x})$ where $\tilde{\rho}(\cdot)$ is a simplified function she estimates. For example, she knows that the profit of a firm depends on prices and interest rate (\mathbf{x}) but is not aware of the firms’ entire trade exposure and neglects the role of some exchange rates (\mathbf{y}). Given \mathbf{x} , let \mathbf{y} have conditional probability $g(\mathbf{y} | \mathbf{x})$. Thus, for each fixed value of \mathbf{x} the financial asset is a lottery in which each arm represents a particular value of \mathbf{y} and is assigned a return $\rho(\mathbf{x}, \mathbf{y})$ with probability $g(\mathbf{y} | \mathbf{x})$. However, the investor attributes the variation of the return for fixed \mathbf{x} to noise and estimates $\tilde{\rho}(\mathbf{x})$ from signals $m(\rho(\mathbf{x}, \mathbf{y}_k)) + \hat{\varepsilon}_k$, $k = 1, \dots, n$. The conditional probability $\tilde{g}(\mathbf{y} | \mathbf{x})$ of observing a signal about the return for $\mathbf{y}_k = \mathbf{y}$ depends on the investor’s sampling; if her sampling is representative, then $\tilde{g} = g$. By Proposition 3, when the number of signals diverges, the investor treats the asset for each \mathbf{x} as if she were an expected-utility maximizer with Bernoulli utility $u(\cdot) = m(\cdot)$ and probability $\tilde{g}(\mathbf{y} | \mathbf{x})$ assigned to each value of y .

4.2 Coarse Decision-Maker

Next, we study a DM who considers distinctions among some but not all lottery arms payoff-relevant. She anticipates that all lottery arms in each element of the partition \mathcal{K} of the set of all arms pay the same reward. That is, she anticipates encountering a lottery from the set $\mathcal{A}_{\mathcal{K}}$ of the lotteries measurable with respect to \mathcal{K} as defined in (1).

We say that the DM’s choice has a *mixed representation* with Bernoulli utility $u(\cdot)$, probabilities $(\rho_i)_{i=1}^I$ and partition \mathcal{K} if the probability that she chooses lottery \mathbf{r} over the safe option s converges to 1 [0] in each decision problem (\mathbf{r}, s) such that

$$\sum_{J \in \mathcal{K}} \rho_J r_J^* > [<] s,$$

where $\rho_J = \sum_{i \in J} \rho_i$ and r_J^* is the certainty equivalent defined by

$$u(r_J^*) = \sum_{i \in J} \frac{\rho_i}{\rho_J} u(r_i)$$

for each $J \in \mathcal{K}$.

Let $J(i)$ be the element of partition \mathcal{K} that contains arm i . Let $p_J = \sum_{i \in J} p_i$ be the overall true probability of the arms $i \in J$. Similarly, $\pi_J = \sum_{i \in J} \pi_i$ is the overall sampling frequency for J .

Proposition 4. *The choice of the coarse DM who forms the ML or Bayesian estimate has a mixed representation with Bernoulli utility $u(\cdot) = m(\cdot)$ and arm probabilities $\rho_i = p_{J(i)} \frac{\pi_i}{\pi_{J(i)}}$ for $i = 1, \dots, I$.*

See Appendix C for the proof. In the limit, the DM chooses as if she was treating the lottery \mathbf{r} as a compound lottery in which each element J of the partition \mathcal{K} constitutes a sub-lottery and these sub-lotteries have probabilities p_J . She behaves as if she first reduced each sub-lottery to its certainty equivalent under the Bernoulli utility $u(\cdot) = m(\cdot)$ and subjective arm probabilities equal to the normalized sampling frequencies. After the reduction, she evaluates the overall lottery in a risk-neutral manner using the true probabilities of each J .

Example (estimating financial returns continued): Unlike in the previous version of this example, the investor does not observe \mathbf{x} (or \mathbf{y}) at the moment of choice. Instead, she observes a signal \mathbf{z} . Conditional on the observed value of \mathbf{z} , the asset is a lottery in which each arm represents a realization of (\mathbf{x}, \mathbf{y}) with associated return $\rho(\mathbf{x}, \mathbf{y})$ and probability $g(\mathbf{x}, \mathbf{y} \mid \mathbf{z})$. Since the investor is unaware of \mathbf{y} 's influence on the return, she forms a coarse counterpart of this lottery in which each arm represents a value of \mathbf{x} , paying $\tilde{\rho}(\mathbf{x})$ with probability $g(\mathbf{x} \mid \mathbf{z}) = \sum_{\mathbf{y}} g(\mathbf{x}, \mathbf{y} \mid \mathbf{z})$. For each value of \mathbf{x} , the investor forms the estimate of the return $\tilde{\rho}(\mathbf{x})$ given the data points $m(\rho(\mathbf{x}, \mathbf{y}_k)) + \hat{\varepsilon}_k$, where \mathbf{y}_k is drawn from $\tilde{g}(\mathbf{y}_k \mid \mathbf{x}, \mathbf{z})$. Again, $\tilde{g}(\mathbf{y} \mid \mathbf{x}, \mathbf{z})$ captures sampling. If sampling is untargeted, then $\tilde{g} = g$. After she forms the MLE $\hat{\rho}_n(\mathbf{x})$ for each value \mathbf{x} , she assigns the expected value $E[\hat{\rho}_n(\mathbf{x}) \mid \mathbf{z}]$ to the asset, where the expectation is with respect to the conditional density $g(\mathbf{x} \mid \mathbf{z})$. By Proposition 4, for each \mathbf{z} this investor values the asset as if she computed the certainty equivalent over $\rho(\mathbf{x}, \mathbf{y}) \mid (\mathbf{x}, \mathbf{z})$ for each (\mathbf{x}, \mathbf{z}) under Bernoulli utility $m(\cdot)$ and subjective probabilities $\tilde{g}(\mathbf{y} \mid \mathbf{x}, \mathbf{z})$, and then computed the risk-neutral value of the reduced lottery under the objective probabilities $g(\mathbf{x} \mid \mathbf{z})$. That is, the investor is risk-neutral with respect to the risk induced by stochastic $\mathbf{x} \mid \mathbf{z}$ that she comprehends but behaves as if she had non-trivial risk-attitudes with respect to the risk induced by stochastic $\mathbf{y} \mid (\mathbf{x}, \mathbf{z})$ that she does not comprehend.

The DM who encounters a lottery $\mathbf{r} \in \mathcal{A}_{\mathcal{K}}$ that she has anticipated learns in a correctly

specified model. The asymptotic results for correctly specified learning of Wald (1949) for ML estimation and of Le Cam (1953) for Bayesian estimation then imply that she correctly learns the encountered lottery as the number of signals diverges. In this case, her perception strategy is irrelevant for her limit choice since she accounts for the encoding and the sampling frequencies when she interprets the perception data. The next corollary of Proposition 4 formalizes this in our framework.

Corollary 1. *For each decision problem (\mathbf{r}, s) such that $\mathbf{r} \in \mathcal{A}_{\mathcal{K}}$ and $\sum_{i=1}^I p_i r_i > [<] s$, the probability that the DM chooses the lottery converges to 1 [0].*

Our predictions of the DM’s risk attitudes more generally depend on the combination of the adaptation experienced, as in Section 3, and her misapprehension of the lottery at the moment of choice. Recall that \mathcal{J} denotes the partition that the DM has employed during adaptation and partition \mathcal{K} specifies the DM’s anticipation of lotteries at the moment of choice; \mathcal{J} and \mathcal{K} may differ. The optimal encoding function is S-shaped regardless of the adaptation partition \mathcal{J} . Hence, we predict risk aversion (loving) for upper (lower) tail rewards with respect to the unanticipated risk under \mathcal{K} . Our results of Section 3 predict relative sampling frequencies only for arms that the DM was distinguishing under adaptation partition \mathcal{J} . If the DM then ignores distinctions between such arms at the moment of choice (for instance because she was framed by a car dealer to ignore the role of the weather), we predict overweighting of small probability events (for instance weather conditions suitable for driving open). If, on the other hand, the DM has ignored distinctions between arms already at the stage of adaptation, our results of Section 3 do not predict relative sampling frequencies, since this DM believed the arms were payoff-equivalent during adaptation. If sampling is representative, then the sampling frequencies coincide with the arms’ objective probabilities. Any targeted sampling, for instance over-sampling salient contingencies, results in choice that assigns disproportional subjective probabilities to the over-sampled arms.¹⁰

4.3 Somewhat Surprising Risk

As a last extension of our model, we analyze a DM who deems risk a priori possible but unlikely. Her perception frictions are comparable in size to the risk that she typically expects to encounter. To this end, we study a limit in which the prior shrinks to the set of riskless lotteries as the amount of perception data diverges. We find perception distortions that are qualitatively similar to those from Subsection 4.1. Additionally, the approach makes

¹⁰Starmer and Sugden (1993) report that a payoff-irrelevant split of an event increases the weight that lab subjects assign to this event. This effect arises for our coarse DM if splitting a contingency leads to its having a larger overall sampling frequency.

predictions about the impact of framing and time pressure on risk-taking. Risk attitudes are attenuated by the anticipation of high risk or by rich perception data.

The DM of this subsection is Bayesian. Her prior density indexed by n is

$$\varrho_n(\mathbf{r}) = \varrho_n^0 \exp\left(-\frac{n}{2\Delta}\sigma^2(\mathbf{r})\right) \quad (9)$$

with support $[\underline{r}, \bar{r}]^I$, where $\sigma^2(\mathbf{r}) = \sum_{i=1}^I p_i (r_i - r)^2$ is the variance of the arm rewards and $r = \sum_i p_i r_i$ is the true lottery value as usual; ϱ_n^0 is the normalization factor. This prior is mostly concentrated on low-risk lotteries. For any fixed n , $\Delta > 0$ parameterizes the level of a priori anticipated risk. The index n has two roles. As n increases, risky lotteries become a priori less likely, approximating then the anticipation of the DM from Subsection 4.1. In addition to risk becoming less likely, the DM observes more data as n increases. She observes, for each arm i , a sequence of $a\pi_i n$ messages equal to $m(r_i)$ perturbed with iid additive standard normal noise, where the $(\pi_i)_{i=1}^I$ continue to denote the sampling frequencies.¹¹ The parameter $a > 0$ captures attention span; the larger a is, the more signals the DM observes for each fixed n . The DM chooses the lottery \mathbf{r} over the safe option s if and only if the Bayesian posterior expected lottery value exceeds s .

To formulate the next result, we define a function $\mathbf{q}^* : [\underline{r}, \bar{r}]^I \rightarrow [\underline{r}, \bar{r}]^I$ as follows:

$$\mathbf{q}^*(\mathbf{r}) = \arg \min_{\mathbf{r}' \in [\underline{r}, \bar{r}]^I} \left\{ \frac{\sigma^2(\mathbf{r}')}{a\Delta} + \sum_{i=1}^I \pi_i (m(r'_i) - m(r_i))^2 \right\}. \quad (10)$$

We impose the regularity condition that the minimizer is unique. We refer to the posterior expectation $\mathbb{E}[\mathbf{r} \mid \mathbf{m}_n] \in [\underline{r}, \bar{r}]^I$ that the DM forms given the vector of the average perturbed messages \mathbf{m}_n as the Bayesian estimate of the lottery and show that it converges to $\mathbf{q}^*(\mathbf{r})$. As n diverges, the average of the perturbed messages generated by an arm converges almost surely to the unperturbed message for this arm. Thus, the posterior log-likelihood of each lottery \mathbf{r}' is approximately $-n/2$ times the objective in (10) within a constant factor. As n diverges, the posterior converges almost surely to an atom on the minimizer $\mathbf{q}^*(\mathbf{r})$.

Proposition 5. *Suppose the DM has encountered lottery \mathbf{r} . The Bayesian estimate of the lottery converges to $\mathbf{q}^*(\mathbf{r})$ in probability as $n \rightarrow \infty$.*

See Appendix D for the proofs for this subsection. The asymptotic estimate $\mathbf{q}^*(\mathbf{r})$ of the lottery \mathbf{r} is a compromise lottery that is not too risky and does not generate messages too far from the true messages. When $a\Delta$ is small, then the DM anticipates little risk and/or

¹¹We again abstract from uncertainty about the number of the messages sampled for each arm and from divisibility issues.

collects little perception data. Her best explanation of her perception data is a lottery that involves little risk: in the limit as $a\Delta \rightarrow 0$, the solution to (10) minimizes Kullback-Leibler divergence from the true lottery among the riskless lotteries, as in Subsection 4.1. When $a\Delta$ is large, then the DM anticipates large risks and/or collects a lot of perception data. Then, her best explanation of the data minimizes Kullback-Leibler divergence from the true lottery among all lotteries, which yields the correct estimate.

Let $q^*(\mathbf{r}) = \sum_{i=1}^I p_i q_i^*(\mathbf{r})$ be the value of the lottery $\mathbf{q}^*(\mathbf{r})$. Proposition 5 implies:

Corollary 2. *Consider a decision problem (\mathbf{r}, s) such that $q^*(\mathbf{r}) > [<] s$. Then, the probability that the DM chooses the lottery [the safe option] approaches 1 as $n \rightarrow \infty$.*

To focus on the effect of the curvature of the encoding function, we set the sampling frequencies equal to the actual probabilities and compare the asymptotic estimated lottery value $q^*(\mathbf{r})$ with the true value r of the lottery \mathbf{r} .¹²

Proposition 6. *Let the encoding function m be twice differentiable. Let $\pi_i = p_i$, and \mathbf{r} be a fixed lottery. The value of its Bayesian estimate almost surely converges to*

$$r + \frac{1}{2} \frac{m''(r)}{m'(r)} \cdot \frac{1 + 4z(r)}{(1 + z(r))^2} \cdot \sigma^2(\mathbf{r}) + o(\sigma^2(\mathbf{r})), \quad (11)$$

as $n \rightarrow \infty$, where $z(r) = a\Delta m'^2(r)$. The factor $\frac{1+4z(r)}{(1+z(r))^2}$ attains values in $(0, 4/3]$ and approaches 1 and 0 as $a\Delta \rightarrow 0$ and $a\Delta \rightarrow \infty$, respectively.

To interpret the result, recall that the risk premium of an expected-utility maximizer with Bernoulli utility u for a lottery \mathbf{r} with small risk is approximately $\frac{1}{2} \frac{u''(r)}{u'(r)} \sigma^2(\mathbf{r})$. The risk premium of our DM is the same for $u(\cdot) = m(\cdot)$ (up to a negligible term), but scaled by the positive factor $\frac{1+4z(r)}{(1+z(r))^2}$. The DM's bias in the valuation of the lottery relative to r arises because the DM deems risk a priori unlikely and therefore concludes that her perceived data are generated by a lottery with a smaller reward variance than the true variance. The underestimation of the variance leads to a mismatch to the perception data and this mismatch is offset by a bias in the estimated mean of the lottery.

The dependence of the risk premium on the parameters Δ and a sheds light on two apparent instabilities of risk preferences pointed out by Rabin (2000) and Kahneman (2011). Kahnemann distinguishes between fast and slow modes of decision-making, where the fast mode favours the risk-attitudes found in prospect theory whereas the slow mode favours

¹²We say that function $f(\mathbf{r})$ is $o(g(\mathbf{r}))$ if $f(\mathbf{r}_k)/g(\mathbf{r}_k) \rightarrow 0$ for any sequence \mathbf{r}_k such that $\sigma(\mathbf{r}_k) \rightarrow 0$. Specifically, a function is $o(\sigma^2)$ if it is negligible relative to σ^2 for lotteries with small σ . The expression $o(\cdot)$ stands for "term of smaller order than".

risk-neutrality.¹³ If the amount of perception data collected by the DM increases with the time available for the decision, then time pressure is captured in our example by a low value of parameter a . In accord with Kahnemann, we find encoding-based risk attitudes when $a \rightarrow 0$. When our DM, who has anticipated little risk, encounters a risky lottery under time pressure, the relatively few data points that she has collected are best explained by an a priori likely low-risk lottery. Which such low-risk lottery is the best fit to the DM’s data depends on the encoding function, thus the curvature of m determines the DM’s risk attitudes. At the other extreme, in the absence of time pressure, when $a \rightarrow \infty$, the DM collects enough data for her prior to be irrelevant. She then learns the lottery and makes the risk-neutral choice.

Rabin (2000) points out that the risk-averse choices observed for small risks imply implausibly high risk aversion for large risks under a stable Bernoulli utility function. In our model, however, risk attitudes depend on the level of a priori anticipated risk. The anticipation of low risk – captured by small Δ here – induces risk attitudes since it makes risky lotteries surprising, and this leads to distortion of the posteriors when a risky lottery is encountered. If, however, the DM anticipates high-risk lotteries – if parameter Δ is large – then the DM’s risk attitudes are attenuated. Risky lotteries become unsurprising and the DM’s posterior expectation approaches the lottery’s true expected value.

5 Summary

We develop a model of constrained optimal perception of gambles in which psychophysical adaptation affects choices. The impact of the perceptual strategy vanishes for rich perception data if the DM encounters a lottery that she has anticipated, but perception-induced risk attitudes arise for risk that the DM has not anticipated. In the latter case, we provide a unified explanation for various well-documented patterns in risky choice: adaptive risk attitudes, S-shaped reward valuation, probability weighting, and the role of stakes and time pressure.

The model makes several novel predictions. For example, explaining the structure of risk to the DM should attenuate her risk attitudes, while increasing the complexity of the environment should strengthen perception-driven behavior. These predictions are broadly in line with the recent experimental findings of Enke and Graeber (2021), who show that probability weighting is more pronounced for subjects who state a higher level of cognitive

¹³Kirchler et al. (2017) show experimentally that time pressure increases risk aversion for gains and risk loving for losses. Relatedly, Porcelli and Delgado (2009) and Cahliková and Cingl (2017) find that stress accentuates risk attitudes in lab choices. But see also Kocher, Pahlke, and Trautmann (2013) who do not find an increase of risk aversion due to time pressure in their design.

uncertainty about the correct action. Other predictions relate to the effect of framing. For example, a DM who is framed to perceive a risky lottery as riskless will rely on sampling frequencies rather than objective probabilities to evaluate the lottery. Manipulation of the sampling frequencies then has a strong impact on choice. A seller offering a risky prospect can make it more attractive if the presentation of the prospect leads to over-sampling of the upside risk. An additional example is the prediction that risk attitudes become more pronounced if the DM samples less perceptual data, which the seller of an insurance contract could exploit by putting the DM under time pressure.

References

- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological review* 61(3), 183.
- Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. *Sensory communication* 1(01).
- Benkert, J.-M. and N. Netzer (2018). Informational requirements of nudging. *Journal of Political Economy* 126(6), 2323–2355.
- Berk, R. H. (1966). Limiting behavior of posterior distributions when the model is incorrect. *The Annals of Mathematical Statistics*, 51–58.
- Bernheim, B. D. and A. Rangel (2009). Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics. *The Quarterly Journal of Economics* 124(1), 51–104.
- Bethge, M., D. Rotermund, and K. Pawelzik (2002). Optimal short-term population coding: when fisher information fails. *Neural computation* 14(10), 2317–2351.
- Birnbaum, Z. W. (1948). On random variables with comparable peakedness. *The Annals of Mathematical Statistics* 19(1), 76–81.
- Cahlíková, J. and L. Cingl (2017). Risk preferences under acute stress. *Experimental Economics* 20(1), 209–236.
- Eliaz, K. and R. Spiegel (2020). A model of competing narratives. *American Economic Review* 110(12), 3786–3816.
- Enke, B. and T. Graeber (2021). Cognitive uncertainty. Mimeo.

- Esponda, I. and D. Pouzo (2016). Berk–nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica* 84(3), 1093–1130.
- Fechner, G. T. (1860). Elements of psychophysics.
- Frick, M., R. Iijima, and Y. Ishii (2021). Welfare comparisons for biased learning.
- Friedman, D. (1989). The S-shaped value function as a constrained optimum. *American Economic Review* 79(5), 1243–1248.
- Frydman, C. and L. J. Jin (2019). Efficient coding and risky choice. *Available at SSRN 3270773*.
- Heidhues, P., B. Köszegi, and P. Strack (2018). Unrealistic expectations and misguided learning. *Econometrica* 86(4), 1159–1214.
- Heng, J. A., M. Woodford, and R. Polania (2020). Efficient sampling and noisy decisions. *Elife* 9, e54962.
- Herold, F. and N. Netzer (2015). Second-best probability weighting. Mimeo.
- Johnson, D. H. and G. C. Orsak (1993). Relation of signal set choice to the performance of optimal non-gaussian detectors. *IEEE Transactions on Communications* 41(9), 1319–1328.
- Juechems, K., J. Balaguer, B. Spitzer, and C. Summerfield (2021). Optimal utility and probability functions for agents with finite computational precision. *Proceedings of the National Academy of Sciences* 118(2).
- Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291.
- Kahneman, D., P. P. Wakker, and R. Sarin (1997). Back to bentham? explorations of experienced utility. *The Quarterly Journal of Economics* 112(2), 375–406.
- Khaw, M. W., Z. Li, and M. Woodford (2018). Cognitive imprecision and small-stakes risk aversion. *The Review of Economic Studies*.
- Kirchler, M., D. Andersson, C. Bonn, M. Johannesson, E. Ø. Sørensen, M. Stefan, G. Tinghög, and D. Västfjäll (2017). The effect of fast and slow decisions on risk taking. *Journal of Risk and Uncertainty* 54(1), 37–59.

- Kocher, M. G., J. Pahlke, and S. T. Trautmann (2013). Tempus fugit: time pressure in risky decisions. *Management Science* 59(10), 2380–2391.
- Koszegi, B. and M. Rabin (2008). Revealed mistakes and revealed preferences. *The foundations of positive and normative economics: a handbook*, 193–209.
- Laughlin, S. (1981). A simple coding procedure enhances a neuron’s information capacity. *Zeitschrift für Naturforschung c* 36(9-10), 910–912.
- Le Cam, L. (1953). On some asymptotic properties of maximum likelihood estimates and related bayes estimates. *Univ. California Pub. Statist.* 1, 277–330.
- Molavi, P. (2019). Macroeconomics with learning and misspecification: A general theory and applications. *Unpublished manuscript*.
- Netzer, N. (2009). Evolution of time preferences and attitudes toward risk. *The American Economic Review* 99(3), 937–955.
- Ok, E. A. (2011). *Real analysis with economic applications*. Princeton University Press.
- Porcelli, A. J. and M. R. Delgado (2009). Acute stress modulates risk taking in financial decision making. *Psychological Science* 20(3), 278–283.
- Purkayastha, S. (1998). Simple proofs of two results on convolutions of unimodal distributions. *Statistics & probability letters* 39(2), 97–100.
- Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica* 68(5), 1281–1292.
- Rayo, L. and G. S. Becker (2007). Evolutionary efficiency and happiness. *Journal of Political Economy* 115(2), 302–337.
- Robson, A. J. (2001). The biological basis of economic behavior. *Journal of Economic Literature* 39(1), 11–33.
- Robson, A. J., L. A. Whitehead, and N. Robalino (2021). Adaptive hedonic utility. Technical report, Working paper.
- Salant, Y. and A. Rubinstein (2008). (a, f): choice with frames. *The Review of Economic Studies* 75(4), 1287–1296.
- Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons.

- Schaffner, J., P. Tobler, T. Hare, and R. Polania (2021). Fitness-maximizing codes in early sensory areas. University of Zurich mimeo.
- Starmer, C. and R. Sugden (1993). Testing for juxtaposition and event-splitting effects. *Journal of risk and uncertainty* 6(3), 235–254.
- Steiner, J. and C. Stewart (2016). Perceiving prospects properly. *The American Economic Review* 106(7), 1601–31.
- Thurstone, L. L. (1927). A law of comparative judgment. *Psychological review* 34(4), 273.
- Wald, A. (1949). Note on the consistency of the maximum likelihood estimate. *The Annals of Mathematical Statistics* 20(4), 595–601.
- Wang, Z., A. A. Stocker, and D. D. Lee (2016). Efficient neural codes that minimize lp reconstruction error. *Neural computation* 28(12), 2656–2686.
- White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica: Journal of the Econometric Society*, 1–25.
- Woodford, M. (2020). Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics* 12, 579–601.

A Asymptotic Loss Characterization

Let $0 < \underline{h} \leq \bar{h} < +\infty$ and $0 < \underline{m}' \leq \bar{m}' < +\infty$ be bounds on the functions h and m' . These bounds exist since the two functions are continuous on a compact interval.

Lemma 1. *Suppose that the encoding function is continuously differentiable and the reward density h is continuous. Let $r_J^* \in (\underline{r}, \bar{r})$ be a realization of the reward, and $q_{J,n}^B = \mathbb{E}[r_J \mid m_{J,n}]$ and $q_{J,n}^{ML} = m^{-1}(m_{J,n})$ its Bayesian and ML estimators. Then,*

$$(i) \quad \sqrt{n} (q_{J,n}^B - q_{J,n}^{ML}) \rightarrow 0 \text{ as } n \rightarrow \infty \text{ (a.s.)},$$

$$(ii) \quad n \text{Var}[r_J \mid m_{J,n}] \rightarrow \frac{1}{\pi_J m'^2(r_J^*)} \text{ as } n \rightarrow \infty \text{ (a.s.)},$$

(iii) *The mean squared error of the ML estimate (rescaled by n) is uniformly bounded:*

$$n \mathbb{E} \left[(r_J - q_{J,n}^{ML})^2 \mid m_{J,n} \right] \leq \frac{\bar{h}}{\underline{h}} \cdot \frac{\bar{m}'}{m'^3} \frac{1}{\pi_J}.$$

Proof. Consider sufficiently large n so that $m_{J,n} \in [\underline{m}, \bar{m}]$. We introduce rescaled error $\hat{\varepsilon}_{J,n} := \sqrt{\pi_J n}(r_J - q_{J,n})$ and derive its conditional density given $m_{J,n}$. Since $m_{J,n} \sim \mathcal{N}(m(r_J^*), 1/(\pi_J n))$, the pdf of $r_J \mid m_{J,n}$ is proportional to

$$h(\tilde{r}) \varphi\left(\sqrt{\pi_J n}(m(\tilde{r}) - m_{J,n})\right)$$

for any $\tilde{r} \in [\underline{r}, \bar{r}]$ and 0 otherwise; recall φ is the standard normal density. Thus, the pdf of $\hat{\varepsilon}_{J,n}$ conditioned on $m_{J,n} = m(q_{J,n}^{ML})$ is

$$h_{J,n}(\tilde{\varepsilon}) = h_{J,n}^0 \cdot h\left(q_{J,n}^{ML} + \frac{\tilde{\varepsilon}}{\sqrt{\pi_J n}}\right) \varphi\left(\sqrt{\pi_J n}\left(m\left(q_{J,n}^{ML} + \frac{\tilde{\varepsilon}}{\sqrt{\pi_J n}}\right) - m(q_{J,n}^{ML})\right)\right)$$

for any $\tilde{\varepsilon} \in [\sqrt{\pi_J n}(\underline{r} - q_{J,n}^{ML}), \sqrt{\pi_J n}(\bar{r} - q_{J,n}^{ML})]$ and 0 otherwise; $h_{J,n}^0$ is the normalization factor. It follows that $h_{J,n}(\tilde{\varepsilon})/h_{J,n}^0$ is dominated by the integrable function $\bar{h} \cdot \varphi(\underline{m}' \cdot \tilde{\varepsilon})$.

Since $|q_{J,n}^{ML} - r_J^*| = |m^{-1}(m_{J,n}) - r_J^*| \leq \frac{1}{\underline{m}'} |m_{J,n} - m(r_J^*)|$ and $m_{J,n} \sim \mathcal{N}(m(r_J^*), 1/(\pi_J n))$, we have that $q_{J,n}^{ML} \rightarrow r_J^*$ (a.s.). Using this, the Mean Value Theorem, and continuity of m' , we get

$$\sqrt{\pi_J n}\left(m\left(q_{J,n}^{ML} + \frac{\tilde{\varepsilon}}{\sqrt{\pi_J n}}\right) - m(q_{J,n}^{ML})\right) \rightarrow m'(r_J^*) \tilde{\varepsilon} \text{ as } n \rightarrow \infty \text{ a.s.},$$

and thus, using the continuity of h , for any $\tilde{\varepsilon}$,

$$\frac{h_{J,n}(\tilde{\varepsilon})}{h_{J,n}^0} \rightarrow h(r_J^*) \varphi(m'(r_J^*) \tilde{\varepsilon}) \text{ as } n \rightarrow \infty \text{ (a.s.)}.$$

Next, we characterize the limit of the normalization factors. By the Dominated Convergence Theorem,

$$\int_{\mathbb{R}} \frac{h_{J,n}(\tilde{\varepsilon})}{h_{J,n}^0} d\tilde{\varepsilon} \rightarrow \frac{1}{h_J^0} \text{ as } n \rightarrow \infty \text{ (a.s.)}, \text{ where } h_J^0 := \left[\int_{\mathbb{R}} h(r_J^*) \varphi(m'(r_J^*) \tilde{\varepsilon}) d\tilde{\varepsilon} \right]^{-1}.$$

Since $\int_{\mathbb{R}} h_{J,n}(\tilde{\varepsilon}) d\tilde{\varepsilon} = 1$ for all n , it follows that $h_{J,n}^0 \rightarrow h_J^0 > 0$ (a.s.). In particular, $h_{J,n}^0$ is bounded. Then, the posterior errors $\hat{\varepsilon}_{J,n} \mid m_{J,n}$ converge in distribution to $\mathcal{N}(0, 1/m'^2(r_J^*))$ (a.s.).

Applying the Dominated Convergence Theorem to the functions $\tilde{\varepsilon} h_{J,n}(\tilde{\varepsilon})$ and $\tilde{\varepsilon}^2 h_{J,n}(\tilde{\varepsilon})$, we conclude that $E[\hat{\varepsilon}_{J,n} \mid m_{J,n}] \rightarrow 0$ and $\text{Var}[\hat{\varepsilon}_{J,n} \mid m_{J,n}] \rightarrow 1/m'^2(r_J^*)$ as $n \rightarrow \infty$ (a.s.). Claims (i) and (ii) follow from deriving r_J from $\hat{\varepsilon}_{J,n} = \sqrt{\pi_J n}(r_J - q_{J,n}^{ML})$.

For Claim (iii), recall that $h_{J,n}(\tilde{\varepsilon})/h_{J,n}^0$ is dominated by the integrable function $\bar{h} \cdot \varphi(\underline{m}' \cdot \tilde{\varepsilon})$ that equals, up to a multiplicative constant, the pdf of $\mathcal{N}(0, 1/\underline{m}'^2)$. Consider a random variable $\hat{\varepsilon}'_{J,n}$ with pdf proportional to $h_{J,n}(\tilde{\varepsilon})/h_{J,n}^0$ on the domain of $\hat{\varepsilon}_{J,n}$, and $\underline{h} \varphi(\bar{m}' \cdot \tilde{\varepsilon})$

outside of the domain. We can establish the following upper bound on the normalization constant $h_{J,n}^0$ of the pdf of the variable $\varepsilon'_{J,n}$,

$$h_{J,n}^0 \leq \left[\int_{\mathbb{R}} \underline{h} \varphi(\overline{m}' \cdot \tilde{\varepsilon}) d\tilde{\varepsilon} \right]^{-1} = \frac{\overline{h}}{\underline{h}} \cdot \frac{\overline{m}'}{\underline{m}'} \cdot \left[\int_{\mathbb{R}} \overline{h} \varphi(\underline{m}' \cdot \tilde{\varepsilon}) d\tilde{\varepsilon} \right]^{-1}.$$

Then,

$$n \mathbb{E} \left[(r_J - q_{J,n}^{ML})^2 \mid m_{J,n} \right] \leq \frac{1}{\pi_J} \mathbb{E} \left[\varepsilon'^2_{J,n} \mid m_{J,n} \right] \leq \frac{1}{\pi_J} h_{J,n}^0 \int_{\mathbb{R}} \tilde{\varepsilon}^2 \cdot \overline{h} \varphi(\underline{m}' \tilde{\varepsilon}) d\tilde{\varepsilon} = \frac{\overline{h}}{\underline{h}} \cdot \frac{\overline{m}'}{\underline{m}'^3} \frac{1}{\pi_J}.$$

□

Corollary 3. *Conditional on a realization of $\mathbf{r}^* \in (\underline{r}, \overline{r})^{|\mathcal{J}|}$,*

$$(i) \sqrt{n} (q_n^B - q_n^{ML}) \rightarrow 0 \text{ as } n \rightarrow \infty \text{ (a.s.)},$$

$$(ii) n \text{Var}[r \mid \mathbf{m}_n] \rightarrow \sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r^*_J)} \text{ as } n \rightarrow \infty \text{ (a.s.)}.$$

$$(iii) n \mathbb{E} \left[(r - q_n^{ML})^2 \mid \mathbf{m}_n \right] \leq \frac{\overline{h}}{\underline{h}} \cdot \frac{\overline{m}'}{\underline{m}'^3} \cdot \sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J}.$$

For the following two lemmas, we abstract from the specific structure of the messages \mathbf{m}_n , let the DM receive a vector of messages $\mathbf{m} = (m_J)_J$ and then form the estimate $q(\mathbf{m})$ of r as a function of \mathbf{m} . Let

$$\ell := \max\{r, s\} - \mathbb{1}_{q > s} r - \mathbb{1}_{q \leq s} s$$

denote the loss from using the estimate q .

We assume that the pdf h_s is continuously differentiable, thus h_s and h'_s are bounded from above; let \overline{h}_s and \overline{h}'_s be the respective bounds. We say that $O(\cdot)$ has uniform bound \overline{h}'_s if $|O(x)/x| \leq \overline{h}'_s$ for all x and any value of \mathbf{r} and \mathbf{m} .¹⁴

Lemma 2. *The expected loss of the estimate q conditioned on \mathbf{m} and \mathbf{r} is*

$$\mathbb{E}[\ell \mid \mathbf{r}, \mathbf{m}] = \frac{1}{2} h_s(q) (r - q)^2 + O((r - q)^3),$$

where the expectation is over s and $O(\cdot)$ has the uniform bound \overline{h}'_s .

Proof. Consider a fixed realization of \mathbf{r} and \mathbf{m} . The loss is $\ell = |r - s|$ if the DM makes the suboptimal choice, which happens if and only if s is between the true lottery value r and its

¹⁴The term $O(\cdot)$ stands for the “term of the order of”.

estimate q . Taking the expectation over the safe option yields (for both $r < q$ and $r > q$)

$$\mathbb{E}[\ell \mid \mathbf{r}, \mathbf{m}] = \int_r^q (\tilde{s} - r) h_s(\tilde{s}) d\tilde{s}.$$

The lemma follows from the approximation $h_s(\tilde{s}) = h_s(q) + O(\tilde{s} - q)$, in which $O(\cdot)$ has the uniform bound \bar{h}'_s ,

$$\begin{aligned} \int_r^q (\tilde{s} - r) h_s(\tilde{s}) d\tilde{s} &= h_s(q) \int_r^q (\tilde{s} - r) d\tilde{s} + \int_r^q (\tilde{s} - r) O(\tilde{s} - q) d\tilde{s} \\ &= \frac{1}{2} h_s(q) (r - q)^2 + O((r - q)^3). \end{aligned}$$

□

Lemma 3. *The expected loss of the estimate q conditioned on \mathbf{m} is*

$$\mathbb{E}[\ell \mid \mathbf{m}] = \frac{1}{2} h_s(q) \sigma^2 + O(\sigma^3), \text{ where } \sigma^2 := \text{Var}[r \mid \mathbf{m}] + (q^B - q)^2,$$

where $O(\cdot)$ has the uniform bound \bar{h}'_s and $q^B = \mathbb{E}[r \mid \mathbf{m}]$.

Proof. This follows from Lemma 2 by taking the expectation over \mathbf{r} :

$$\begin{aligned} \mathbb{E}[\ell \mid \mathbf{m}] &= \mathbb{E} \left[\mathbb{E}[\ell \mid \mathbf{r}, \mathbf{m}] \mid \mathbf{m} \right] \\ &= \mathbb{E} \left[\frac{1}{2} h_s(q) (r - q)^2 + O((r - q)^3) \mid \mathbf{m} \right], \end{aligned}$$

where $O(\cdot)$ has the uniform bound \bar{h}'_s . Since $|O((r - q)^3)| \leq \bar{h}'_s |r - q|^3$,

$$\left| \mathbb{E} [O((r - q)^3) \mid \mathbf{m}] \right| \leq \mathbb{E} \left[\bar{h}'_s \left((r - q)^2 \right)^{3/2} \mid \mathbf{m} \right] \leq \bar{h}'_s \mathbb{E} [(r - q)^2 \mid \mathbf{m}]^{3/2},$$

where we have used Jensen's inequality in the second step.

We conclude with

$$\begin{aligned} \mathbb{E} [(r - q)^2 \mid \mathbf{m}] &= \mathbb{E} \left[\left((r - q^B) + (q^B - q) \right)^2 \mid \mathbf{m} \right] \\ &= \mathbb{E} \left[(r - q^B)^2 + 2(r - q^B)(q^B - q) + (q^B - q)^2 \mid \mathbf{m} \right] \\ &= \text{Var}[r \mid \mathbf{m}] + (q^B - q)^2 = \sigma^2. \end{aligned}$$

□

Proof of Proposition 1. Let $\ell_n^z = \max\{r, s\} - \mathbb{1}_{q_n^z > s} r - \mathbb{1}_{q_n^z \leq s} s$ be the loss of the estimator

q_n^z , $z \in \{B, ML\}$. For a given realization of the lottery \mathbf{r}^* with value r^* , we prove that the expected loss conditioned on \mathbf{m}_n satisfies

$$n \mathbb{E}[\ell_n^z \mid \mathbf{m}_n] \rightarrow \frac{1}{2} h_s(r^*) \sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J^*)} \text{ as } n \rightarrow \infty \text{ (a.s.)}, \quad (12)$$

where the expectation is over s and q_n^z is a function of \mathbf{m}_n .

Lemma 3 applied to \mathbf{m}_n and q_n^z implies

$$n \mathbb{E}[\ell_n^z \mid \mathbf{m}_n] = \frac{1}{2} h_s(q_n^z) n \sigma_n^2(z) + n O(\sigma_n^3(z)), \text{ where } \sigma_n^2(z) := \text{Var}[r \mid \mathbf{m}_n] + (q_n^B - q_n^z)^2. \quad (13)$$

Corollary 3 implies that $n(q_n^B - q_n^z)^2 \rightarrow 0$ (a.s.) (this holds trivially for the Bayesian estimator). Further, Claim (ii) of Corollary 3 implies

$$n \sigma_n^2(z) \rightarrow \sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J^*)} \text{ as } n \rightarrow \infty \text{ (a.s.)}.$$

Thus, $\sigma_n(z) \rightarrow 0$ as $n \rightarrow \infty$ (a.s.); and so $|n O(\sigma_n^3(z))| \leq n \bar{h}'_s \sigma_n^3(z) = \bar{h}'_s \cdot n \sigma_n^2(z) \cdot \sigma_n(z) \rightarrow 0$ as $n \rightarrow \infty$ (a.s.). Substituting back into (13) and taking into account that $q_n^z \rightarrow r^*$ as $n \rightarrow \infty$ (a.s.), we obtain (12).

The proposition follows from taking expectation over \mathbf{r} and applying the Dominated Convergence Theorem. In particular, (13) implies that $n \mathbb{E}[\ell_n^z \mid \mathbf{m}_n]$ has integrable bound:

$$|n \mathbb{E}[\ell_n^z \mid \mathbf{m}_n]| \leq \frac{1}{2} \bar{h}_s \cdot n \sigma_n^2(ML) + \bar{h}'_s \cdot n \sigma_n^2(ML) \cdot \sigma_n(ML) \leq \frac{1}{2} \bar{h}_s \bar{\Sigma} + \bar{h}'_s \bar{\Sigma} \cdot \max\{\bar{\Sigma}, 1\},$$

where $\bar{\Sigma} = \frac{\bar{h}}{\bar{h}} \cdot \frac{\bar{m}'}{\bar{m}'^3} \cdot \sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J}$ is the uniform bound from Claim (iii) of Corollary 3. \square

B Optimal Perception

B.1 Notes on Unimodal Symmetric Random Variables

We use the results of this subsection in the proof of Proposition 2.

Definition 1. *A continuous random variable is unimodal and symmetric around 0 if its density function $h(x)$ is strictly decreasing on the positive part of its domain and $h(x) = h(-x)$ for all $x \in \mathbb{R}$.*

This property is preserved by summation: the sum of unimodal and symmetric random variables is unimodal and symmetric, see e.g. Purkayastha (1998).

Definition 2 (Birnbaum (1948)). *Let X and Y be two unimodal random variables symmetric around 0. We say that X is more peaked than Y if $P(|X| < \alpha) > P(|Y| < \alpha)$ (unless the right-hand side is 1) for all $\alpha > 0$.*

Equivalently, for two unimodal symmetric random variables, X is more peaked than Y whenever the cdf of X is greater than the cdf of Y at any $\alpha > 0$ from the support of Y .

For the next two lemmas, let X_0, X_1, \dots, X_I be independent real-valued continuous random variables that are unimodal and symmetric around 0, where X_1, \dots, X_I are identically distributed and the distribution of X_0 may be distinct from that of X_i , $i > 0$. Denote by h the pdf of each of the iid variables X_1, \dots, X_I . Let $(p_1, \dots, p_I) \in \Delta(\{1, \dots, I\})$, and $X := \sum_{i=1}^I p_i X_i$.

Lemma 4. *The random variable $X_i \mid (X = X_0)$, $i = 1, \dots, I$, is unimodal and symmetric around 0.*

Proof. Since unimodality together with symmetry is preserved by affine combinations, the variable $X_{-i} := \frac{1}{p_i}(X_0 - \sum_{k \neq i} p_k X_k)$ is unimodal and symmetric around 0. Denote by h_{-i} the pdf of X_{-i} . Then $X_i \mid (X = X_0)$ is identical to $X_i \mid (X_i = X_{-i})$, and so its pdf is, up to a normalization constant, $h(x_i)h_{-i}(x_i)$, which is unimodal and symmetric around 0, as those properties are preserved when taking product of pdfs. \square

Lemma 5. *The random variable $X_i \mid (X = X_0)$ is more peaked than $X_j \mid (X = X_0)$ if and only if $p_i > p_j$.*

Proof. Without loss of generality, assume $\{i, j\} = \{1, 2\}$ (that is, either $i = 1$ and $j = 2$ or $i = 2$ and $j = 1$). Define $X_{-12} := X_0 - \sum_{k=3}^I p_k X_k$ (if $I = 2$, then $X_{-12} = X_0$) and let h_{-12} be its pdf. This is a unimodal random variable symmetric around 0. The random variable $X_i \mid (X = X_0)$ is identical to $X_i \mid (p_i X_i + p_j X_j = X_{-12})$ and so its pdf equals

$$h_i(x_i) = \frac{\int_{\mathbb{R}} h_{-12}(p_1 x_1 + p_2 x_2) h(x_1) h(x_2) dx_j}{\mathbb{E}[h_{-12}(p_1 X_1 + p_2 X_2)]},$$

where the expectation, which is with respect to X_1 and X_2 , is independent of i . Thus, for any $\alpha > 0$,

$$P(|X_1| < \alpha \mid X = X_0) = \frac{\iint_{(-\alpha, \alpha) \times \mathbb{R}} h_{-12}(p_1 x_1 + p_2 x_2) h(x_1) h(x_2) dx_1 dx_2}{\mathbb{E}[h_{-12}(p_1 X_1 + p_2 X_2)]}$$

$$P(|X_2| < \alpha \mid X = X_0) = \frac{\iint_{(-\alpha, \alpha) \times \mathbb{R}} h_{-12}(p_1 x_2 + p_2 x_1) h(x_1) h(x_2) dx_1 dx_2}{\mathbb{E}[h_{-12}(p_1 X_1 + p_2 X_2)]},$$

where we used that $P(|X_1| < \alpha | X = X_0)$ and $P(|X_2| < \alpha | X = X_0)$ are both (up to the same normalization constant) integrals of the same function $(x_1, x_2) \mapsto h_{-12}(p_1x_1 + p_2x_2)h(x_1)h(x_2)$, but the first is over the region $[-\alpha, \alpha] \times \mathbb{R}$, and the second is over $\mathbb{R} \times [-\alpha, \alpha]$. This is equivalent to integrating both over the same region but switching the roles of x_1 and x_2 . Then,

$$\begin{aligned} & (P(|X_1| < \alpha | X = X_0) - P(|X_2| < \alpha | X = X_0)) \cdot E[h_{-12}(p_1X_1 + p_2X_2)] = \\ & \iint_{(-\alpha, \alpha) \times \mathbb{R}} \left(h_{-12}(p_1x_1 + p_2x_2) - h_{-12}(p_1x_2 + p_2x_1) \right) h(x_1)h(x_2) dx_1 dx_2 = \\ & \iint_{(-\alpha, \alpha) \times (\mathbb{R} \setminus (-\alpha, \alpha))} \left(h_{-12}(p_1x_1 + p_2x_2) - h_{-12}(p_1x_2 + p_2x_1) \right) h(x_1)h(x_2) dx_1 dx_2 = \\ & 2 \iint_{(-\alpha, \alpha) \times [\alpha, +\infty)} \left(h_{-12}(p_1x_1 + p_2x_2) - h_{-12}(p_1x_2 + p_2x_1) \right) h(x_1)h(x_2) dx_1 dx_2, \end{aligned}$$

where we used that both integrals cancel each other out on the region $(-\alpha, \alpha) \times (-\alpha, \alpha)$, and that h and h_{-12} are symmetric around 0.

Suppose that $p_2 > p_1$, and consider any $(x_1, x_2) \in (-\alpha, \alpha) \times [\alpha, +\infty)$. It follows from the identity

$$p_1x_1 + p_2x_2 = (p_1x_2 + p_2x_1) + (p_2 - p_1)(x_2 - x_1)$$

that

$$p_1x_1 + p_2x_2 > p_1x_2 + p_2x_1,$$

where the left-hand side (LHS) is always positive. The right-hand side (RHS) is either positive or negative, but smaller in absolute value than the LHS. Indeed, if the RHS is negative, then $x_1 < 0$, and

$$|p_1x_2 + p_2x_1| = -p_1x_2 + p_2|x_1| = -p_1|x_1| + p_2x_2 - (p_1 + p_2)(x_2 - |x_1|) < -p_1|x_1| + p_2x_2.$$

Thus,

$$|p_1x_1 + p_2x_2| > |p_1x_2 + p_2x_1|,$$

and due to the symmetry and unimodality of h_{-12} ,

$$h_{-12}(p_1x_1 + p_2x_2) < h_{-12}(p_1x_2 + p_2x_1),$$

unless both are zero. It follows that $X_2 | (X = X_0)$ is more peaked than $X_1 | (X = X_0)$, as needed. \square

Lemma 6. *Let the function f be continuous, symmetric around 0 and increasing on \mathbb{R}_+ , and let X_1, X_2 be unimodal continuous random variables that are symmetric around 0 and have bounded support. Then $\mathbb{E}[f(X_1)] < \mathbb{E}[f(X_2)]$ whenever X_1 is more peaked than X_2 .*

Proof. Denote by $h_i(x)$ and $H_i(x)$ the pdf and cdf of X_i , $i = 1, 2$. Then,

$$\begin{aligned} \frac{1}{2} \mathbb{E}[f(X_i)] &= \int_0^\infty f(x)h_i(x)dx \\ &= \left[f(x)(H_i(x) - 1) \right]_0^{+\infty} - \int_0^\infty (H_i(x) - 1)df(x) \\ &= \frac{1}{2}f(0) + \int_0^\infty (1 - H_i(x))df(x), \end{aligned}$$

where we have used integration by parts for Stieltjes integral, see e.g. Ok (2011). If X_1 is more peaked than X_2 , then $1 - H_1(x) < 1 - H_2(x)$ unless both are zero for all $x > 0$. It follows that $\mathbb{E}[f(X_1)] < \mathbb{E}[f(X_2)]$. \square

B.2 Proof of Proposition 2

Proof of Proposition 2. The objective (4) of the information-processing problem is a functional

$$\mathcal{L}(m'(\cdot), (\pi_J)_{J \in \mathcal{J}}) = \mathbb{E} \left[\sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J)} \mid r = s \right].$$

Since $\frac{p_J^2}{\pi_J m'^2(r_J)}$ is convex with respect to $(m'(r_J), \pi_J)$, the functional \mathcal{L} is convex:

$$\alpha \mathcal{L}(m'_1(\cdot), (\pi_{1,J})_J) + (1-\alpha) \mathcal{L}(m'_2(\cdot), (\pi_{2,J})_J) > \mathcal{L}(\alpha m'_1(\cdot) + (1-\alpha)m'_2(\cdot), (\alpha\pi_{1,J} + (1-\alpha)\pi_{2,J})_J),$$

for each $\alpha \in (0, 1)$ and any two perception strategies. Thus, the first-order conditions are sufficient for a global minimum of the information-processing problem.

Since the objective (4) is strictly monotone in $m'(\cdot)$, the constraint (5) is binding at the optimum. The Lagrangian of the constrained optimization problem (4)-(6) is

$$\begin{aligned} &\mathbb{E} \left[\sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J)} \mid r = s \right] + \lambda \left(\int_{\underline{r}}^{\bar{r}} m'(\tilde{r})d\tilde{r} - (\bar{m} - \underline{m}) \right) + \mu \left(\sum_{J \in \mathcal{J}} \pi_J - 1 \right) = \\ &\sum_{J \in \mathcal{J}} \int_{\underline{r}}^{\bar{r}} \frac{p_J^2}{\pi_J m'^2(\tilde{r}_J)} h_J(\tilde{r}_J) d\tilde{r}_J + \lambda \left(\int_{\underline{r}}^{\bar{r}} m'(\tilde{r})d\tilde{r} - (\bar{m} - \underline{m}) \right) + \mu \left(\sum_{J \in \mathcal{J}} \pi_J - 1 \right), \end{aligned}$$

where λ and μ are the Lagrange multipliers for (5) and (6), respectively. The first-order

condition (7) then follows by summing the derivatives w.r.t. $m'(\tilde{r})$, $\tilde{r} \in [\underline{r}, \bar{r}]$, of all the integrands in the last inline expression. Expressing $m'(\tilde{r})$ from (7) gives (8). Statement 1 of the proposition follows because by Lemma 4, each h_J is unimodal with the same mode as the unconditional reward density h .

Additionally, m' is symmetric around r_m since each h_J is symmetric around r_m . Further, m' is continuous since each h_J is continuous: since h_s is continuous on a compact interval, it is uniformly continuous, and thus the function $\tilde{r} \rightarrow \mathbb{E}[h_s(r) \mid r_J = \tilde{r}]$ is continuous; thus h_J is continuous.

The first-order condition of the information-processing problem with respect to π_J is for each $J \in \mathcal{J}$,

$$\left(\frac{p_J}{\pi_J}\right)^2 \mathbb{E}\left[\frac{1}{m'^2(r_J)} \mid r = s\right] = \mu. \quad (14)$$

Suppose $p_J < p_{J'}$. For Statement 2 it suffices to show that

$$\mathbb{E}\left[\frac{1}{m'^2(r_J)} \mid r = s\right] > \mathbb{E}\left[\frac{1}{m'^2(r_{J'})} \mid r = s\right]. \quad (15)$$

This indeed holds since by Lemma 5, $r_{J'} \mid (r = s)$ is more peaked than $r_J \mid (r = s)$ and the inequality (15) follows from Lemma 6 and from the fact that $1/m'^2(r)$ is continuous and symmetric around r_m and increasing above r_m . \square

B.3 Extension

We discuss here an extension of Proposition 2 to a setting in which the DM does not know the payoff-relevant partition \mathcal{J} and the probabilities $\mathbf{p} = (p_J)_{J \in \mathcal{J}}$ at the point of optimization of the encoding function but believes that $(\mathcal{J}, \mathbf{p})$ is drawn from a density $g(\mathcal{J}, \mathbf{p})$. The timing is as follows: first, the DM chooses the encoding function. Afterwards, \mathcal{J} and $\mathbf{p} \in \Delta(\mathcal{J})$ are realized, the DM observes these, and chooses sampling frequencies $(\pi_J(\mathcal{J}, \mathbf{p}))_{J \in \mathcal{J}}$. Finally, the DM observes the safe option s and the sequence of perturbed messages generated according to the chosen encoding function and the sampling frequencies, conditional on the encountered reward vector $\mathbf{r} \in [\underline{r}, \bar{r}]^{|\mathcal{J}|}$. As in Section 3, the rewards are iid from a symmetric unimodal reward density h and independent of \mathcal{J} and \mathbf{p} . The safe option s is drawn from a symmetric unimodal density h_s . The information-processing problem for this setting is as follows:

$$\min_{m'(\cdot) > 0, (\pi_J(\cdot))_{J \in \mathcal{J}} > 0} \mathbb{E}\left[\sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J(\mathcal{J}, \mathbf{p}) m'^2(r_J)} \mid \sum_{J \in \mathcal{J}} p_J r_J = s\right] \quad (16)$$

$$\text{s.t.:} \quad \int_{\underline{r}}^{\bar{r}} m'(\tilde{r}) d\tilde{r} \leq \bar{m} - \underline{m}$$

$$\sum_{J \in \mathcal{J}} \pi_J(\mathcal{J}, \mathbf{p}) = 1 \text{ for all } \mathcal{J} \text{ and } \mathbf{p} \in \Delta(\mathcal{J}).$$

The expectation in (16) is with respect to \mathcal{J} , \mathbf{p} , $(r_J)_{J \in \mathcal{J}}$ and s .

Proposition 2 extends to this setting. That is, its statement 1 holds for the optimal encoding function: it is S-shaped. Statement 2 holds for the optimal sampling frequencies $(\pi_J(\mathcal{J}, \mathbf{p}))_{J \in \mathcal{J}}$ for each realized $(\mathcal{J}, \mathbf{p})$: low-probability arms are over-sampled.

To see that statement 1 extends, observe that $m'(\cdot)$ minimizes the Lagrangian

$$\mathbb{E} \left[\mathbb{E} \left[\sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J(\mathcal{J}, \mathbf{p}) m'^2(r_J)} \mid \sum_{J \in \mathcal{J}} p_J r_J = s; \mathcal{J}, \mathbf{p} \mid \sum_{J \in \mathcal{J}} p_J r_J = s \right] \right. \\ \left. + \lambda \left(\int_{\underline{r}}^{\bar{r}} m'(r) dr - (\bar{m} - \underline{m}) \right) + \mathbb{E} \left[\mu(\mathcal{J}, \mathbf{p}) \left(\sum_{J \in \mathcal{J}} \pi_J(\mathcal{J}, \mathbf{p}) - 1 \right) \mid \sum_{J \in \mathcal{J}} p_J r_J = s \right] \right],$$

where the outer expectation in the first line is over $(\mathcal{J}, \mathbf{p})$ and the inner expectation is over \mathbf{r} and s ; the whole expectation equals the objective in (16) by the law of iterated expectations. The expectation in the second line is over $(\mathcal{J}, \mathbf{p})$.

The same steps as in the proof of Proposition 2 imply that $m'(\tilde{r})$ satisfies for each \tilde{r} the analogue of condition (7):

$$2 \mathbb{E} \left[\sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J(\mathcal{J}, \mathbf{p}) m'^3(\tilde{r})} h_J(\tilde{r}; \mathcal{J}, \mathbf{p}) \mid \sum_{J \in \mathcal{J}} p_J r_J = s \right] = \lambda,$$

where $h_J(r_J; \mathcal{J}, \mathbf{p})$ is the density of r_J conditional on the tie between the lottery value $\sum_{J' \in \mathcal{J}} p_{J'} r_{J'}$ and s for the given $(\mathcal{J}, \mathbf{p})$. The expectation is over $(\mathcal{J}, \mathbf{p})$.

Therefore, the analogue of (8) holds:

$$m'(\tilde{r}) \propto \left(\mathbb{E} \left[\sum_{J \in \mathcal{J}} \frac{p_J^2}{\pi_J(\mathcal{J}, \mathbf{p})} h_J(\tilde{r}; \mathcal{J}, \mathbf{p}) \mid \sum_{J \in \mathcal{J}} p_J r_J = s \right] \right)^{\frac{1}{3}}.$$

Since by Lemma 4, $h_J(\cdot; \mathcal{J}, \mathbf{p})$ is unimodal with maximum at the mode r_m of the ex ante reward density for all $(\mathcal{J}, \mathbf{p})$, m' is hump-shaped with maximum at r_m and hence m is S-shaped with its inflection point at r_m .

The proof of Statement 2 of Proposition 2 extends for each realization of $(\mathcal{J}, \mathbf{p})$ since the

analogue of the first-order condition (14) for $\pi_J(\mathcal{J}; \mathbf{p})$ continues to hold for each $(\mathcal{J}, \mathbf{p})$:

$$\frac{p_J^2}{\pi_J^2(\mathcal{J}, \mathbf{p})} \mathbb{E} \left[\frac{1}{m'^2(r_J)} \mid \sum_{J' \in \mathcal{J}} r_{J'} = s; \mathcal{J}, \mathbf{p} \right] = \mu(\mathcal{J}, \mathbf{p})$$

for each $J \in \mathcal{J}$. Continuity, unimodality and symmetry of $m'(\cdot)$ are unaffected by the studied extension. Thus, again, if $p_{J_1} < p_{J_2}$, then $r_{J_2} \mid (\sum_{J' \in \mathcal{J}} p_{J'} r_{J'} = s)$ is more peaked than $r_{J_1} \mid (\sum_{J' \in \mathcal{J}} p_{J'} r_{J'} = s)$ and the result follows from Lemma 6.

C Proofs of Propositions 3 and 4

Proposition 3 follows from Proposition 4 for $\mathcal{K} = \{\{1, \dots, I\}\}$.

Proof of Proposition 4. Let $f_{\mathbf{r}}(x)$ be the signal density conditional on the encountered lottery \mathbf{r} . That is, for signal $x = (\hat{m}, i)$, $f_{\mathbf{r}}(x) = \pi_i \varphi(\hat{m} - m(r_i))$ where φ is the standard normal density. Kullback-Leibler divergence of the signal densities for any two lotteries \mathbf{r}, \mathbf{r}' is

$$\begin{aligned} D_{\text{KL}}(f_{\mathbf{r}} \parallel f_{\mathbf{r}'}) &= \int_{\mathbb{R} \times \{1, \dots, I\}} f_{\mathbf{r}}(x) \ln \frac{f_{\mathbf{r}}(x)}{f_{\mathbf{r}'}(x)} dx \\ &= \sum_{i=1}^I \int_{\mathbb{R}} \pi_i \varphi(\hat{m} - m(r_i)) \ln \frac{\pi_i \varphi(\hat{m} - m(r_i))}{\pi_i \varphi(\hat{m} - m(r'_i))} d\hat{m} \\ &= \sum_{i=1}^I \pi_i \int_{\mathbb{R}} \varphi(\hat{m} - m(r_i)) \ln \frac{\varphi(\hat{m} - m(r_i))}{\varphi(\hat{m} - m(r'_i))} d\hat{m} \\ &= \sum_{i=1}^I \pi_i D_{\text{KL}}(\varphi_{m(r_i)} \parallel \varphi_{m(r'_i)}) \\ &= \frac{1}{2} \sum_{i=1}^I \pi_i (m(r_i) - m(r'_i))^2. \end{aligned}$$

where $\varphi_m(\hat{m}) = \varphi(\hat{m} - m)$ is the density of the perturbed message \hat{m} conditional on the unperturbed message m . The last equality follows from the fact that the Kullback-Leibler divergence of two Gaussian densities with means μ_1, μ_2 and variances equal to 1 is $(\mu_1 - \mu_2)^2 / 2$, e.g. Johnson and Orsak (1993).

Let

$$\mathbf{q} = \arg \min_{\mathbf{r}' \in \mathcal{A}_{\mathcal{K}}} D_{\text{KL}}(f_{\mathbf{r}} \parallel f_{\mathbf{r}'}) = \arg \min_{\mathbf{r}' \in \mathcal{A}_{\mathcal{K}}} \sum_{i=1}^I \pi_i (m(r_i) - m(r'_i))^2.$$

This minimizer $\mathbf{q} = (q_i)_i$ is unique and satisfies for each arm $i = 1, \dots, I$,

$$\begin{aligned} m(q_i) &= \arg \min_{m \in [\underline{m}, \bar{m}]} \sum_{j \in J(i)} \pi_j (m(r_j) - m)^2 \\ &= \sum_{j \in J(i)} \frac{\pi_j}{\pi_{J(i)}} m(r_j), \end{aligned}$$

where $J(i)$ is the element of the partition \mathcal{K} that contains i .

The estimated lottery value q_n^z , $z \in \{ML, B\}$, almost surely converges to $\sum_{i=1}^I p_i q_i$. For the ML estimate, this follows from White (1982) who proves that the MLE almost surely converges to the minimizer of the Kullback-Leibler divergence (provided it is unique). For the Bayesian estimate, the result follows from Berk (1966) who proves that the posterior belief almost surely converges in probability to the atom on the minimizer of the Kullback-Leibler divergence (again, provided the minimizer is unique). \square

D Proofs for Subsection 4.3

We use the next lemma in the proof of Proposition 5.

Lemma 7. *Let $\psi_n(\mathbf{x}) : [\underline{r}, \bar{r}]^I \rightarrow \mathbb{R}$ be a sequence of continuous functions uniformly converging to a function $\psi(\mathbf{x})$ which has a unique minimizer \mathbf{x}^* . Then, the random variable X_n with pdf equal to $\alpha_n \exp(-n\psi_n(\mathbf{x}))$, where α_n is the normalization factor, converges to \mathbf{x}^* in probability as $n \rightarrow \infty$.*

Proof. We need to prove that for every $\delta > 0$, the probability $P(X_n \in B_\delta) \rightarrow 1$ as $n \rightarrow \infty$, where B_δ is the open Euclidean δ -ball centered at \mathbf{x}^* . Fix $\delta > 0$ and define

$$d = \min_{\mathbf{x} \in [\underline{r}, \bar{r}]^I \setminus B_\delta} \{\psi(\mathbf{x}) - \psi(\mathbf{x}^*)\}.$$

The minimum exists as ψ is continuous and the set $[\underline{r}, \bar{r}]^I \setminus B_\delta$ is closed. Additionally, $d > 0$ since \mathbf{x}^* is the unique minimizer of ψ on $[\underline{r}, \bar{r}]^I$.

Because the convergence $\psi_n \rightarrow \psi$ is uniform, for any $d' > 0$ there exists $n_{d'} \in \mathbb{N}$ such that $|\psi_n(\mathbf{x}) - \psi(\mathbf{x})| < d'$ for all $\mathbf{x} \in [\underline{r}, \bar{r}]^I$ and $n \geq n_{d'}$. Consider $n \geq n_{d/4}$. Because $\psi_n(\mathbf{x}) \geq \psi(\mathbf{x}) - \frac{d}{4} \geq \psi(\mathbf{x}^*) + \frac{3d}{4}$ for \mathbf{x} outside of the ball B_δ , the probability density of X_n is at most $\alpha_n \exp(-n\psi(\mathbf{x}^*) - \frac{3d}{4}n)$. This implies,

$$P(X_n \notin B_\delta) \leq \tilde{\alpha}_n \exp\left(-\frac{3d}{4}n\right) (\bar{r} - \underline{r})^I, \quad \text{where } \tilde{\alpha}_n := \alpha_n \exp(-n\psi(\mathbf{x}^*)). \quad (17)$$

We conclude by establishing an upper bound for $\tilde{\alpha}_n$. Given $\delta > 0$, let $\delta' > 0$ be such that $\psi(\mathbf{x}) \leq \psi(\mathbf{x}^*) + d/4$ for all $\mathbf{x} \in B_{\delta'} \cap [\underline{r}, \bar{r}]^I$. Existence of such δ' follows from the continuity of ψ . Then, $\psi_n(\mathbf{x}) \leq \psi(\mathbf{x}) + \frac{d}{4} \leq \psi(\mathbf{x}^*) + \frac{d}{2}$ for all $\mathbf{x} \in B_{\delta'} \cap [\underline{r}, \bar{r}]^I$ and $n > n_{d/4}$. Thus the probability density of X_n is at least $\tilde{\alpha}_n \exp(-\frac{d}{2}n)$ on this set. It follows that,

$$1 \geq P(X_n \in B_{\delta'}) \geq \tilde{\alpha}_n \exp\left(-\frac{d}{2}n\right) b',$$

where $b' > 0$ is the volume of the set $B_{\delta'} \cap [\underline{r}, \bar{r}]^I$. Substituting the implied upper bound on $\tilde{\alpha}_n$ into (17) gives

$$P(X_n \notin B_{\delta}) \leq \exp\left(-\frac{d}{4}n\right) \frac{(\bar{r} - \underline{r})^I}{b'}.$$

Since the right-hand vanishes as $n \rightarrow \infty$, the claim follows. \square

Proof of Proposition 5. Let $\mathbf{m}_n = (m_{i,n})_{i=1}^I$ be the vector of the averages of $a\pi n$ perturbed messages received for each arm i . Since the encoding errors are standard normal, $m_{i,n} \mid r_i \sim \mathcal{N}\left(m(r_i), \frac{1}{a\pi n}\right)$. By Bayes' Rule, the posterior density of each lottery $\mathbf{r}' \in [\underline{r}, \bar{r}]^I$, is for given \mathbf{m}_n , proportional to

$$\varrho_n(\mathbf{r}') \prod_{i=1}^I \varphi\left((m_{i,n} - m(r'_i))\sqrt{a\pi n}\right) \propto \exp\left(-n\psi(\mathbf{r}'; \mathbf{m}_n)\right),$$

where \propto denotes equality modulo normalization and

$$\psi(\mathbf{r}'; \mathbf{m}) := \frac{1}{2} \sum_{i=1}^I \left(\frac{\sigma(\mathbf{r}')}{\Delta} + a\pi_i (m(r'_i) - m_i)^2 \right).$$

The first inline equality follows from the specification of the prior ϱ_n in (9).

Since $m_{i,n} \rightarrow m(r_i)$ (a.s.), $\psi(\mathbf{r}'; \mathbf{m}_n)$ converges to $\psi(\mathbf{r}'; (m(r_i))_i)$, uniformly in \mathbf{r}' . Additionally, $\psi(\mathbf{r}'; (m(r_i))_i)$ as a function of \mathbf{r}' has the unique minimizer $\mathbf{q}^*(\mathbf{r})$ by assumption. Lemma 7 implies that the posterior formed given \mathbf{m}_n converges in probability to $\mathbf{q}^*(\mathbf{r})$. Since the support of the rewards is bounded, convergence in probability implies convergence in expected value, and thus the Bayesian estimate $\mathbb{E}[\mathbf{r} \mid \mathbf{m}_n]$ converges to $\mathbf{q}^*(\mathbf{r})$. \square

Proof of Proposition 6. By Proposition 5, the Bayesian estimate of \mathbf{r} converges to $\mathbf{q}^*(\mathbf{r})$. We write $\mathbf{q}^* = (q_i^*)_{i=1}^I$ as an abbreviation for $\mathbf{q}^*(\mathbf{r})$ and let $q^* = \sum_i p_i q_i^*$. The first-order condition applied to the minimization in (10) implies,

$$(q_i^* - q^*) + a\Delta(m(q_i^*) - m(r_i))m'(q_i^*) = 0, \tag{18}$$

for all $i = 1, \dots, I$, where we have used that $\pi_i = p_i$ and $\sum_i^I p_i (q_i^* - q^*) = q^* - q^* = 0$. We write σ^2 for $\sigma^2(\mathbf{r})$ and $\sigma^{*2} := \sum_{i=1}^I p_i (q_i^* - q^*)^2$ for the variance of \mathbf{q}^* .

We will prove the following claims (see Footnote 12 for the definition of the $o(\cdot)$ convention.)

Claim 1: Any function that is $o(r_i - r)$ or $o(q_i^* - r)$ is also $o(\sigma)$.

Claim 2: $q^* = r + o(\sigma)$.

Claim 3: $\sigma^{*2} = \frac{z(r)^2}{(1+z(r))^2} \sigma^2 + o(\sigma^2)$.

Claim 4: $q^* = r + \frac{1}{2} \frac{m''(r)}{m'(r)} \left(\sigma^2 + \left(\frac{2}{z(r)} - 1 \right) \sigma^{*2} \right) + o(\sigma^2)$.

To prove Claim 1, we provide a bound on the distance of r_i and r_i' from r . It follows from definition of σ^2 that $(r_i - r)^2 \leq \sigma^2/p_i$, and thus $|r_i - r| \leq \sigma/\sqrt{p_i}$. Therefore, any function that is $o(r_i - r)$ is also $o(\sigma)$. Bounding $|q_i^* - r|$ is complicated by the fact that \mathbf{q}^* is defined implicitly. We first establish a bound on $|q^* - r|$. Define \underline{m}' and \overline{m}' to be the minimum and the maximum of $m'(\cdot)$ on $[r, \bar{r}]$, respectively, and let $\underline{z} = a\Delta\underline{m}'^2$, $\bar{z} = a\Delta\overline{m}'^2$. We have $0 < \underline{m}' \leq \overline{m}' < +\infty$ and $0 < \underline{z} \leq \bar{z} < +\infty$ since $m'(\cdot)$ is continuous and strictly positive on the closed interval $[r, \bar{r}]$.

For fixed values of \mathbf{r} and \mathbf{q}^* define $z_i \in \mathbb{R}$ by

$$a\Delta m'(q_i^*) (m(q_i^*) - m(r_i)) = (q_i^* - r_i) z_i$$

whenever $q_i^* \neq r_i$, and $z_i := a\Delta m'^2(r_i)$ otherwise. It follows from its definition that $z_i \geq \underline{z}$ for all i . Then, equation (18) can be written as

$$0 = (q_i^* - q^*) + (q_i^* - r_i)z_i = (1 + z_i)(q_i^* - q^*) - (r_i - q^*)z_i,$$

and thus,

$$q_i^* - q^* = \frac{z_i}{1+z_i}(r_i - q^*) = \frac{z_i}{1+z_i}(r_i - r) + \frac{z_i}{1+z_i}(r - q^*). \quad (19)$$

Summing up the last equation weighted by p_i over i gives

$$0 = \sum_{i=1}^I \left(p_i \frac{z_i}{1+z_i} (r_i - r) \right) + (r - q^*) \sum_{i=1}^I \left(p_i \frac{z_i}{1+z_i} \right),$$

in which $0 < \frac{\underline{z}}{1+\underline{z}} \leq \frac{z_i}{1+z_i} < 1$. The triangle inequality implies

$$|q^* - r| \leq \frac{1+\bar{z}}{\underline{z}} \sum_{i=1}^I p_i |r_i - r| \leq \frac{1+\bar{z}}{\underline{z}} \sigma \sum_{i=1}^I \sqrt{p_i} \leq \frac{1+\bar{z}}{\underline{z}} I \sigma.$$

Returning to equation (19),

$$|q_i^* - r| \leq \frac{z_i}{1+z_i}|r_i - r| + \frac{z_i}{1+z_i}|r - q^*| + |q^* - r| < |r_i - r| + 2|r - q^*| \leq \left(p_i^{-1/2} + 2\frac{1+\underline{z}}{\underline{z}}I\right)\sigma.$$

We conclude that $|q_i^* - r| \leq \left(p_i^{-1/2} + 2\frac{1+\underline{z}}{\underline{z}}I\right)\sigma$ for any $\mathbf{r} \in [\underline{r}, \bar{r}]^I$, and thus any function that is $o(q_i^* - r)$ is also $o(\sigma)$. This establishes Claim 1.

We will prove the remaining claims by taking first- and second-order approximations of the first-order condition (18) for $\sigma > 0$ small. Since $m(\cdot)$ is twice differentiable, the functions m and m' can be expressed using first-order Taylor approximations around r :

$$\begin{aligned} m(r_i) &= m(r) + m'(r)(r_i - r) + o(\sigma), \\ m(q_i^*) &= m(r) + m'(r)(q_i^* - r) + o(\sigma), \\ m'(q_i^*) &= m'(r) + m''(r)(q_i^* - r) + o(\sigma), \end{aligned}$$

where we used Claim 1 to replace $o(r_i - r)$ and $o(q_i^* - r)$ by $o(\sigma)$. Equation (18) then implies

$$\begin{aligned} 0 &= (q_i^* - q^*) + a\Delta \left(m'(r)(q_i^* - r_i) + o(\sigma) \right) \left(m'(r) + m''(r)(q_i^* - r) + o(\sigma) \right) \\ &= (q_i^* - q^*) + a\Delta m'^2(r)(q_i^* - r_i) + o(\sigma), \end{aligned}$$

where we used that $(q_i^* - r_i)(q_i^* - r) = o(\sigma)$. The last inline equation can be written as

$$0 = (q_i^* - q^*) + z(r)(q_i^* - r_i) + o(\sigma). \quad (20)$$

Summing up these equations weighted by p_i , we get $0 = z(r)(q^* - r) + o(\sigma)$. Thus $|q^* - r| \leq \frac{1}{\underline{z}}o(\sigma)$, as needed for Claim 2.

We rewrite (20) as

$$(1 + z(r))(q_i^* - q^*) = z(r)(r_i - r) + z(r)(r - q^*) + o(\sigma) = z(r)(r_i - r) + o(\sigma),$$

where the second equality follows from Claim 2. Squaring both sides of the equation and summing up the equations weighted by p_i , we get

$$(1 + z(r))^2\sigma^{*2} = z^2(r)\sigma^2 + o(\sigma^2),$$

where we used that $z(r) \leq \bar{z}$ and thus $z(r)(r_i - r)o(\sigma)$ is $o(\sigma^2)$. Claim 3 follows.

To prove Claim 4, we use the second-order Taylor approximation of $m(\cdot)$ around r :

$$\begin{aligned} m(q_i^*) &= m(r) + m'(r)(q_i^* - r) + \frac{1}{2}m''(r)(q_i^* - r)^2 + o(\sigma^2) \\ m(r_i) &= m(r) + m'(r)(r_i - r) + \frac{1}{2}m''(r)(r_i - r)^2 + o(\sigma^2). \end{aligned}$$

This implies the second-order approximation of the equation (18),

$$\begin{aligned} 0 &= (q_i^* - q^*) + a\Delta \left(m'(r)(q_i^* - r_i) + \frac{1}{2}m''(r)\left((q_i^* - r)^2 - (r_i - r)^2\right) + o(\sigma^2) \right) \\ &\quad \cdot \left(m'(r) + m''(r)(q_i^* - r) + o(\sigma) \right), \end{aligned}$$

which we rewrite as

$$0 = (q_i^* - q^*) + z(r) \left((q_i^* - r_i) + \frac{1}{2} \frac{m''(r)}{m'(r)} \left((q_i^* - r)^2 - (r_i - r)^2 \right) \right) \left(1 + \frac{m''(r)}{m'(r)} (q_i^* - r) \right) + o(\sigma^2).$$

Summing up these equations weighted by p_i and dividing by $z(r)$, we arrive at

$$0 = (q^* - r) - \frac{1}{2} \frac{m''(r)}{m'(r)} \left(\sigma^2 - \sigma^{*2} + 2 \sum_{i=1}^I p_i (r_i - q_i^*) (q_i^* - r) \right) + o(\sigma^2). \quad (21)$$

Expressing $q_i^* - r_i$ from (20) allows us to write

$$\sum_{i=1}^I p_i (r_i - q_i^*) (q_i^* - r) = \frac{1}{z(r)} \sum_{i=1}^I p_i (q_i^* - r)^2 + o(\sigma^2) = \frac{1}{z(r)} \sigma^{*2} + o(\sigma^2),$$

where we used that $r = q^* + o(\sigma)$ for the second equality. Substituting the last inline equation back into (21) completes the proof of Claim 4.

Finally, substituting for σ^{*2} from Claim 3 into the expression from Claim 4 gives

$$\begin{aligned} q^* &= r + \frac{1}{2} \frac{m''(r)}{m'(r)} \left(1 + \left(\frac{2}{z(r)} - 1 \right) \frac{z(r)^2}{(1+z(r))^2} \right) \sigma^2 + o(\sigma^2) \\ &= r + \frac{1}{2} \frac{m''(r)}{m'(r)} \left(1 + \frac{2z(r) - z(r)^2}{(1+z(r))^2} \right) \sigma^2 + o(\sigma^2), \end{aligned}$$

and using $1 + \frac{2z(r) - z(r)^2}{(1+z(r))^2} = \frac{1+4z(r)}{(1+z(r))^2}$, we obtain (11), concluding the proof. \square