Endogenous Risk Attitudes*

Nick Netzer
University of Zurich

Arthur Robson
Simon Fraser University

Jakub Steiner
University of Zurich and Cerge-Ei

Pavel Kocourek
Cerge-Ei

May 6, 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 DM fully understands the decision problem. If, however, the DM 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 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 the SSHRC grant 435-2013-0426 (Robson) and by the 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 endogenous risk attitudes and probability weighting as the joint 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 at 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 akin to those from prospect theory when the DM applies a simplifying model to the encountered risk. 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 safe option. She knows the probabilities of the lottery arms, observes the value of the safe option but faces a friction in processing of the lottery rewards. She learns about the reward vector by sampling signals (from her own memory, experience of others, 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 makes the a posteriori optimal choice. We characterize the ex ante optimal perception strategy, and derive the implications of any perception strategy for behavior.

Choice over the perception strategies is a specific form of the 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. Since she reads the needle position with noise, 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 tractability and by the fact that biases that survive this limit are likely to be large even when stakes of the decision problems are large relative to the perception frictions. We first prove that the expected loss from misperception, relative to the choice under complete information, is approximately the mean squared error in perception of the lottery value, integrated across all decision problems in which the lottery value ties with the safe option. The conditioning on ties arises endogenously. The accuracy of perception has instrumental value for choice, and choice is trivial except where the values of two options are nearly equal, given information is nearly complete.

We then show for unimodal symmetric reward densities that an S-shaped encoding function and over-sampling of small-probability lottery arms jointly minimize the mean squared error over ties. 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 the 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 non-marginal decision problems. The DM with an S-shaped encoding function therefore struggles to comprehend the low-probability contingencies in relatively many marginal decision problems, since these involve tail rewards relatively often. It is optimal to compensate for this by allocating disproportionate attention to unlikely arms by over-sampling them. 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. 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. Assume that the needle trembles due to the stochasticity of the input and the engineer observes the distribution of the needle position. If she correctly understands that the input is stochastic, then she inverts each observed needle position to obtain the corresponding input value, thus 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 concludes that a 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 at 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.

This result holds irrespective of whether the perception strategy was chosen optimally or not. Combined with our previous insights on optimal perception, it provides an explanation for the adaptive risk attitudes of prospect theory. 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.

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. She might also have been framed to believe that the decision problem involves less risk than it does. Alternatively, she might not be aware that she is omitting an explanatory variable from her econometric model of the reward.

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 encoding function 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. 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.

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.[1] A large literature in brain sciences and psychology views perception as information processing via a limited channel and studies the optimal

[1] See Woodford (2020) for a review of psychophysics from an economic perspective.

encoding of stimuli for a given channel capacity (see Attneave (1954) and Barlow (1961) for early contributions). 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.

Kahneman and Tversky (1979) rely on analogies to the adaptation of sensory perception to rationalize their S-shaped value function. Within 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 e.g. Bethge, Rotermund, and Pawelzik (2002) and Wang, Stocker, and Lee (2016)). Economics can help here by providing a microfoundation 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 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]

Two recent economic papers study risk attitudes stemming from reward encoding in pres-

[2] Herold 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.

ence of non-vanishing encoding noise. Khaw, Li, and Woodford (2018) show theoretically and verify experimentally that logarithmic stochastic encoding which is then optimally decoded generates risk attitudes in an effect akin to reversion to the mean. Frydman and Jin (2019) allow for endogenous encoding of the lottery reward and show both theoretically and experimentally that this encoding adapts to the distribution of the decision problems and the adaptation affects choice.

In contrast to these papers, we focus on the limit of vanishing encoding noise. 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 encoding and behavior. While the impact of encoding on behavior vanishes when the decoding model is well specified (as in the two papers cited above), the encoding implications for behavior remain substantial if misspecified decoding oversimplifies the structure of the risk.

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 the 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, Kőszegi, and Strack (2018) characterize a vicious circle of overconfident learning, Molavi (2019) studies the macroeconomic consequences of misspecification, Frick, Iijima, and Ishii (2021) rank the short- and long-run costs of various forms of misspecification and Eliaz and Spiegler (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 of arms $i = 1, \dots, I$, $I \geq 1$, where each arm i has a positive probability p_i and pays reward $r_i \in [\underline{r}, \overline{r}]$ where $\underline{r} < \overline{r}$ are arbitrary bounds. For the sake of simplicity, we fix the vector of the arm probabilities and let the DM observe it frictionlessly. The lottery rewards and the safe option value are generated randomly. The DM observes the exact value of the safe option; we focus on a friction in perception of the lottery rewards. We let $\mathbf{r} = (r_i)_i \in [\underline{r}, \overline{r}]^I$ 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 *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. 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 can be an appropriate concave function of the monetary prizes if the DM chooses monetary prospects and money has diminishing returns to personal production.

The DM does not know the lottery $\mathbf{r}$ and estimates it from a sequence of n signals, where each signal is a monotone transformation of one of the arm rewards perturbed with an additive noise: the DM observes signals $x_k = (\hat{m}_k, i_k)$, $k = 1, \ldots, n$. We refer to the first element of the signal, $\hat{m}_k$, as to the *perturbed message*. The second element, 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 the 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 iid standard normal. Each sampled arm i_k is one of the lottery arms $i = 1, \ldots, I$, iid across k with positive probabilities π_i independently from the preceding signal history. The function $m : [\underline{r}, \overline{r}] \longrightarrow [\underline{m}, \overline{m}]$ is a bijection, strictly increasing and continuously differentiable; we refer to it as to the *encoding function*. We dub π_i to be *sampling frequencies* and refer to $(m(\cdot), (\pi_i)_i)$ as to the *perception strategy*. The size of the sample, n, is exogenous.

After she has observed the signals x_k, $k = 1 \ldots n$, 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 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 $\mathcal{A} \subseteq [\underline{r}, \overline{r}]^I$ of lotteries she anticipates, concludes that she has encountered the lottery

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

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

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

We study decision-makers who employ simplifying models of risk in the spirit of the *small world* from Savage (1954). That is, 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

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

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 any reward vector from $[\underline{r}, \overline{r}]^I$.

3 Optimal Perception in a Small World

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 the lottery arms that the DM deems as payoff-equivalent: $r_i = r_j$ for all arms $i, j \in J$, $J \in \mathcal{J}$. Since the distinction between arms in each J is subjectively redundant for this DM, 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 for the small world consists of the increasing encoding function $m(\cdot)$ and interior sampling frequencies $(\pi_J)_J \in \Delta(\mathcal{J})$. In Section 4, we will allow for the possibility that the DM's small-world model is, in fact, a misspecified model of the grand world; either because the world has become more complex or because it was more complex than the DM thought to start with.

The DM optimizes her perception strategy ex ante for a given distribution of the decision problems. Specifically, the rewards r_J, $J \in \mathcal{J}$ 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 $[\underline{r}, \overline{r}]$.[4] We characterize the expected loss for general perception

[4]Since s may have a distinct density from that of r_J, the safe option may be interpreted as a value of

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 the sampled signals to be large and abstract in this section from uncertainty over the number of the perturbed messages sampled for each arm and from divisibility issues. That is, we suppose the number of the 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, the DM believes that $m_{J,n} - m(r_J)$ 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.

We first consider how the DM forms the maximum-likelihood estimate (MLE) and then note that the Bayesian estimate differs from the MLE only by a negligible term when n is large. The DM forms the MLE $\mathbf{q}_n = (q_{J,n})_J \in [\underline{r}, \overline{r}]^{|\mathcal{J}|}$ of the lottery $\mathbf{r}$. Since the signals are conditionally independent across arms J, $q_{J,n}$ depends only the data collected on the arm J. If $m_{J,n} \in (\underline{m}, \overline{m})$, then the MLE of the unperturbed message $m(r_J)$ is the average perturbed message $m_{J,n}$ and thus,[5]

$$q_{J,n} = m^{-1}(m_{J,n}) \,. \tag{2}$$

To derive how the perception strategy affects the error of the estimates $q_{J,n}$ for large n, we write $m_{J,n} = m(r_J) + \frac{\varepsilon_J}{\sqrt{\pi_J n}}$ where ε_J is $\mathcal{N}(0,1)$.[6] Using the first-order Taylor approximation for $m^{-1}(\cdot)$ around $m(r_J)$, we get for each realization of ε_J that

$$q_{J,n} - r_J = \frac{\varepsilon_J}{m'(r_J)\sqrt{\pi_J n}} + O\left(\frac{1}{n}\right). \tag{3}$$

Hence, the error of the MLE is of order $\frac{1}{\sqrt{n}}$ for each realization of ε_J and it is mitigated by a steep slope of the encoding function and by a high sampling frequency.

The estimate in (2) is also, except for a higher-order term, the Bayesian estimate of r_J of the DM who a priori believes that all the rewards are drawn iid from density h. More precisely, we show in Appendix A that $\mathrm{E}[r_J \mid m_{J,n}] = m^{-1}(m_{J,n}) + O(1/n)$ for each realization of the average perturbed message $m_{J,n} \in (\underline{m}, \overline{m})$. That is, the ML and Bayesian

alternative lottery, etc.

[5]If $m_{J,n} < \underline{m}$ or $m_{J,n} > \overline{m}$, then the MLE estimate of r_J is $\underline{r}$ and $\overline{r}$, respectively. Note, that $m_{J,n} \in (\underline{m}, \overline{m})$ with probability approaching 1 as n diverges.

[6]That is, we are coupling the random error terms $m_{J,n} - m(r_J)$ across distinct values of n. We are using the fact that the joint distribution of the error terms across n is not specified and hence we can choose the error terms $m_{J,n} - m(r_J) = \frac{\varepsilon_J}{\sqrt{\pi_J n}}$ to be generated by a single random variable ε_J for all n. See e.g. Lindvall (2002) on coupling.

estimates are the same up to the leading term of order $1/\sqrt{n}$. Further, letting $r = \sum_J p_J r_J$ and $q_n^{ML} = \sum_J p_J q_{J,n}$ be the true overall lottery value and its MLE, respectively, the error $q_n^{ML} - r$ is also of order $1/\sqrt{n}$. Subsequently, q_n refers to q_n^{ML}.

Let

$$L(n) = \mathrm{E}\left[\max\{r, s\} - \mathbb{1}_{q_n > s} r - \mathbb{1}_{q_n \leq s} s\right]$$

be the ex ante expected loss relative to the choice under complete information. The expectation is over s and the true and estimated lottery values r and q_n.

Proposition 1. *The loss satisfies*

$$\lim_{n\to\infty} nL(n) = \frac{1}{2}\,\mathrm{E}\left[\sum_{J\in\mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J)} h_s(r)\right], \tag{4}$$

where the expectation is with respect to $\mathbf{r} = (r_J)_{J\in\mathcal{J}}$.

The limit loss characterization in (4) has an intuitive interpretation. It is the mean squared error $n\,\mathrm{E}(q_n - r)^2$ in the perception of the lottery value (rescaled by n) integrated across 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 quickly vanishes with increasing n and small perception errors distort choice only in the decision problems in which an approximate tie, $r \approx s$, arises. In the limit, the set of the 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 as r and q_n, 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. When n is large, and hence the error is small, this occurs with approximate probability $h_s(r)|q_n - r|$. Conditional on the choice being distorted like this, the expected loss is approximately $|q_n - r|/2$ since s is approximately uniformly distributed between r and q_n. Hence the overall loss over all s is approximately $h_s(r)\left(q_n - r\right)^2/2$.

The mean squared error of perception depends on the realization of the lottery rewards. Since by (3), the distribution of $q_{J,n}$ is approximately $\mathcal{N}\left(r_J, \frac{1}{\pi_J n m'^2(r_J)}\right)$, the expected loss for a given lottery is approximately

$$\frac{1}{2n} h_s(r) \sum_{J\in\mathcal{J}} \frac{p_J^2}{\pi_J m'^2\left(r_J\right)},$$

which is the half of the MSE of perceived lottery value multiplied by the likelihood $h_s(r)$

that s ties with r. The overall loss in (4) integrates these contributions across all lotteries.

Motivated by the limit loss characterization from (4) we define the *information-processing problem* as follows.

$$\min_{m'(\cdot)>0,(\pi_J)_J>0} \quad \int_{[\underline{r},\overline{r}]^{|\mathcal{J}|}} \sum_{J\in\mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J)} h_s(r) \prod_{J\in\mathcal{J}} (h(r_J)\, dr_J) \tag{5}$$

$$\text{s.t.:} \quad \int_{\underline{r}}^{\overline{r}} m'(\tilde{r}) d\tilde{r} \leq \overline{m} - \underline{m} \tag{6}$$

$$\sum_{J\in\mathcal{J}} \pi_J = 1. \tag{7}$$

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 (6) is required by the finite range of the encoding function – the encoding function cannot be steep everywhere. Constraint (7) together with the restriction to positive sampling frequencies requires $(\pi_J)_J$ to be a probability distribution over $J \in \mathcal{J}$; the DM must also treat sampling frequencies as a scarce resource. A special case arises when the lottery has one arm; then, optimization over the sampling frequencies is trivial.

3.2 Optimization

Reward density $f(x)$ is *unimodal* and *symmetric* around $r^* = (\underline{r} + \overline{r})/2$ if it is strictly decreasing on $(r^*, \overline{r}]$ and $f(r^* + x) = f(r^* - x)$ whenever $|x| \leq (\overline{r} - \underline{r}+)/2$. 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 unimodal and symmetric around r^*, then*

1. *the optimal encoding function m is S-shaped: It is convex below and concave above r^*, 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. 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. See the extension on p. 36 in Appendix D.

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

$$h_J\left(r_J\right) \quad = \quad \frac{h\left(r_J\right)\int_{[\underline{r},\overline{r}]^{|\mathcal{J}|-1}} h_s(r)\prod_{J'\neq J}\left(h(r_{J'})dr_{J'}\right)}{\mathrm{E}[h_s(r)]}$$

be the density of the reward r_J at the arm J conditional on the event of tie between the lottery value and the safe option, $r = s$. The first-order condition for the optimal slope of the encoding function $m'(\tilde{r})$ is then shown to be

$$\sum_{J\in\mathcal{J}} \frac{p_J^2}{\pi_J m'^3(\tilde{r})} h_J\left(\tilde{r}\right) = \lambda \tag{8}$$

for each reward value $\tilde{r}$, where λ is the shadow price of the constraint (6) multiplied by a constant factor. The left-hand side of (8) is proportional to 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 (i) 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 (ii) 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 (6) implies that, at the optimum, the marginal benefit of a slope increase is equal across all reward values $\tilde{r}$.

We show in in Appendix C.1 that the density of the reward conditional on tie is, for each arm, unimodal with the same mode as the unconditional reward density. The first-order condition (8) then implies that the optimal slope, $m'(\tilde{r})$, is proportional to a monotone transformation of a sum of unimodal functions that all have their maxima at the unconditional reward mode,

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

establishing Claim 1 of the Proposition.

Let us now turn to Claim 2. We show in Appendix using an argument based on diminishing returns to sampling, 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 tie because a marginal change of the sampling frequency affects choice only at ties. Conditional on a tie, the density of reward r_J at each arm J is concentrated towards the unconditional modal reward since these modal rewards lead to ties relatively often. This effect is, however, heterogenous across arms. The condition $\sum_J p_J r_J = s$ is relatively uninformative about the low-probability rewards and hence the posterior distributions of these are relatively spread-out at ties. 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 optimally compensates for the expected errors in measurement at ties by over-sampling the low-probability arms.[7]

Optimal over-sampling arises from our microfoundation of the optimization objective (5). 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 certain 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 tie is proportional to $h^2(r)$ for $|\mathcal{J}| = 1$ and the result then follows from the first-order condition (8).[8]

Asymptotically, the perception strategy has no implications for behavior in the absence of further frictions because 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 small-world 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 to perception data exhibit risk attitudes dictated by properties of the perception strategy.

4 Behavior

The implications of the perception strategy for behavior depend on the DM's anticipation of the risk. Consider a vivid example from Savage (1954). The DM is considering the purchase

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

[8] The perception friction assumed in Netzer differs from that assumed here in technical details. Netzer studies the limit of increasingly fine discretizations of the reward space, whereas we take the limit of vanishing additive encoding noise. Like Netzer, we can dispose with the assumption of the Gaussian rewards when $|\mathcal{J}| = 1$ since the conditional density $r \mid (r = s)$ is simple to analyze in this case.

of a convertible car. Not purchasing the convertible guarantees a certain payoff of s whereas 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 $k = 1, \ldots, n$, 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, or the experience of her peers, 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 a 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. By White's (1982) result on asymptotic misspecified estimation applied to Gaussian additive errors, the single message that maximizes the likelihood of the observed data 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 small world in which all the lotteries were measurable with respect to the coarsest partition $\mathcal{J} = \{\{1, 2\}\}$ of the set of the arms $\{1, 2\}$. As outlined in the previous section, they both then optimized their encoding functions in this environment. Afterwards, their environments became more complex so they currently encounter lotteries where $r_1 \neq r_2$. Although neither DM can adjust her encoding function in the short run, 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 ad-

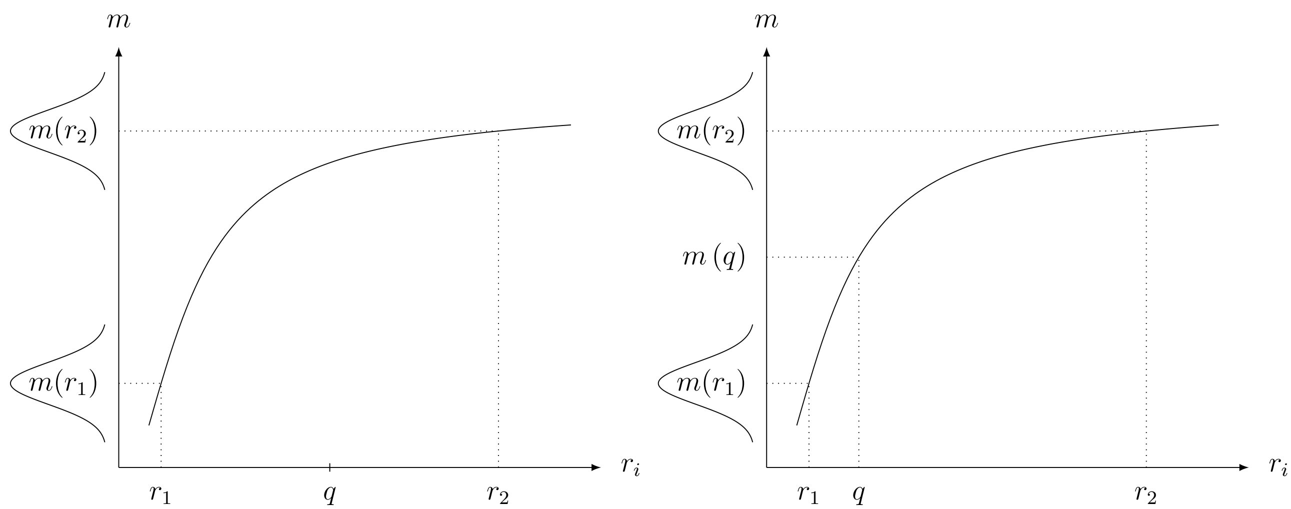

Figure 1: Example: Asymptotic estimated lottery value q of the fine (left) and the coarse (right) DMs.

justment and continues to anticipate riskless lotteries only. Or, both DMs evolved in a risky environment with partition $\mathcal{J} = \{\{1\}, \{2\}\}$ and optimized their encoding functions for such binary lotteries. Afterwards, the coarse DM was (incorrectly) assured that her next lottery will be riskless, while the fine DM was not told this. Alternatively, both DM's know they may 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 to apply the coarse procedure to each arm separately and then compute 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 DM has some but only partial awareness of the risk she faces. Subsection 4.3 focuses on a DM who anticipates some risk but believes that large differences between rewards across the lottery arms are a priori unlikely. If such a DM observes data generated by a risky lottery she then underestimates the degree to which the arms rewards differ. This misjudgement generates non-trivial risk-attitudes akin to but somewhat milder than those of the DM who anticipates no risk.

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

$$\mathcal{A} = \left\{ \mathbf{r} \in [\underline{r}, \overline{r}]^I : r_i = r_j \text{ for all arms } i, j \right\}. \tag{9}$$

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

To describe her behavior, we say that the DM's choice is represented by 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[9] that the DM chooses the lottery $\mathbf{r}$ converges to 1 [0] as $n \to \infty$.

Proposition 3. *When the DM anticipates a riskless lottery, then her 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 $i = 1, \ldots, I$.*

This proposition follows from the result on misspecified estimation by White (1982). He lets an agent observe n iid signals from a signal density and form the MLE from a set of hypothesized signal densities that may fail to include the true density. White proves that the MLE almost surely converges to the minimizer of the Kullback-Leibler divergence from the true signal density as n diverges.

To apply White's result 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 searches for a single unperturbed message m_n such that, with the added noise, it maximizes the likelihood of the observed data. White's result implies that $m^* = \lim_{n\to\infty} m_n$ almost surely minimizes Kullback-Leibler divergence from the true data-generating distribution to the signal distribution given m^*. For Gaussian errors, this implies $m^* = \sum_i \pi_i m(r_i)$ almost surely. Thus, the DM's estimate of the lottery value almost surely converges to the "certainty equivalent" $m^{-1}\left(\sum_{i=1}^I \pi_i m(r_i)\right)$.

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

The behavior of the DM who anticipates a riskless lottery is governed by the sampling frequencies rather than by the true arm probabilities. Indeed, the DM believes 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.

Proposition 3 extends to Bayesian decision-makers by the result of Berk (1966) who shows that, under general conditions, the Bayesian posterior converges to an atom that coincides with the asymptotic misspecified MLE of White as number of observed signals diverges.

Example (omitted-variable error): The DM chooses between a known safe payoff s and a reward $\rho(\mathbf{x}, \mathbf{y})$ whose value depends on vectors of inputs $\mathbf{x}$ and $\mathbf{y}$. For concreteness, we refer to ρ as a production function; it is unknown to the DM. She employs a misspecified model: she neglects the role of the input $\mathbf{y}$, so believing that the reward is $\tilde{\rho}(\mathbf{x})$ where $\tilde{\rho}(\cdot)$ is a simplified production function she estimates from data. Given $\mathbf{x}$, let the input $\mathbf{y}$ have conditional probability $g(\mathbf{y} \mid \mathbf{x})$ in the DM's environment. Thus, given the DM's observation of $\mathbf{x}$, the true reward is a lottery in which each arm represents a particular value of $\mathbf{y}$ and is assigned a reward $\rho(\mathbf{x}, \mathbf{y})$ and probability $g(\mathbf{y} \mid \mathbf{x})$.[10] However, the DM believes the lottery is riskless and estimates $\tilde{\rho}(\mathbf{x})$ from signals $m\left(\rho\left(\mathbf{x}, \mathbf{y}_k\right)\right) + \hat{\varepsilon}_k$, $k = 1, \ldots, n$. The conditional probability $\tilde{g}(\mathbf{y} \mid \mathbf{x})$ of observing the input $\mathbf{y}_k = \mathbf{y}$ depends on the DM's sampling; if her sampling is representative, then $\tilde{g} = g$. By Proposition 3, when the number of the signals diverges, this DM, who is unaware of the input $\mathbf{y}$, chooses for each observed $\mathbf{x}$ as if she were an expected-utility maximizer with Bernoulli utility $u(\cdot) = m(\cdot)$ and probability $\tilde{g}(\mathbf{y} \mid \mathbf{x})$ assigned to each value of y.

4.2 Coarse Decision-Maker

Next, we consider 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 $\{1, \ldots, I\}$ 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] pointwise in each decision problem $(\mathbf{r}, s)$ such that

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

[10]To keep the number of the arms finite, assume that $\mathbf{y}$ has a finite support.

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

$$u\left(r_J^*\right) = \sum_{i \in J} \frac{\rho_i}{\rho_J} u\left(r_i\right) \tag{10}$$

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

Let $J(i)$ be the element of the partition $\mathcal{K}$ that contains the 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 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 each arm $i = 1, \ldots, I$.*

In the limit, the DM makes choice 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. She then evaluates the overall lottery in a risk-neutral manner using the true probabilities of each J.

Example (omitted-variable error continued): As before, the reward $\rho(\mathbf{x}, \mathbf{y})$ depends on inputs $\mathbf{x}$ and $\mathbf{y}$ and the DM estimates the misspecified production function $\tilde{\rho}(\mathbf{x})$. Unlike in the previous version of this example, the DM 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 reward is a lottery in which each arm represents a realization of $(\mathbf{x}, \mathbf{y})$ with associated reward $\rho(\mathbf{x}, \mathbf{y})$ and probability $g(\mathbf{x}, \mathbf{y} \mid \mathbf{z})$.

Since the DM is unaware of $\mathbf{y}$'s influence on the reward, she forms a coarse counterpart of this lottery in which each arm represents a value of $\mathbf{x}$, paying a reward $\tilde{\rho}(\mathbf{x})$ 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 DM forms the estimate of the reward $\tilde{\rho}(\mathbf{x})$ 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 $\mathrm{E}\left[\hat{\rho}_n(\mathbf{x}) \mid \mathbf{z}\right]$ to the lottery where the expectation is with respect to the conditional density $g(\mathbf{x} \mid \mathbf{z})$. By Proposition 4, for each $\mathbf{z}$, this DM values the reward 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, this DM 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 $\mathbf{r} \in \mathcal{A}_{\mathcal{K}}$ she has anticipated learns in a correctly specified model. The asymptotic results for correctly specified learning of Wald (1949) for maximum-likelihood and of Le Cam (1953) for Bayesian estimation then imply that she 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 model.

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 $\mathcal{K}$.

Consider the DM with the finest adaptation partition $\mathcal{J} = \{\{1\}, \ldots, \{I\}\}$ who has concluded at the moment of choice that she faces a riskless lottery, $\mathcal{K} = \{\{1, \ldots, I\}\}$. Then, the results of Section 3 yield the optimal sampling frequencies for each arm $i = 1 \ldots, I$ and Proposition 4 represents this DM's choice with an EUT choice rule under subjective beliefs equal to these sampling frequencies.

If, on the other hand, $\mathcal{J} = \mathcal{K} = \{\{1, \ldots, I\}\}$, then the DM has adapted for riskless lotteries and anticipates a riskless lottery at the moment of decision. Her choice also has an EUT representation with the subjective probability of each arm equal to its sampling frequency. But in this case, the optimal-adaptation result does not predict the sampling frequencies since this DM believed all arms were payoff-equivalent during adaptation. If the sampling is representative, then the sampling frequencies coincide with the arms' objective probabilities. Any targeted sampling, for instance over-sampling salient contingencies, will result in an EUT representation of choice that assigns disproportional subjective probabilities to the over-sampled arms.[11]

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

4.3 Somewhat Surprising Risk

In the final version of the model, we analyze a DM who deems risk a priori possible but unlikely. To do this, we formulate an example and study the limit in which both the precision of the prior information and the sample size of the perception data diverge. We find perception distortions that are qualitatively similar to those from Subsection 4.1. Additionally, the example makes predictions on impact of framing and time pressure on risk-taking. Risk attitudes are attenuated by the anticipation of high risk or by a high volume of perception data.

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

$$\varrho_n(\mathbf{r}) = \rho_n^0 \exp\left(-\frac{n}{2\Delta}\sigma^2(\mathbf{r})\right) \tag{11}$$

with support $[\underline{r}, \overline{r}]^I$, where $\sigma^2(\mathbf{r}) = \sum_{i=1}^I p_i(r_i - r)^2$ is variance of the arm rewards and ϱ_n^0 is a normalization factor. This prior is concentrated on low-risk lotteries. For any fixed n, it becomes more dispersed as Δ increases, and thus Δ parameterizes the level of a priori anticipated risk. As n increases, risky lotteries become a priori unlikely, approximating then the anticipation of the DM from Subsection 4.1.

The index n has two roles. As n increases, then, in addition to risk becoming a priori unlikely, the DM observes more data. 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 $(\pi_i)_{i=1}^I$ continue to denote the sampling frequencies.[12] 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 $\mathbf{q}^* : [\underline{r}, \overline{r}]^I \longrightarrow [\underline{r}, \overline{r}]^I$:

$$\mathbf{q}^*(\mathbf{r}) = \arg\min_{\mathbf{r}' \in [\underline{r}, \overline{r}]^I} \left\{ \frac{\sigma^2(\mathbf{r}')}{\Delta} + \sum_{i=1}^I a\pi_i \left(m\left(r_i'\right) - m(r_i)\right)^2 \right\}. \tag{12}$$

We impose the regularity condition that the minimizer is unique.

We refer to the posterior expectation $\mathrm{E}[\mathbf{r} \mid \mathbf{m}_n] \in [\underline{r}, \overline{r}]^I$ that the DM forms given the vector of the average perturbed messages $\mathbf{m}_n$ as to the Bayesian estimate of the lottery.

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* $n \to \infty$.

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

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 $-\frac{n}{2}$ times the objective in (12) within a constant factor. As n diverges, the posterior converges almost surely to an atom on the minimizer $\mathbf{q}^*(\mathbf{r})$. The asymptotic estimate of the lottery is a compromise lottery that is not too risky and does not generate messages too far from the true messages.

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* $n \to \infty$*.*

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 value $q^*(\mathbf{r})$ of the Bayesian estimate with the true value r of the lottery $\mathbf{r}$. We say that function $f(\mathbf{r})$ is $o(g(\mathbf{r}))$ if $f(\mathbf{r}_k)/g(\mathbf{r}_k) \to 0$ for any sequence $\mathbf{r}_k$ such that $\sigma(\mathbf{r}_k) \to 0$. Specifically, a function is $o(\sigma^2)$ if it is negligible relative to σ^2 for lotteries with small σ.[13]

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

$$r + \frac{1}{2}\frac{m''(r)}{m'(r)} \cdot \frac{1 + 4z(r)}{(1 + z(r))^2} \cdot \sigma^2(\mathbf{r}) + o(\sigma^2(\mathbf{r})), \tag{13}$$

as $n \to \infty$*, where* $z(r) = a\Delta m'^2(r)$ *and* $r = \sum_i p_i r_i$ *is the true lottery value. The factor* $\frac{1+4z(r)}{(1+z(r))^2}$ *attains values in* $(0, 4/3]$ *and approaches 1 and 0 as* $a\Delta \to 0$ *and* $a\Delta \to \infty$*, respectively.*

To interpret the result, recall that the risk premium for the lottery $\mathbf{r}$ with small risk of an expected-utility maximizer with Bernoulli utility u is approximately $\frac{1}{2}\frac{u''(r)}{u'(r)}\sigma^2(\mathbf{r})$. The risk premium of our DM is the same, 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 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.

Dependence of the risk premium on 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

[13] The expression $o(\cdot)$ stands for "terms of smaller order than".

the risk-attitudes found in prospect theory whereas the slow mode favours risk-neutrality.[14] 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 model by a low value of parameter a. In accord with Kahnemann, we find that $a \to 0$ induces risk attitudes. 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, so the curvature of m determines the DM's risk attitudes. At the other extreme, in the absence of time pressure, when $a \to \infty$, 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 an inconsistency of risk attitudes across tasks with small and large levels of risk: the risk-averse choices observed for small risks imply implausibly high risk aversion for large risks with 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 – $\Delta \to 0$ here – induces risk attitudes since it makes risky lotteries surprising and this leads the distortion of the posteriors when a risky lottery is encountered. If, however, the DM is framed to anticipate high-risk lotteries – if our parameter $\Delta \to \infty$ – 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 inspired by neuroscience 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, an S-shaped reward valuation, probability weighting, and the role of stakes and time pressure.

The model makes several novel predictions about the effect of framing. For example, explaining the structure of risk to the DM may attenuate her risk attitudes. Further, if the DM conceptualizes a risky lottery as riskless, then manipulation of her sampling frequencies has a strong impact on choice. For instance, a seller offering a risky prospect can make it

[14] Kirchler et al. (2017) show experimentally that time pressure increases risk aversion for gains and risk loving for losses. Relatedly, Porcelli and Delgado (2009) and Cahlíková 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.

more attractive if the presentation of the prospect leads to over-sampling of the upside risk. Conversely, the seller of an insurance contract can make the contract more attractive if he prevents the DM from sampling a lot of perceptual data, for instance by putting her 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. Spiegler (2020). A model of competing narratives. *American Economic Review 110*(12), 3786–3816.

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. Kőszegi, and P. Strack (2018). Unrealistic expectations and misguided learning. *Econometrica 86*(4), 1159–1214.

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.

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.

Lindvall, T. (2002). *Lectures on the coupling method*. Courier Corporation.

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.

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 Bayesian Estimate

note: Pavel plans to edit this. Distinct density for s hasn't been implemented here yet.

Consider average perturbed message of value $m_{J,n} \in (\underline{m}, \overline{m})$. Recall that φ denotes the standard normal density. Since the average perturbed message $m_{J,n} \mid r_J \sim \mathcal{N}\left(m(r_J), \frac{1}{\pi_J n}\right)$, Bayes' law implies

$$\mathrm{E}\left[r_J \mid m_{J,n}\right] =$$

$$\frac{\int_{-\underline{r}}^{\overline{r}} \tilde{r} h\left(\tilde{r}\right) \sqrt{\pi_J n} \varphi\left(\left(m_{J,n} - m\left(\tilde{r}\right)\right)\sqrt{\pi_J n}\right) d\tilde{r}}{\int_{-\underline{r}}^{\overline{r}} h\left(\tilde{r}\right) \sqrt{\pi_J n} \varphi\left(\left(m_{J,n} - m\left(\tilde{r}\right)\right)\sqrt{\pi_J n}\right) d\tilde{r}} =$$

$$\frac{\int_{\sqrt{\pi_J n}\left(m_{J,n}-\overline{m}\right)}^{\sqrt{\pi_J n}\left(m_{J,n}-\underline{m}\right)} m^{-1}\left(m_{J,n} - \frac{\varepsilon}{\sqrt{\pi_J n}}\right) h\left(m^{-1}\left(m_{J,n} - \frac{\varepsilon}{\sqrt{\pi_J n}}\right)\right) \varphi\left(\varepsilon\right) \frac{1}{m'\left(m^{-1}\left(m_{J,n} - \frac{\varepsilon}{\sqrt{\pi_J n}}\right)\right)} d\varepsilon}{\int_{\sqrt{\pi_J n}\left(m_{J,n}-\overline{m}\right)}^{\sqrt{\pi_J n}\left(m_{J,n}-\underline{m}\right)} h\left(m^{-1}\left(m_{J,n} - \frac{\varepsilon}{\sqrt{\pi_J n}}\right)\right) \varphi\left(\varepsilon\right) \frac{1}{m'\left(m^{-1}\left(m_{J,n} - \frac{\varepsilon}{\sqrt{\pi_J n}}\right)\right)} d\varepsilon} =$$

$$\frac{\int_{\sqrt{\pi_J n}\left(m_{J,n}-\overline{m}\right)}^{\sqrt{\pi_J n}\left(m_{J,n}-\underline{m}\right)} \left(m^{-1}\left(m_{J,n}\right) h\left(m^{-1}\left(m_{J,n}\right)\right) \frac{1}{m'\left(m^{-1}\left(m_{J,n}\right)\right)} + \nu_1 \frac{\varepsilon}{\sqrt{\pi_J n}} + o\left(\frac{1}{n}\right)\right) \varphi\left(\varepsilon\right) d\varepsilon}{\int_{\sqrt{\pi_J n}\left(m_{J,n}-\overline{m}\right)}^{\sqrt{\pi_J n}\left(m_{J,n}-\underline{m}\right)} \left(h\left(m^{-1}\left(m_{J,n}\right)\right) \frac{1}{m'\left(m^{-1}\left(m_{J,n}\right)\right)} + d_1 \frac{\varepsilon}{\sqrt{\pi_J n}} + o\left(\frac{1}{n}\right)\right) \varphi\left(\varepsilon\right) d\varepsilon} =$$

$$\frac{\int_{-\infty}^{\infty} \left(m^{-1}\left(m_{J,n}\right) h\left(m^{-1}\left(m_{J,n}\right)\right) \frac{1}{m'\left(m^{-1}\left(m_{J,n}\right)\right)} + \nu_1 \frac{\varepsilon}{\sqrt{\pi_J n}} + o\left(\frac{1}{n}\right)\right) \varphi\left(\varepsilon\right) d\varepsilon + o\left(\frac{1}{n}\right)}{\int_{-\infty}^{\infty} \left(h\left(m^{-1}\left(m_{J,n}\right)\right) \frac{1}{m'\left(m^{-1}\left(m_{J,n}\right)\right)} + d_1 \frac{\varepsilon}{\sqrt{\pi_J n}} + o\left(\frac{1}{n}\right)\right) \varphi\left(\varepsilon\right) d\varepsilon + o\left(\frac{1}{n}\right)} =$$

$$\frac{m^{-1}\left(m_{J,n}\right) h\left(m^{-1}\left(m_{J,n}\right)\right) \frac{1}{m'\left(m^{-1}\left(m_{J,n}\right)\right)} + o\left(\frac{1}{n}\right)}{h\left(m^{-1}\left(m_{J,n}\right)\right) \frac{1}{m'\left(m^{-1}\left(m_{J,n}\right)\right)} + o\left(\frac{1}{n}\right)} =$$

$$m^{-1}\left(m_{J,n}\right) + o\left(\frac{1}{n}\right).$$

The second equation follows from substitution $\varepsilon = \left(m_{J,n} - m\left(\tilde{r}\right)\right)\sqrt{\pi_J n}$. For the third equation, we have used the first-order Taylor approximation of

$$\nu(\eta) = m^{-1}\left(m_{J,n} - \eta\right) h\left(m^{-1}\left(m_{J,n} - \eta\right)\right) \frac{1}{m'\left(m^{-1}\left(m_{J,n} - \eta\right)\right)}$$

around 0, where $\nu_1 = \nu'(0)$ in the numerator, and the first-order Taylor approximation of

$$d(\eta) = h\left(m^{-1}\left(m_{J,n} - \eta\right)\right) \frac{1}{m'\left(m^{-1}\left(m_{J,n} - \eta\right)\right)}$$

around 0, where $d_1 = d'(0)$ in the denominator. The fourth equation follows from the fact that $\int_{\text{const.}\sqrt{n}}^{+\infty} \varphi(\varepsilon) d\varepsilon$, $\int_{\text{const.}\sqrt{n}}^{+\infty} \varepsilon\varphi(\varepsilon) d\varepsilon$, $\int_{-\infty}^{-\text{const.}\sqrt{n}} \varphi(\varepsilon) d\varepsilon$, $\int_{-\infty}^{-\text{const.}\sqrt{n}} \varepsilon\varphi(\varepsilon) d\varepsilon$ are $o(1/n)$ for any positive const.

B Asymptotic Loss Characterization

Let us extend the inverse encoding function m^{-1} outside of the interval $[\underline{m}, \overline{m}]$ by setting $m^{-1}\left(m_{J,n}\right) = \underline{r}$ for $m_{J,n} < \underline{m}$ and $m^{-1}\left(m_{J,n}\right) = \overline{r}$ for $m_{J,n} > \overline{m}$. Recall that the ML estimate of the lottery value of the DM who has observed n signals is

$$\begin{aligned} q_n &= \sum_{J \in \mathcal{J}} p_J q_{J,n} \\ &= \sum_{J \in \mathcal{J}} p_J m^{-1}\left(m_{J,n}\right) \\ &= \sum_{J \in \mathcal{J}} p_J m^{-1}\left(m(r_J) + \frac{\varepsilon_J}{\sqrt{\pi_J n}}\right), \end{aligned}$$

where ε_J, $J \in \mathcal{J}$, are iid standard normal draws. We let $\varepsilon = (\varepsilon_J)_{J \in \mathcal{J}}$ and define the DM's rescaled squared error $e_n(\mathbf{r}, \varepsilon) = n(q_n - r)^2$; it depends on $\mathbf{r}$ and ε via q_n and r.

Lemma 1. *There exists a constant $\overline{e}$ such that $e_n(\mathbf{r}, \varepsilon) \leq \overline{e} \sum_J \varepsilon_J^2$ for all n, $\mathbf{r}$ and ε.*

Proof. Follows from the mean value theorem and from the fact that m^{-1} has a bounded derivative (since m' is positive and continuous on a compact interval). □

DM's loss as a function of $\mathbf{r}$, ε and s is

$$\hat{\ell}_n(\mathbf{r}, \varepsilon, s) = \max\{r, s\} - \mathbb{1}_{q_n > s} r - \mathbb{1}_{q_n \leq s} s;$$

it depends on $\mathbf{r}$ and ε via q_n and r. We introduce substitution $s = r + \frac{\sigma}{\sqrt{n}}$ and let

$$\ell_n(\mathbf{r}, \varepsilon, \sigma) = \sqrt{n}\hat{\ell}_n\left(\mathbf{r}, \varepsilon, r + \frac{\sigma}{\sqrt{n}}\right).$$

Let

$$\ell^*(\mathbf{r}, \varepsilon, \sigma) = \begin{cases} \sigma & \text{if } 0 \le \sigma \le \sum_{J\in\mathcal{J}} p_J \frac{\varepsilon_J}{\sqrt{\pi_J} m'(r_J)}, \\ -\sigma & \text{if } 0 \ge \sigma \ge \sum_{J\in\mathcal{J}} p_J \frac{\varepsilon_J}{\sqrt{\pi_J} m'(r_J)}, \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 2. *If m is differentiable, then $\lim_{n\to\infty} \ell_n(\mathbf{r}, \varepsilon, \sigma) = \ell^*(\mathbf{r}, \varepsilon, \sigma)$ almost everywhere.*

Proof. Choice of the DM differs from the optimal choice under complete information if and only if s attains value in between r and q_n. For each r, q_n, and s in between r and q_n, the loss of the DM relative to the complete-information choice is $|s - r|$ and it is 0 otherwise. Therefore,

$$\begin{aligned} \ell_n(\mathbf{r}, \varepsilon, \sigma) &= \sqrt{n}\hat{\ell}_n\left(\mathbf{r}, \varepsilon, r + \frac{\sigma}{\sqrt{n}}\right) \\ &= \begin{cases} \sigma & \text{if } 0 \le \sigma \le \sum_{J\in\mathcal{J}} p_J \left(m^{-1}\left(m(r_J) + \frac{\varepsilon_J}{\sqrt{\pi_J n}}\right) - r_J\right)\sqrt{n}, \\ -\sigma & \text{if } 0 \ge \sigma \ge \sum_{J\in\mathcal{J}} p_J \left(m^{-1}\left(m(r_J) + \frac{\varepsilon_J}{\sqrt{\pi_J n}}\right) - r_J\right)\sqrt{n}, \\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

and the right-hand side converges pointwise to $\ell^*(\mathbf{r}, \varepsilon, \sigma)$ almost everywhere.[15] □

Proof of Proposition 1.

$$\begin{aligned} nL(n) &= \int_{[\underline{r},\overline{r}]^{|\mathcal{J}|+1}\times\mathbb{R}^{|\mathcal{J}|}} n\hat{\ell}_n(\mathbf{r}, \varepsilon, s) h_s(s) \prod_J (h(r_J)\varphi(\varepsilon_J)) ds \prod_J (dr_J d\varepsilon_J) \\ &= \int_{[\underline{r},\overline{r}]^{|\mathcal{J}|+1}\times\mathbb{R}^{|\mathcal{J}|}} \sqrt{n}\ell_n(\mathbf{r}, \varepsilon, \sigma) h_s\left(r + \frac{\sigma}{\sqrt{n}}\right) \prod_J (h(r_J)\varphi(\varepsilon_J)) \frac{d\sigma}{\sqrt{n}} \prod_J (dr_J d\varepsilon_J) \\ &= \int_{[\underline{r},\overline{r}]^{|\mathcal{J}|+1}\times\mathbb{R}^{|\mathcal{J}|}} \ell_n(\mathbf{r}, \varepsilon, \sigma) h_s\left(r + \frac{\sigma}{\sqrt{n}}\right) \prod_J (h(r_J)\varphi(\varepsilon_J)) d\sigma \prod_J (dr_J d\varepsilon_J), \end{aligned}$$

where the second equality follows from substitution $s = r + \frac{\sigma}{\sqrt{n}}$.

[15] The convergence claim holds for all points $(\mathbf{r}, \varepsilon, \sigma)$ for which $\sigma \ne \sum_J p_J \frac{\varepsilon_J}{\sqrt{\pi_J} m'(r_J)}$.

In order to apply the Dominated Convergence Theorem, we show that the last integrand is bounded as follows:

$$\ell_n(\mathbf{r}, \varepsilon, \sigma)\, h_s\left(r + \frac{\sigma}{\sqrt{n}}\right) \prod_J \left(h(r_J)\varphi(\varepsilon_J)\right) \le \overline{\ell}(\mathbf{r}, \varepsilon, \sigma)\, \overline{h} \prod_J \left(h(r_J)\varphi(\varepsilon_J)\right) \tag{14}$$

where $\overline{h}$ is an upper bound on the density h_s and

$$\overline{\ell}(\mathbf{r}, \varepsilon, \sigma) = \begin{cases} |\sigma| & \text{if } |\sigma| \le \sqrt{\overline{e}\sum_J \varepsilon_J^2}, \\ 0 & \text{otherwise.} \end{cases}$$

This is because $\sqrt{\overline{e}\sum_J \varepsilon_J^2}$ is an upper bound on the perception error (rescaled by $\sqrt{n}$),

$$\left|\sum_{J\in\mathcal{J}} p_J \left(m^{-1}\left(m(r_J) + \frac{\varepsilon_J}{\sqrt{\pi_J n}}\right) - r_J\right)\sqrt{n}\right| \le \sqrt{\overline{e}\sum_J \varepsilon_J^2}.$$

Since

$$\begin{aligned}
&\int_{[\underline{r},\overline{r}]^{|\mathcal{J}|+1}\times\mathbb{R}^{|\mathcal{J}|}} \overline{\ell}(\mathbf{r}, \varepsilon, \sigma)\, \overline{h} \prod_J (h(r_J)\varphi(\varepsilon_J)) d\sigma \prod_J (dr_J d\varepsilon_J) \\
= \;& \overline{h} \int_{[\underline{r},\overline{r}]^{|\mathcal{J}|}\times\mathbb{R}^{|\mathcal{J}|}} \int_{-\sqrt{\overline{e}\sum_J \varepsilon_J^2}}^{\sqrt{\overline{e}\sum_J \varepsilon_J^2}} |\sigma| d\sigma \prod_J (h(r_J)\varphi(\varepsilon_J)) \prod_J (dr_J d\varepsilon_J) \\
= \;& \overline{h} \int_{[\underline{r},\overline{r}]^{|\mathcal{J}|}\times\mathbb{R}^{|\mathcal{J}|}} \overline{e}\left(\sum_J \varepsilon_J^2\right) \prod_J (h(r_J)\varphi(\varepsilon_J)) \prod_J (dr_J d\varepsilon_J) \\
= \;& \overline{h}\overline{e}|\mathcal{J}|,
\end{aligned}$$

the bound from the right-hand side of (14) is integrable.

Hence, the Dominated Convergence Theorem, Lemma 2 and continuity of h_s imply

$$\begin{aligned}\lim_{n\to\infty} nL(n) &= \int_{[\underline{r},\overline{r}]^{|\mathcal{J}|+1}\times\mathbb{R}^{|\mathcal{J}|}} \ell^*(\mathbf{r},\varepsilon,\sigma)\, h_s(r) \prod_J (h(r_J)\varphi(\varepsilon_J)) d\sigma \prod_J (dr_J d\varepsilon_J) \\ &= \mathrm{E}\left[\int_0^{\sum_{J\in\mathcal{J}} p_J \frac{\varepsilon_J}{\sqrt{\pi_J} m'(r_J)}} \sigma h_s(r) d\sigma\right] \\ &= \mathrm{E}\left[\frac{1}{2}\left(\sum_{J\in\mathcal{J}} p_J \frac{\varepsilon_J}{\sqrt{\pi_J} m'(r_J)}\right)^2 h_s(r)\right] \\ &= \mathrm{E}\left[\frac{1}{2}\sum_{J\in\mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J)} h_s(r)\right],\end{aligned}$$

where the first two expectations are over $\mathbf{r}$ and ε and the last expectation is with respect to $\mathbf{r}$. The last step follows from the fact that ε_J are iid standard normal random variables and thus $\mathrm{E}\,\varepsilon_J^2 = 1$ and $\mathrm{E}\,[\varepsilon_J \varepsilon_{J'}] = 0$ for all $J \neq J'$. □

C Optimal Perception

C.1 Notes on Peakedness of Conditioned Random Variables

This subsection contains the lemmas that we use 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 its domain and $h(x) = h(-x)$ for all $x \in \mathbb{R}$.*

This property is preserved by summation: sum of unimodal and symmetric random variables is unimodal and symmetric, 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, X is more peaked than Y whenever the CDF of X is greater than the CDF of Y at any α from the support of Y.

Lemma 3. *Let $X_0, X_1, \ldots, X_I$ be independent real-valued continuous random variables that are unimodal and symmetric around 0, where $X_1, \ldots, X_I$ are identically distributed and distribution of X_0 may be distinct. Let $(p_1, \ldots, p_I) \in \Delta(\{1, \ldots, I\})$, and $\overline{X} := \sum_{i=1}^{I} p_i X_i$. Then $X_i \mid \left(\overline{X} = X_0\right)$ is unimodal and symmetric around 0.*

Proof. Denote h the PDF of each of the iid variables $X_1, \dots, X_I$. Since the unimodality together with symmetry is preserved by taking 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 h_{-i} the PDF of X_{-i}. Then $X_i \mid \left(\overline{X} = X_0\right)$ is identical to $X_i \mid (X_i = X_{-i})$, and so its PDF is up to the 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. □

Lemma 4. *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 distribution of X_0 may be distinct. Let $(p_1, \dots, p_I) \in \Delta(\{1, \dots, I\})$, and $\overline{X} := \sum_{i=1}^{I} p_i X_i$. Then $X_i \mid \left(\overline{X} = X_0\right)$ is more peaked than $X_j \mid (X = X_0)$ if and only if $p_i > p_j$.*

Proof. Let h dento the PDF of each of the iid variables $X_1, \dots, X_I$. Without loss of generality, assume $\{i, j\} = \{1, 2\}$ (that is, either $i = 1$ and $j = 2$ or $i = 2$ and $j = 1$). Define $X_{-12} := X_0 - \sum_{k=3}^{n} p_k X_k$ (if $n = 2$, then $X_{-12} = X_0$), and let h_{-12} be its PDF. This is a unimodal random variable symmetric around 0. The conditioned random variable $X_i \mid \left(\overline{X} = X_0\right)$ is identical to $X_i \mid (p_i X_i + p_j X_j = X_{-12})$ and so its PDF equals

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

where $\mathrm{E}[h_{-12}(p_1 X_1 + p_2 X_2)]$ is the expectation with respect to X_1 and X_2 and it is independent of i. Thus, for any $\alpha \geq 0$,

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

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

where we used that both $P\left(|X_1| < \alpha \mid \overline{X} = X_0\right)$ and $P\left(|X_2| < \alpha \mid \overline{X} = X_0\right)$ are both (up to a same normalization constant) integrals of the same function $(x_1, x_2) \mapsto h_{-12}(p_1 x_1 + p_2 x_2) h(x_1) h(x_2)$, but the first is over the region $[-\alpha, \alpha] \times \mathbb{R}$, and the second is over $\mathbb{R} \times [-\alpha, \alpha]$; and that is equivalent to integrating both over the same region but switching the roles of x_1

and x_2. Then,

$$\Big(P\left(|X_1| < \alpha \mid \overline{X} = X_0\right) - P\left(|X_2| < \alpha \mid \overline{X} = X_0\right)\Big) \cdot \mathrm{E}[h_{-12}(p_1 X_1 + p_2 X_2)] =$$

$$\iint_{[-\alpha,\alpha]\times\mathbb{R}} \Big(h_{-12}(p_1 x_1 + p_2 x_2) - h_{-12}(p_1 x_2 + p_2 x_1)\Big) h(x_1) h(x_2) dx_1 dx_2 =$$

$$\iint_{[-\alpha,\alpha]\times\left(\mathbb{R}\setminus[-\alpha,\alpha]\right)} \Big(h_{-12}(p_1 x_1 + p_2 x_2) - h_{-12}(p_1 x_2 + p_2 x_1)\Big) h(x_1) h(x_2) dx_1 dx_2 =$$

$$2 \iint_{[-\alpha,\alpha]\times(\alpha,+\infty)} \Big(h_{-12}(p_1 x_1 + p_2 x_2) - h_{-12}(p_1 x_2 + p_2 x_1)\Big) h(x_1) h(x_2) dx_1 dx_2,$$

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

$$p_1 x_1 + p_2 x_2 = (p_1 x_2 + p_2 x_1) + (p_2 - p_1)(x_2 - x_1)$$

that

$$p_1 x_1 + p_2 x_2 > p_1 x_2 + p_2 x_1,$$

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

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

Thus,

$$|p_1 x_1 + p_2 x_2| > |p_1 x_2 + p_2 x_1|,$$

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

$$h_{-12}(p_1 x_1 + p_2 x_2) < h_{-12}(p_1 x_2 + p_2 x_1),$$

unless both are zero. It follows that $X_2 \mid \left(\overline{X} = X_0\right)$ is more peaked than $X_1 \mid \left(\overline{X} = X_0\right)$, as needed. □

Lemma 5. *Let f be a bounded, increasing, differentiable function on $\mathbb{R}_+$, and X_1, X_2 be unimodal continuous random variables symmetric around* 0*. Then* $\mathrm{E}[f(|X_1|)] < \mathrm{E}[f(|X_2|)]$ *whenever X_1 is more centered than X_2.*

Proof. Denote $h_i(x)$ and $H_i(x)$ the PDF and CDF of X_i, respectively. Then,

$$\begin{aligned}
\tfrac{1}{2}\,\mathrm{E}[f(|X_i|)] &= \int_0^\infty f(x)h_i(x)dx \\
&= \Big[f(x)(H_i(x)-1)\Big]_0^{+\infty} - \int_0^\infty f'(x)(H_i(x)-1)dx \\
&= \tfrac{1}{2}f(0) + \int_0^\infty f'(x)(1-H_i(x))dx.
\end{aligned}$$

If X_1 is more centered than X_2, then $1 - H_1(x) < 1 - H_2(x)$ unless both are zero for all $x > 0$. It follows that $\mathrm{E}[f(|X_1|)] < \mathrm{E}[f(|X_2|)]$. □

C.2 Proof of Proposition 2

Proof of Proposition 2. The objective (5) of the information-processing problem is a functional

$$\mathcal{L}\left(m'(\cdot), (\pi_J)_{J\in\mathcal{J}}\right) = \int_{[\underline{r},\overline{r}]^{|\mathcal{J}|}} \sum_{J\in\mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J)} h_s(r) \prod_{J\in\mathcal{J}} \left(h\left(r_J\right)dr_J\right).$$

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:

$$\alpha\mathcal{L}\left(m_1'(\cdot), (\pi_{1,J})_J\right) + (1-\alpha)\mathcal{L}\left(m_2'(\cdot), (\pi_{2,J})_J\right) > \mathcal{L}\left(\alpha m_1'(\cdot) + (1-\alpha)m_2'(\cdot), (\alpha\pi_{1,J} + (1-\alpha)\pi_{2,J})_J\right),$$

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.

Let us first fix any given interior sampling frequencies $(\pi_J)_J$, possibly suboptimal ones, and let us optimize with respect to $m'(\cdot)$. We show that the first-order condition of the information-processing problem with respect to $m'(\tilde{r})$ is, for all $\tilde{r}$, equivalent to (8) up to a multiplicative factor. The Lagrangian for minimization of the objective (5) with respect to $m'(\cdot) > 0$ under constraint (6) is

$$\int_{[\underline{r},\overline{r}]^{|\mathcal{J}|}} \sum_{J\in\mathcal{J}} \frac{p_J^2}{\pi_J m'^2(r_J)} h_s(r) \prod_{J\in\mathcal{J}} \left(h\left(r_J\right)dr_J\right) + \tilde{\lambda}\left(\overline{m} - \underline{m} - \int_{\underline{r}}^{\overline{r}} m'(\tilde{r})d\tilde{r}\right),$$

where $\tilde{\lambda}$ is the shadow price of (6). The first-order condition with respect to $m'(\tilde{r})$ is, for each $\tilde{r}$,

$$\int_{[\underline{r},\overline{r}]^{|\mathcal{J}|-1}} \sum_{J\in\mathcal{J}} \left(\frac{d}{dm'(r_J)} \frac{p_J^2}{\pi_J m'^2(r_J)}\bigg|_{r_J=\tilde{r}} h(\tilde{r})\, h_s(r) \times \prod_{J'\neq J} \left(h(r_{J'})\, dr_{J'}\right)\right) = \tilde{\lambda}$$

$$-\sum_{J\in\mathcal{J}} \left(\frac{2p_J^2}{\pi_J m'^3(\tilde{r})} h(\tilde{r}) \times \int_{[\underline{r},\overline{r}]^{|\mathcal{J}|-1}} h_s(r) \prod_{J'\neq J} \left(h(r_{J'})\, dr_{J'}\right)\right) = \tilde{\lambda}.$$

Multiplying the last inline equation by $-1/2\,\mathrm{E}[h_s(r)]$ gives (8) as needed. Condition (8) implies,

$$m'(\tilde{r}) \propto \left(\sum_{J\in\mathcal{J}} \frac{p_J^2}{\pi_J} h_J(\tilde{r})\right)^{\frac{1}{3}}, \tag{15}$$

where the right-hand side is positive for all $\tilde{r} \in [\underline{r}, \overline{r}]$ as needed. Statement 1 of the proposition follows because by Lemma 3, each h_J is unimodal with the same mode as the unconditional reward density h.

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

$$-\int_{[\underline{r},\overline{r}]^{|\mathcal{J}|}} \frac{p_J^2}{\pi_J^2 m'^2(r_J)} h_s(r) \prod_{J'\in\mathcal{J}} \left(h(r_{J'})\, dr_{J'}\right) = \tilde{\mu},$$

where $\tilde{\mu}$ is the shadow price of the constraint (7). Dividing both sides by $-\mathrm{E}[h_s(r)]$ gives

$$\int_{[\underline{r},\overline{r}]^{|\mathcal{J}|}} \frac{p_J^2}{\pi_J^2 m'^2(r_J)} \frac{h_s(r) \prod_{J'\in\mathcal{J}} \left(h(r_{J'})\, dr_{J'}\right)}{\mathrm{E}[h_s(r)]} = \frac{-\tilde{\mu}}{\mathrm{E}[h_s(r)]}$$

$$\int_{[\underline{r},\overline{r}]} \frac{p_J^2}{\pi_J^2 m'^2(r_J)} h_J(r_J)\, dr_J = \frac{-\tilde{\mu}}{\mathrm{E}[h_s(r)]}$$

$$\mathrm{E}\left[\frac{p_J^2}{\pi_J^2 m'^2(r_J)} \mid r = s\right] = \mu, \tag{16}$$

where $\mu = -\tilde{\mu}/\,\mathrm{E}[h_s(r)]$ and the expectations are with respect to $\mathbf{r}$.

Suppose $p_J < p_{J'}$. By condition (16), to prove Statement 2 it suffices to show that

$$\mathrm{E}\left[\frac{1}{m'^2(r_J)} \mid r = s\right] > \mathrm{E}\left[\frac{1}{m'^2(r_{J'})} \mid r = s\right]. \tag{17}$$

This indeed holds since by Lemma 4, $r_{J'} \mid (r = s)$ is more peaked than $r_J \mid (r = s)$ and the inequality (17) follows from Lemma 5 and from the fact that m' is decreasing above r^*. □

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 $\mathbf{p}$ realize, the DM observes these, and chooses sampling frequencies $(\pi_J(\mathcal{J}, \mathbf{p}))_{J\in\mathcal{J}}$. Finally, the DM observes the sequence of perturbed messages generated according to the chosen encoding function and the sampling frequencies, conditional on the encountered reward vector $\mathbf{r} \in [\underline{r}, \overline{r}]^{|\mathcal{J}|}$. As in Section 3, the rewards are iid from a unimodal reward density h and independent of $\mathcal{J}$ and $\mathbf{p}$. The information-processing problem for this setting is as follows:

$$\min_{m'(\cdot)>0,(\pi_J(\cdot))_J>0} \quad \mathrm{E}\left[\int_{[\underline{r},\overline{r}]^{|\mathcal{J}|}} \sum_{J\in\mathcal{J}} \frac{p_J^2}{\pi_J(\mathcal{J},\mathbf{p})m'^2(r_J)} h_s(r) \prod_{J\in\mathcal{J}} \left(h\left(r_J\right)dr_J\right)\right] \tag{18}$$

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

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

The expectation in (18) is with respect to $(\mathcal{J}, \mathbf{p})$ drawn from g.

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

$$\mathrm{E}\left[\int_{[\underline{r},\overline{r}]^{|\mathcal{J}|}} \sum_{J\in\mathcal{J}} \frac{p_J^2}{\pi_J(\mathcal{J},\mathbf{p})m'^2(r_J)} h_s(r) \prod_{J\in\mathcal{J}} \left(h\left(r_J\right)dr_J\right)\right] + \tilde{\lambda}\left(\overline{m} - \underline{m} - \int_{\underline{r}}^{\overline{r}} m'(r)dr\right),$$

where the expectation is with respect to $(\mathcal{J}, \mathbf{p})$. The same steps as in the proof of Proposition 2 imply that $m'(\tilde{r})$ satisfies for each $\tilde{r}$ the analogue of the condition (8):

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

where λ equals the shadow price $\tilde{\lambda}$ up to a constant factor and $h_J\left(r_J;\mathcal{J},\mathbf{p}\right)$ 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 $(\mathcal{J}, \mathbf{p})$.

Therefore, the analogue of (15) holds:

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

Since by Lemma 3, $h_J(\cdot;\mathcal{J},\mathbf{p})$ is unimodal with maximum at the mode of the ex ante reward density, m' is hump-shaped with maximum at r^*.

The proof of Statement 2 of Proposition 2 extends for each realization of $(\mathcal{J},\mathbf{p})$ since the first-order condition (16) for $\pi_J(\mathcal{J};\mathbf{p})$ continues to hold for each $(\mathcal{J},\mathbf{p})$:

$$\mathrm{E}\left[\frac{p_J^2}{\pi_J^2(\mathcal{J},\mathbf{p})m'^2(r_J)} \mid r=s\right] = \mu(\mathcal{J},\mathbf{p}),$$

properties of the conditional densities $h_J(\cdot;\mathcal{J},\mathbf{p})$ are unaffected by the studied extension and the proof relies of the unimodality of $m'(\cdot)$, which continues to hold.

D Proofs of Propositions 3 and 4

Proposition 3 is a special case of Proposition 4. For expositional reasons, we first prove Proposition 3 separately and then generalize the proof to establish Proposition 4.

Proof of Proposition 3. 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 distributions for any two lotteries $\mathbf{r},\mathbf{r}'$ is

$$\begin{aligned}
D_{\mathrm{KL}}(f_{\mathbf{r}}\parallel f_{\mathbf{r}'}) &= \int f_{\mathbf{r}}(x)\ln\frac{f_{\mathbf{r}}(x)}{f_{\mathbf{r}'}(x)}dx \\
&= \sum_{i=1}^{I}\int_{\mathbb{R}}\pi_i\varphi(\hat{m}-m(r_i))\ln\frac{\pi_i\varphi(\hat{m}-m(r_i))}{\pi_i\varphi(\hat{m}-m(r_i'))}d\hat{m} \\
&= \sum_{i=1}^{I}\pi_i\int_{\mathbb{R}}\varphi(\hat{m}-m(r_i))\ln\frac{\varphi(\hat{m}-m(r_i))}{\varphi(\hat{m}-m(r_i'))}d\hat{m} \\
&= \sum_{i=1}^{I}\pi_i D_{\mathrm{KL}}\left(\varphi_{m(r_i)}\parallel\varphi_{m(r_i')}\right)
\end{aligned}$$

where $\varphi_m(\hat{m})=\varphi(\hat{m}-m)$ is the density of the perturbed message $\hat{m}$ conditional on the unperturbed message m. Since the Kullback-Leibler divergence of two Gaussian densities with means μ_1, μ_2 and variances equal to 1 is $(\mu_1 - \mu_2)^2/2$, see e.g. Johnson and Orsak (1993), we get for the Gaussian errors,

$$D_{\text{KL}}(f_{\mathbf{r}} \parallel f_{\mathbf{r}'}) = \frac{1}{2}\sum_{i=1}^{I} \pi_i \left(m(r_i) - m(r'_i)\right)^2 . \tag{19}$$

Let $\mathbf{q}_n \in \mathcal{A}$ be the maximum-likelihood estimator of the encountered lottery $\mathbf{r}$, given the sequence of signals $x_1, \ldots, x_n$. White (1982), Theorem 2.2, implies that $\mathbf{q}_n$ almost surely converges to

$$\mathbf{q} = \arg\min_{\mathbf{r}' \in \mathcal{A}} D_{\text{KL}}(f_{\mathbf{r}} \parallel f_{\mathbf{r}'}), \tag{20}$$

where $\mathcal{A}$ is the set of the riskless lotteries as defined in (9).[16] We have from (19) that the unique minimizer $\mathbf{q}$ estimates each arm reward as

$$q = \arg\min_{q' \in \mathbb{R}} \sum_{i=1}^{I} \pi_i \left(m(r_i) - m(q')\right)^2,$$

and thus $m(q) = \sum_{i=1}^{I} \pi_i m(r_i)$. Therefore, $q > [<]\, s$ if and only if $\sum_{i=1}^{I} \pi_i m(r_i) > [<]\, m(s)$, as needed. □

The next proof is a minor extension of Proof of Proposition 3.

Proof of Proposition 4. Recall that $f_{\mathbf{r}}(\hat{m}, i) = \pi_i \varphi(\hat{m} - m(r_i))$ is the signal density conditional on lottery $\mathbf{r}$. Using the assumption of additive Gaussian noise, we have derived in the proof of Proposition 3 that the Kullback-Leibler divergence of the signal densities for any two lotteries $\mathbf{r}, \mathbf{r}' \in \mathbb{R}^I$ is

$$D_{\text{KL}}(f_{\mathbf{r}} \parallel f_{\mathbf{r}'}) = \frac{1}{2}\sum_{i=1}^{I} \pi_i \left(m(r_i) - m(r'_i)\right)^2 .$$

Let $\mathbf{q}_n \in \mathcal{A}_{\mathcal{K}}$ be the maximum-likelihood estimator of the lottery, given the sequence of signals $x_1, \ldots, x_n$. White (1982), Theorem 2.2, implies that $\mathbf{q}_n$ almost surely converges to[17]

$$\mathbf{q} = \arg\min_{\mathbf{r}' \in \mathcal{A}_{\mathcal{K}}} D_{\text{KL}}(f_{\mathbf{r}} \parallel f_{\mathbf{r}'}) = \arg\min_{\mathbf{r}' \in \mathcal{A}_{\mathcal{K}}} \sum_{i=1}^{I} \pi_i \left(m(r_i) - m(r'_i)\right)^2 . \tag{21}$$

[16] White's conditions are satisfied: the minimizer in (20) is unique; $\mathrm{E}[\log f_{\mathbf{r}}(x_k)]$ exists where the expectation is with respect to the signal density $f_{\mathbf{r}}(x_k)$ for the realized lottery $\mathbf{r}$. Finally, there exists a bound on $\log |f_{\mathbf{r}'}(x_k)|$ across all $\mathbf{r}' \in \mathcal{A}$ that is integrable with respect to the true signal density $f_{\mathbf{r}}(x_k)$ since the interval of the unperturbed messages is bounded.

[17] White's conditions are satisfied for the same reasons as those mentioned in Footnote 16.

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

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

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

Therefore, $q_i = q_J$ for all arms $i \in J$, where q_J is the certainty equivalent as defined in (10) for $u(\cdot) = m(\cdot)$ and $\rho_i = p_J \frac{\pi_i}{\pi_J}$. Thus, the DM chooses the lottery with the probability almost surely converging to 1 [0] if $\sum_{J \in \mathcal{J}} p_J q_J > [<]\, s$, as needed. □

E Proofs for Subsection 4.3

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

Lemma 6. *Let $\psi_n(\mathbf{x}) : [\underline{r}, \overline{r}]^I \longrightarrow \mathbb{R}$ 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 probability density function equal to $\alpha_n \exp(-n\psi_n(\mathbf{x}))$, where α_n is the normalization factor, converges to $\mathbf{x}^*$ in probability as $n \to \infty$.*

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

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

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

Since the convergence $\psi_n \to \psi$ is uniform, for every $d > 0$, there exists $n_d \in \mathbb{N}$ such that $|\psi_n(\mathbf{x}) - \psi(\mathbf{x})| < d$ for all $\mathbf{x} \in [\underline{r}, \overline{r}]^I$ and $n \geq n_d$. Consider a fixed $\delta > 0$ and any $n \geq n_{d_\delta/4}$. Because $\psi_n(\mathbf{x}) \geq \psi(\mathbf{x}) - \frac{d_\delta}{4} \geq \psi(\mathbf{x}^*) + \frac{3d_\delta}{4}$ for $\mathbf{x}$ outside of the ball B_δ, the probability density of X_n is no more than $\alpha_n \exp\left(-n\psi(\mathbf{x}^*) - \frac{3d_\delta}{4} n\right)$. This implies,

$$P(X_n \notin B_\delta) < \widetilde{\alpha}_n \exp\left(-\frac{3d_\delta}{4} n\right) (\overline{r} - \underline{r})^I, \qquad \text{where } \widetilde{\alpha}_n := \alpha_n \exp(-n\psi(\mathbf{x}^*)). \tag{22}$$

We conclude by establishing an upper bound for $\widetilde{\alpha}_n$. Let $\delta' > 0$ be such that $\psi(\mathbf{x}) \leq \psi(\mathbf{x}^*) + d_{\delta/4}$ for all $\mathbf{x} \in B_{\delta'} \cap [\underline{r}, \overline{r}]^I$. Existence of such δ' follows from the continuity of ψ.

Then, $\psi_n(\mathbf{x}) \le \psi(\mathbf{x}) + \frac{d_\delta}{4} \le \psi(\mathbf{x}^*) + \frac{d_\delta}{2}$ for all $\mathbf{x} \in B_{\delta'} \cap [\underline{r}, \overline{r}]^I$ and thus the probability density of X_n is at least $\widetilde{\alpha}_n \exp\left(-\frac{d_\delta}{2} n\right)$. It follows that,

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

where $b' > 0$ is the volume of the set $B_{\delta'} \cap [\underline{r}, \overline{r}]^I$. Substituting the implied upper bound on $\widetilde{\alpha}_n$ into (22) gives

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

Since the right-hand vanishes as $n \to \infty$, the claim follows. □

Proof of Proposition 5. Let $\mathbf{m}_n = (m_{i,n})_{i=1}^I$ be the vector of the averages of $a\pi n$ signals received for each arm i. Since signal errors are standard normal, $m_{i,n} \mid r_i \sim \mathcal{N}\left(m(r_i), \frac{1}{a\pi_i n}\right)$. The DM whose vector of the average signals is $\mathbf{m}_n$ attaches posterior density (the symbol $\sim$ means equality modulo normalization)

$$\varrho_n\left(\mathbf{r}' \mid \mathbf{m}_n\right) \sim \varrho_n\left(\mathbf{r}'\right) \prod_{i=1}^{I} \varphi\Big(\big(m_{n,i} - m(r_i')\big)\sqrt{a\pi_i n}\Big) \sim \exp\big(- n\psi(\mathbf{r}'; \mathbf{m}_n)\big)$$

to each lottery $\mathbf{r}' \in [\underline{r}, \overline{r}]^I$, where

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

The first inline equality follows from the Bayes' law (recall that φ denotes the standard normal density). The second equality follows from the specification of the prior ϱ_n in (11).

Since $m_{n,i}$ converges to $m(r_i)$ (almost surely), $\psi(\mathbf{r}'; \mathbf{m}_n)$ converges to $\psi\big(\mathbf{r}'; (m(r_i))_i\big)$, uniformly across $\mathbf{r}'$. Additionally, $\psi\big(\mathbf{r}'; (m(r_i))_i\big)$ as a function of $\mathbf{r}'$ has the unique minimizer $\mathbf{q}^*(\mathbf{r})$ by assumption. By Lemma 6, it follows that the DM's posterior formed after observing $\mathbf{m}_n$ converges in probability to $\mathbf{q}^*(\mathbf{r})$. However, as the support of the rewards is bounded, the convergence in probability implies the convergence in expected value, and thus the Bayesian estimate $\mathrm{E}\left[\mathbf{r} \mid \mathbf{m}_n\right]$ converges to $\mathbf{q}^*(\mathbf{r})$. □

Proof of Proposition 6. By Proposition 5, the Bayesian estimate of $\mathbf{r}$ converges to $\mathbf{q}^*(\mathbf{r})$. We write shortly $\mathbf{q}^* = (q_i^*)_{i=1}^I$ for $\mathbf{q}^*(\mathbf{r})$ and let $q^* = \sum_i p_i q_i^*$. The first-order condition applied to the objective in (12) implies,

$$(q_i^* - q^*) + a\Delta\big(m\left(q_i^*\right) - m(r_i)\big)m'\left(q_i^*\right) = 0, \tag{23}$$

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

We will prove the following claims:

Claim 1: Any function that is $o(r_i - r)$ or $o\left(q_i^* - r\right)$ 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: $q^* = r + \frac{1}{2}\frac{m''(r)}{m'(r)}\left(\sigma^2 + \left(\frac{2}{z(r)} - 1\right)\sigma^{*2}\right) + o(\sigma^2)$.

The proposition follows from substituting for σ^{*2} from Claim 3 into the formula for q^* in Claim 4.

Claim 1 will allow us to apply approximation techniques. To prove it, we bound 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 $[\underline{r}, \overline{r}]$, respectively, and let $\underline{z} = a\Delta\underline{m}'^2$, $\overline{z} = a\Delta\overline{m}'^2$. Clearly, $\underline{m}' > 0$ and $\underline{z} > 0$, since $m'(\cdot)$ is continuous and strictly positive on the closed interval $[\underline{r}, \overline{r}]$.

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

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

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 (23) 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,

$$q_i^* - q^* = \tfrac{z_i}{1+z_i}(r_i - q^*) = \tfrac{z_i}{1+z_i}(r_i - r) + \tfrac{z_i}{1+z_i}(r - q^*). \tag{24}$$

Summing up the last equation weighted by p_i across i gives

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

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

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

Returning to the equation (24),

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

We conclude that $|q_i^* - r| \leq \left(p_i^{-1/2} + 2\frac{1+\underline{z}}{\underline{z}} I\right)\sigma$ for any $\mathbf{r} \in [\underline{r}, \overline{r}]^I$, 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 (23) for $\sigma > 0$ small. Since the function $m(\cdot)$ has continuous derivatives up to the third order, and its domain is the closed interval $[\underline{r}, \overline{r}]$, each of its derivatives is a bounded function. Thus, 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 (23) then implies

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

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

$$0 = (q_i^* - q^*) + z(r)(q_i^* - r_i) + o(\sigma). \tag{25}$$

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 can rewrite (25) as

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

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 $z(r) \le \overline{z}$ 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:

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

This implies the second-order approximation of the equation (23),

$$\begin{aligned} 0 = (q_i^* - q^*) + a\Delta\Big(m'(r)\big(q_i^* - r_i\big) + \tfrac{1}{2}m''(r)\Big((q_i^* - r)^2 - (r_i - r)^2\Big) + o(\sigma^2)\Big) \\ \cdot \Big(m'(r) + m''(r)\big(q_i^* - r\big) + o(\sigma)\Big), \end{aligned}$$

which can be written as

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

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

$$0 = (q^* - r) - \tfrac{1}{2}\tfrac{m''(r)}{m'(r)}\left(\sigma^2 - \sigma^{*2} + 2\sum_{i=1}^{I} p_i\big(r_i - q_i^*\big)\big(q_i^* - r\big)\right) + o(\sigma^2). \tag{26}$$

Expressing $q_i^* - r_i$ from (25) allows us to write

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

where we used that $r = q^* + o(\sigma)$ for the second equality. Substituting the last inline equation back into (26) completes the proof of Claim 4.

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

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

and using that $1+\frac{2z(r)-z(r)^2}{(1+z(r))^2}=\frac{1+4z(r)}{(1+z(r))^2}$, we obtain (13), concluding the proof. $\square$