Risk Perception*

Nick Netzer
University of Zurich

Arthur Robson
Simon Fraser University

Jakub Steiner
University of Zurich, Cerge-Ei and CTS

Pavel Kocourek
University of Duisburg-Essen

January 29, 2024

Abstract

In a model inspired by neuroscience, we study choice between lotteries as a process of encoding and decoding noisy perceptual signals. The implications of this process for behavior depend on the decision-maker's understanding of the risk. The encoding strategy does not influence choice in the limit as perception noise vanishes when the decision-maker correctly understands the decision problem during decoding. If, however, the decision-maker underrates the complexity of the decision problem, then the encoding strategy generates behavioral risk attitudes even for vanishing perception noise. We show that constrained optimal perception encodes lottery rewards using an S-shaped encoding function and over-samples low-probability events. Taken together, the model can explain adaptive risk attitudes and probability weighting as in prospect theory. Additionally, it predicts that risk attitudes are influenced by the anticipation of risk, time pressure, experience, salience, and availability heuristics.

*We have benefited from comments of Olivier Compte, Cary Frydman, Philippe Jehiel, Franz Ostrizek, Antonio Rosato, Ryan Webb, and various seminar audiences. Giorgi Chavchanidze and Pavel Ilinov provided excellent research assistantship. An earlier version of this paper was circulated under the title "Endogenous Risk Attitudes." This work was supported by the Alfred P. Sloan Foundation and the NOMIS Foundation (Netzer), by SSHRC grant 435-2013-0426 (Robson), by ERC grant 770652 and GAČR 24-10145S (Steiner), and by DFG grant 450175676 (E. Kováč, Kocourek).

1 Introduction

Choice under risk exhibits inconsistencies with expected utility theory. Some of these revolve around the fact that behavior cannot be represented by fixed risk preferences. We know at least since Kahneman and Tversky (1979) that risk attitudes adapt to the environment. In a similar vein, Rabin's (2000) paradox implies that choices over small and large risks are represented by distinct utility functions. Risk attitudes are further modulated by factors such as the salience of rewards (Bordalo et al., 2012), time pressure (Kirchler et al., 2017), and experience (Ert and Haruvy, 2017). Additional inconsistencies involve the overweighting of small relative to large probability events (Kahneman and Tversky, 1979) and the importance of availability of an event in memory for the assessment of its probability (Tversky and Kahneman, 1973).

In recent work, Oprea (2023) has shown that some of these regularities may have little to do with risk. In a laboratory experiment, Oprea replaces lotteries by their "deterministic mirrors," which are described exactly like the original lotteries but pay the lotteries' expected value and hence involve no risk. The subjects' behavior is essentially identical when they choose between risky lotteries to when they choose between the respective deterministic mirrors. Oprea concludes that apparent risk attitudes may be due to the complexity of aggregation of the lotteries' rewards and not due to genuine risk.

In this paper, we develop a model of perception and aggregation that is compatible with the findings outlined above and, in addition, makes novel predictions. The model is inspired by literature on noisy coding in neuroscience and economics.

A risk-neutral decision-maker (DM) chooses between a lottery and a safe option. For illustration, consider the vivid example from Savage (1954) in which the lottery represents purchase of a convertible car, the enjoyment of which depends on the random weather conditions. The DM knows the probabilities of the possible states of the world (the different weather conditions), observes the value of the safe option (the price of the car), but faces a friction in perceiving the rewards of the lottery in the different states (the weather-dependent enjoyments). She learns about the rewards by sampling signals. Each signal is a reward of the lottery in some state encoded via a non-linear encoding function that maps rewards to their mental representations (e.g. neural firing rates), then perturbed by additive Gaussian noise. Depending on the application, these perturbed signals can be cues extracted from the visual description of the decision problem, own experiences retrieved from memory, or experiences provided by others (e.g. the car dealer). When making the decision, the DM estimates the value of the lottery by decoding the collection of signals that she obtained and chooses the lottery if its estimated value is larger than the outside option.

We make two distinct contributions. In the first part of the paper, we treat the encoding strategy of the DM as exogenous and show that decoding of the perception data via a misspecified statistical model is an important mechanism by which encoding has behavioral implications, even when noise in perception is small. In the second part of the paper, we derive optimal encoding by jointly optimizing perception of lottery payoffs and probabilities.

To understand our first contribution, suppose that the encoding function, which maps actual rewards into mental representations, is exogenously given. Suppose furthermore that the frequencies with which the DM samples signals from each state of the world are exogenous and do not necessarily coincide with the true probabilities of the states. Finally, suppose that the DM uses a very simple decoding procedure: she just averages all collected signals and then applies the inverted coding function to determine a value of the lottery. Then, when the numbers of signals which the DM collects becomes large, her behavior converges to the behavior of an expected utility maximizer with Bernoulli utility function equal to her encoding function and subjective probabilities equal to her sampling frequencies. We thus establish a tight connection between the elements of the encoding strategy (the encoding function and the sampling frequencies) and the elements of expected utility theory (the Bernoulli utility function and subjective probabilities).

The described behavior is a result of how the rewards of the lottery are encoded and decoded in the aggregation process, and not a consequence of the riskiness of these rewards. In line with the findings of Oprea (2023), we predict behavioral "risk attitudes" for choice between deterministic mirrors of lotteries, which are not risky but still complex objects that the DM must aggregate. From a welfare perspective such behavioral risk attitudes reflect mistakes rather than genuine risk preferences.

The simple model makes several additional predictions that are consistent with empirical regularities. First, since behavioral risk attitudes depend on the properties of encoding, they will adapt if the DM's encoding function adapts to the environment (Kahneman and Tversky, 1979; Frydman and Jin, 2022). Similarly, since behavior depends on sampling frequencies rather than true probabilities, subjective probability weights may differ from objective probabilities (Kahneman and Tversky, 1979) and can furthermore be affected by the salience of states (Bordalo et al., 2012) or availability in memory (Tversky and Kahneman, 1973). Our procedural choice model thus provides a microfoundation for these phenomena.

The simple decoding procedure described above is a special case of maximum likelihood estimation using 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. If the DM anticipates a riskless lottery that pays the same reward in all states, then averaging all the signals and applying the inverted coding function gives the maximum likelihood estimate of the value of the riskless lottery. The DM's statistical model is misspecified if the lottery's rewards differ across the states, resulting in a biased estimate of the lottery's value. This bias generates the behavioral risk attitudes. If, on the other hand, the DM anticipates that the lottery pays different rewards in different states, her maximum likelihood estimate averages the signals and applies the inverted coding function state-by-state, leading to an unbiased estimate of the lottery's value and to risk-neutral behavior in the limit of rich perception data. With a general coarse partitional model of the state space, we show that maximum likelihood estimation leads to risk-neutral choices whenever the DM faces risk that she comprehends (i.e., which is measurable with respect to her partition) but implies perception-driven risk attitudes for those elements of the risk that she does not comprehend. Our proof relies on results by White (1982) for misspecified maximum likelihood estimation. Analogous results hold for Bayesian estimation (Berk, 1966).

There are various reasons why the DM might employ a coarse model of the world. She might have evolved in a simple environment and the complexity of the environment has then increased, making previously payoff-irrelevant contingencies relevant, without the DM adapting to the change. For example, a financial expert when analysing a new financial asset may estimate its expected performance in a model that fails to include all variables that are relevant to the asset's return. Alternatively, the DM might have been framed to believe that the decision problem involves less risk than it does (by a car dealer, for instance). Finally, and in line with findings by Oprea (2023) for laboratory subjects, the DM may deliberately choose to use a simplified decoding procedure to economize on cognitive resources.

The partitional decoding procedure predicts risk attitudes to be more pronounced with a coarser partition. To the extent that more experience with a decision environment implies a better understanding of which contingencies matter, and thus a finer and more appropriate partition, behavior is predicted to become more risk-neutral with experience (Ert and Haruvy, 2017).

We then study a variant of the model where the DM anticipates some risk but finds large risks unlikely. We formalize this by letting the DM form a Bayesian estimate of the lottery value and taking a joint limit in which the number of perception signals grows and the prior belief of the DM gradually concentrates on the set of riskless lotteries. The DM therefore possesses a comparable amount of both a priori and perceptual information. Changing the relative rates of divergence allows us to vary the relative influence of these two sources of information on choice. We find perception-driven risk attitudes akin to those of the DM discussed above. Choice becomes more risk-neutral when the DM anticipates larger risk a priori, consistent with the finding that framing a decision problem as one which features high risk dampens risk attitudes (Rabin, 2000). Choice also becomes more risk-neutral when the DM collects more data, predicting that risk attitudes are intensified under time pressure when decisions have to rely more on prior information and less on perceptual signals (Kirchler et al., 2017).

The previous arguments all treated the encoding strategy as exogenous. In the second part of the paper, we propose one possible way in which encoding may have been shaped. For this second contribution of our paper, we assume that the DM decodes her signals correctly in a well-specified model at the time when optimization of the encoding strategy took place (which might have been during evolutionary times, so the assumption does not imply that the DM is still decoding correctly today). Choice of the encoding strategy is a specific form of an attention allocation problem. The DM is akin to an engineer who measures a physical input by reading off the position of a needle on a meter (Robson, 2001). The engineer can choose the measurement function that maps the physical input to the needle position. If the needle position has a stochastic component, then the engineer can increase the precision of her measurement for a specific range of inputs by making the measurement function steep in this range. Our DM can increase the precision of her reward perception for a specific range of rewards by making the encoding function steep in this range. Since the range of possible mental representations is finite, the encoding function cannot be steep everywhere. Further, our DM can allocate attention to a specific state of the world by sampling it frequently, but this comes at the cost of sampling other states less frequently.

We analyze the limit of rich perception data, which enables us to obtain tractable results. These limit results serve as an approximation for the more realistic case of non-vanishing but small noise relative to the stakes of the decision problem. We prove that the expected loss from noisy perception, compared to choice under complete information, is approximately the mean squared error in the estimate of the lottery value, integrated over all decision problems that the DM may face in her environment in which the lottery value ties with the safe option. The conditioning on ties arises endogenously. Accuracy of perception has instrumental value for choice, and choice is trivial except where the values of two options are nearly equal, given that information is nearly complete.

We then derive the encoding strategy that minimizes the mean squared error conditional on ties. For the plausible case of unimodal symmetric reward densities, we show that an S-shaped encoding function and over-sampling of low-probability states are jointly optimal. The DM chooses the encoding function to be steep near the modal rewards and flatter towards the tails of the reward distribution. She thus perceives the reward values that are typical for her environment relatively precisely, at the expense of precision at the tail rewards. Conditioning on ties induces a statistical association between tail rewards and low-probability states. Since tail rewards in high-probability states typically result in very attractive or unattractive lotteries rather than in ties, tail rewards arise relatively more often in low-probability states once we condition on the event of a tie. The DM with an S-shaped encoding function therefore relatively often struggles to estimate the rewards in low-probability states, and it is optimal to compensate for this by over-sampling such states.

Our results in the second part of the paper generalize earlier findings in the literature, which studied choice between simple, one-dimensional objects (Robson, 2001; Netzer, 2009). Here, we provide a microfoundation for an objective rooted in choice, and we derive results for choice between complex, multi-dimensional objects. Combined with the results in the first paper of the paper, our findings provide a possible explanation for prospect theory type behavior in situation with and without risk.

2 Decoding

We introduce a model of perception and aggregation of lottery rewards based on noisy encoding and decoding of signals about these rewards. Throughout this section, we treat the encoding as fixed and focus on the behavioral implications of decoding. We make predictions that are in line with the findings of Oprea (2023) and the literatures on prospect theory (Kahneman and Tversky, 1979), Rabin's paradox (Rabin, 2000), salience (Bordalo et al., 2012), time pressure (Kirchler et al., 2017), experience (Ert and Haruvy, 2017), and availability heuristics (Tversky and Kahneman, 1973).

2.1 A Simple Decoding Procedure

There is a set of states of the world $i \in \{1, \ldots, I\}$, where each state i has a fixed positive probability p_i. The DM chooses between a safe option of value s and a lottery that pays a reward $r_i \in [\underline{r}, \overline{r}]$ in each state i, where $\underline{r} < \overline{r}$ are arbitrary bounds. We let $\mathbf{r} = (r_i)_i \in [\underline{r}, \overline{r}]^I$ denote the tuple of rewards and refer to it as the *lottery*. The pair $(\mathbf{r}, s)$ is the *decision problem.*

The DM observes the value of the safe option and the probabilities but faces frictions in the perception of the lottery rewards. She receives a sequence of n signals, where each signal is a monotone transformation of one of the rewards, perturbed with additive noise. Signals are given by $x_k = (\hat{m}_k, i_k)$, $k = 1, \ldots, n$, where the size n of the sample is exogenous. We refer to the first component, $\hat{m}_k$, as the *perturbed message.* The second component, i_k, indicates the state that the k'th signal pertains to. Each perturbed message is generated by encoding the reward r_{i_k} in state i_k into *unperturbed message* $m\left(r_{i_k}\right)$ and then perturbing it to $\hat{m}_k = m\left(r_{i_k}\right) + \hat{\varepsilon}_k$, where the noise term $\hat{\varepsilon}_k$ is independently and identically distributed (iid) standard normal.[1] The sampled state i_k is one of the states $i = 1, \ldots, I$, iid with positive probabilities π_i. The function $m : [\underline{r}, \overline{r}] \longrightarrow [\underline{m}, \overline{m}]$ is assumed to be continuously differentiable and strictly increasing with $m' > 0$ throughout the paper. The assumption that the DM encodes rewards into messages from a finite range and that this encoding is noisy follows a long tradition in the psychometric literature (see e.g. the discussion in Frydman and Jin, 2022). We refer to m as the *encoding function* and to $(\pi_i)_i$ as the *sampling frequencies*. The pair $(m, (\pi_i)_i)$ of encoding function and sampling frequencies is the *encoding strategy*. Throughout this section, we treat the encoding strategy as exogenous.

In Savage (1954)'s example, the DM is contemplating the purchase of a convertible car for price s. The purchase is a lottery because the enjoyments of driving a convertible car depend on the random weather. The states of the world represent the possible weather conditions: rainy or sunny. The perturbed messages $\hat{m}_k$ are noisy signals about the enjoyment r_{i_k} of driving the car in weather i_k, which the DM retrieves from memory and other sources. Signals can encode own past experiences, stories of peers, or information provided by the car dealer, etc. The shape of the encoding function determines the neural intensity with which these experiences are represented. The sampling frequencies are the proportions of experiences under different weather conditions that the DM processes, which are not necessarily representative (Tversky and Kahneman, 1973).

In other applications, where lottery rewards and probabilities are presented as numbers, as in Oprea (2023)'s experiment or in financial applications, noisy perception occurs in the process of visual inspection of the options (Schaffner et al., 2023). Furthermore, memory is again involved when processing the consequences of different events such as winning or losing (Ludvig et al., 2015). More generally, in line with a large literature (e.g. Robson, 2001; Netzer, 2009; Frydman and Jin, 2022), our approach assumes that the DM does not use an algebraic calculation mechanism to make the choice but instead processes stated numbers like any other percept (see Khaw et al., 2021, for a careful discussion and motivation of this assumption).

We now describe a procedure, which we call the *simple decoding procedure*, that the DM may employ to decode the collected data: she averages all the perturbed messages and then applies the inverted encoding function to obtain an estimated lottery value $m^{-1}(\sum_{k=1}^{n} \hat{m}_k/n)$. She chooses the lottery if $m^{-1}(\sum_{k=1}^{n} \hat{m}_k/n) > s$ and the outside option otherwise. The following observation describes the behavioral implications of this simple procedure.

[1] We use the Gaussian assumption mainly because it will yield a tractable form of the Kullback-Leibler divergence in the next subsection.

Observation 1. *With the simple decoding procedure, the probability that the DM chooses the lottery in problem* $(\mathbf{r}, s)$ *converges almost surely to* 1 (0) *as* $n \to \infty$ *if*

$$\sum_{i=1}^{I} \pi_i m(r_i) > (<)\, m(s).$$

The observation shows that the DM behaves in the limit as if the encoding function m were her Bernoulli utility function and the sampling frequencies π_i were her subjective probabilities. By the strong law of large numbers, the average empirical message $\sum_{k=1}^{n} \hat{m}_k/n$ almost surely converges to $\sum_{i=1}^{I} \pi_i m(r_i)$ as $n \to \infty$, and the value that the DM places on the lottery therefore converges to $m^{-1}(\sum_{i=1}^{I} \pi_i m(r_i))$, which is the certainty equivalent of the lottery for Bernoulli utility m and probabilities π_i. The result establishes a tight connection between the elements of the encoding strategy (encoding function, sampling frequencies) and the elements of expected utility theory (Bernoulli utility, subjective probabilities).

The simple decoding procedure has several implications that differ from those of expected utility theory and are in line with empirically documented patterns. First, the behavior in Observation 1 arises as a consequence of the simplistic reward aggregation procedure, rather than due to risk preferences. The behavior is the result of aggregation of a complex object and would arise equally if the lottery was replaced by its "deterministic mirror," an option that is described in the same complex way as the lottery but pays the lottery's expected value with certainty, in line with the findings by Oprea (2023). Thus, choice over lotteries may not reveal the DM's genuine risk preferences.

Second, behavior is affected by the subjective sampling frequencies π_i rather than by the objective probabilities p_i. Tversky and Kahneman (1973) have pointed out that subjective probability weights are influenced by the availability of the respective events in memory and do not necessarily correspond to the true frequencies of these events. In Bordalo et al. (2012), salience of a state shifts subjective probabilities in favor of that state. These authors derive their salience weights from the set of options that the DM considers, while we treat the sampling frequencies as exogenous in this section, but the simple decoding procedure provides a microfoundation for salience-driven behavior based on non-representative sampling.

Third, since behavior depends on the encoding function and the sampling frequencies, rather than on an exogenous Bernoulli utility function and the true probabilities, the DM's behavior will be affected by changes in the encoding function (e.g. through environmental adaptation, see Frydman and Jin, 2022; Schaffner et al., 2023) and by changes in the sampling frequencies (e.g. by manipulation of her focus through marketing interventions, see Koszegi and Szeidl, 2013).

2.2 Maximum Likelihood Decoding

We now generalize the simple decoding procedure and show that it and its generalizations describe a DM who forms a maximum likelihood estimate of the lottery value using a coarse model of the state space. The DM is endowed with a compact set $\mathcal{A} \subseteq [\underline{r}, \overline{r}]^I$ of lotteries she deems possible and, given a sequence of n signals as in the previous subsection, 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 the likelihood of the observed signals, where φ is the standard normal density.[2] Given the estimate $\mathbf{q}_n = (q_{1n}, \ldots, q_{In})$, she estimates the value of the lottery to be $q_n = \sum_i p_i q_{in}$ and chooses the lottery if $q_n > s$ and the outside option otherwise. The risk-neutrality with respect to rewards embodied in this rule is an implicit assumption on the units of measurement in which the rewards are expressed. For instance, the rewards might be an appropriate concave function of monetary prizes if the DM chooses among monetary lotteries and money has diminishing returns.[3]

We consider a DM who employs a possibly simplifying model of the risk in the spirit of the *small world* of Savage (1954). The DM anticipates, rightly or wrongly, distinctions among some of the states of the world to be payoff-irrelevant. Let $\mathcal{P}$ be a partition of the set of all the states $\{1, \ldots, I\}$. The partition captures the DM's coarse model of the state space similarly to Jehiel (2005) as follows. The DM anticipates that $r_i = r_j$ for all pairs of states $i, j \in J$ that belong to a same element J of the partition $\mathcal{P}$. That is, she anticipates lotteries from a set

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

If $\mathcal{P} = \{\{1, \ldots, I\}\}$ is the coarsest partition, then the DM anticipates only degenerate lotteries that pay the same reward in all states. In Savage's original example, the coarse DM believes that the convertible car will lead to "definite and sure enjoyments", so she

[2]The maximum-likelihood estimate exists since $\mathcal{A}$ is compact. It is unique for the specifications that we consider in the following.

[3]If, for example, rewards are $r_i = f(w_i)$, where w_i is a monetary prize and f is a fitness function, and the DM measures the rewards r_i by applying a non-linear encoding function $m(r_i)$, then the simple procedure from the previous subsection predicts that the DM behaves as expected utility maximizer with Bernoulli utility $u(w) = m(f(w))$. If she encounters a lottery with rewards in the concave part of the encoding function, her choices would make her appear to be more risk-averse than the fitness function f would suggest. For monetary lotteries with relatively small stakes, where a linear function f appears particularly plausible, all non-degenerate risk attitudes are entirely driven by the aggregation frictions.

anticipates a degenerate lottery $(r, \ldots, r)$. If, on the other extreme, $\mathcal{P} = \{\{1\}, \ldots, \{I\}\}$ is the finest partition, then the DM anticipates that any reward tuple in $\mathcal{A}_{\mathcal{P}} = [\underline{r}, \overline{r}]^I$ is possible. Such fine DM knows that the weather is payoff-relevant and hence anticipates that the purchase of the convertible car leads to weather-dependent rewards.

We treat the partition $\mathcal{P}$ as exogenous. There are, however, various paths that could have led the DM to the adoption of a partition, and in particular a partition that is too coarse to measure lotteries she actually encounters. The DM could have evolved in a simple environment in which all lotteries were measurable with respect to a relatively coarse partition. Afterwards, the environment became more complex, so that she currently encounters lotteries that involve more risk, but the DM has not made an adjustment and continues to model the world as relatively riskless. It is plausible that real-world DMs are sometimes not aware of all contingencies that affect their payoffs, thus effectively omitting relevant variables from their model of the risk. Alternatively, the DM may have been assured (incorrectly and possibly by a strategically interested party such as a car dealer) that her next lottery will be relatively riskless. Finally, the DM may know that she encounters a risky lottery but applies a coarse estimation procedure due to its simplicity, as it requires the estimation of a smaller number of distinct rewards. Subjects in Oprea (2023)'s experiment explicitly report being aware that they use a simplifying procedure when aggregating the value of a lottery. Our approach is also compatible with the idea that the DM's partition changes over time. An unexperienced agent may first use a coarse partition but become increasingly fine in subsequent choices.

Our next result characterizes the behavior of the DM who forms a maximum likelihood estimate using a given partition and a given encoding strategy.

Proposition 1. *With maximum likelihood decoding, the probability that the DM chooses the lottery in problem* $(\mathbf{r}, s)$ *converges almost surely to* 1 (0) *as* $n \to \infty$ *if*

$$\sum_{J \in \mathcal{P}} p_J r_J^* > (<)\, s,$$

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

$$m\left(r_J^*\right) = \sum_{i \in J} \frac{\pi_i}{\sum_{j \in J} \pi_j} m\left(r_i\right)$$

for each element of the partition $J \in \mathcal{P}$.

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

We prove the proposition in Appendix A.1. The proof relies on a result about misspecified maximum-likelihood estimation by White (1982). White lets an agent observe iid signals from a signal density and form the maximum likelihood estimate from a set of hypothesized signal densities that may fail to include the true density. He proves that the estimate almost surely converges to the minimizer of the Kullback-Leibler divergence from the true signal density as the number of signals diverges. To apply this result in our setting, consider a DM who encounters a lottery $\mathbf{r}$. She observes the empirical distribution of approximately $\pi_i n$ perturbed messages drawn iid from $\mathcal{N}(m(r_i), 1)$ for each state i. Since the DM anticipates a lottery which pays the same reward in all states $i \in J$, she forms an estimate of a single unperturbed message for each $J \in \mathcal{P}$, a perturbation of which has generated the observed data. For Gaussian errors, this estimate is the arithmetic average of the perturbed messages for J, denoted $\hat{m}_{Jn}$, which almost surely converges to $\sum_{i \in J} \frac{\pi_i}{\sum_{j \in J} \pi_j} m(r_i)$. Thus, the DM's estimate of the reward of J converges to the certainty equivalent r_J^* defined in the proposition. Across elements J of the partition, the DM's anticipation of distinct rewards implies that her maximum likelihood estimate aggregates the values r_J^* in a risk-neutral manner.

The simple decoding procedure studied in the previous subsection corresponds to the special case of a DM who anticipates no risk at all and uses the coarsest partition $\mathcal{P} = \{\{1, \ldots, I\}\}$. This DM explains her perception data by the same reward across all states. The maximum likelihood estimate of that reward is given by the average of all observed perturbed messages mapped through the inverse of the encoding function. The behavior of the DM is therefore governed by the sampling frequencies rather than by the true probabilities. Indeed, this DM believes that the true probabilities are payoff-irrelevant. In contrast, the sampling frequencies govern the proportions of her data generated for each state and hence her estimate of the encoded riskless reward that she thinks she has encountered.

The other extreme is a DM who uses the finest partition $\mathcal{P} = \{\{1\}, \ldots, \{I\}\}$ and therefore always behaves in a risk-neutral manner based on correct probabilities. More generally, whenever the DM encounters a lottery $\mathbf{r} \in \mathcal{A}_{\mathcal{P}}$ that she has anticipated, she learns in a well-specified model. Asymptotic results for well-specified maximum-likelihood estimation by Wald (1949) imply that she correctly learns the encountered lottery as the number of signals diverges, implying that her encoding strategy becomes irrelevant and she chooses in a risk-neutral way.

Between these two extremes, our model gives rise to interesting comparative statics. In addition to being compatible with Oprea (2023), we predict that behavior is risk-neutral for risk that the DM anticipates in her environment, while the encoding strategy generates non-trivial risk attitudes for risk that the DM does not anticipate. All the behavioral patterns discussed in the context of the simple decoding procedure in Subsection 2.1 are modulated by the degree of understanding of the risk and thus by the DM's experience with the situation, with more experience generating more risk-neutral and less manipulable behavior. This is in line with experimental results showing that experience with a decision-problem tends to shift risk-preferences towards risk-neutrality (Bradbury et al., 2015; Ert and Haruvy, 2017; Charness et al., 2023).

Our results contrast with Savage (1954). In his discussion of small world models, Savage argues that a coarse representation of the complex grand world does not necessarily distort behavior. In his argument, Savage assumes that the DM learns the correct reward average for each element of the coarse state space partition. Our approach departs from Savage in that we explicitly model the process of learning. 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. We thus add the observation that long-run biases in estimates may asymptotically survive in a natural coarse learning process.

We conclude this subsection by two remarks. First, Proposition 1 has an immediate extension to choice between multiple lotteries. The DM's estimate of the value of any lottery $\mathbf{r}$ converges to $\sum_{J\in\mathcal{P}} p_J r_J^*$ and therefore, in the limit, she chooses the one that maximizes this expression among all feasible lotteries. Second, analogous results hold when the DM forms a Bayesian estimate of the lottery using a prior with support $\mathcal{A}_{\mathcal{P}}$, rather than a maximum likelihood estimate. The results of Berk (1966) for misspecified Bayesian estimation imply that the Bayesian estimate also converges to the minimizer of the Kullback-Leibler divergence from the true signal distribution, generating the same limiting behavior of the DM as with the maximum likelihood estimate. The two different estimation approaches deliver the same result because the effect of the prior becomes negligible in the limit with many signals.

2.3 Decoding with Additional Prior Information

The distinction between anticipated and unanticipated lotteries in the previous section is dichotomous. In this subsection, we pursue a more continuous approach that allows for additional comparative statics. We study, in a specific example, a Bayesian DM who possesses a comparable amount of both a priori and perceptual information, allowing both these sources to influence the choice. The DM is well-specified but has a prior belief leaning towards riskless lotteries. For the sake of comparability with the setting from the previous subsection, we consider a sequence of information structures in which both the prior concentrates on the set of riskless lotteries and the quantity of perception data diverges. Changing the relative rates of divergence allows us to modulate the relative influence of a priori versus perceptual information.

We set the prior belief of the DM over lotteries $\mathbf{r}$, indexed by n, to the density

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

with support $[\underline{r}, \overline{r}]^I$, where $\sigma^2(\mathbf{r}) = \sum_{i=1}^I p_i(r_i - \sum_j p_j r_j)^2$ is the variance of the states' rewards and ϱ_n^0 is the normalization factor. For large n, this prior approximates a prior over a riskless value uniformly distributed on $[\underline{r}, \overline{r}]$. The parameter $\Delta > 0$ specifies the level of a priori anticipated risk. The larger Δ is, the more risk the DM anticipates for any given n.[4] Additionally to indexing the prior, we let n control the volume of the DM's perception data. Akin to our model in the previous subsections, for each state i the DM observes a sequence of approximately $a\pi_i n$ messages equal to $m(r_i)$ perturbed with iid additive standard normal noise, where m and $(\pi_i)_i$ continue to denote the exogenous encoding strategy.[5] The parameter $a > 0$ represents attention span. A larger a implies that the DM observes more signals for every fixed n.

The DM chooses the lottery $\mathbf{r}$ over the safe option s if and only if the posterior expected lottery value exceeds s. The new parameters Δ and a jointly determine how much the DM's prior and the perceptual information affect her posterior expectation. If $a\Delta$ is small, then the DM's prior anticipation of a relatively riskless lottery dominates the data obtained through perception. Conversely, if $a\Delta$ is large, the rich perception data dominate the relatively dispersed prior belief.

The result of Berk (1966), who characterizes the Bayesian posterior asymptotically for expanding data volume keeping the prior fixed, does not apply directly to our setting where the prior varies alongside the data volume. However, we adapt his result and show that the Bayesian posterior concentrates on a compromise lottery that balances the unexpected elements observed in the data (akin to Berk's conclusion) with the surprises stemming from our progressively focused prior.

[4] For large n, one can think of this prior as first drawing a common value for rewards in all states uniformly from $[\underline{r}, \overline{r}]$ and then perturbing each reward with a Gaussian shock with variance proportional to Δ^2.

[5] Since we take the number n of signals to be large, from now on we abstract from uncertainty over the number of perturbed messages sampled for each state and from divisibility issues, that is, we suppose that the average of the sampled messages for state i is drawn from $\mathcal{N}\big(m(r_i), 1/(a\pi_i n)\big)$.

To formulate our result, we define a function $\mathbf{q}^* : [\underline{r}, \overline{r}]^I \longrightarrow [\underline{r}, \overline{r}]^I$ by

$$\mathbf{q}^*(\mathbf{r}) = \underset{\mathbf{r}' \in [\underline{r}, \overline{r}]^I}{\arg\min} \left\{ \frac{\sigma^2(\mathbf{r}')}{a\Delta} + \sum_{i=1}^{I} \pi_i \left(m\left(r_i'\right) - m(r_i) \right)^2 \right\} \tag{3}$$

and impose the regularity condition that the minimizer is unique for the given $\mathbf{r}$ of interest, which holds generically. We will show that the DM's posterior expected lottery converges to $\mathbf{q}^*(\mathbf{r})$ almost surely as $n \to \infty$. The asymptotic estimate $\mathbf{q}^*(\mathbf{r}) = (q_1^*(\mathbf{r}), \ldots, q_I^*(\mathbf{r}))$ is a compromise lottery that is not too risky and does not generate messages too far from the true messages. When $a\Delta$ is small, a main concern in (3) is the minimization of $\sigma^2(\mathbf{r}')$, and hence $\mathbf{q}^*(\mathbf{r})$ will involve little risk. In the limit as $a\Delta \to 0$, the solution minimizes Kullback-Leibler divergence from the true lottery (the last term on the right side of (3)) among the riskless lotteries. When $a\Delta$ is large, the main concern in (3) is the minimization of the discrepancy between the messages, which in the limit as $a\Delta \to \infty$ yields the correct lottery $\mathbf{q}^*(\mathbf{r}) = \mathbf{r}$.

Proposition 2. *With Bayesian decoding and concentrated prior beliefs, the probability that the DM chooses the lottery in problem* $(\mathbf{r}, s)$ *converges almost surely to* 1 (0) *as* $n \to \infty$ *if*

$$\sum_{i=1}^{I} p_i q_i^*(\mathbf{r}) > (<)\, s.$$

See Appendix A.2 for the proof. If the DM anticipates relatively large risk and/or collects a lot of perception data ($a\Delta$ large), the relevance of the prior information is diminished and the DM learns the true value of the lottery, becoming risk-neutral like the well-specified DM in Subsection 2.2. If the DM anticipates relatively little risk and/or collects little perception data ($a\Delta$ small), prior information remains influential and the DM's posterior mirrors that of the coarse DM, leading her to maximize Bernoulli utility equal to her encoding function and to use sampling frequencies as subjective probability weights. Furthermore, the result yields additional comparative statics in between these two extremes. As the volume of perception data grows, the DM's choice progressively shifts from the risk attitudes shaped by her encoding strategy to risk-neutrality. Similarly, risk attitudes are attenuated by the anticipation of high risk.[6]

To illustrate these comparative statics further, we compute the risk premium defined in the standard manner as the excess return needed to compensate the DM for accepting a

[6] Our prediction that anticipation of risk makes the DM less risk-averse is reminiscent of Köszegi and Rabin (2007) but is based on a different mechanism. In Köszegi and Rabin (2007), anticipation is modelled by a reference lottery and behavioral risk-attitudes are generated by aversion to losses relative to that reference lottery. In our model, anticipation is modelled by prior beliefs about the decision problem and behavioral risk-attitudes are generated by biases in the estimation of surprising lotteries.

risk. As common in expected utility theory, we compute the risk premium for a lottery that involves little risk, in that the variance $\sigma^2(\mathbf{r})$ of the rewards across states is small. To further facilitate comparison with expected utility theory, we set the sampling frequencies equal to the actual probabilities.[7]

Proposition 3. *Let the encoding function m be twice differentiable and assume $\pi_i = p_i$. Given a lottery $\mathbf{r}$, the expected value of its Bayesian estimate converges almost surely to*

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

as $n \to \infty$, where $z(r) = a\Delta m'^2(r)$.

See Appendix A.3 for the proof. To interpret the result, recall that the risk premium of an expected utility maximizer with Bernoulli utility u for a lottery $\mathbf{r}$ with small risk is given by $\frac{1}{2}\frac{u''(r)}{u'(r)}\sigma^2(\mathbf{r})$. The risk premium of our DM is the same for $u(\cdot) = m(\cdot)$ but scaled by the positive factor $\frac{1+4z(r)}{(1+z(r))^2}$ that depends on the parameters Δ and a. This factor approaches 1 and 0 as $a\Delta \to 0$ and $a\Delta \to \infty$, respectively.

The bias in the valuation of the lottery – expressed here as the risk premium – arises as follows. The DM who encounters a risky lottery faces a conflict between the perception data and the prior information and resolves this conflict by concluding that she has encountered a compromise lottery that is more risky than a priori anticipated but less risky than indicated by the perception data. This underestimation of the reward variance leads to a mismatch with the perception data. To minimize the mismatch, the curvature of the encoding function requires the DM to form a bias in the estimated average reward. The resulting bias depends on the curvature of the encoding function (captured by the term $\frac{1}{2}\frac{m''}{m'}$) and on the extent of the underestimation of the reward variance (captured by the factor $\frac{1+4z}{(1+z)^2}$).[8]

The dependence of the risk premium on the parameter a sheds light on the apparent instability of risk preferences pointed out by Kahneman (2011). Kahnemann distinguishes between fast and slow decision-making, where the fast mode favours the risk-attitudes found in prospect theory whereas the slow mode favours risk-neutrality.[9] In accord with Kahnemann, we find encoding-based risk attitudes when $a \to 0$, corresponding to a fast decision

[7]Otherwise, the effect of distorted sampling dominates the effect of the curvature of the encoding function, because the first effect is of the order of $\sigma(\mathbf{r})$ while the second effect is only of the order of $\sigma^2(\mathbf{r})$. In the following, the expression $o(\cdot)$ stands for "term of smaller order than." Specifically, 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 $g(\mathbf{r}_k) \to 0$.

[8]We remark that this latter factor is not monotone in z and larger than one for small z.

[9]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.

that is based primarily on prior beliefs, without much consideration of the data about the specific decision problem. In slow decisions, corresponding to $a \to \infty$, the DM collects enough data for her prior to become irrelevant and makes the risk-neutral choice.

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

3 Optimal Encoding

So far we have treated the encoding strategy $(m, (\pi_i)_i)$ as exogenous. In this section, we propose one way of endogenizing it using optimality arguments.

The encoding strategy determines how the DM allocates attention. An increase of the sampling frequency of a state increases the DM's attention to its reward but reduces attention to the rewards in other states. Similarly, making the encoding function steeper in a neighborhood of a reward value increases attention and reduces perception error in this neighborhood, but entails decreased attention and increased perception error elsewhere, due to the finite range of the encoding function. Therefore, the encoding strategy needs to adapt to the prevailing environment if it is to allocate attention efficiently. We will show that, for the limit of many signals, the DM's loss caused by perception errors equals the mean squared error of her estimate of the lottery value, averaged over the pivotal decision problems where perception matters even when the data volume is large. We then optimize the encoding strategy. The loss-minimizing solution generalizes results in the literature (Robson, 2001; Netzer, 2009; Woodford, 2012; Payzan-LeNestour and Woodford, 2021) and, under reasonable assumptions, entails an S-shaped encoding function and oversampling of low-probability states.

3.1 Objective

We fix a state space partition $\mathcal{P}$ and assume that this partition correctly describes the environment at the point of optimization of the encoding function, i.e., all lotteries are measurable with respect to $\mathcal{P}$. The encoding strategy is therefore optimized in a statistical model with well-specified decoding of the perception data. We interpret this optimization process as evolutionary selection.

Since the distinction between states within each $J \in \mathcal{P}$ is redundant, we treat J as an index of a state, refer to the rewards in states $i \in J$ simply as r_J, and model the whole lottery $\mathbf{r} = (r_J)_{J\in\mathcal{P}}$ as having $|\mathcal{P}|$ rewards, each with probability $p_J = \sum_{i\in J} p_i$. An encoding strategy consists of the encoding function $m(\cdot)$ and positive sampling frequencies $(\pi_J)_J$.[10] The encoding strategy is optimized ex ante for a given distribution of decision problems. Specifically, the rewards r_J are assumed to be 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 $[\underline{r}, \overline{r}]$.[11]

Let $r = \sum_J p_J r_J$ denote the value of lottery $\mathbf{r}$ and $q_n = \sum_J p_J q_{Jn}$ its maximum likelihood estimate, where each $q_{Jn} = m^{-1}(\hat{m}_{Jn})$ is the maximum-likelihood estimate of r_J given the average perturbed message $\hat{m}_{Jn}$. The DM's ex ante expected loss relative to choice under complete information is

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

where the expectation is over the estimate q_n and the decision problem $(\mathbf{r}, s)$.

Proposition 4. *The expected loss satisfies*

$$\lim_{n\to\infty} nL(n) = \frac{1}{2}\,\mathrm{E}\left[h_s(r)\sum_J \frac{p_J^2}{\pi_J m'^2(r_J)}\right], \tag{5}$$

where the expectation is with respect to $\mathbf{r}$*.*

See Appendix A.4 for the proof. The limit loss characterization in (5) has an intuitive interpretation. It is the DM's mean squared error (MSE) in the perception of the lottery value integrated over all decision problems in which the true lottery value r ties with the outside option value s (multiplied by $n/2$). The conditioning on the tie arises because the likelihood of large perception errors vanishes quickly with increasing n. Asymptotically, only small perception errors contribute significantly to the loss, and these distort choice only in decision problems in which an approximate tie arises. In the limit, the set of decision problems in which perception errors have nontrivial loss consequences approaches the set of

[10] Assuming positive sampling frequencies is without loss, because in the limit when the number of signals grows it is optimal to gather at least some information about the rewards in each state.

[11] Since s can have a distinct density from that of r_J, the safe option may for example capture, in reduced form, the choice of an alternative lottery with each of its rewards drawn from h.

problems with exact ties.[12] To understand the relevance of the MSE for loss, fix the nearby true and perceived lottery values. The perception error distorts choice and causes loss if and only if the safe option s falls in between these two values. Hence, a mistaken choice arises with a probability proportional to the size of the perception error. The loss associated with such a mistake is also proportional to the error size, making the expected loss proportional to the MSE.

The precision of the DM's perception varies across lotteries. Observe that, approximately for large n, the DM measures reward r_J with MSE equal to $1/(\pi_J n m'^2(r_J))$. This is because she observes the encoded value $m(r_J)$ with MSE $1/(\pi_J n)$, and the map between the encoded and the true value can be locally linearized for large n, so that the local slope $m'(r_J)$ of the encoding function determines the precision of the estimate. The sum in (5) is thus the MSE of the perception of the value of lottery $\mathbf{r}$, scaled up by n.

Motivated by the asymptotic loss characterization, we now fix a large n, use (5) to approximate the expected loss, and solve for the encoding strategy which minimizes the loss. Formally, we define the *information-processing problem* as the minimization of the expected MSE conditional on ties:

$$\min_{m'(\cdot)>0,(\pi_J)_J>0} \quad \mathrm{E}\left[\sum_J \frac{p_J^2}{\pi_J m'^2(r_J)} \mid r = s\right] \tag{6}$$

$$\text{s.t.:} \quad \int_{\underline{r}}^{\overline{r}} m'(\tilde{r})d\tilde{r} \leq \overline{m} - \underline{m} \tag{7}$$

$$\sum_J \pi_J = 1. \tag{8}$$

The objective in (6) equals the asymptotic loss characterized in (5), up to a factor that is independent of the encoding strategy.[13] We let the DM control the derivative $m'(\cdot)$. Constraint (7) is implied by the finite range of the encoding function – the encoding function

[12]A related marginal argument was used in Steiner and Stewart (2016) in their analysis of the optional perception of probabilities.

[13]The conditional expectation in (6) is

$$\mathrm{E}\left[\sum_J \frac{p_J^2}{\pi_J m'^2(r_J)} \mid r = s\right] = \frac{\mathrm{E}\left[\sum_J \frac{p_J^2}{\pi_J m'^2(r_J)} h_s(r)\right]}{\mathrm{E}[h_s(r)]}.$$

It coincides with the asymptotic loss from (5) up to the ex ante likelihood of a tie, $\mathrm{E}[h_s(r)]$, and the scaling factor $2/n$. A special case in which conditioning on ties can be ignored is when s is uniformly distributed, because conditional and unconditional MSE are then identical up to a constant. Earlier work has assumed the minimization of unconditional MSE (e.g. Woodford, 2012).

cannot be steep everywhere. Constraint (8) requires $(\pi_J)_J$ to be a probability distribution – the DM must also treat sampling frequencies as a scarce resource. We say that an encoding strategy $(m(\cdot), (\pi_J)_J)$ is optimal if $(m'(\cdot), (\pi_J)_J)$ solves the information-processing problem.

3.2 Optimization

We say that a density $f(x)$ on $[\underline{r}, \overline{r}]$ is *unimodal* and *symmetric* around the mode $r_m = (\underline{r} + \overline{r})/2$ if it is strictly decreasing on $(r_m, \overline{r}]$ and $f(r_m + y) = f(r_m - y)$ for all y.[14]

Proposition 5. *If the densities h and h_s are unimodal and symmetric, then the optimal encoding strategy has the following properties:*

1. *The encoding function is S-shaped: strictly convex below and strictly concave above r_m. Additionally, it is continuously differentiable.*
2. *The DM oversamples low-probability states. For any two states J, J' such that $p_J < p_{J'}$, it holds that $\frac{\pi_J}{p_J} > \frac{\pi_{J'}}{p_{J'}}$.*

The proof in Appendix A.5 derives first-order conditions of the information-processing problem for general distributions and then exploits unimodality and symmetry to establish the above properties.

The first statement of the proposition extends an existing result in the literature for choice over two riskless rewards. The core intuition is that to minimize the loss, the optimal encoding function is steep in the range of rewards that often occur in decision problems with ties. The novel challenge arising in our setting is to prove that the reward in each state conditional on a tie inherits the unimodality and symmetry property of its unconditional distribution, which then implies the S-shape of the encoding function. The solution restricted to riskless lotteries with a single state coincides with the optimal encoding from Netzer (2009) (see Lemma 4 in Appendix A.5 for details).

The second statement of the proposition has no counterpart in the existing literature. It is driven by our microfoundation of the objective of the information-processing problem. While rewards are assumed to be iid across the states unconditionally, conditional on a tie they are no longer identically distributed. The tie condition $\sum_J p_J r_J = s$ is relatively uninformative about rewards in low-probability states, and hence the conditional reward distributions for

[14] Symmetry is sufficient but not necessary for the statement of Proposition 5. Our proof exploits that symmetry combined with unimodality of random variables is preserved by summation. This implies that the distribution of each reward r_J conditional on a tie is unimodal. Unimodality in absence of symmetry is generally not preserved by summation. Note that, if the safe option is the value of an alternative lottery with rewards drawn from h, as discussed in footnote 11, then unimodality and symmetry of h implies unimodality and symmetry of h_s.

the low-probability states are more spread-out compared to the high-probability states. In simple words, since a tail reward in a high-probability state makes the lottery value extreme and not likely to result in a tie with the outside option, once the DM conditions on the pivotal event of a tie, she expects to encounter more tail rewards in low- than in high-probability states. Because the optimal encoding function is relatively flat at such tail rewards, the DM measures the rewards of the low-probability states relatively poorly. Sampling according to the true probabilities would, therefore, leave the DM relatively poorly informed about the low-probability states, and hence the marginal benefit of an additional signal would be larger for those than for the other states. As a consequence, the DM finds it optimal to oversample the low-probability states relative to proportional sampling.[15] In particular, when there are two states, then $\pi_J > p_J$ for the state with probability $p_J < 1/2$ and vice versa for the high-probability state. Had the DM minimized the unconditional MSE, the effect would not arise. By taking the instrumental perspective that focuses on the payoff consequences of perception errors in choice problems, we obtain an objective that conditions on ties and induces nontrivial sampling frequencies as the optimal adaptation.

4 Related Literature

We build on a rich literature in neuroscience and economics, to which we make two distinct contributions. First, we clarify the role of misspecification for behavioral consequences of any perception strategy when stakes are large relative to perception frictions. Second, we jointly optimize both encoding of the lottery rewards and their sampling frequencies.

For the first contribution, we apply the statistical results of Berk (1966) and White (1982) on asymptotic outcomes of misspecified Bayesian and maximum-likelihood estimation. The concept of Berk-Nash equilibrium in Esponda and Pouzo (2016) is defined as a fixed point of misspecified learning. This has motivated a renewed interest in misspecification across economics. Heidhues et al. (2018) characterize a vicious circle of overconfident learning, Molavi (2019) studies the macroeconomic consequences of misspecification, Frick et al. (2021) rank the short- and long-run costs of various forms of misspecification, and Eliaz and Spiegler (2020) focus on political-economy consequences of misspecification. We study the interplay of encoding and misspecified decoding of rewards, thereby connecting perception to classical representations using a Bernoulli utility function and subjective probability weights.

[15] In Woodford (2012), a decision-maker collects information about multi-attribute objects through an information channel and allocates finite channel capacity across the different attributes. It is then optimal to allocate more capacity to attributes with higher prior variance. In our analysis, the more spread-out reward distributions in low-probability states arise endogenously in the pivotal decision problems, and the (optimal) S-shaped value function is crucial for making oversampling optimal.

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

Our second contribution derives ultimately from psychophysics, a field that originated in Fechner's (1860) study of stochastic perceptual comparisons based on Weber's data.[16] A large literature in brain sciences and psychology views perception as information processing via a limited channel and studies the optimal encoding of stimuli for a given channel capacity. Laughlin (1981) derives and tests the hypothesis that optimal encoding under an information-theoretic objective encodes random stimuli with neural activities proportional to their cumulative distribution value. This implies S-shaped encoding for unimodal densities.[17]

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.[18] Economics can help here by providing microfoundations for the most appropriate optimization objective for perceptions related to choice. Robson (2001) has studied encoding of rewards that minimizes the probability of making a wrong choice and has shown that, in the limit of vanishing perception frictions, the optimal encoding function likewise coincides with the cumulative distribution function of rewards in the decision environment. Netzer (2009) has studied maximization of the expected chosen reward, an objective rooted in the instrumental approach of economics to information. The optimal encoding function still tracks the cumulative distribution function but is flattened. Schaffner et al. (2023) report that the optimal encoding function as in Netzer provides a better fit to neural data than do encodings derived under competing objectives.

These models study choices over riskless prizes and thus 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 that literature (see Rayo and Becker, 2007).[19] We extend Netzer's instrumental approach to choices over

[16] Woodford (2020) provides a review of psychophysics from an economics perspective.

[17] See Attneave (1954) and Barlow (1961) for early contributions and Heng et al. (2020) for recent work.

[18] See e.g. Bethge et al. (2002) and Wang et al. (2016).

[19] The optimal hedonic utility function of Rayo and Becker (2007) is a step function. They provide an

gambles, finding a connection to one of the above reduced-form objectives.[20] That is, in the limit with rich perception data, maximization of the expected chosen reward is equivalent to minimization of the expected mean squared error in the perceived lottery value, where the expectation is over all decision problems with a tie. This conditioning on ties not only generates the better fit of the optimal encoding function documented by Schaffner et al. (2023) but is also crucial for the result of optimal oversampling of low-probability contingencies. Oversampling would not arise under reduced-form objectives that maximize unconditional measures of precision.[21]

Some recent papers study risk attitudes stemming from reward encoding in the presence of noise. Khaw et al. (2021) show theoretically and verify experimentally that exogenous logarithmic stochastic encoding and Bayesian decoding generates risk attitudes in an effect akin to reversion to the mean, compatible with the paradox of Rabin (2000). Vieider (2023) proposes a model in which probabilities are also encoded in an exogenous logarithmic way and establishes a connection to stochastic prospect theory. Frydman and Jin (2022) and Juechems et al. (2021) allow for optimal encoding of the lottery reward and show both theoretically and experimentally that this encoding adapts to the distribution of the decision problems and that the adaptation affects choice. Like ours, these papers are in principle compatible with the findings of Oprea (2023). Relative to these papers, we analyze optimal encoding of rewards alongside optimal treatment of probabilities. We also differ in the proposed source of behavioral distortions. The discussed models assume well-specified learning, and thus they approximate the frictionless benchmark when noise becomes small. We focus on the limit of small encoding noise right away. This focus uncovers a novel connection between coding and behavior. While the impact of coding on behavior must necessarily vanish when the decoding model is well-specified, as in the previous literature, the implications for behavior remain substantial if the cognitive model used for decoding oversimplifies the structure of the risk.

---

extension in which this function becomes s-shaped. Robson et al. (2023) is a dynamic version of Robson (2001) and Netzer (2009) that captures low-rationality, real-time adaptation of a hedonic utility function used to make ultimately deterministic choices. Friedman (1989) is an early approach dealing with gambles.

[20]Our model differs from Robson (2001) and Netzer (2009) concerning the perception friction. Those papers model frictions as minimal just noticeable differences, while here we rely on the modeling framework of Thurstone (1927) who hypothesized that perception is a Gaussian perturbation of an encoded stimulus. Payzan-LeNestour and Woodford (2021) have shown that the Gaussian approach yields the same limiting results as in Robson (2001) and Netzer (2009).

[21]Herold and Netzer (2023) 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. Lieder et al. (2017) argue that a contingency should be oversampled if it has extreme payoff consequences and decisions are based on a small sample. The present paper derives both S-shaped encoding and low-probability over-sampling in a joint optimization.

5 Conclusion

We develop a model of perception and aggregation of lotteries (or other complex options). Our model has an *encoding* stage, where the DM generates signals about the lottery rewards using an encoding function and sampling frequencies that jointly determine how she allocates her attention, and a *decoding* stage, where the DM estimates the value of the lottery based on the generated signals using a more or less sophisticated estimation procedure.

We first show that the behavioral impact of encoding vanishes for rich perception data if the DM encounters a lottery that she has anticipated and that she therefore decodes using a well-specified estimation procedure. On the other hand, encoding-induced behavioral risk attitudes arise for lotteries that the DM has not anticipated and that she therefore decodes using a misspecified estimation procedure. In the latter case, our model provides a unified explanation for multiple well-documented empirical patterns: adaptive risk preferences, different risk attitudes for small and large risks, probability weighting, availability heuristics, the roles of salience, time pressure and experience, and behavioral risk attitudes toward non-risky options. Second, we derive properties of optimal encoding in a model with well-specified decoding and show that the optimal encoding strategy typically exhibits S-shaped reward encoding and oversampling of low-probability states.

An interesting question that our paper leaves open is about the properties of optimal encoding when decoding is anticipated to be misspecified. Even though evolutionary processes are typically not forward looking, there may be benefits to a robust encoding solution that performs well when the DM acts in changing environments for which she is misspecified. When considering the limit of rich perception data, like we did in this paper, the large mistakes from misspecification will always outweigh the small mistakes from perception errors, and this constitutes a force towards linear encoding functions (Rustichini et al., 2017) which mitigate the misspecification bias. A more suitable framework for the study of optimal encoding with misspecified decoding may be a Bayesian model with prior information like in our Subsection 2.3, where perception errors and misspecification errors remain of comparable size. Optimal perception of lotteries in the presence of large encoding noise and with anticipation of one's own decoding frictions remains a hard open problem that may require new conceptual breakthroughs.

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.

Bordalo, P., N. Gennaioli, and A. Shleifer (2012). Salience theory of choice under risk. *The Quarterly Journal of Economics 127*(3), 1243–1285.

Bradbury, M., T. Hens, and S. Zeisberger (2015). Improving investment decisions with simulated experience. *Review of Finance 19*, 1019–1052.

Cahlíková, J. and L. Cingl (2017). Risk preferences under acute stress. *Experimental Economics 20*(1), 209–236.

Charness, G., N. Chemaya, and D. Trujano-Ochoa (2023). Learning your own risk preferences. *Journal of Risk and Uncertainty 67*, 1–19.

Eliaz, K. and R. Spiegler (2020). A model of competing narratives. *American Economic Review 110*(12), 3786–3816.

Ert, E. and E. Haruvy (2017). Revisiting risk aversion: Can risk preferences change with experience? *Economics Letters 151*, 91–95.

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 (2022). Efficient coding and risky choice. *Quarterly Journal of Economics 137*(1), 161—-213.

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

Heng, J. A., M. Woodford, and R. Polania (2020). Efficient sampling and noisy decisions. *Elife 9*, e54962.

Herold, F. and N. Netzer (2023). Second-best probability weighting. *Games and Economic Behavior* (138), 112–125.

Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory 123*(2), 81–104.

Johnson, D. H. and G. C. Orsak (1993). Relation of signal set choice to the performance of optimal non-gaussian detectors. *IEEE Transactions on Communications 41*(9), 1319–1328.

Juechems, K., J. Balaguer, B. Spitzer, and C. Summerfield (2021). Optimal utility and probability functions for agents with finite computational precision. *Proceedings of the National Academy of Sciences 118*(2).

Kahneman, D. (2011). *Thinking, fast and slow.* Macmillan.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–291.

Kahneman, D., P. P. Wakker, and R. Sarin (1997). Back to bentham? explorations of experienced utility. *The Quarterly Journal of Economics 112*(2), 375–406.

Khaw, M. W., Z. Li, and M. Woodford (2021). Cognitive imprecision and small-stakes risk aversion. *The Review of Economic Studies 88*, 1979—-2013.

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.

Köszegi, B. and M. Rabin (2007). Reference-dependent risk attitudes. *American Economic Review 97*, 1047–1073.

Koszegi, B. and M. Rabin (2008). Revealed mistakes and revealed preferences. *The foundations of positive and normative economics: a handbook*, 193–209.

Koszegi, B. and A. Szeidl (2013). A model of focusing in economic choice. *Quarterly Journal of Economics 128*, 53—-104.

Laughlin, S. (1981). A simple coding procedure enhances a neuron's information capacity. *Zeitschrift für Naturforschung c 36*(9-10), 910–912.

Lieder, F., T. Griffiths, and M. Hsu (2017). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review 125*, 1—-32.

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

Ludvig, E., C. Madan, and M. Spetch (2015). Priming memories of past wins induces risk seeking. *Journal of Experimental Psychology: General 144*, 24–29.

Molavi, P. (2019). Macroeconomics with learning and misspecification: A general theory and applications. *Unpublished manuscript*.

Netzer, N. (2009). Evolution of time preferences and attitudes toward risk. *The American Economic Review 99*(3), 937–955.

Ok, E. A. (2011). *Real analysis with economic applications*. Princeton University Press.

Oprea, R. (2023). Simplicity equivalents. Mimeo.

Payzan-LeNestour, E. and M. Woodford (2021). Outlier blindness: A neurobiological foundation for neglect of financial risk. *Journal of Financial Economics* (in press).

Porcelli, A. J. and M. R. Delgado (2009). Acute stress modulates risk taking in financial decision making. *Psychological Science 20*(3), 278–283.

Purkayastha, S. (1998). Simple proofs of two results on convolutions of unimodal distributions. *Statistics & probability letters 39*(2), 97–100.

Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica 68*(5), 1281–1292.

Rayo, L. and G. S. Becker (2007). Evolutionary efficiency and happiness. *Journal of Political Economy 115*(2), 302–337.

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

Robson, A. J., L. A. Whitehead, and N. Robalino (2023). Adaptive utility. *Journal of Economic Behavior and Organization 211*, 60–81.

Rustichini, A., K. Conen, X. Cai, and C. Padoa-Schioppa (2017). Optimal coding and neuronal adaptation in economic decisions. *Nature Communications 8*, 1–14.

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., S. D. Bao, P. N. Tobler, T. A. Hare, and R. Polania (2023). Sensory perception relies on fitness-maximizing codes. *Nature Human Behavior 7*, 1135—-1151.

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.

Tversky, A. and D. Kahneman (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology 5*, 207–232.

Vieider, F. (2023). Decisions under uncertainty as bayesian inference on choice options. *Management Science*.

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. (2012). Prospect theory as efficient perceptual distortion. *American Economic Review, Papers & Proceedings 102*(3), 41–46.

Woodford, M. (2020). Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics 12*, 579–601.

A Proofs

A.1 Proof of Proposition 1

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. The Kullback-Leibler divergence of the signal densities for any two lotteries $\mathbf{r}, \mathbf{r}'$ is

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

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

Let

$$
\mathbf{r}^* = \operatorname*{arg\,min}_{\mathbf{r}' \in \mathcal{A}_{\mathcal{P}}} D_{\mathrm{KL}}\left(f_{\mathbf{r}} \parallel f_{\mathbf{r}'}\right) = \operatorname*{arg\,min}_{\mathbf{r}' \in \mathcal{A}_{\mathcal{P}}} \sum_{i=1}^{I} \pi_i \left(m\left(r_i\right) - m\left(r_i'\right)\right)^2 .
$$

This minimizer $\mathbf{r}^* = (r_i^*)_i$ is unique and satisfies, for each state $i = 1, \ldots, I$,

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

where $J(i)$ is the element of the partition $\mathcal{P}$ that contains i and $\pi_{J(i)} = \sum_{j \in J(i)} \pi_j$. The estimated lottery value q_n almost surely converges to $\sum_{i=1}^{I} p_i r_i^*$, which follows from White (1982) who proves that $\mathbf{q}_n$ almost surely converges to the minimizer of the Kullback-Leibler divergence (provided the minimizer is unique).

A.2 Proof of Proposition 2

We first state a lemma that will be useful for the proof of Proposition 2.

Lemma 1. *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 PDF equal to $\alpha_n \exp(-n\psi_n(\mathbf{x}))$, where α_n is the normalization factor, converges to $\mathbf{x}^*$ in probability as $n \to \infty$.*

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

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

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

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

$$P(X_n \notin B_\delta) \leq \tilde{\alpha}_n \exp\left(-\frac{3d}{4} n\right) (\overline{r} - \underline{r})^I, \qquad \text{where } \tilde{\alpha}_n := \alpha_n \exp(-n\psi(\mathbf{x}^*)). \tag{9}$$

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

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

where $b' > 0$ is the volume of the set $B_{\delta'} \cap [\underline{r}, \overline{r}]^I$. Substituting the implied upper bound on

$\tilde{\alpha}_n$ into (9) gives

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

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

We now prove Proposition 2. Let $\hat{\mathbf{m}}_n = (\hat{m}_{in})_{i=1}^I$ be the tuple of the averages of $a\pi_i n$ perturbed messages received for each state i. Since the encoding errors are standard normal, this tuple of averages is a sufficient statistic for the Bayesian estimation, and we have $\hat{m}_{in} \sim \mathcal{N}\left(m(r_i), \frac{1}{a\pi_i n}\right)$. By Bayes' Rule, the posterior density of each lottery $\mathbf{r'} \in [\underline{r}, \overline{r}]^I$ is, for given $\hat{\mathbf{m}}_n$, proportional to

$$\varrho_n(\mathbf{r'}) \prod_{i=1}^{I} \varphi\Big(\big(\hat{m}_{in} - m(r'_i)\big)\sqrt{a\pi_i n}\Big) \propto \exp\big(-n\psi(\mathbf{r'}; \hat{\mathbf{m}}_n)\big),$$

where $\propto$ denotes equality modulo normalization and

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

Throughout this paragraph, consider a fixed realization of the sequence $(\hat{\mathbf{m}}_n)_n$ such that $\hat{m}_{in} \to m(r_i)$ for all i. Then, $\psi(\mathbf{r'}; \hat{\mathbf{m}}_n)$ converges to $\psi\big(\mathbf{r'}; (m(r_i))_i\big)$, uniformly in $\mathbf{r'}$. 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. Lemma 1 implies that the posterior formed given $\hat{\mathbf{m}}_n$ converges in probability to $\mathbf{q}^*(\mathbf{r})$. Since the support of the rewards is bounded, convergence in probability implies convergence in expected value, and thus the Bayesian estimate $\mathrm{E}\left[\hat{\mathbf{r}} \mid \hat{\mathbf{m}}_n\right] \in [\underline{r}, \overline{r}]^I$ converges to $\mathbf{q}^*(\mathbf{r})$. Since $\hat{m}_{in} \to m(r_i)$ almost surely, we conclude that $\mathrm{E}[\sum_{i=1}^I p_i \hat{r}_i \mid \hat{\mathbf{m}}_n] \in [\underline{r}, \overline{r}]$ converges to $\sum_{i=1}^I p_i q_i^*(\mathbf{r})$ almost surely. Here, $\hat{\mathbf{r}}$ and $\hat{r}_i$ stand for random variables and $\mathbf{r}$ and r_i are their realizations.

A.3 Proof of Proposition 3

By Proposition 2, the Bayesian estimate of $\mathbf{r}$ converges to $\mathbf{q}^*(\mathbf{r})$ almost surely. We write $\mathbf{q}^* = (q_i^*)_{i=1}^I$ as an abbreviation for $\mathbf{q}^*(\mathbf{r})$ and let $q^* = \sum_i p_i q_i^*$. The first-order condition of the minimization in (3) implies

$$(q_i^* - q^*) + a\Delta\big(m(q_i^*) - m(r_i)\big)m'(q_i^*) = 0 \tag{10}$$

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 σ^2 for the true reward variance $\sigma^2(\mathbf{r})$ and $\sigma^{*2} := \sum_{i=1}^I p_i\left(q_i^* - q^*\right)^2$ for the estimated variance. We will prove the following claims (see Footnote 7 for the definition of the "order smaller than" convention $o(\cdot)$):

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

To prove Claim 1, we provide a bound on the distance of r_i and q_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 $\underline{z} = a\Delta\underline{m}'^2$, $\overline{z} = a\Delta\overline{m}'^2$. We have $0 < \underline{m}' \leq \overline{m}' < +\infty$ and $0 < \underline{z} \leq \overline{z} < +\infty$ 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 (10) 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,

$$q_i^* - q^* = \tfrac{z_i}{1+z_i}(r_i - q^*) = \tfrac{z_i}{1+z_i}(r_i - r) + \tfrac{z_i}{1+z_i}(r - q^*). \tag{11}$$

Summing up the last equation weighted by p_i over 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 $0 < \frac{\underline{z}}{1+\underline{z}} \leq \frac{z_i}{1+z_i} < 1$. The 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 equation (11),

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

We conclude that $|q_i^* - r| \leq \left(p_i^{-1/2} + 2\frac{1+\overline{z}}{\underline{z}}I\right)\sigma$ 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 (10) for $\sigma > 0$ small. Since $m(\cdot)$ is twice differentiable, the functions m and m' can be expressed using first-order Taylor approximations around r:

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

where we used Claim 1 to replace $o(r_i - r)$ and $o(q_i^* - r)$ by $o(\sigma)$. Equation (10) 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

$$0 = (q_i^* - q^*) + z(r)(q_i^* - r_i) + o(\sigma). \tag{12}$$

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

We rewrite (12) 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 $z(r) \leq \overline{z}$ 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 equation (10),

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

which we rewrite 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

$$0 = (q^* - r) - \tfrac{1}{2}\tfrac{m''(r)}{m'(r)}\left(\sigma^2 - \sigma^{*2} + 2\sum_{i=1}^{I} p_i\big(r_i - q_i^*\big)\big(q_i^* - r\big)\right) + o(\sigma^2). \tag{13}$$

Expressing $q_i^* - r_i$ from (12) 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 (13) 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 $1 + \frac{2z(r)-z(r)^2}{(1+z(r))^2} = \frac{1+4z(r)}{(1+z(r))^2}$, we obtain (4), concluding the proof.

A.4 Proof of Proposition 4

The encoding error $\hat{m}_{Jn} - m(r_J)$ is drawn from $\mathcal{N}(0, 1/(\pi_J n))$. For each n, we set $\hat{m}_{Jn} - m(r_J) := \varepsilon_J/\sqrt{\pi_J n}$, where $\varepsilon_J \sim \mathcal{N}(0,1)$ is an error factor common across all n, and independent across J. This choice of correlation of the errors across n is without loss of generality since it does not affect the expected loss for each n (see e.g. Lindvall (2002) for this technique known in probability theory as coupling).

We extend the inverse encoding function m^{-1} outside of the interval $(\underline{m}, \overline{m})$ by setting $m^{-1}(\hat{m}_{Jn}) = \underline{r}$ for $\hat{m}_{Jn} \leq \underline{m}$ and $m^{-1}(\hat{m}_{Jn}) = \overline{r}$ for $\hat{m}_{Jn} \geq \overline{m}$. This allows us to express the ML estimate of the lottery value as

$$\begin{aligned} q_n &= \sum_J p_J q_{Jn} \\ &= \sum_J p_J m^{-1}\left(m(r_J) + \frac{\varepsilon_J}{\sqrt{\pi_J n}}\right). \end{aligned}$$

We start the proof of Proposition 4 with a lemma that we will use below for an application of the Dominated Convergence theorem. The lemma establishes an integrable bound on the rescaled error of the estimated lottery value. Let $\varepsilon := (\varepsilon_J)_J$.

Lemma 2. *There exists a function $\overline{e}(\varepsilon)$ such that $|\sqrt{n}(q_n - r)| \leq \overline{e}(\varepsilon)$ for all $\mathbf{r}$, ε and n, and $\mathrm{E}\,\overline{e}^2(\varepsilon) < \infty$.*

Proof. Let $\underline{m'} > 0$ be a lower bound for $m'(r)$ on $[\underline{r}, \overline{r}]$, which exists since m' is positive and continuous. Observe a bound on the estimation error for the reward r_J,

$$|q_{Jn} - r_J| \leq \frac{|\varepsilon_J|}{\underline{m'}\sqrt{\pi_J n}}, \tag{14}$$

which holds uniformly for all J, r_J and ε_J. To see that (14) holds, note that $q_{Jn} - r_J$ and ε_J have a same sign since m^{-1} is monotone and $q_{Jn} = r_J$ if $\varepsilon_J = 0$. Consider positive $q_{Jn} - r_J$ (the negative case is analogous). If $m(r_J) + \frac{\varepsilon_J}{\sqrt{\pi_J n}} \leq \overline{m}$, then (14) follows from $\partial m^{-1}(\cdot) \leq 1/\underline{m'}$. If $m(r_J) + \frac{\varepsilon_J}{\sqrt{\pi_J n}} > \overline{m}$, then $q_{Jn} = \overline{r}$ and thus $q_{Jn}(r_J, \varepsilon_J) - r_J = q_{Jn}(r_J, \varepsilon'_J) - r_J$, where $\varepsilon'_J \in (0, \varepsilon_J)$ is defined by $m(r_J) + \frac{\varepsilon'_J}{\sqrt{\pi_J n}} = \overline{m}$. Then, $q_{Jn}(r_J, \varepsilon_J) - r_J = q_{Jn}(r_J, \varepsilon'_J) - r_J \leq \frac{|\varepsilon'_J|}{\underline{m'}\sqrt{\pi_J n}} \leq \frac{|\varepsilon_J|}{\underline{m'}\sqrt{\pi_J n}}$, as needed. We can thus define $\overline{e}(\varepsilon) := \sum_J p_J \frac{|\varepsilon_J|}{\underline{m'}\sqrt{\pi_J}}$. □

For a given $\mathbf{r}$, ε and s, we denote the DM's loss by

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

which depends on $\mathbf{r}$ and ε via q_n and r. We introduce substitution $s = r + \frac{\sigma}{\sqrt{n}}$ and denote the rescaled loss as $\ell_n(\mathbf{r}, \varepsilon, \sigma) := \sqrt{n}\tilde{\ell}_n\left(\mathbf{r}, \varepsilon, r + \frac{\sigma}{\sqrt{n}}\right)$. Let

$$\ell^*(\mathbf{r}, \varepsilon, \sigma) = \begin{cases} \sigma & \text{if } 0 \leq \sigma \leq \sum_J p_J \frac{\varepsilon_J}{\sqrt{\pi_J} m'(r_J)}, \\ -\sigma & \text{if } 0 \geq \sigma \geq \sum_J p_J \frac{\varepsilon_J}{\sqrt{\pi_J} m'(r_J)}, \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 3. $\lim_{n\to\infty} \ell_n(\mathbf{r}, \varepsilon, \sigma) = \ell^*(\mathbf{r}, \varepsilon, \sigma)$ *almost everywhere.*

Proof. Choice of the DM differs from the optimal choice under complete information if and only if s attains a value in between r and q_n. In such cases, the loss of the DM relative to the complete-information choice is $|s - r|$. Therefore,

$$\ell_n(\mathbf{r}, \varepsilon, \sigma) = \begin{cases} \sigma & \text{if } 0 \leq \sigma \leq \sum_J p_J \left(m^{-1}\left(m(r_J) + \frac{\varepsilon_J}{\sqrt{\pi_J n}}\right) - r_J\right)\sqrt{n}, \\ -\sigma & \text{if } 0 \geq \sigma \geq \sum_J p_J \left(m^{-1}\left(m(r_J) + \frac{\varepsilon_J}{\sqrt{\pi_J n}}\right) - r_J\right)\sqrt{n}, \\ 0 & \text{otherwise.} \end{cases}$$

The right side converges pointwise to $\ell^*(\mathbf{r}, \varepsilon, \sigma)$, because

$$\begin{aligned} \lim_{n\to\infty}\left(m^{-1}\left(m(r_J) + \frac{\varepsilon_J}{\sqrt{\pi_J n}}\right) - r_J\right)\sqrt{n} &= \partial m^{-1}(m(r_J))\frac{\varepsilon_J}{\sqrt{\pi_J}} \\ &= \frac{\varepsilon_J}{m'(r_J)\sqrt{\pi_J}}. \end{aligned}$$

$\square$

To prove Proposition 4, observe that

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

where we applied substitution $s = r + \frac{\sigma}{\sqrt{n}}$. To apply the Dominated Convergence Theorem, we note that the last integrand is bounded as follows:

$$0 \leq \ell_n(\mathbf{r}, \varepsilon, \sigma)\, h_s\left(r + \frac{\sigma}{\sqrt{n}}\right)\prod_J(h(r_J)\varphi(\varepsilon_J)) \leq \overline{\ell}(\mathbf{r}, \varepsilon, \sigma)\,\overline{h}_s\prod_J(h(r_J)\varphi(\varepsilon_J)), \tag{15}$$

where $\overline{h}_s$ is an upper bound on the density h_s,[22]

$$\overline{\ell}\left(\mathbf{r},\varepsilon,\sigma\right)=\begin{cases}|\sigma| & \text{if } |\sigma|\leq\overline{e}(\varepsilon)\text{ and } r+\frac{\sigma}{\sqrt{n}}\in[\underline{r},\overline{r}],\\ 0 & \text{otherwise,}\end{cases}$$

and $\overline{e}(\varepsilon)$ is the bound from Lemma 2. Since $\overline{e}(\varepsilon)$ is an upper bound on the size of the error of the DM's estimate, $\overline{\ell}$ expands, relative to ℓ_n, the set of $(\mathbf{r},\varepsilon,\sigma)$ for which the erroneous choice occurs.

The upper bound in (15) is integrable as needed for the use of the Dominated Convergence Theorem:

$$\begin{aligned}&\int\overline{\ell}\left(\mathbf{r},\varepsilon,\sigma\right)\overline{h}_s\prod_J(h(r_J)\varphi(\varepsilon_J))d\sigma\prod_J(dr_Jd\varepsilon_J)\\ \leq\;&\overline{h}_s\int\int_{-\overline{e}(\varepsilon)}^{\overline{e}(\varepsilon)}|\sigma|d\sigma\prod_J(h(r_J)\varphi(\varepsilon_J))\prod_J(dr_Jd\varepsilon_J)\\ =\;&\overline{h}_s\int\overline{e}^2(\varepsilon)\prod_J(h(r_J)\varphi(\varepsilon_J))\prod_J(dr_Jd\varepsilon_J)\\ =\;&\overline{h}_s\,\mathrm{E}\,\overline{e}^2\left(\varepsilon\right),\end{aligned}$$

where the expectation in the last line is finite, as needed, by Lemma 2.

Hence, the Dominated Convergence Theorem, Lemma 3, and continuity of h_s imply

$$\begin{aligned}\lim_{n\to\infty}nL(n)&=\int\ell^*\left(\mathbf{r},\varepsilon,\sigma\right)h_s(r)\prod_J(h(r_J)\varphi(\varepsilon_J))d\sigma\prod_J(dr_Jd\varepsilon_J)\\ &=\mathrm{E}\left[\int_0^{\sum_J p_J\frac{\varepsilon_J}{\sqrt{\pi_J}m'(r_J)}}\sigma h_s(r)d\sigma\right]\\ &=\mathrm{E}\left[\frac{1}{2}\left(\sum_J p_J\frac{\varepsilon_J}{\sqrt{\pi_J}m'(r_J)}\right)^2h_s(r)\right]\\ &=\mathrm{E}\left[\frac{1}{2}\sum_J\frac{p_J^2}{\pi_Jm'^2(r_J)}h_s(r)\right],\end{aligned}$$

where the first two expectations are over $\mathbf{r}$ and ε and the last expectation is over $\mathbf{r}$. The last

[22] The bound on h_s exists since h_s continuous and the support is compact.

step follows from the fact that ε_J are iid standard normal and thus $\mathrm{E}\,\varepsilon_J^2 = 1$ and $\mathrm{E}\left[\varepsilon_J \varepsilon_{J'}\right] = 0$ for all $J \neq J'$.

A.5 Proof of Proposition 5

The next lemma states the first-order conditions to the information-processing problem without imposing unimodality and symmetry on the reward densities. To state the result, define

$$h_J\left(\tilde{r}\right) = \frac{h(\tilde{r})\,\mathrm{E}[h_s(r)|r_J = \tilde{r}]}{\mathrm{E}[h_s(r)]} = \frac{\int h_s(r)h(\tilde{r})\prod_{J' \neq J} h(r_{J'})dr_{J'}}{\int h_s(r)\prod_{J'} h(r_{J'})dr_{J'}},$$

which is the density of reward r_J in state J conditional on a tie $r = s$.

Lemma 4. *The information-processing problem has a unique optimal encoding strategy. This optimal strategy has the following properties:*

1. *The encoding function satisfies, for all $\tilde{r} \in [\underline{r}, \overline{r}]$,*

$$m'(\tilde{r}) = m_0 \cdot \left(\sum_J \frac{p_J^2}{\pi_J} h_J\left(\tilde{r}\right)\right)^{\frac{1}{3}}, \tag{16}$$

   where $m_0 \in \mathbb{R}_+$ is a normalization factor chosen such that $\int_{\underline{r}}^{\overline{r}} m'(\tilde{r})d\tilde{r} = \overline{m} - \underline{m}$.

2. *The sampling frequencies satisfy, for all $J, J' \in \mathcal{P}$,*

$$\left(\frac{p_J}{\pi_J}\right)^2 \mathrm{E}\left[\frac{1}{m'^2(r_J)} \mid r = s\right] = \left(\frac{p_{J'}}{\pi_{J'}}\right)^2 \mathrm{E}\left[\frac{1}{m'^2(r_{J'})} \mid r = s\right], \tag{17}$$

   where the expectations are over r_J and $r_{J'}$ with respect to the densities h_J and $h_{J'}$, respectively.

Proof. The objective of the information-processing problem is a functional

$$\mathcal{L}\left(m'(\cdot), (\pi_J)_J\right) = \mathrm{E}\left[\sum_J \frac{p_J^2}{\pi_J m'^2(r_J)} \mid r = s\right].$$

Since $\frac{p_J^2}{\pi_J m'^2(r_J)}$ is convex with respect to $\left(m'\left(r_J\right), \pi_J\right)$, the functional $\mathcal{L}$ is convex. Thus, considering that the constraints are linear, the first-order conditions are sufficient for a global minimum of the information-processing problem. Since the objective (6) is strictly decreasing in $m'(\cdot)$, the constraint (7) is binding. The Lagrangian of the constrained optimization problem (6)-(8) is

$$\sum_J \mathrm{E}\left[\frac{p_J^2}{\pi_J m'^2(r_J)} \mid r=s\right] + \lambda\left(\int_{\underline{r}}^{\overline{r}} m'(\tilde{r})d\tilde{r} - (\overline{m}-\underline{m})\right) + \mu\left(\sum_J \pi_J - 1\right) \quad =$$

$$\sum_J \int_{\underline{r}}^{\overline{r}} \frac{p_J^2}{\pi_J m'^2(\tilde{r}_J)} h_J(\tilde{r}_J)d\tilde{r}_J + \lambda\left(\int_{\underline{r}}^{\overline{r}} m'(\tilde{r})d\tilde{r} - (\overline{m}-\underline{m})\right) + \mu\left(\sum_J \pi_J - 1\right),$$

where λ and μ are the Lagrange multipliers for (7) and (8), respectively. For any $\tilde{r} \in [\underline{r}, \overline{r}]$, summing the derivatives w.r.t. $m'(\tilde{r})$ of all the integrands in the last inline expression gives the first-order condition

$$2\sum_J \frac{p_J^2}{\pi_J m'^3(\tilde{r})} h_J(\tilde{r}) = \lambda. \tag{18}$$

Expressing $m'(\tilde{r})$ from (18) gives (16). The first-order condition with respect to each π_J is

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

which implies (17). $\square$

Observe that the optimal m' is continuous since each h_J as defined above is continuous: h is continuous and, since h_s is continuous on a compact interval, it is uniformly continuous, and thus the function $\tilde{r} \mapsto \mathrm{E}[h_s(r) \mid r_J = \tilde{r}]$ is continuous as well.

Observe also that, when the DM compares two riskless rewards drawn independently from the same density h, the first statement of the lemma replicates the optimal encoding result from Netzer (2009). In this case, (16) implies that $m'(r) \propto h^{\frac{2}{3}}(r)$.

We next state three auxiliary lemmas about unimodal and symmetric random variables.

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

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

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

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

For the next two lemmas, let $X_0, X_1, \ldots, X_I$ be independent real-valued continuous random variables that are unimodal and symmetric around 0, where $X_1, \ldots, X_I$ are identically distributed while the distribution of X_0 may be distinct. Denote by h the density of each of the variables $X_1, \ldots, X_I$ and let h_0 be the density of X_0. Let $(p_1, \ldots, p_I) \in \Delta(\{1, \ldots, I\})$ and $X := \sum_{i=1}^{I} p_i X_i$. We define the density of $(X_1, \ldots, X_I)|(X = X_0)$ to be $h(X_1) \times \cdots \times h(X_I) \times h_0(X)$, up to normalization, and we define $X_i|(X = X_0)$ by marginalizing it.

Lemma 5. *The random variable $X_i \mid (X = X_0)$, $i = 1, \ldots, I$, is unimodal and symmetric around 0.*

Proof. Unimodality together with symmetry is preserved by multiplication by a constant and by summation, so the variable $X_{-i} := \frac{1}{p_i}(X_0 - \sum_{k \neq i} p_k X_k)$ is unimodal and symmetric around 0. Denote by h_{-i} the density of X_{-i}. Then $X_i \mid (X = X_0)$ is identical to $X_i \mid (X_i = X_{-i})$, and so its density is, up to a normalization constant, $h(x_i)h_{-i}(x_i)$. This function is unimodal and symmetric around 0, as needed. □

Lemma 6. *The random variable $X_i \mid (X = X_0)$ is more peaked than $X_j \mid (X = X_0)$ if and only if $p_i > p_j$.*

Proof. Without loss of generality, assume $\{i, j\} = \{1, 2\}$ (that is, either $i = 1$ and $j = 2$ or $i = 2$ and $j = 1$). Define $X_{-12} := X_0 - \sum_{k=3}^{I} p_k X_k$ (if $I = 2$, then $X_{-12} = X_0$) and let h_{-12} be its density. This is a unimodal random variable symmetric around 0. The random variable $X_i \mid (X = X_0)$ is identical to $X_i \mid (p_i X_i + p_j X_j = X_{-12})$ and so its density 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 the expectation, which is with respect to X_1 and X_2, is independent of i. Thus, for any $\alpha > 0$,

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

and

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

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

$$\left(P\left(|X_1|<\alpha \mid X=X_0\right)-P\left(|X_2|<\alpha \mid X=X_0\right)\right)\cdot \mathrm{E}[h_{-12}(p_1X_1+p_2X_2)]=$$

$$\iint_{(-\alpha,\alpha)\times\mathbb{R}}\Big(h_{-12}(p_1x_1+p_2x_2)-h_{-12}(p_1x_2+p_2x_1)\Big)h(x_1)h(x_2)dx_1dx_2=$$

$$\iint_{(-\alpha,\alpha)\times\left(\mathbb{R}\setminus(-\alpha,\alpha)\right)}\Big(h_{-12}(p_1x_1+p_2x_2)-h_{-12}(p_1x_2+p_2x_1)\Big)h(x_1)h(x_2)dx_1dx_2=$$

$$2\iint_{(-\alpha,\alpha)\times[\alpha,+\infty)}\Big(h_{-12}(p_1x_1+p_2x_2)-h_{-12}(p_1x_2+p_2x_1)\Big)h(x_1)h(x_2)dx_1dx_2,$$

where we used that the integral is 0 on the region $(-\alpha,\alpha)\times(-\alpha,\alpha)$ and that h and h_{-12} are symmetric around 0.

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

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

that

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

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

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

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

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

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

Lemma 7. *Let a function f be continuous, symmetric ($f(x)=f(-x)$) and increasing on $\mathbb{R}_+$, and let X_1,X_2 be unimodal continuous random variables that are symmetric around* 0 *and have bounded support. Then* $\mathrm{E}[f(X_1)] < \mathrm{E}[f(X_2)]$ *whenever* X_1 *is more peaked than* X_2*.*

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

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

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

We now prove Proposition 5. Statement 1 follows from (16) because, by Lemma 5, each conditional reward density h_J is unimodal with the same mode as that of the unconditional reward density h. Additionally, m' is symmetric around r_m since each h_J is symmetric around r_m. Now consider Statement 2. Suppose $p_J < p_{J'}$. By (17) it suffices to show that

$$\mathrm{E}\left[\frac{1}{m'^2(r_J)} \mid r = s\right] > \mathrm{E}\left[\frac{1}{m'^2(r_{J'})} \mid r = s\right]. \tag{19}$$

This indeed holds since, by Lemma 6, $r_{J'} \mid (r = s)$ is more peaked than $r_J \mid (r = s)$, and the inequality (19) then follows from Lemma 7 and the fact that $1/m'^2(r)$ is continuous and symmetric around r_m and increasing above r_m.