Rational Misspecification: Framework and Applications*

Ran Eilat[†] Kfir Eliaz[‡]

October 8, 2024

Abstract

This paper proposes a framework for assessing whether misspecified decision makers would be willing to pay for information that can potentially make them less misspecified. We introduce a prior-free approach, based on "constrained" maximal regret, to derive an upper bound on the subjective assessment of potential gains from acquiring a more accurate model. The constraint stems from the information currently available to the decision maker. We apply our approach to three prominent models of misspecified beliefs: coarse reasoning, sampling equilibria, and causal misperceptions.

Keywords: Misspecified beliefs; regret; value of information

JEL Classification Codes: D83; D84; D90

*We thank Ran Spiegler and Heidi Thysen for valuable comments.
[†]Department of Economics, Ben-Gurion University of the Negev, eilatr@bgu.ac.il
[‡]School of Economics, Tel Aviv University, kfire@tauex.tau.ac.il

1 Introduction

The literature on non-rational expectations, where decision-makers have misspecified beliefs about the steady-state mapping from their actions to consequences, offers several ways to model misspecifications and a variety of solution concepts for analyzing single-person and interactive decision-making under the various misspecifications. Notable examples include analogy-based expectation equilibrium (Jehiel, 2005; Jehiel, 2022), sampling (Osborne and Rubinstein, 1998; Salant and Cherry, 2020), causal misperceptions (Spiegler, 2016; Spiegler, 2020), cursed equilibrium (Eyster and Rabin, 2005; Cohen and Li, 2023), Berk-Nash (Esponda and Pouzo, 2016), and behavioral equilibrium (Esponda, 2008).

This literature typically takes as given the decision-makers' particular form of misspecification. For instance, a decision-maker may form beliefs about the mapping from actions to consequences based on a limited sample of the consequences that resulted from each action in the past (Osborne and Rubinstein, 1998). Or, a decision maker may believe in a particular, not necessarily correct, causal model that explains the relationship of some relevant variables (Spiegler, 2016). However, in these cases and others, a natural question arises: how does the decision-maker arrive at these particular forms of misspecifications to begin with? And, in particular, if a decision-maker suspects that his model of the environment is not perfectly accurate, why does he not attempt to reduce his error by acquiring more knowledge?

One possible answer is that more data is simply not available – if there exists only a single sample for the past consequence of each of the decision maker's actions, then he must make a decision based solely on this information. Another answer might be that the decision-maker is just completely unaware of his misspecification. While this may be true in certain cases, oftentimes reality is more nuanced. Indeed, in many situations decision-makers are aware that they do not *fully* understand the relationship between actions and consequences in the environment they operate in, yet they still do not engage in improving their model, even when data can be collected.

A possible explanation for this behavior is that acquiring more data may be costly. In this case, if the costs exceed the benefits, it is a "rational" choice for the decision-maker to remain misspecified. However, this raises another conceptual question: how can a misspecified decision-maker compute the benefit of becoming less misspecified?[1] Our goal in this paper is to propose a framework that takes a first step towards addressing this question and provides an approach to analyzing *rational misspecification*.

To illustrate the challenge, consider a two-player Bayesian game where a player forms

[1] In the context of analogy-based expectation equilibrium, Jehiel (2022) highlighted the difficulty of employing a cost-benefit analysis to endogenize a decision-maker's analogy partition, stating that "it is not clear how players would have the correct understanding about how their choices of analogy partitions translate into true payoffs."

expectations about his opponent's behavior using coarse data on the opponent's strategy, as in analogy-based expectation equilibrium (Jehiel, 2005). Specifically, suppose the set of types is $[0,1]$, and a player knows the marginal distribution over his opponent's actions and the marginal distribution over the opponent's types, but does not know the joint distribution of types and actions. Now, suppose this player is presented with the opportunity to refine his data by learning the marginal distributions over his opponent's actions for opponent types in $[0,0.5]$ and for opponent types in $[0.5,1]$. Clearly, this knowledge can only improve the player's decision. But how can the player quantify the extent of this improvement?

Although the basic idea that individuals weigh costs and benefits when deciding whether to become more informed appears also in the literature on costly information acquisition, there is an inherent conceptual difference between the two problems. The distinction lies in the difficulty of the misspecified decision-maker to quantify the benefits of acquiring additional data. For example, in the rational inattention literature (Sims, 2003; Maćkowiak et al., 2023) an uninformed decision-maker has to decide which information structure to acquire. To evaluate the expected benefit of any given information structure, the decision-maker relies on one of the framework's primitives – the true prior distribution over the states – to form beliefs about the possible consequences of learning. In contrast, a misspecified decision-maker operates with an erroneous model of the steady-state and needs to form beliefs about the expected gain from employing a more precise model. Short of exogenously imposing some arbitrary prior beliefs on the set of correct models, there is no primitive of the environment to guide the decision-maker in forming these beliefs.

In the absence of an objective prior beliefs on the set of models, there are myriad ways to form beliefs about what one might learn from a acquiring more data. In this paper we propose an *upper bound* on the subjective assessment of the expected gain from learning a more accurate model. This upper bound is computed by finding a new stochastic mapping from actions to consequences that satisfies the following properties: (i) it is consistent with the partial but correct information on the true mapping derived from the decision-maker's current misspecified model, and (ii) it maximizes the difference between the expected payoff the decision-maker could achieve from the more accurate model and the expected payoff he would obtain if he remained with the action he planned to take given his current model, with the expectation taken with respect to the new stochastic mapping. We refer to this difference as the *maximal regret* from not acquiring the more accurate model, and interpret it as the decision-maker's maximal willingness-to-pay for reducing his misspecification. Thus, when the cost of reducing the misspecification exceeds this maximal willingness-to-pay, the decision-maker can be said to be rationally-misspecified.

Our approach is motivated by the economic literature on decision-making without priors (in particular, Bergemann and Schlag (2008) and Bergemann and Schlag (2011)) and by the common phenomenon of "fear of missing out" or FOMO (see Milyavskaya et al., 2018;Laurence and Temple, 2023). The idea is that when a decision-maker is presented with the opportunity to acquire new knowledge that could potentially make him better off, he considers what he might be giving up if he forgoes that opportunity. Specifically, the new knowledge might prompt him to take a different action and obtain a significantly higher payoff; it might also make him realize that the action he was planning to take without the new knowledge would result in a low payoff. The greater the difference between these two potential payoffs, the more valuable the new knowledge becomes. The worst-case, in terms of what the decision-maker would lose by not acquiring the new knowledge, is represented by the maximal value this difference in payoffs can take.

When assessing his potential regret, the decision-maker does not consider the entire set of (stochastic) mappings from actions to consequences. Rather, he restricts his attention only to those mappings that are *consistent* with the partial information he already possesses. To illustrate this, recall the above example of the player with coarse data on his opponent. Suppose the decision-maker already knows that the steady-state marginal distribution over his opponent's actions is uniform when the opponent's type is in $[0,1]$, and he contemplates learning the conditional marginal distributions over actions when the opponent's type is in the intervals $[0,0.5]$ and $[0.5,1]$. In this case, when computing his maximal regret, the decision-maker considers only those conditional marginal distributions that are consistent with the overall distribution of his opponents actions being uniform on $[0,1]$.

As a different example, consider a decision-maker who only knows the correlation between two pairs of variables, (x,y) and (y,z). If this decision-maker can learn the true joint distribution over all three variables, consistency requires that this distribution must align with the pairwise correlations he already knows. Thus, a decision-maker's maximal regret from not reducing his misspecification is constrained by this consistency requirement.

To demonstrate our approach we apply it to three well-studied models of decision-making with misspecified beliefs: *coarse reasoning* (as captured by Jehiel (2005) notion of Analogy-Based-Expectations-Equilibrium or ABEE), *sampling* (as captured by Osborne and Rubinstein (1998)'s notion of S_k equilibrium) and *causal misperceptions* (as captured by Spiegler (2016)'s Bayesian networks framework). We provide a detailed explanation of each model in the corresponding section below. Here, we offer a brief summary of our analysis of these models.

In our first application we consider a classic adverse selection setting where an informed seller and an uninformed buyer trade using a simple double auction. Initially, the buyer is "fully coarse" in the sense that he knows the marginal distribution over the seller's type and the equilibrium marginal distribution over the seller's ask, and believes the two are independent.

The buyer is presented with the opportunity to refine his knowledge by partitioning the seller's types and acquiring the respective margins over each cell in the partition. We characterize the regret-maximizing partition.

Our second application analyzes sampling equilibria. We characterize a necessary and sufficient condition for when a player would "rationally" decide to sample each of his actions only once for any positive cost of a second sample. We illustrate this condition for binary action games that were the focus of Salant and Cherry (2020).

Finally, we consider a continuum of decision makers who believe in a misspecified causal model of the effect of a costly action on a personal outcome. For some range of costs there exists a mixed equilibrium in which only a subset of the population chooses the rational-expectations action (but they choose it for the wrong reason). In this equilibrium, the highest willingness-to-pay for learning the true model depends on the chosen action. In particular, if the cost of the action is below some threshold, the highest willingness-to-pay is higher for agents who choose the rational-expectations action. We then consider a simple competitive market for consultants who reveal the true model, and show that in equilibrium only agents who choose the rational-expectations action hire consultants.

These applications showcase the potential of our approach to "rationalize" decision makers with misspecified beliefs, and to endogenize the level of misspecification. The core idea of our approach is that by focusing on the "worst (constrained) case" of *not* treating the potential misspecification of one's beliefs, we can derive conditions for when decision-makers would choose to remain misspecified without introducing arbitrary prior beliefs. Of course, we are not claiming that our max-regret approach is the only plausible one. The literature on ambiguity may offer alternative approaches for conducting cost-benefit analysis without priors. Our hope is that this paper may enourage further research in this direction.

The remainder of the paper is organized as follows. Related literature is discussed immediately below. Section 2 formally introduces our approach. The next three sections analyze the three applications of our approach: coarse reasoning in Section 3, sampling in Section 4 and causal misperceptions in Section 5. Section 6 concludes.

Related literature. A number of alternative approaches have been proposed to endogenize decision-makers' misspecified beliefs. Gonçalves (2023) considers normal form games where each player is endowed with some exogenous prior over the other players' mixed strategies and decides sequentially whether to sample costly signals about these strategies. When a player decides to stop sampling, he chooses a best response to his updated beliefs given all the signals he observed. The paper proposes a new solution concept according to which a strategy profile is an equilibrium if each element in the profile is a best response to an optimal sequential sampling strategy, where the samples are taken from the equilibrium strategy profile. The one feature that is common to our approach and that of Gonçalves (2023), is an attempt to endogenize players' misspecified beliefs through a rational decision to acquire only partial information. In contrast to Gonçalves (2023) we do not impose an exogenous prior belief about the equilibrium. Instead, we adopt the approach of decision-making without priors to propose an upper bound on a player's willingness-to-pay to reduce his misspecification.

Heller and Winter (2020) take a different approach to justify the persistence of some forms of misspecified beliefs. They propose to view a pair consisting of an action profile and a profile of "belief distortion functions" (functions that take as input other players' actions and outputs a misspecified belief about them) as an equilibrium if (i) each player's action is a best response to the beliefs induced by applying his distortion function to the equilibrium actions of the other players, and (ii) if a player were to unilaterally adopt some other distortion function, then he would be worse off in some action profile where each player best responds to his distortion function applied to this new action profile.

He and Libgober (2023) propose an evolutionary approach to define a notion of "stable misspecifications". The authors consider a setting where a population of players are randomly matched every period and choose best responses to a possibly misspecified mapping from action profiles to outcomes. One mapping is deemed stable relative to another, if the former yields a weakly higher average (over the different stage games) equilibrium payoff than the latter when the population share of the stable model is close to one.

In the context of ABEE, Jehiel and Weber (2024) propose a framework for endogenizing analogy partitions with a given number of cells. In their setting, a player faces an opponent in a game, which is randomly drawn from set of games that have the same action set. A player partitions the set of games into K (exogenously given) cells and best responds to a belief that aggregates the opponent's strategies across all the games in each cell of the partition. A pair of strategies and a pair of partitions with K cells is an equilibrium if (i) the strategies form an ABEE given the partitions, and (ii) given the strategies, the partitions satisfy a property, which can be interpreted as minimizing prediction errors.

There is also a literature that takes a learning approach to justifying persistent misspecification (notable examples include Cho and Kasa (2015) and more recently, Ba (2024)). This literature considers a decision-maker who starts with a parametric model that maps actions to distributions over outcomes. The decision-maker is endowed with a misspecified prior belief over the true parameters (i.e., the true parameters are not in the distribution's support). Each period the decision-maker does two things: (1) he best replies to his model, updating his beliefs given the realized outcomes, and (2) he compares his current model to some alternative model by conducting some statistical test, and switches if the test result passes some threshold. The works in this literature characterize which misspecified model will persist in the long run.

A completely different approach to endogenize misspecified beliefs is to consider an interested third party that strategically provides a decision-maker with a (possibly misspecified) model of the steady-state in order to persuade him to choose a particular action. Some recent examples include Eliaz and Spiegler (2020); Eliaz et al. (2021c,a,b); Schwartzstein and Sunderam (2021) and Aina (2024).

The problem of evaluating the impact of new information in the absence of objective priors naturally comes up in decision-making under ambiguity: A decision-maker, who does not know the true distribution over some states, may encounter opportunities to acquire information that will reduce his ambiguity. As in our setting, here too, what the decision-maker thinks he might learn, may be constrained by his current knowledge.

Li (2020) considers a situation where a decision-maker facing ambiguity over some set of states can receive information that only a subset of these states is relevant. However, the decision-maker also faces ambiguity over which of these subsets will actually materialize. Li (2020) provides an axiomatic model in which the decision-maker evaluates the information in two steps: First, for each possible realized event, he computes his payoff according to his ambiguity-aversion over states (e.g., according to max-min expected utility); second, he evaluates the ambiguity over which subset will materialize according to the same model of ambiguity-aversion (e.g., max-min expected utility). Li (2020) shows that this representation can lead to a negative value of information. Kops and Pasichnichenko (2023) and Shishkin and Ortoleva (2023) conduct experimental tests on the relation between ambiguity-aversion and negative value of information and find mixed evidence. In contrast, models of belief-misspecifications, which are the focus of this paper, do not provide guidance on how the decision-maker may evaluate information that reduces his misspecification. Consequently, a two-step recursive procedure, like the one proposed by Li (2020), cannot be applied. This is where our approach, which bounds the willingness to pay for information with maximal regret, proves useful.

2 Framework

In this section, we introduce our framework, which provides a rationale for why a player might choose to remain misspecified and not seek new knowledge about the environment in which he operates. The core idea is that rational misspecification may occur when the costs of learning exceed the potential benefits. The challenge in applying this idea lies in finding an effective method to measure the benefit of learning. Indeed, as discussed earlier, unlike in other settings where the benefits of learning can be naturally computed using the model's primitives (e.g., prior beliefs over states), in the context of rational misspecification assuming that the player has prior beliefs about the potential models he might learn raises a conceptual difficulty. To address this issue, we propose to quantify the player's maximum willingness to pay for knowledge by using a notion of maximal regret.

We present our framework in four steps. First, we define the objective (or "true") environment in which the player operates. This environment is known to the modeler but not to the player. Next, we introduce the concept of a misspecified 'type' and explain how a player's type affects his decisions. We then define a player's regret from *not* adopting an alternative model. Based on this, we derive an upper bound on the player's willingness to pay for data that can lead him to adopt a new model, which is unknown at the time of acquiring the data. Finally, we say that a player rationally decides to remain misspecified if, among all models consistent with his type, there is no model for which the player's willingness to pay exceeds the cost.

The objective environment. A player has to choose an action from a compact set A. Each action is stochastically mapped to a consequence from a set Y via a function $g : A \to \Delta(Y)$.[2] The stochastic nature of this mapping could be due to an unknown state of nature, or because the consequence also depends on an unknown action by another player. We refer to g as the *true model of the environment*. The player's preferences are defined over $A \times Y$ and are represented by a bounded and continuous utility function $u : A \times Y \to \mathbb{R}$.

Misspecified models. The player does not know g. Instead, he works according to a (potentially) misspecified model which we represent by his *type*. Let Θ denote the set of types. Each type $\theta \in \Theta$ possesses a subjective model $g_\theta : A \to \Delta(Y)$ from actions to consequences which guides his choice of actions.[3] A model g_θ is considered *misspecified* if it differs from g. We assume that θ encapsulates all relevant information the player has about the environment. Thus, the optimal action for a player of type θ, denoted by a_θ, is given by:[4]

$$a_\theta = \arg\max_{a \in A} \int_{y \in Y} u(a, y) \mathrm{d}g_\theta(y \mid a) \tag{1}$$

where $g_\theta(\cdot \mid a)$ is the probability measure over consequences in Y that is generated by $g_\theta(a)$.

For example, a type θ may know the true marginal distribution over consequences in the population, i.e. $h(y) = \sum_{\theta \in \Theta} \pi(\theta) g_\theta(y \mid a_\theta)$ for all y, where $\pi(\theta)$ denotes the frequency of type $\theta \in \Theta$ in the population. Another example is that a type θ may know the support of $g(a)$, i.e. the set of possible consequences for each of his actions, but not the relative probabilities of these consequences that his action entails.

[2]We assume that Y is a Polish space, and denote by $\Delta(Y)$ the set of probability distributions over Y endowed with the weak* topology. The function g is assumed to be both measurable and continuous.

[3]We assume that the function g_θ is both measurable and continuous.

[4]Existence of a maximum in Eq. (1) follows from the compactness of A and the continuity of u and g_θ. For simplicity, we assume that this maximizer is unique. In Section 5 we relax this assumption.

Regret from *not* adopting a new model. Consider a player of type θ who possesses a subjective model g_θ which guides his choice of actions. Now suppose this player encounters an alternative model $\tilde{g} : A \to \Delta(Y)$. The player is uncertain whether $\tilde{g}$ is the correct model, yet recognizes that ignoring it could result in *regret*. We define regret as the difference between the expected payoff from the optimal action under the new model $\tilde{g}$, and the expected payoff from the original optimal action a_θ, with expectations about the relationship between actions and consequences are calculated according to $\tilde{g}$. Formally, the regret of type θ from continuing to operate under g_θ instead of adopting $\tilde{g}$ is given by:[5]

$$R_\theta(\tilde{g}) = \max_{a \in A} \int_{y \in Y} u(a, y) \mathrm{d}\tilde{g}(y \mid a) - \int_{y \in Y} u(a_\theta, y) \mathrm{d}\tilde{g}(y \mid a_\theta). \tag{Regret}$$

where $\tilde{g}(\cdot \mid a)$ is the probability measure over consequences in Y that is generated by $\tilde{g}(a)$.

This approach to quantifying the player's regret is inspired by the common phenomenon of FOMO. That is, a player is concerned that if he were to pass on the opportunity to act according to a new model, he would not only miss out the chance to earn a significantly high payoff, but also could realize that his current action is truly suboptimal.

Rational misspecification. Suppose a player of type θ is offered the opportunity to pay a cost c to obtain *data* about the environment that would lead him to revise his model. We represent such data by a set of potential new models G_θ that are all *consistent* with the information encoded in the player's type θ. The notion of consistency is context dependent and will be defined precisely for each of the applications in the subsequent sections. For now, we keep the definition abstract and illustrate it with the following example.

Consider a type θ who knows only the following information for each of his actions: (i) the feasible outcomes, and (ii) a single observation of an outcome from a previous instance when the action was taken. Suppose this type chooses an action according to some model g_θ that is consistent with this information. Now, suppose θ is offered access to new data that provides an additional observation of a past outcome for each action. This new data may lead θ to adopt a new model. Crucially, the set of possible new models, namely G_θ, includes only those consistent with the original information θ has. Specifically, for each action, models in G_θ must assign a positive probability to the outcome that was originally observed.

Before acquiring the data, the player does not know what his revised model will be and he holds no prior beliefs about the possible models in G_θ. However, given our definition of the player's regret from not adopting a model, the player would certainly decline the offer if the cost c exceeds the upper bound on potential regret across all models in G_θ. This upper bound,

[5] When there are more than one optimal actions according to g_θ, the regret may also depend on which of these actions were chosen. We demonstrate this in Section 5.

which we denote by $R^*(G_\theta)$, is given by:

$$R^*(G_\theta) = \sup_{\tilde{g} \in G_\theta} R_\theta(\tilde{g}). \tag{WTP}$$

We interpret $R^*(G_\theta)$ as an upper bound on what a player will pay for the data given by G_θ. To streamline the exposition, we will henceforth simply refer to $R^*(G_\theta)$ as the player's *willingness-to-pay (WTP)* for G_θ. Thus, if $c > R^*_\theta$ the player would pass on the opportunity to acquire the data and learn about a new model from actions to consequences. In this case, we say that the player is *rationally misspecified*.

Finally, we say that data represented by a set of models G_θ is *more valuable* than data represented by a set of models G'_θ if $R^*(G_\theta) > R^*\left(G'_\theta\right)$. I.e., if the WTP for G_θ is above the WTP for G'_θ. In the subsequent sections we explore whether data sets which are in some natural sense "more informative" are indeed more valuable.

In the next sections we apply this framework to three forms of misspecifications that have been analyzed in the literature: coarse reasoning, sampling and causal misperceptions. In each of these settings we explain how they map to the primitives we defined in this section.

3 Coarse reasoning

We start by applying our framework to misspecifications arising from coarse reasoning. By this we mean that a player forms beliefs about the mapping from actions to consequences based on coarse data on the equilibrium joint distribution over the action profile and states of nature. For instance, a player may have access only to the marginal distribution over another player's action and the marginal distribution over states, without further data about the joint distribution. In this case, the player extrapolates and fills in the missing data to form a subjective belief about the joint distribution over actions and states.

The literature offers several approaches to how a player extrapolates from his coarse data. We adopt the Analogy Based Expectations Equilibrium (ABEE) approach originally proposed in Jehiel (2005) and later extended to Bayesian games in Jehiel and Koessler (2008). Rather than present this framework in its full generality, we apply it to a classic adverse selection setting. This setting was analyzed under various approaches to coarse reasoning: "cursed equilibrium" in Eyster and Rabin (2005), "behavioral equilibrium" in Esponda (2008) and ABEE in Jehiel and Koessler (2008) and Spiegler (2011). Our analysis would continue to hold in other market settings that generate the same initial information held by the decision-maker (see also footnote 8).

The setting. A seller owns an object with privately observed quality ϕ, which is distributed according to a distribution F supported on $[0,1]$ with density f. We assume that F is such that $\phi + F(\phi)/f(\phi)$ is increasing in ϕ. In the mechanism-design terminology, this means that the seller's "virtual value" is increasing in his type.

Trade proceeds according to the following double auction protocol. The seller submits an ask $x \in [0,1]$; the buyer submits a bid $p \in [0,1]$; trade takes place at price p if $p \geq x$. The value of the object for the buyer is $v(\phi, b)$, where $v(\cdot,\cdot)$ is increasing in both parameters, and $v(\phi, b) \geq \phi$ for all $\phi \in [0,1]$, ensuring there are always gains from trade. The parameter $b \in \mathbb{R}^+$ captures the gains from trade in this setting. The seller's payoff is 0 if there is no trade, and $p - \phi$ otherwise. The buyer's payoff is 0 if there is no trade, and $v(\phi, b) - p$, otherwise.

For expositional purposes, it is convenient to have a stark rational expectations benchmark in which the the market collapses due to adverse selection. We therefore make the following assumption:[6]

$$\mathbb{E}_{\phi \sim F}\left[v(\phi, b) \,|\, \phi < p\right] < p \qquad \forall p \in (0,1] \tag{2}$$

In words, Eq. (2) implies that, for any price p, the buyer's expected value from trading with a seller who agrees to sell at price p is less than p.

It is straightforward to show that for a seller with quality ϕ, submitting an ask $x = \phi$ is a dominant strategy. Thus, in what follows, we focus on the buyer's problem.

ABEE. We follow Spiegler (2011) in describing the notion of ABEE in the present context. Let $\sigma : [0,1] \to \Delta([0,1])$ be a seller's strategy that maps each quality to a distribution over prices. As noted above, given the trading rule, the seller will use the deterministic mapping σ^* which equates the ask to the quality. A buyer with rational expectations would choose a bid that maximizes his expected payoff given a perfect perception of σ^*. In contrast, a misspecified buyer will maximize his expected payoff with respect to a coarse representation of σ^*. This representation takes the following form. The buyer is endowed with an *analogy partition* $\mathscr{C} = (C_1, \ldots, C_K)$ of $[0,1]$, where each cell C_k is an interval. Let $C(\phi)$ denote the cell containing the seller's quality ϕ. The buyer's coarse representation of σ^* is a mixed strategy $\sigma^{\mathscr{C}}$ such that for every seller quality ϕ, the strategy $\sigma^{\mathscr{C}}$ mimics the price distribution in the entire cell $C(\phi)$ in the sense that for all $x \in [0,1]$:[7]

$$\Pr[\sigma^{\mathscr{C}}(\phi) \leq x] = \Pr[\sigma^*(t) \leq x \mid t \in C(\phi)]$$

[6]For a simple example that satisfies this condition, see Example 1 below.

[7]As Spiegler (2011) explains, a possible interpretation of this notion of misspecification is the following: "when the buyer enters the market, he has access to records of all the ask prices that were previously submitted by the seller (or his previous incarnations), but he does not have acess to the records of the valuations that lay behind these ask prices. Following the "Occam's razor" principle, the buyer adopts the simplest theory that is consistent with the historical records, where simplicity here means that the theory is not allowed to depend on unobserveable variables as long as it is consistent with the data."

Mapping the setting to our framework. To cast the setting within our framework, we proceed in two steps. First, we describe how a misspecified buyer computes the distribution over the possible consequences of submitting a bid (i.e. a model $g : A \to \Delta(Y)$). This computation is performed given that the buyer is misspecified in the ABEE sense and has certain beliefs about the marginal distributions over the set of seller's qualities and the marginal distribution of asks based on his analogy partition. Next, we explain how these marginal distributions are determined for the misspecified buyer, both when he possesses the initial analogy partition and when he considers which models are consistent with new data he might acquire.

Fix an analogy partition $\mathscr{C} = (C_1, \ldots, C_K)$ of $[0,1]$. Suppose the buyer believes that the marginal distribution over seller's quality is given by H_ϕ, which admits density h_ϕ. The buyer also believes that the marginal distribution over asks, conditional on the seller's quality being in the cell C_k, is given by H_x^k. These marginals can come either from the buyer's initial misspecified model, or these could be what he "imagines" the marginals would be when he is offered the opportunity to refine his original partition to $\mathscr{C}$.

For example, suppose the seller's quality is uniformly distributed over $[0,1]$ and he plays his dominant strategy. If the buyer is initially fully coarse, i.e. his analogy partition $\mathscr{C} = (C_1 = [0,1])$ has a single cell, then $H_\phi = H_x^1 = \text{U}[0,1]$. If the buyer initially has the analogy partition $\mathscr{C} = (C_1 = [0,1/2], C_2 = [1/2,1])$, then $H_\phi = \text{U}[0,1]$, $H_x^1 = \text{U}[0,1/2]$ and $H_x^1 = \text{U}[1/2,1]$.

Trade occurs at the bid price, whenever it is higher than the ask. Thus, each bid price p induces a probability distribution over the set of consequences $Y = \{\varnothing\} \cup [0,1]$, where the outcome $\{\varnothing\}$ is interpreted as "no-trade" and any outcome $\phi \in [0,1]$ is interpreted as trade with a seller of quality ϕ. Therefore, for any bid price p, the buyer computes the (ex-ante) probability for no trade as follows:

$$\Pr(\{\varnothing\} \mid p) = \sum_{k=1}^{K} \Pr\left(\phi \in C_k\right) \Pr\left(x > p | \phi \in C_k\right) = \sum_{k=1}^{K} \left(H_\phi(\overline{C}_k) - H_\phi(\underline{C}_k)\right) \cdot \left(1 - H_x^k(p)\right), \tag{3}$$

where $\overline{C}_k$ and $\underline{C}_k$ denote the upper and lower boundaries of the cell C_k, respectively. For each k, the first multiplier on the right-hand side in Eq. (3) is the probability that the seller's quality ϕ is in the cell C_k, and the second multiplier is the probability that the ask is greater than the bid, conditional on the seller's quality being in the cell C_k.

For each bid price p, the buyer can also compute the probability of trade with any set of qualities $\Phi \subseteq [0,1]$. Under ABEE, his misspecification leads him to compute this probability using the marginals as follows:

$$\Pr(\textit{trade with sellers in } \Phi \mid p) = \sum_{k=1}^{K} \Pr\left(\phi \in C_k\right) \cdot \Pr\left(x < p \mid \phi \in C_k\right) \cdot \Pr\left(\phi \in \Phi \mid \phi \in C_k\right)$$

$$= \sum_{k=1}^{K} \left(H_x^k(p) \cdot \int_{\phi \in \Phi \cap C_k} h_\phi(z)\,dz \right). \tag{4}$$

Note that the player's misspecification is captured by his use of $\Pr(\phi \in \Phi \mid \phi \in C_k)$ instead of $\Pr(\phi \in \Phi \mid \phi \in C_k,\ x < p)$, as he ignores the fact that the seller's ask and quality are dependent.

A buyer's type is a partition of the interval $[0,1]$, and represents the buyer's analogy partition at the outset, before potentially acquiring new data. For simplicity, we restrict our attention to partitions with countably many elements. A type $\theta = \mathscr{C}^\theta = (C_1^\theta, \ldots, C_K^\theta)$ correctly perceives the marginal distribution of the seller's quality, and the marginal distribution of the seller's asks, conditional on the seller's quality being within any of the cells. Consequently, the buyer's type θ determines the marginal distributions $(H_\phi^\theta, H_x^{\theta,1}, \ldots, H_x^{\theta,K})$ in which the buyer believes as follows:

$$H_\phi^\theta(\phi) = F(\phi) \qquad \forall \phi \in [0,1] \quad \text{and} \tag{5}$$

$$H_x^{\theta,k}(x) = \frac{F(x) - F(\underline{C}_k^\theta)}{F(\overline{C}_k^\theta) - F(\underline{C}_k^\theta)} \qquad \forall k \in \{1, \ldots, K\} \text{ and } \forall x \in C_k^\theta. \tag{6}$$

Thus, the misspecified model g_θ of type θ is determined by Eqs. (3) and (4), which are computed based on the marginal distributions in Eqs. (5) and (6).[8]

Now, suppose the buyer is offered the opportunity to get access to a new partition $C = (C_1, ..., C_M)$, which is a refinement of $\mathscr{C}$. The set of mappings that are feasible for θ (the set G_θ) includes all the mappings that are induced according to the marginals $(H_\phi, H_x^1, \ldots, H_x^M)$, which satisfy the following:

$$H_\phi(\phi) = H_\phi^\theta(\phi) \qquad \forall \phi \in [0,1] \quad \text{and} \tag{7}$$

$$H_x^{\theta,k}(x) = \sum_{\ell : C_\ell \subseteq C_k} \left(F(\overline{C}_\ell) - F(\underline{C}_\ell) \right) H_x^\ell(x) \qquad \forall k \in \{1, \ldots, K\} \text{ and } \forall x \in C_k. \tag{8}$$

In words, Eq. (7) means that the marginal over the seller's quality is consistent with what the buyer's knowledge prior to obtaining the new partition. Equation (8) guarantees that for cell that was refined, the new marginals over asks aggregate to the coarser marginal that the buyer started with. Note these constraints still leave the buyer with substantial freedom in imagining what the marginal over the asks may be in the new (refined) partition. In Section 3.2 we illustrate how to operationalize these constraints.

[8] Our analysis depends on the information encoded in the buyer's type θ, i.e. the marginals on the quality and on the asks as specified by Eqs. (5) and (6). In principle, these could be generated by a different market setting than the double auction we described above. That is, our analysis would continue to hold for any joint distribution as long as it induces these marginals.

3.1 Two polar benchmarks

To illustrate the impact of the buyer's misspecification we begin by comparing the case of a correctly specified buyer with the case of a fully coarse one.

Rational expectations. Under rational expectations, the buyer has correct beliefs about the joint distribution of the object's quality and the seller's ask (i.e., he knows they are perfectly correlated, because it is a dominant strategy for the seller to submit an ask that is equal to the quality). Hence, the buyer's problem is given by:

$$\max_p F(p) \cdot \left(\mathbb{E}_{\phi \sim F}\left[v\left(\phi, b\right) \mid \phi < p\right] - p\right)$$

By our assumption in Eq. (2), the optimal solution is obtained at $p = 0$. Thus, there is no trade in equilibrium.

Full coarseness. Next, consider a buyer whose analogy partition consists of a single cell, i.e., $\theta = \left(C_1^\theta = [0,1]\right)$. This buyer only knows the marginal distribution over the seller's quality and the overall marginal distribution over the seller's ask. We refer to this buyer as being *fully coarse*. According to Eqs. (5) and (6), we have $H_\phi^\theta = H_x^{\theta,1} = F$.

The problem of a fully coarse buyer is given by:

$$\max_p F(p) \cdot \left(\mathbb{E}_{\phi \sim F}\left[v\left(\phi, b\right)\right] - p\right).$$

The optimal price satisfies:

$$\mathbb{E}_{\phi \sim F}\left[v\left(\phi, b\right)\right] = p + \frac{F(p)}{f(p)}. \tag{9}$$

There exists a unique price that solves this equation. We denote this solution by p_0.

3.2 The willingness to pay of a fully coarse buyer

Suppose that a fully coarse buyer has the opportunity to refine his data by paying a fee to add a cell to his partition. Specifically, the buyer can acquire the analogy partition ($C_1 = [0,t], C_2 = [t,1]$) for some $t \in (0,1)$. This enables the buyer to refine his data and learn the marginal distribution over asks when the seller's quality is in $[0,t]$, denoted H_x^1, and the marginal over seller's ask when the quality is in $[t,1]$, denoted H_x^2.

Before obtaining the new data, the buyer does not know what the marginal distributions H_x^1 and H_x^2 might be. However, by Eqs. (6) and (8), he knows that:

$$F(t) \cdot H_x^1(x) + (1 - F(t)) \cdot H_x^2(x) = F(x) \quad \forall x \tag{10}$$

Denote by $W(p, H_x^1, H_x^2)$ the buyer's (misspecified) expected payoff from a bid p when the marginals over the seller's ask are given by H_x^1 and H_x^2. By Eqs. (4)-(7) we can compute this expected payoff as follows:

$$
\begin{aligned}
W(p, H_x^1, H_x^2) &= \Pr(\textit{trade with sellers in } C_1 \mid p) \cdot \mathbb{E}_{\phi \sim H\phi}\left[v(\phi, b) - p \mid \phi \in C_1\right] \\
&\quad + \Pr(\textit{trade with sellers in } C_2 \mid p) \cdot \mathbb{E}_{\phi \sim H\phi}\left[v(\phi, b) - p \mid \phi \in C_2\right] \\
&= F(t) \cdot H_x^1(p) \cdot (V_1 - p) + (1 - F(t)) \cdot H_x^2(p) \cdot (V_2 - p), \qquad (11)
\end{aligned}
$$

where $V_k \equiv \mathbb{E}_{\phi \sim F}(v(\phi, b) \mid \phi \in C_k)$ denotes the expected value of $v\left(\phi, b\right)$, conditional on the seller's quality being in C_k. Note that the buyer's misspecification is captured by his use of the expected values V_1 and V_2, which do not condition on the event that trade takes place, thereby ignoring the fact that the object's quality and the seller's ask are dependent.

The buyer's WTP is derived by solving the following maximization problem:

$$
\max_{H_x^1, H_x^2} \left\{ \left(\max_p W(p, H_x^1, H_x^2) \right) - W(p_0, H_x^1, H_x^2) \right\} \qquad (12)
$$

subject to Eq. (10). The first component, $W(p, H_x^1, H_x^2)$, represents the expected payoff the buyer can obtain with the new data, where the maximum reflects the fact that the buyer can choose the optimal price based on this new data. The second component, $W(p_0, H_x^1, H_x^2)$, represents the expected payoff the buyer would receive if he does not refine his data (thus sticking with his original bid), with the expectation computed with respect to the refined data. Our first result characterizes the marginals, H_x^1 and H_x^2, and the new bid p_1 that these marginals induce:

Proposition 1. *Let p_1 be the price that solves the buyer's WTP given by Eq. (12), subject to the constraint in Eq. (10). This price satisfies* one *of the following conditions:*

i. The price p_1 satisfies $V_1 = p_1 + \frac{F(p_1)}{f(p_1)}$, provided that $F(p_0) - F(p_1) \leq F(t)$. Otherwise, it satisfies $F(p_0) - F(p_1) = F(t)$.

ii. The price p_1 satisfies $V_2 = p_1 + \frac{F(p_1)}{f(p_1)}$, provided that $F(p_1) - F(p_0) \leq 1 - F(t)$. Otherwise, it satisfies $F(p_1) - F(p_0) = 1 - F(t)$.

Furthermore, the marginal distributions H_x^1 and H_x^2 that solve the buyer's TWP given by Eq. (12), subject to the constraint in Eq. (10), satisfy the following condition: If the price p_1 is determined by (i) above, then $H_x^2(p_1) = H_x^2(p_0)$; if the price p_1 is determined by (ii) above, then $H_x^1(p_1) = H_x^1(p_0)$.

To interpret the above conditions recall that a buyer faces two potential sources of regret from not refining his data: The new data might have allowed the buyer to make a more precise

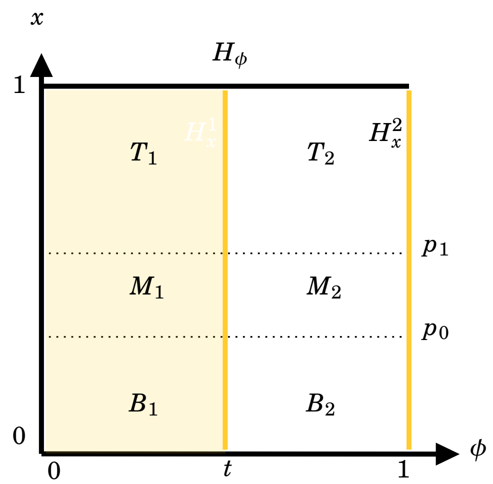

Figure 1: Feasible signals

bid and potentially secure better quality, or the new data might have led the buyer to realize that his original bid was too high relative to the quality. Proposition 1 states that the buyer experiences maximal regret if he focuses on only one of these sources and imagines the worse-case scenario (in terms of regret) for that source.

To see this, note that the new price p_1 can be either above or below the original price p_0. If $p_1 > p_0$, then p_1 will result in trade with higher quality sellers relative to p_0 if sellers who submit asks between the p_0 and p_1 are more likely to belong to the top cell in the partition, i.e. $[t,1]$. Thus, by not refining his data, the buyer would miss an opportunity to trade with "good" sellers and secure a high payoff. Indeed, $H_x^1(p_1) = H_x^1(p_0)$ implies that the regret-maximizing distribution assigns a probability of zero to the event that a seller with an object whose quality is in the interval $[0,t]$ submits an ask between p_0 and p_1. In other words, increasing the bid attracts only sellers of high quality.

On the other hand, if $p_1 < p_0$, then the original price p_0 was too high relative to the quality purchased. The loss from keeping the original high price is accentuated if the sellers who agree to trade at p_0 but not at p_1 are more likely to belong to the bottom cell of the partition, i.e. $[0,t]$. Indeed, $H_x^2(p_1) = H_x^2(p_0)$ implies that the regret-maximizing distribution assigns a probability of zero to the event that a seller with an object whose quality is in the interval $[t,1]$ submits an ask between p_1 and p_0.

The main idea behind the proof is to look for a *joint* distribution H over asks (x) and qualities (ϕ) that maximizes the objective function in Eq. (12) and satisfies the marginal constraint in Eq. (10). To give a rough outline of the proof, we focus on the case where the price p in the maximization problem is restricted to be greater than p_0. In this case, we can partition the space $[0,1] \times [0,1]$ of asks and qualities, over which the joint distribution H is defined, into six regions as depicted in Figure 1.

With this partition of the space, we can rewrite the objective function in Eq. (12) in terms of the probability mass that H assigns to the areas M_1 and M_2. The marginal constraints, (5) and (10), pin down the sum of probability masses along each row and column of regions in Figure 1. This allows us to show that, since $V_2 > V_1$, the distribution H that maximizes the buyer's WTP assigns as much probability as possible to M_2 "at the expense" of the probability mass assigned to M_1. Consequently, the objective function greately simplifies, and the solution follows from maximizing it.

The following example demonstrates how Proposition 1 can be operationalized:

Example 1: Suppose the seller's quality ϕ is uniformly distributed on $[0,1]$, i.e., $F(\phi)=\phi$, and $v(\phi,b)=\phi b$ for some $b\in(1,2)$. Solving Eq. (9) reveals that a coarse buyer sets the price $p_0=\frac{b}{4}$. Now, suppose the buyer can refine his data according to the partition $\left([0,\frac{1}{2}],[\frac{1}{2},1]\right)$. What is the buyer's WTP for this refined data?

Computation shows that $V_1=\frac{b}{4}$ and $V_2=\frac{3b}{4}$. Therefore, by Proposition 1, the bid p_1 that maximizes Eq. (12) subject to Eq. (10) satisfies either $V_2=2p_1$ or $V_1=2p_1$. To determine the buyer's WTP for the refined data, the two cases have to be compared.

If $V_2=2p_1$, then $p_1=\frac{3b}{8}$. Indeed, since $F\left(\frac{3b}{8}\right)-F\left(\frac{b}{4}\right)=\frac{b}{8}<\frac{1}{2}=1-F\left(\frac{1}{2}\right)$, the price $p_1=\frac{3b}{8}$ is a candidate for being a maximizer in the computation of the buyer's regret. From the proof of Proposition 1, we know that the buyer's regret from not acquiring the new data in this case is given by $F(p_1)(V_2-p_1)-F(p_0)(V_2-p_0)$, which equals $\frac{b^2}{64}$.

If $V_1=2p_1$, then $p_1=\frac{b}{8}$. Here too, because $F\left(\frac{b}{4}\right)-F\left(\frac{b}{8}\right)=\frac{b}{8}<\frac{1}{2}=1-F\left(\frac{1}{2}\right)$, it follows that $p_1=\frac{b}{8}$ is a candidate for being a maximizer in the computation of the buyer's regret. A similar computation shows that, in this case as well, the buyer's regret from not acquiring the new data is $\frac{b^2}{64}$. It is worth noting that the equality in the regret computed in both cases follows from the specific setup and does not hold in general.

To conclude, the buyer's WTP for refining his data and obtaining a new partition is $\frac{b^2}{64}$. □

Proposition 1 took the refined partition offered to the buyer as given. A natural question that arises is: what (two-cell) partition maximizes the buyer's WTP for data?

To answer this question, we introduce the following notation. Given an analogy partition $([0,t],[t,1])$ defined by a cutoff t, let $p_1^*(t)$ denote the price that solves the buyer's WTP to learn this analogy partition, as characterized by Proposition 1. We then have that:

Proposition 2. *The cutoff t^* that generates the partition $([0,t^*],[t^*,1])$ for which the buyer's WTP is maximal satisfies*

i. $F(p_0)-F(p_1^*(t^*))=F(t^*)$ *if $p_0>p_1^*$, and*

ii. $F(p_1^*(t^*))-F(p_0)=1-F(t^*)$ *if $p_0\le p_1^*$ otherwise.*

In terms of the explanation provided after Proposition 1 above, Proposition 2 implies that when $p_1^* > p_0$, the regret-maximizing distribution H assigns zero mass to regions M_1, T_2, and B_2 depicted in Figure 1. This suggests that the buyer believes that if he maintains the original bid p_0, he will receive only quality below t, whereas the new bid will yield quality above t. An analogous intuition applies for the case that $p_1^* < p_0$.

Thus, if the cost c of a new partition exceeds the WTP associated with the partition characterized by Proposition 2, then the buyer would choose to remain misspecified with his coarsest partition. The next example illustrates this for a simple specification.

Example 2. Consider the specification in Example 1. With these parameters, the partition that maximizes the buyer's WTP splits the interval $[0,1]$ at $t^* = \frac{4}{4+b}$.[9] Given our restriction on b, this implies that $t^* > \frac{2}{3}$.

To see how we arrived at this conclusion, focus on the case where the solution satisfies (ii) in Proposition 1. By the proof of Proposition 1, the buyer's WTP is given by $F(p_1)(V_2 - p_1) - F(p_0)(V_2 - p_0)$. Recall that, given the parameters of the example, $p_0 = \frac{b}{4}$. Additionally, note that $V_2 = \frac{(1+t)b}{2} = 2p_1$, so $p_1 = \frac{1}{4}(1+t)b$. Hence, by Proposition 2, the threshold t^* that maximizes the buyer's WTP satisfies

$$\frac{1}{4}(1+t^*)b - \frac{b}{4} = 1 - t^*$$

yielding $t^* = 4/(4+b)$ and a buyer's WTP of $b^2/(b+4)^2$ Indeed, since $b < 2$, this is higher than $b^2/64$, which is the buyer's WTP for the partition $([0,\frac{1}{2}],[\frac{1}{2},1])$, as shown in Example 1. The arguments for the case where the solution satisfies case (i) in Proposition 1 are analogous.

To conclude, with the parameters of Example 1, if the cost c of obtaining the new data exceeds $\frac{b^2}{(b+4)^2}$, the buyer would always prefer to remain fully coarse. □

We conclude this section by considering a fully coarse buyer who is offered two analogy partitions, where one is a refinement of the other. Clearly, the finer partition reveals more information about the "true" mapping from bids to consequences. Our next result verifies that the buyer's willingness to pay for the finer partition is higher.

Proposition 3. *Suppose a fully coarse buyer is offered two data sets that correspond to a pair of analogy partitions, where one is a refinement of the other. Then the finer partition is more valuable.*

The proof shows that the set of models that is associated with the finer partition is a superset of the set of models that is associated with the coarser one.

[9] An additional partition that yields the same WTP has a threshold at $t = \frac{b}{b+4}$. For brevity, we will focus on the first partition in this example.

4 Sampling

We next illustrate our notion of rational misspecification in a strategic environment where players' beliefs about the mapping from actions to consequences is based on a sampling procedure proposed by Osborne and Rubinstein (1998). The idea is that each player knows the possible payoffs he can obtain from each action but is unaware of the distribution of these payoffs, conditional on the action. A player may not even be aware that that his is engaged in a game. Thus, his belief about the distribution of payoffs conditional on one action, may be independent of his belief about the distribution conditional on another action.

To decide which action to choose, players apply the following procedure. For each action a, a player draws m independent samples of consequences (i.e., action profiles in which his action is a) from the steady-state distribution. The player then associates with each action the average payoff obtained in the sample, and chooses the action with the highest average payoff.

If in steady state each action is played with some probability, then the average payoff associated with each action is a random variable. Focusing on two-player symmetric games, a steady-state or a sampling equilibrium is a probability distribution over actions that satisfies the following (fixed point) property: the probability that action a is played is equal to the probability that this action is associated with the highest average payoff in the sample.

Consider an environment in which the first sample that each player draws is free, and suppose that prior to taking an action, each player has the option to pay c to obtain a second sample from each action. Would he be willing to do so? The challenge in answering this question lies in the player's lack of prior beliefs about the outcomes of a second sample, which prevents him from computing the expected benefit of the second sample. We address this challenge by using our framework from Section 2 to compute an *upper bound* on the player's willingness to pay for the second sample. We do this by calculating the player's maximal regret from *not* taking the second sample, and comparing this to the cost c.

The setting. Assume that each player has a finite set $A = (a_1, \ldots, a_K)$ of actions. The player obtains the payoff $u(a,b)$ when she chooses the action $a \in A$ and the other player chooses $b \in A$. To simplify the notation, we assume that every pair of actions generate a distinct payoff.[10] A strategy for a player is a distribution α over actions in A.

At the outset, each player samples one observation from his opponent's response to each action in A. Thus, a sample s can be described by an action profile $(r_1,...,r_K)$, with the interpretation that $r_k \in A$ is the response of the opponent when action a_k was sampled. Given s, let

$$(\overline{a}(s), \overline{r}(s)) := (a_k, r_k) : u(a_k, r_k) > u(a_{k'}, r_{k'}) \quad \forall k, k' \in \{1,...,K\} \tag{13}$$

[10] It is straightforward to relax this assumption, but this require cumbersome notation.

In words, $(\overline{a}(s), \overline{r}(s))$ is the sampled outcome with the highest payoff, which leads the player to choose $\overline{a}(s)$.

A single sample (or S(1)) equilibrium is then defined as follows:

Definition. *An $S(1)$ equilibrium is a distribution over actions $\alpha \in \Delta(A)$ such that for all $a_k \in A$ we have that*

$$\alpha(a_k) = \Pr(s \mid \overline{a}(s) = a_k),$$

where the probability on the right hand side is computed under the assumption that the opponent chooses his actions according to the distribution α.

To illustrate this equilibrium notion, consider two players who play the following symmetric game, which appears as Example 1 in Osborne and Rubinstein (1998) (we depict only the row player's payoff):

	b_1	b_2
a_1	2	4
a_2	3	1

Table 1. Row player's payoff in a symmetric game

In an $S1$ equilibrium, the probability α that a player chooses a_1 is equal to the probability that he either samples $((a_1, b_1), (a_2, b_2))$, or any sample with (a_1, b_2). These are the cases where the payoff from a_1 is greater than the payoff from a_2. Thus, in an S1 equilibrium, $\alpha = \alpha(1-\alpha) + (1-\alpha)$.[11]

Following Osborne and Rubinstein (1998), if a player takes a second sample, he will associate with each action a_k the average payoffs across the two samples.[12] In order to decide whether to pay for a second sample, a player forms a belief over the possible outcomes of the second sample (i.e., for each action a, what outcome (a, b) will realize). We assume that this belief must be consistent with the outcome of the first sample in the 'maximum likelihood' sense:

Assumption 1 (Maximum Likelihood Consistency). *Given an initial sample $s = (r_1, \ldots, r_K)$, a belief over the realization of the second sample is permissible, if for every action a_k, the probability the belief assigns to the outcome (a_k, r_k) is at least the probability it assigns to any other outcome (a_k, b).*

[11] The first term on the R.H.S. is the probability of drawing b_1 when a_1 is played and b_2 when a_2 is played, given that the opponent chooses his first and second action with probability α and $1-\alpha$, respectively. The second term is the probability that the opponent played b_2 when a_1 was played.

[12] For binary action games, Salant and Cherry (2020) allow for more general ways to summarize the two samples.

Put differently, a permissible belief is one where the first sample's outcome is the most probable. Given a permissible belief, a player associates with each action the expected average payoff across the two samples.

To compute the maximal regret from not taking a second sample, a player selects from all admissible beliefs about the outcome of the second sample, the one that maximizes the difference between the expected average payoff from some new action and the expected average payoff from the action he would take with only one sample. If the maximal regret, computed according to this belief, is less than the cost c, then purchasing the second sample is not considered beneficial, and so no player will purchase the second sample.

Mapping the setting to our framework. The "real" mapping g faced by a player in an $S1$ equilibrium is given by

$$g(a_k)[b_k] = \alpha[b_k]$$

where $g(a_k)[b_k]$ denotes the probability that the distribution $g(a_k)$ assigns to the outcome (a_k, b_k). A type is a sample $\theta = s = (r_1, ..., r_K)$. The misspecified mapping g_θ of type θ is:

$$g_\theta(a_k)[r_k] = 1$$

In words, when a player of type θ plays a_k his (misspecified) mapping assigns probability 1 to the outcome r_k. The set of *premissible mappings* G_θ for type $\theta = (r_1, ..., r_K)$ are all the mappings $\tilde{g}$ that satisfy:

$$\tilde{g}(a_k)[r_k] \geq \tilde{g}(a_k)[b] \text{ for all } a_k \in A \text{ and all } b \in A \tag{ML}$$

In words, a mapping is premissible if, following any action a_k, the probability it assigns to the outcome r_k (i.e., the outcome from the first sample) is the greatest.

4.1 Binary action games

Following Salant and Cherry (2020) we examine two-player symmetric binary action games. We begin with the following observation:

Claim 1. *In the example depicted in Table 1, no player is willing to pay for a second sample.*

To see why, consider first a type $\theta = (b_1, b_2)$. This is a type who would choose a_1 if he does not take a second sample. A second sample is worthwhile only if it leads to a different action. To change his action, the new sample will have to be (b_1, b_1). What is the highest likelihood for this sample given Assumption 1? This assumption implies that in the second sample, the likelihood of drawing b_1 when sampling a_1 is *at least* 0.5, and the likelihood of drawing b_1 when sampling a_2 is *at most* 0.5. It follows that the permissible belief that maximizes the likelihood of observing $s_2 = (b_1, b_1)$ in the second sample assigns probability 1 to observing b_1 when sampling a_1 and probability 0.5 of observing b_1 when sampling a_2. Thus, type θ's highest permissible expected average payoff from a_2 across the two samples is equal to

$$\frac{1+(0.5\cdot 1+0.5\cdot 3)}{2}=1.5$$

The expression on the L.H.S. is the average of the payoff from a_2 in the initial sample (which was equal to 1) and the expected payoff from the second sample, where the expectation according to the belief described above. Since the lowest permissible expected payoff from a_1 is 2, it follows that the player expects to continue choosing a_1 even after taking a second sample. Therefore, this type will not be willing to pay for another sample.

Consider next a type $\theta = (b_1, b_1)$, who would choose a_2 if he does not acquire a second sample. To maximize the chances that a second sample will lead to a change of action, the belief should put as much weight as possible on observing b_2 when sampling each of the actions. By the Assupmtion 1, and because the outcome of the first sample was b_1 for each action, the maximal probability that can be put on observing b_2 for each action is 0.5. It follows that the maximal permissible expected average payoff from a_1 is

$$\frac{2+(0.5\cdot 4+0.5\cdot 2)}{2}=2.5,$$

while the lowest permissible expected average payoff from a_2 is

$$\frac{3+(0.5\cdot 3+0.5\cdot 1)}{2}=2.5.$$

Since at the regret maximizing belief it is still optimal to continue to choose a_2, this type will not be willing to pay anything for a second sample.

Similar computations show that in this example, the regret-maximizing permissible beliefs for the remaining types cannot change their actions. Therefore, these types will also decline a second sample for any positive fee.

The reasoning illustrated above can be applied to a broader class of binary action games. In particular, we will focus next on the class that exhibits the following payoff structure:

	b_1	b_2
a_1	$\overline{x}$	$\underline{x}$
a_2	$\overline{y}$	$\underline{y}$

Table 2. A class of symmetric binary action games

where $\overline{x} > \underline{x}$ and $\overline{y} > \underline{y}$. Without loss of generality, let $\overline{x} > \overline{y}$. For non-triviality, we assume no action is dominated, and thus $\underline{y} > \underline{x}$.[13] This payoff structure is characterized by the highest payoff from each action being attained for a fixed action of the opponent (note that the game depicted in Table 2 does not belong to this class). Our next result provides the necessary and sufficient conditions under which a player would not be willing to pay for a second sample.

Proposition 4. *In the class of games characterized by Table 2, no player would pay any positive amount for a second sample if and only if*

$$\overline{x} - \overline{y} \geq \frac{\overline{y} - \underline{x}}{3} \text{ and } \underline{y} - \underline{x} \geq \frac{\overline{x} - \underline{y}}{3}. \tag{14}$$

The proof, which appears in the Appendix, computes the maximal regret for all possible types and shows that the conditions in Eq. (14) are sufficient for the maximal regret to be non-positive for all types and necessary for two of the types. The proof leverages Assumption 1 which implies that, in binary action games, the permissible beliefs that maximize regret either put probability one on one of the outcomes or equal probabilities on each of the outcomes.

The remaining class of games can be analyzed using similar arguments to those used in proof of Proposition 4. We omit the characterization of the necessary and sufficient conditions ensuring that no player would pay any positive amount for a second sample because they are less transparent.

4.2 Multiple action games

We next extend the analysis to symmetric games with multiple actions for which we derive necessary and sufficient condition for declining a second sample for any positive fee. This condition guarantees that the maximal regret from not taking a second sample is zero. An equivalent condition states that the worst permissible expected average payoff from the original action (the one that would be chosen with the single initial sample) is at least as large as the best permissible expected average payoff from taking a different action. When this holds, a player realizes that no new data will induce him to change his action, hence, he will not be willing to pay for it.

Proposition 5. *Given an initial sample $s = (r_1, \ldots, r_K)$, let $\overline{a} = \overline{a}(s)$ and $\overline{r} = \overline{r}(s)$, where $\overline{a}(\cdot)$ and $\overline{r}(\cdot)$ are defined in Eq. (13). A player does not want to pay for a second sample for any positive cost if and only if*

$$\frac{1}{2}\left[u(\overline{a},\overline{r}) + \min_{R\subseteq A} \sum_{r\in R\cup\{\overline{r}\}} \frac{1}{|R|+1} u(\overline{a},r)\right] \geq \max_{a_k\neq\overline{a}} \frac{1}{2}\left[u(a_k,r_k) + \max_{R\subseteq A} \sum_{r\in R\cup\{r_k\}} \frac{1}{|R|+1} u(a_k,r)\right] \tag{15}$$

[13] Our assumptions on payoffs imply that the payoffs along the diagonal can be ranked as $\overline{x} > \underline{y}$ and $\overline{y} > \underline{x}$.

The left-hand side of Eq. (15) represents the *minimal* average payoff across the two samples that a player expects to get if he continues to choose the action that he would have chosen with the initial sample ($\overline{a}$), where the expectation is taken with respect to a permissible belief. Specifically, the first term in the brackets is the payoff observed in the initial sample from $\overline{a}$. The second term represents the minimal expected payoff from $\overline{a}$ in the second sample. The key step in the proof is to show that this expectation is minimized for a permissible belief which is uniform over some subset of responses (that includes the initially sampled response, $\overline{r}$). Indeed, in the binary action case, the distributions were either degenerate or $50-50$.

The right-hand side of Eq. (15) represents the *maximal* average payoff across the two samples that a player expected to get if he were to change his action, where the expectation is taken with respect to a permissible belief. To derive this payoff, we first need to choose an action that is different from $\overline{a}$. This is captured by the external maximum operator outside the brackets. The first term in the brackets is the payoff from the new action a_k that was observed in the initial sample. The second term represents the maximal expected payoff from a_k in the second sample. Here again, this expectation is maximized for a permissible belief which is uniform over some subset of responses (that includes the initially sampled response, r_k).

We conclude this section by showing that the WTP increases with the number of samples that are offered to a player with one initial sample. While intuitive, this is not an immediately obvious result since sampling more times changes the set of permissible models that a player can imagine. For example, it is not apriori clear how the set of models that are permissible for three additional samples relate to those that are permissible with only two additional samples. We view the next result as evidence supporting the use of our notion of WTP for models.

Proposition 6. *Consider a player who sampled once. Then sampling M more times is more valuable than sampling $M-1$ more times.*

The proof demonstrates that the set of permissible models that correspond to taking M additional samples subsumes the set of feasible models that correspond to $M-1$ additional samples. To show this, we consider a feasible model that corresponds to $M-1$ additional samples along with its implied distribution over the opponent's actions. We then adjust this distribution by shifting probability mass towards the opponent's action that was observed in the initial sample. This only relaxes the ML constraints and yields a permissible model for M additional samples that replicates the same mapping from actions to expected average payoffs as the model for $M-1$ additional samples.

5 Causal misperceptions

In this section we apply our approach to misspecifications that arise from causal misperceptions. These arise when individuals misinterpret spurious correlations between variables as causal relations. To model such misperceptions, we adopt the framework introduced by Spiegler (2016). This framework borrows tools from the Bayesian networks literature to represent individuals' subjective causal perceptions by directed acyclic graphs (DAGs). We consider agents with a misspecified causal model who can pay some price to learn the true model. We apply our approach to derive these agents' WTP for this knowledge. We then use this setting to illustrate that a market for information may exclude the misspecified agents who can gain from learning the true model.

The setting. We demonstrate our approach in the context of the "dieter's dilemma" analyzed in Spiegler (2016). This is a convenient setting for introducing the Bayesian network framework of causal misperceptions (for the general framework, see Spiegler (2016) and Spiegler (2020)). There is a unit measure of agents, who each needs to decide whether to take a dietary supplement. We let $a = 1$ denote the action of taking the supplement and $a = 0$ the action of not taking it. An agent cares about his health h, which is either good ($h = 1$) or bad ($h = 0$). His health h and supplement consumption a are potentially correlated with the level of some chemical c in his blood, which can be either normal ($c = 0$) or not ($c = 1$). An agent's payoff is equal to $h - ka$, where k is some positive constant.

The relations between these three variables are governed by a given data generating process, which can be represented by a long-run joint distribution p over (a, c, h).[14] Let $p(a)$ and $p(h)$ denote the marginal distributions over the action and health outcome. Let $p(c \mid a, h)$ be the conditional probability of c, given a and h. The true objective joint distribution over (a, c, h) can be factorized as follows:

$$p(a, c, h) = p(a)p(h)p(c \mid a, h) \tag{16}$$

where $p(a)$ is determined endogenously by the agents' behavior according to the equilibrium notion defined below. We assume that the true distribution p satisfies the following: for all $a, h \in \{0, 1\}$,

$$p(h \mid a) = p(h) = \frac{1}{2}$$

and

$$p(c = 1 \mid a, h) = (1 - a)(1 - h).$$

[14] One can think of this steady-state distribution as a giant excel sheet with infinite rows and three columns, one for each variable. Each entry in this sheet is a particular realization of the three variables, and the empirical frequencies are interpreted as the long-run probabilities.

That is, in reality an agent's health state is independent of supplementation and his blood chemical level is abnormal if and only if he's unhealthy and did not consume the supplement. Thus, given p, the rational decision is $a = 0$.

The factorization in Eq. (16) can be depicted by a DAG where each node represents a variable, and where a link from one variable to another means that the former causes the latter. That is, the DAG Γ associated with p is

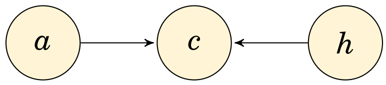

We refer to Γ as the true DAG and to p as the true data generating process. Thus, according to p, the action a does not affect h, which is determined exogenously by nature. In addition, the chemical level c is just a symptom that is jointly affected by the action and the person's health.

Given some DAG Γ', any joint distribution p' on (a,c,h) that is consistent with Γ' can be factorized as follows:

$$p'(a,c,h) = p'(a)p'(c \mid \Gamma'(c))p'(h \mid \Gamma'(h))$$

where $\Gamma(c)$ and $\Gamma(h)$ are the variables that cause c and h, respectively (i.e., those variables with a link going into c and into h). If one of these variables, say $x \in \{c,h\}$, is exogenous (i.e., there are no links going into x), then $\Gamma(x) = \emptyset$ and $p'(x \mid \Gamma'(x)) = p'(x)$.

Suppose the agents have a misspecified belief about the data generating process. Specifically, they believe that the joint distribution over (a,c,h) is consistent with the DAG Γ' given by:

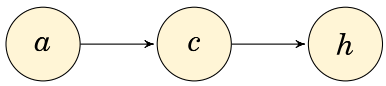

I.e., the agents falsely believe that supplementation causes the health outcome via its effect on the chemical level. Thus, they believe that the joint distribution p' that is consistent with Γ' is factorized as follows:

$$p'(a,c,h) = p(a)p(c \mid a)p(h \mid c).$$

In deciding which action to take, the agents do not have access to the true $p(h \mid a)$, but rather, they must derive this quantity using their subjective model. That is, agents have access only to the quantities in their factorization formula, $p(a)$, $p(c \mid a)$ and $p(h \mid c)$. They then derive $p'(h = 1 \mid a)$ by calculating

$$p(c = 0 \mid a)p(h = 1 \mid c = 0) + p(c = 1 \mid a)p(h = 1 \mid c = 1).$$

Note that the misspecification lies in that the agent uses $p(h = 1 \mid c = 0)$ (which is an expectation over all values of a) instead of $p(h = 1 \mid c = 0, a)$. Given our assumptions on p, these calculations yield

$$p'(h = 1 \mid a) = \frac{1}{(2-\alpha)(1+\alpha)}$$

where α is the (endogenous) long-run/steady-state probability of choosing $a = 1$ (or equivalently, the steady-state proportion of agents who choose $a = 1$).

Mapping the setting to our framework. To describe this setting in the language of our framework, let $A = \{0,1\}$ and $Y = \{0,1\}$ (where $y = h$ here). The true mapping from actions to consequences $g(a)$ is a uniform distribution over Y for all $a \in A$. Thus, the optimal action under rational expectations is $a = 0$. The set of types Θ is given by the set of possible DAGs over the variables a, c and y such that a is an ancestral node. For $\theta = \Gamma'$, the agents' misspecified mapping g_θ is given by

$$g_\theta(a)[1] = \begin{cases} \frac{1}{2(1+\alpha)} & if \quad a = 0 \\ \frac{1}{1+\alpha} & if \quad a = 1. \end{cases}$$

Note that the steady-state frequency of taking the supplement affects the agents' misspecified belief about the supplement's effect on health, which in turn affects the steady-state frequency of taking the supplement. Such feedback effect is common in models of misspecified beliefs.

Personal equilibrium. Because the agents strategy affects their mapping from actions to consequences, their decisions are determined as a fixed point rather than as a solution to a maximization problem. Spiegler (2016) refers to this fixed point as a *personal equilibrium*. To present this definition, denote by $\mathbb{E}_{\tilde{g}} u(a, y)$ the expectation of $u(a, y)$ with respect to some mapping $\tilde{g} : A \to \Delta(Y)$.

Definition. *Given a (possibly misspecified) mapping $\tilde{g} : A \to \Delta(Y)$, a probability $\alpha \in (0,1)$ is an ε-**personal equilibrium**, if whenever $\alpha > \varepsilon$ (respectively, $1-\alpha > \varepsilon$) then $\mathbb{E}_{\tilde{g}} u(1, y) \geq \mathbb{E}_{\tilde{g}} u(0, y)$ (respectively, $\mathbb{E}_{\tilde{g}} u(1, y) \leq \mathbb{E}_{\tilde{g}} u(0, y)$). A probability $\alpha \in [0,1]$ is a **personal equilibrium** if it is the limit of ε-personal equilibria as $\varepsilon \to 0$.*

Equipped with this definition, the following can be shown:

Proposition 7. *[Spiegler (2016)] Suppose the agents' misspecified mapping is g_θ. If $k \in (\frac{1}{4}, \frac{1}{2})$ there is a personal equilibrium in which the proportion of agents who take the supplement is given by*

$$\alpha = \frac{1-2k}{2k}. \tag{17}$$

The willingness to pay for the true model. Suppose the agents could obtain access to the true $p(h \mid a)$ for a price. What is their WTP for this information in the above personal equilibrium? According to our approach, this amount is given by the maximal expected regret from *not* obtaining this information. To derive this, we consider the joint distribution q over (a,c,h) that maximizes the difference between the expected payoff from the action that will be taken under q, and the expected payoff from the agent's equilibrium action, where: (i) this expectation is computed according to q, and (ii) the joint distribution q must be consistent with the information that the agents use with their misspecified model:

$$
\begin{aligned}
&q(h) = \frac{1}{2} \\
&q(c=0 \mid a=1) = 1 \quad , \quad q(c=0 \mid a=0) = \frac{1}{2} \\
&q(h=1 \mid c=1) = 0 \quad , \quad q(h=1 \mid c=0) = 2k.
\end{aligned}
$$

These constraints leave only one degree of freedom in the specification of q. Denoting $\beta := q(a=1, c=0, h=0)$, the values of $q(a,c,h)$ are determined by β and α (where α is given by Eq. (17)) as follows:

	$a=0$		$a=1$	
	$c=0$	$c=1$	$c=0$	$c=1$
$h=0$	$\frac{\alpha}{2}-\beta$	$\frac{1-\alpha}{2}$	β	0
$h=1$	$\frac{1}{2}-\alpha+\beta$	0	$\alpha-\beta$	0

Table 3. the constrained joint distribution q

We denote by Q the set of probability distributions that are consistent with Table 3.

In contrast to the applications presented in Sections 3 and 4, here, the agent's regret depends not only on his type, but also on his equilibrium action. Let $R_a(q)$ denote the expected regret of agents who choose a and believe that q is the joint distribution over (a,c,h). Let $q_a^* \in \arg\max_{q \in Q} R_a(q)$ and $R_a^* = R_a(q_a^*)$.[15] We denote by β_a^* the value of β that corresponds to q_a^*. Our next result characterizes these values:

Proposition 8. *The regret-maximizing distribution* q_a^* *satisfies* $\beta_0^* = \max\{0, \frac{1-3k}{2k}\}$ *and* $\beta_1^* = \frac{1-2k}{4k}$.

To see the intuition for this result, consider the agents who choose $a=0$ in equilibrium. The distribution q that maximizes the regret of these agents assigns a high probability to

[15] Since $\beta \in [0,1]$, this guarantees that a maximizer q_a^* exists.

the outcome $h = 0$ when $a = 0$ and a high probability to $h = 1$ when $a = 1$. Both objectives are attained when β is minimal, subject to the constraint that all the entries in Table 3 are are non-negative. We show that the effective constraint are $\frac{1}{2} - \alpha + \beta \geq 0$. This yields that $\beta_0^* = \max\{0, \alpha - \frac{1}{2}\}$, which translates to the condition in the proposition after substituting for α. Similarly, for agents who choose $a = 1$ in equilibrium, the distribution q that maximizes the regret of these agents assigns a high probability to the outcome $h = 1$ when $a = 0$ and a high probability to $h = 0$ when $a = 1$. Both objectives are attained when β is maximal, provided that all the entries in Table 3 are non-negative. The effective constraint in this case is $\frac{\alpha}{2} - \beta \geq 0$. Hence, $\beta_1^* = \frac{\alpha}{2}$, which translates to the condition in the proposition

Proposition 8 implies a threshold price for learning the true causal model, above which, all agents will "rationally" decide to remain misspecified. This price is computed using the value of R_a^* (which is derived in the proof of Proposition 8. This threshold is described in the following corollary:

Corollary 1. *No agent will pay to learn the true model if the price of learning is higher than*

$$\begin{array}{ccc} \frac{2k^2}{1-2k} & if & \frac{1}{4} < k \leq \frac{1}{3}, \\ \frac{2k(1-2k)}{4k-1} & if & \frac{1}{3} < k \leq \frac{3}{8}, \\ k & if & \frac{3}{8} < k < \frac{1}{2}. \end{array}$$

Corollary 1 implies the following observation. Suppose there are consultants who can reveal the true model to the agents. However, the capacity of the consultants is limited, so that they can serve at most a mass of $\lambda < 1 - \alpha$ of agents. In this market for consultants, the equilibrium prices will be determined by the WTP of the marginal agent (as defined above). Will this market for consultants eliminate the misspecification and lead agents not to consume the supplement? The next observation shows that the answer is negative when the supplement's cost is below some threshold.

Corollary 2. *If $k \in (\frac{1}{4}, \frac{3}{8})$, then in the equilibrium of the market for consultants, only the agents who choose the rational action (not taking the supplement) will pay for consultants.*

This follows directly from the proof of Proposition 8 and Corollary 1. The WTP of agents who take the supplement is equal to k, and this is lower than the WTP of agents who do not take the supplement when $k < \frac{3}{8}$. Thus, there is a range of parameter values for which the only agents who buy information are those who cannot benefit from it. Note that when the agents who originally chose the rational action pay to learn $p(h \mid a)$, they would realize they were choosing the right action so that the distribution over actions will remain unchanged.

Our next result characterizes the causal models that are consistent with the joint distributions that maximize the agents' regret.

Proposition 9. (i) *If* $k \geq \frac{1}{3}$, *then* q_0^* *is consistent with a causal model represented by the DAG*

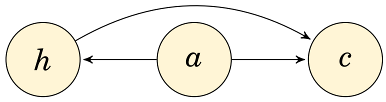

(ii) *if* $k < \frac{1}{3}$, *then* q_0^* *is consistent with a causal model represented by the DAG*

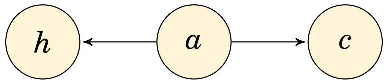

(iii) q_1^* *is consistent with the true casual model represented by the DAG*

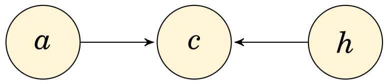

Thus, for an agent who chooses the irrational action $a = 1$, the WTP to learn the true model is obtained for a belief that this model is the one given by the objectively true DAG (i.e., the one where the action has no effect on the outcome and both affect the chemical level). Intuitively, such a model would make the agent realize that choosing $a = 1$ is a mistake.

On the other hand, an agent who is actually choosing the rational action (but for the wrong reason) is willing to pay the most to learn the true model, when he believes that this model would make him realize that the action affects both the outcome and the chemical level (and for high k, that the health outcome also affects the chemical level). Intuitively, such a model would make him realize that he should take the supplement, and hence, he would forego a high payoff if he didn't learn the truth.

6 Conclusion

This paper proposes a theoretical framework for studying how players with misspecified beliefs may confront opportunities for attenuating their misspecifications. Our approach circumvents the need to impose arbitrary prior beliefs on the set of all possible models. Instead, we derive an upper bound on the players willingness-to-pay for refining their model, where the bound is taken over all *feasible* models that the player may envisage given his current knowledge. The player is rationally misspecified when the upper bound exceeds the cost of acquiring new data. The three applications we study demonstrate how our framework can be accommodated by very different approaches to modeling belief misspecification.

7 Proofs

Proof of Proposition 1

Instead of maximizing Eq. (12) over the domain of marginals H_x^1 and H_x^2 that satisfy the constraint in Eq. (10), we solve the equivalent problem of maximizing the objective in Eq. (12) over the domain of joint distributions over $[0,1]\times[0,1]$, whose marginals satisfy Eqs. (5) and (10). A solution to this problem is a pair, $\left(p_1^*, H^*\right)$, of a price and a joint distribution over the quality (ϕ) and seller's ask (x).

We start by considering the case where $p_1^* \geq p_0$ (recall that p_0 is the price that solves Eq. (9)). To solve this problem, it is useful to partition the space $[0,1]\times[0,1]$ into the following six subsets, as depicted in Figure 1:

$$B_i \equiv \left\{(\phi, x) \mid \phi \in C_i \wedge x \in (0, p_0)\right\}$$
$$M_i \equiv \left\{(\phi, x) \mid \phi \in C_i \wedge x \in (p_0, p)\right\}$$
$$T_i \equiv \left\{(\phi, x) \mid \phi \in C_i \wedge x \in (p, 1)\right\}$$

For any joint distirubtion H over $[0,1]\times[0,1]$, and for each $i = 1,2$, denote by $\mu_H(B_i), \mu_H(M_i)$, and $\mu_H(T_i)$ the probability mass of the sets B_i, M_i, and T_i, respectively, according to H. Hence, the conditional marginal distributions induced by H can be computed as follows:

$$H_x^i(p_0) = \frac{\mu_H(B_i)}{\mu_H(B_i) + \mu_H(M_i) + \mu_H(T_i)} \text{ and } H_x^i(p) = \frac{\mu_H(B_i) + \mu_H(M_i)}{\mu_H(B_i) + \mu_H(M_i) + \mu_H(T_i)}.$$

Moreover, in terms of these probability masses, the marginal constraint (5) implies that:

$$\mu_H(B_1) + \mu_H(M_1) + \mu_H(T_1) = F(t) \quad \text{and} \quad \mu_H(B_2) + \mu_H(M_2) + \mu_H(T_2) = 1 - F(t), \tag{18}$$

whereas the marginal constraint (10) implies that:

$$\mu_H(B_1) + \mu_H(B_2) = F(p_0) \quad \text{and} \quad \mu_H(M_1) + \mu_H(M_2) = F(p) - F(p_0). \tag{19}$$

Let $\mathscr{H}$ denote the set of joint distributions over qualities (ϕ) and asks (x), whose marginals satisfy Eqs. (18) and (19). We can then rewrite the maximization problem in Eq. (12), restricted to the case that $p_1^* \geq p_0$, as follows:

$$\max_{p \geq p_0} \left[\max_{H \in \mathscr{H}} \left(V_1 \mu_H(M_1) + V_2 \mu_H(M_2)\right) - pF(p)\right] + F(p_0)p_0. \tag{20}$$

Recall that, by definition, $V_2 > V_1$. Moreover, given any price $p \geq p_0$, the constraint in Eq.

(19) implies that the sum $\mu_H(M_1)+\mu_H(M_2)$ is fixed. Hence, the distribution H^* assigns as much probability as possible to $\mu_H(M_2)$ "at the expense" of the probability mass assigned to $\mu_H(M_1)$.

Consequently, the optimal solution $\left(p_1^*, H^*\right)$ cannot satisfy $F\left(p_1^*\right)-F(p_0)>1-F(t)$. To see this, suppose that $F\left(p_1^*\right)-F(p_0)>1-F(t)$, and note that this implies $p_1^*>p_0$. Moreover, the constraints in Eqs. (18) and (19) imply that the highest probability mass that H^* can assign to the region M_2 cannot exceed $1-F(t)$, and therefore $\mu_{H^*}(M_2)=1-F(t)$ and $\mu_{H^*}(M_1)=(F(p_1^*)-F(p_0))-(1-F(t))$.[16] Plugging these into Eq. (20) we obtain that $p^*=\arg\max_{p\geq p_0} F(p)(V_1-p)$. However, our assumption that $p+F(p)/f(p)$ is strictly increaseing in p implies that the derivative $\frac{d}{dp}[F(p)(V_1-p)]=f(p)\left[V_1-\left(p+\frac{F(p)}{f(p)}\right)\right]$ is negative for all $p\geq p_0$. This is a contradiction for the optimal price p^* being strictly above p_0.

Therefore, the solution (p_1^*, H^*) must satisfy $F(p_1^*)-F(p_0)\leq 1-F(t)$. The fact that $V_2>V_1$ and the constraints (18) and (19) then imply that $\mu_{H^*}(M_1)=0$ and $\mu_{H^*}(M_2)=F(p^*)-F(p_0)$.[17] Substituting into Eq. (20), we obtain:

$$p_1^* \in \arg\max_{p\geq p_0} F(p)(V_2-p)$$
$$\text{subject to } F(p)-F(p_0)\leq 1-F(t)$$

Our assumption that $p+\frac{F(p)}{f(p)}$ is strictly increasing in p implies that the function $F(p)(V_2-p)$ has a unique extremum on $[p_0,1]$, which occurs at $\tilde{p}$ that satisfies $V_2=\tilde{p}+\frac{F(\tilde{p})}{f(\tilde{p})}$. Furthermore, this extremum is a maximum, and $\tilde{p}\geq p_0$. Therefore, $p_1^*=\tilde{p}$, provided that $F(\tilde{p})-F(p_0)\leq 1-F(t)$. Otherwise, p_1^* is equal to the price p that solves $F(p)-F(p_0)=1-F(t)$.

The proof in the case of $p_1^*<p_0$ is analogous and is therefore omitted.

Proof of Proposition 2

Suppose that the partition $((0,t),(t,1))$ maximizes the buyer's WTP. Let (p_1^*, H^*) be the solution to the maximization problem in Eq. (12), subject to the constraint in Eq. (10). Suppose, by contradiction, that $F(p_1^*)-F(p_0)<1-F(t)$. By Proposition 1, this implies that $V_2=p_1^*+\frac{F(p_1^*)}{f(p_1^*)}$, and therefore $p_1^*>p_0$.

Because H^* is determined optimally given p_1^*, we know from the proof of Proposition 1 that the buyer's maximal WTP, as presented in Eq. (20), can be written as follows:

$$V_2\left(F(p_1^*)-F(p_0)\right)-p_1^*F(p_1^*)+F(p_0)p_0. \tag{21}$$

[16] Additionally, $\mu_{H^*}(T_2)=\mu_{H^*}(B_2)=0$, $\mu_{H^*}(T_1)=1-F(p_1^*)$ and $\mu_{H^*}(B_1)=F(p_0^*)$

[17] Additionally, $\mu_{H^*}(T_2)$ and $\mu_{H^*}(B_2)$ satisfy $0\leq\mu_{H^*}(T_2)\leq 1-F(p^*)$, $0<\mu_{H^*}(B_2)<F(p_0)$, and $\mu_{H^*}(T_2)+\mu_{H^*}(B_2)=(1-F(t))-(F(p^*)-F(p_0))$.

Recall that $V_2 \equiv \mathbb{E}_{\phi\sim F}\left(v(\phi,b) \mid \phi \in C_2\right)$, where $C_2 = (t,1)$ is the second element in the analogy partition. Therefore, by slightly increasing the boundary between the partition elements from t to $t^+ > t$, so that $C_2 = (t^+,1)$, we increase the value of V_2. Holding (p_1^*, H^*) fixed, this change only increases the expression in Eq. (21), while the inequality $F(p_1^*) - F(p_0) < 1 - F(t^+)$ still holds. Note, however, that this new value of Eq. (21), which is computed when (p_1^*, H^*) is held fixed, is only a lower bound to the buyer's willingness to pay for the new analogy partition, which is computed by solving the maximization problem in Eq. (12), subject to the constraint in Eq. (10), for the analogy partition $\left((0,t^+),(t^+,1)\right)$. This contradicts $((0,t),(t,1))$ being the partition that maximizes the buyer's willingness to pay for knowledge. The proof of the second case is analogous and is omitted. ■

Proof of Proposition 3

Denote the set of potential models that are associated with the coarser and finer partitions by G_θ and G'_θ, respectively.

A buyer's model g is induced by a partition of $[0,1]$ into cells, and a belief on how the seller picks an ask in each cell. Because the buyer is fully coarse, he knows only the marginal over asks on the entire interval of seller types. This restricts the distribution over asks, conditional on the seller type being in a given cell (specifically, the unconditional distribution over asks has to be equal to the marginal distribution over asks that is known to the fully coarse buyer).

Fix a model g in G_θ. This model corresponds to a collection of conditional distributions over asks $\left(H_x^1,\dots,H_x^K\right)$, one for each cell of the coarser partition. This model is also an element in G'_θ. To see this, consider some cell C_k in the coarser partition, which is a union of a set of cells $C'_{k1},\dots,C'_{kj}$, for some $j>1$ in the finer partition. Assign for each seller type in $C'_{k1},\dots,C'_{kj}$ the conditional distribution over asks H_x^k. Clearly, the unconditional distribution over the finer partition is identical to the unconditional distribution over the coarser partition. This collection of conditional distributions of asks over the finer partition induces the same mapping from actions to consequences as g. It follows that $G_\theta \subset G'_\theta$. Hence, $R^*(G'_\theta) > R^*(G_\theta)$. ■

Proof of Proposition 4

We show that if condition (14) holds, then the regret-maximizing permissible belief of each player type θ will lead him to stick with his one-sample action. Conversely, if this condition is violated, then at least one type would be willing to pay some positive amount for a second sample.

Consider first type $\theta = (b_1, b_1)$. This type would choose a_1 if he does not acquire a second sample. The permissible belief that maximizes the difference between the expected average payoff from a_2 and the expected average payoff from a_1 has the following property: It puts probability one on observing b_1 when sampling a_2 again, and probability 0.5 on observing b_2 when sampling a_1 again. This means that the maximal regret is

$$\overline{y} - \left[(0.5)(\overline{x}) + (0.5)(\frac{\overline{x} + \underline{x}}{2}) \right]$$

Thus, the condition $\overline{x} - \overline{y} \geq \frac{\overline{y} - \underline{x}}{3}$ is necessary and sufficient for the maximal regret to be non-positive making this type unwilling to pay any positive amount for a second sample.

Consider next the type $\theta = (b_1, b_2)$. This type would also choose a_1 if he does not acquire a second sample. The permissible belief that maximizes the difference between the expected average payoff from a_2 and the expected average payoff from a_1 has the following property: It puts probability 0.5 on observing b_1 when sampling a_2 again, and probability 0.5 on observing b_2 when sampling a_1 again. This means that the maximal regret is

$$\left[(0.5)(\underline{y}) + (0.5)(\frac{\overline{y} + \underline{y}}{2}) \right] - \left[(0.5)(\overline{x}) + (0.5)(\frac{\overline{x} + \underline{x}}{2}) \right]$$

This expression is non-positive if $\overline{x} - \underline{y} \geq \frac{\overline{y} - \underline{x}}{3}$, which is implied by the condition $\overline{x} - \overline{y} \geq \frac{\overline{y} - \underline{x}}{3}$.

Now consider the type $\theta = (b_2, b_1)$. This type would choose a_2 if he does not acquire a second sample. The permissible belief that maximizes the difference between the expected average payoff from a_1 and the expected average payoff from a_2 has the following property: It puts probability 0.5 on observing b_1 when sampling a_2 again, and probability 0.5 on observing b_2 when sampling a_1 again. This means that the maximal regret is

$$\left[(0.5)(\underline{x}) + (0.5)(\frac{\overline{x} + \underline{x}}{2}) \right] - \left[(0.5)(\overline{y}) + (0.5)(\frac{\overline{y} + \underline{y}}{2}) \right]$$

This expression is non-positive if $\overline{y} - \underline{x} \geq \frac{\overline{x} - \underline{y}}{3}$, which is implied by the condition $\underline{y} - \underline{x} \geq \frac{\overline{x} - \underline{y}}{3}$.

Finally, consider type $\theta = (b_2, b_2)$. This type would choose a_2 if he does not acquire a second sample. The permissible belief that maximizes the difference between the expected average payoff from a_1 and the expected average payoff from a_2 has the following property: It puts probability 1 on observing b_2 when sampling a_2 again, and probability 0.5 on observing b_2 when sampling a_1 again. This means that the maximal regret is

$$\left[(0.5)(\underline{x}) + (0.5)(\frac{\overline{x} + \underline{x}}{2}) \right] - \underline{y}$$

This expression is non-positive if and only if condition $\underline{y} - \underline{x} \geq \frac{\overline{x} - \underline{y}}{3}$ holds. ■

Proof of Proposition 5

A player does not want to pay for a second sample for any positive cost if and only if his maximal regret is zero. To find the necessary and sufficient condition for this, we first derive the minimal permissible expected average payoff from $\overline{a}$, the original action that would be chosen with the initial sample. The key step is to show that this minimal payoff is obtained with a uniform distribution over some subset of the other player's actions.

To show this, let $\tilde{g}(\overline{a})$ denote the permissible belief over the other player's action that minimizes the expected average payoff from $\overline{a}$, and let $B \subseteq A$ denote its support. Clearly, $u(\overline{a}, b) \leq u(\overline{a}, \overline{b})$ for every $b \in B$. Without loss of generality suppose $B = \{b_1, ..., b_K\}$ and $u(\overline{a}, b_K) \geq u(\overline{a}, b_{K-1}) \geq \cdots \geq u(\overline{a}, b_1)$. First, note that it has to be the case that

$$\tilde{g}(\overline{a})[b_{K-1}] = \tilde{g}(\overline{a})[b_{K-2}] = \cdots = \tilde{g}(\overline{a})[b_1] = \tilde{g}(\overline{a})[\overline{b}]$$

First, by the (ML) constraint, $\tilde{g}(\overline{a})[b_k] \leq \tilde{g}(\overline{a})[\overline{b}]$ for every $k \leq K$. Second, if $\tilde{g}(\overline{a})[b_k] > 0$ for some $k \leq K$, then $\tilde{g}(\overline{a})[b_\ell] = \tilde{g}(\overline{a})[\overline{b}]$ for all $\ell < k$. Otherwise, a belief that assigns slightly lower probability to b_k and slightly higher probability to b_ℓ would lower the expected average payoff without violating the (ML) constraint.

If $\tilde{g}(\overline{a})[b_K] = \tilde{g}(\overline{a})[\overline{b}]$, then we are done: we have shown that $\tilde{g}(\overline{a})$ is uniform. Otherwise, $0 < \tilde{g}(\overline{a})[b_K] < \tilde{g}(\overline{a})[\overline{b}]$. In this case, consider a belief g' such that $g'[b_K] = \tilde{g}(\overline{a})[b_K] + \varepsilon$, $g'[b_k] = \tilde{g}(\overline{a})[b_k] - \frac{\varepsilon}{K}$ for all $k < K$, and $g'[\overline{b}] = \tilde{g}(\overline{a})[\overline{b}] - \frac{\varepsilon}{K}$. Because the expected average payoff is linear in the probabilities on the other player's actions, it follows that if g' generates a lower (higher) expected average payoff for some "small" ε compared with $\tilde{g}(\overline{a})$, then it does so also for "large" ε. By the optimality of $\tilde{g}(\overline{a})$, it follows that both $\tilde{g}(\overline{a})$ and g' generate the same expected average payoff. By choosing $\varepsilon = -\tilde{g}(\overline{a})[b_K]$ we obtain a probability distribution g' that is uniform over $\{b_1, ..., b_{K-1}\} \cup \{\overline{b}\}$.

We now derive the maximal permissible expected average payoff from some $a \neq \overline{a}$. For each such a, we can compute the maximal permissible expected average payoff using the same argument we used for deriving the minimal permissible expected average payoff from $\overline{a}$. In particular, this maximal payoff is obtained with a uniform distribution over some subset of the other player's actions. It remains to maximize over all $a \neq \overline{a}$. This yields the R.H.S. of (15). ■

Proof of Proposition 6

We prove this by showing that the set of feasible models that correspond to M additional samples subsumes the set of feasible models that correspond to $M-1$ additional samples.

Let g_{M-1} denote a feasible model that corresponds to $M-1$ additional samples. Recall that

g maps each action to an expected average payoff, where the expectation is taken with respect to a distribution over the other player's actions. This distribution satisfies our ML constraint. Thus, the expected average payoff from action a_1 is given by:

$$W_1 = \frac{1}{M} u(a_1, b_1) + \frac{M-1}{M} \mathbb{E}_D u(a_1, b_k) \tag{22}$$

where D is some distribution over b_k that satisfies the ML constraint.

Consider a model g_M that corresponds to taking M more samples, and which uses the same distribution D. Such a model is feasible. Suppose $W_1 > u(a_1, b_1)$. Then

$$\frac{1}{M+1} u(a_1, b_1) + \frac{M}{M+1} \mathbb{E}_D u(a_1, b_k) > W_1 \tag{23}$$

By continuity, there exists a distribution D' that is obtained from D by shifting probability mass to b_1 and decreasing probability mass over all other b_k proportionally such that

$$\frac{1}{M+1} u(a_1, b_1) + \frac{M}{M+1} \mathbb{E}_{D'} u(a_1, b_k) = W_1 \tag{24}$$

Since D' increases the probability mass on the opponent's action that was originally samples, it only relaxes the ML constraint, and is therefore, feasible. A similar argument can be applied if $W_1 < u(a_1, b_1)$.

By performing the same procedure for each action a_k we obtain a model g_M that replicates the same mapping from actions to expected average payoffs as g_{M-1}. ■

Proof of Proposition 8

For agents who choose $a = 0$, we have

$$R_0^* = \max_q [q(h=1 \mid a=1) - k - q(h=1 \mid a=0)]$$

which reduces to

$$R_0^* = \max_\beta \left(\frac{\alpha - \beta}{\alpha} - k - \frac{\frac{1}{2} - \alpha + \beta}{1 - \alpha} \right)$$

Since the R.H.S. *decreases* with β, regret is maximized for the *minimal* β that satisfies $\frac{1}{2} - \alpha + \beta \geq 0$. This yields that $\beta_0^* = \max\{0, \alpha - \frac{1}{2}\}$ and implies that

$$R_0^* = \begin{cases} \frac{2k(1-2k)}{4k-1} & if \quad k \geq \frac{1}{3} \\ \frac{2k^2}{1-2k} & if \quad k < \frac{1}{3} \end{cases} \tag{25}$$

(note that both values are positive since $k < \frac{1}{2}$).

For agents who choose $a = 1$, we have

$$R_1^* = \max_q [q(h=1 \mid a=0) - q(h=1 \mid a=0) - k]$$

which reduces to

$$R_1^* = \max_\beta \left(\frac{\frac{1}{2} - \alpha + \beta}{1-\alpha} - \frac{\alpha - \beta}{\alpha} + k \right)$$

Since the R.H.S. *increases* with β, regret is maximized for the *maximal* β that satisfies $\frac{\alpha}{2} - \beta \geq 0$ and $\frac{1}{2} - \alpha + \beta \leq 1$. Since $\frac{\alpha}{2} < \alpha + \frac{1}{2}$, the solution is $\beta = \frac{\alpha}{2}$. It follows that $R_1^*(k) = k$. ■

Proof of Proposition 9

Proof of part (i). If $k \geq \frac{1}{3}$, then q_0^* satisfies

$$q_0^*(h=1 \mid a=1) = 1 > \frac{3k-1}{4k-1} = q_0^*(h=1 \mid a=0)$$

and

$$q_0^*(c=1 \mid a=0, h=0) = \frac{4k-1}{2k}$$

while

$$q_0^*(c=1 \mid a=0, h=1) = q_0^*(c=1 \mid a=1, h=0) = 0$$

Proof of part (ii). If $k < \frac{1}{3}$, then

$$\begin{aligned} &q_0^*(h=1 \mid a=1) = \frac{1}{2} > 0 = q_0^*(h=1 \mid a=0), \\ &q_0^*(c=1 \mid a=1, h=0) = q_0^*(c=1 \mid a=1, h=1) = 0, \text{ and} \\ &q_0^*(c=1 \mid a=0, h=0) = q_0^*(c=1 \mid a=0) = \frac{1}{2}. \end{aligned}$$

(note that the event $(a = 0 \wedge h = 1)$ has zero probability).

Proof of part (iii). Note that

$$q_1^*(h=1 \mid a=1) = q_1^*(h=1 \mid a=0) = \frac{1}{2},$$

and that

$$q_1^*(c=1 \mid a, h) = (1-a)(1-h).$$

■

Bibliography

Aina, Chiara. 2024. "Tailored Stories." *Working paper*.

Ba, Cuimin. 2024. "Robust Misspecified Models and Paradigm Shift." *Working paper*.

Bergemann, Dirk, and Karl H. Schlag. 2008. "Pricing without Priors." *Journal of the European Economic Association* 6 (2-3): 560–569. 10.1162/JEEA.2008.6.2-3.560.

Bergemann, Dirk, and Karl H. Schlag. 2011. "Robust monopoly pricing." *Journal of Economic Theory* 146 (6): 2527–2543. 10.1016/j.jet.2011.10.018.

Cho, In-Koo, and Kenneth Kasa. 2015. "Learning and Model Validation." *The Review of Economic Studies* 82 (1): 45–82.

Cohen, Shani, and Shengwu Li. 2023. "Sequential Cursed Equilibrium."

Eliaz, Kfir, and Ran Spiegler. 2020. "A Model of Competing Narratives." *American Economic Review* 110 (12): 3786–3816. 10.1257/aer.20180624.

Eliaz, Kfir, Ran Spiegler, and Heidi C. Thysen. 2021a. "Persuasion with endogenous misspecified beliefs." *European Economic Review* 134 103712. 10.1016/j.euroecorev.2021.103712.

Eliaz, Kfir, Ran Spiegler, and Heidi C. Thysen. 2021b. "Strategic interpretations." *Journal of Economic Theory* 192 105192. 10.1016/j.jet.2021.105192.

Eliaz, Kfir, Ran Spiegler, and Yair Weiss. 2021c. "Cheating with Models." *American Economic Review: Insights* 3 (4): 417–434. 10.1257/aeri.20200635.

Esponda, Ignacio. 2008. "Behavioral Equilibrium in Economies with Adverse Selection." *American Economic Review* 98 (4): 1269–1291. 10.1257/aer.98.4.1269.

Esponda, Ignacio, and Demian Pouzo. 2016. "Berk–Nash Equilibrium: A Framework for Modeling Agents With Misspecified Models." *Econometrica* 84 (3): 1093–1130. 10.3982/ECTA12609.

Eyster, Erik, and Matthew Rabin. 2005. "Cursed Equilibrium." *Econometrica* 73 (5): 1623–1672. 10.1111/j.1468-0262.2005.00631.x.

Gonçalves, Duarte. 2023. "Sequential Sampling Equilibrium." *Working paper*.

He, Kevin, and Jonathan Libgober. 2023. "Evolutionarily Stable (Mis)specifications: Theory and Applications." *Working paper*.

Heller, Yuval, and Eyal Winter. 2020. "Biased-Belief Equilibrium." *American Economic Journal: Microeconomics* 12 (2): 1–40. 10.1257/mic.20180319.

Jehiel, Philippe. 2005. "Analogy-Based Expectation Equilibrium." *Journal of Economic Theory* 123 (2): 81–104. doi.org/10.1016/j.jet.2003.12.003.

Jehiel, Philippe. 2022. "Analogy-based expectation equilibrium and related concepts: Theory, applications, and beyond." *Working paper*.

Jehiel, Philippe, and Frédéric Koessler. 2008. "Revisiting games of incomplete information with analogy-based expectations." *Games and Economic Behavior* 62 (2): 533–557. 10.1016/j.geb.2007.06.006.

Jehiel, Philippe, and Giacomo Weber. 2024. "Endogenous Clustering and Analogy-Based Expectation Equilibrium." *Working paper*.

Kops, Christopher, and Illia Pasichnichenko. 2023. "Testing negative value of information and ambiguity aversion." *Journal of Economic Theory* 213 105730.

Laurence, Emily, and Nicholas Temple. 2023. "The Psychology Behind FOMO." *Forbes Health*, `https://www.forbes.com/health/mind/the-psychology-behind-fomo/`.

Li, Jian. 2020. "Preferences for partial information and ambiguity." *Theoretical Economics* 15 (3): 1059–1094.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. 2023. "Rational Inattention: A Review." *Journal of Economic Literature* 61 (1): 226–273. 10.1257/jel.20211284.

Milyavskaya, Marina, Mark Saffran, Nadia Hope, and Richard Koestner. 2018. "Fear of missing out: prevalence, dynamics, and consequences of experiencing FOMO." *Motivation and Emotion* 42 (5): 725–737. 10.1007/s11031-018-9683-5.

Osborne, Martin, and Ariel Rubinstein. 1998. "Games with Procedurally Rational Players." *American Economic Review* 88 (4): 834–847.

Salant, Yuval, and Josh Cherry. 2020. "Statistical Inference in Games." *Econometrica* 88 (4): 1725–1752. https://doi.org/10.3982/ECTA17105.

Schwartzstein, Joshua, and Adi Sunderam. 2021. "Using Models to Persuade." *American Economic Review* 111 (1): 276–323. 10.1257/aer.20191074.

Shishkin, Denis, and Pietro Ortoleva. 2023. "Ambiguous information and dilation: An experiment." *Journal of Economic Theory* 208 105610.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics* 50 (3): 665–690. 10.1016/S0304-3932(03)00029-1.

Spiegler, Ran. 2011. *Bounded Rationality and Industrial Organization*. Oxford: Oxford University Press, . 10.1093/acprof:oso/9780195398717.001.0001.

Spiegler, Ran. 2016. "Bayesian networks and boundedly rational expectations." *The Quarterly Journal of Economics* 131 (3): 1243–1290. 10.1093/qje/qjw017.

Spiegler, Ran. 2020. "Behavioral Implications of Causal Misperceptions." *Annual Review of Economics* 12 (Volume 12, 2020): 81–106. https://doi.org/10.1146/annurev-economics-072219-111921.