IDENTIFYING ASSUMPTIONS AND RESEARCH DYNAMICS

ANDREW ELLIS AND RAN SPIEGLER

Abstract. A representative researcher has repeated opportunities for empirical research. To process findings, she must impose an "identifying assumption." She conducts research when the assumption is sufficiently plausible (taking into account both current beliefs and the quality of the opportunity), and updates beliefs as if the assumption were perfectly valid. We study the dynamics of this learning process. While the rate of research cannot always increase over time, research slowdown is possible. We characterize environments in which the rate is constant. Long-run beliefs can exhibit history-dependence and "incredible certitude." We apply the model to stylized examples of empirical methodologies.

1. Introduction

How should scientific research be conducted? The standard model of rational behavior assumes that agents undergo Bayesian learning (following de Finetti (1974/5) and Savage (1954)). Bayesianism is thus a natural benchmark if one thinks of scientific researchers as rational. A Bayesian researcher holds a probabilistic prior belief regarding a research question, accumulates evidence (such as controlled experiments or observational data), and updates her prior in light of that evidence via Bayes' rule. Bayesian researchers would eventually learn the truth (or at least some part of it), provided that the evidence is informative and their prior does not rule it out.[1]

Empirical research in economics often instead interprets evidence through the lens of *identifying assumptions*. These are explicit hypotheses regarding the relationship

Date: 14 January 2025.
Ellis: LSE, a.ellis@lse.ac.uk. Spiegler: Tel Aviv University and UCL, rani@tauex.tau.ac.il. Spiegler acknowledges financial support from ISF grant no. 320/21. We thank Daron Acemoglu, Tim Christensen, In-Koo Cho, Martin Cripps, Tuval Danenberg, Andrew Gelman, Duarte Goncalves, Charles Manski, Ignacio Esponda, Jesse Shapiro, Kate Smith, Zihan Jia, and audiences at Bar-Ilan University, BRIC, BU, NYU, Maryland, a Stony Brook workshop on bounded rationality and learning, and an LSE/UCL theory conference for helpful comments.

[1]For philosophy-of-science discussions of Bayesianism as a normative theory for scientific learning, see Howson and Urbach (2006) and Gelman and Shalizi (2013).

between the evidence and the research question that enable researchers to draw clear-cut conclusions from their observations. For instance, consider learning about the causal effect of class size on students' test performance. To draw causal inferences from an observational dataset that links these two variables, researchers need to form beliefs about how students were assigned to classes. A typical identifying assumption in this context is that the assignment was random, which enables a sharp causal interpretation of students' performance differences across class sizes.

Researchers do not mindlessly impose these assumptions. They debate over their plausibility for the dataset in question, and adopt them only when deemed sufficiently plausible. This judgment may evolve over time. For example, as researchers accumulate evidence about the effect of class size on student performance, they develop more precise beliefs about this effect, which may lead them to become more (or less) stringent in admitting an identifying assumption that is only approximately valid. The ongoing, case-by-case evaluation of an assumption's plausibility distinguishes this assumption-based learning process from Bayesianism. There, assumptions may inform the prior belief but then cease to play any subsequent role.

Assumption-based learning circumvents some practical difficulties with implementing the strict Bayesian recipe. Identifying assumptions reduce the dimensionality of the uncertainty that the research community needs to process and communicate. They are also typically simple to describe. These features simplify the process of belief revision. For example, if the research community accepts the random-assignment assumption in the example, then its subjective beliefs about the assignment process do not affect the updating process. Identifying assumptions also facilitate reaching a long-run consensus in the research community, since according to them, the evidence alone determines the long-run answer to the research question.

Despite these advantages, if researchers make assumptions for practical convenience but later take them too seriously, then they effectively engage in misspecified learning. As a result, they may come to hold incorrect and (by the nature of identifying assumptions) strongly held beliefs. That is, the process may lead to what Manski (2020) called "incredible certitude."

In this paper, we model a research process that treats identifying assumptions as necessary for conducting research but that otherwise conforms to Bayesian learning. A representative researcher, a stand-in for the relevant research community, wants to determine some of (or all of) the values of a collection of fixed parameters. (e.g., the effect of class size on students' test performance). She faces a sequence of research designs of random quality, given by *i.i.d.* context parameters (e.g., the extent to which assignment of students to classes in a dataset is random). Both fixed and context parameters directly affect the data-generating process of the study, if it is carried out.

The community accepts the findings only when they are conducted under a sufficiently plausible identifying assumption (e.g., that student assignment is perfectly random). The assumption fixes the value of the context parameters in such a way that if the study were independently repeated many times, the results would produce a definitive answer to the research question. If the assumption is not sufficiently plausible, then the study is ignored (or not carried out in the first place), and the researcher waits for the next opportunity. When the study is carried out, the identifying assumption informs how its results are incorporated into updated beliefs.

The community evaluates an assumption's plausibility by comparing its beliefs regarding the distribution of variables under the actual and assumed values of the context parameters. If the former distribution differs too much from the latter, the researcher deems the assumption implausible and passes over the opportunity to conduct research. If the difference is small, she deems the assumption plausible, conducts the study, observes its result as determined by the true data-generating process, and updates her beliefs as if the assumption held *exactly*. We quantify the difference using an *f-divergence* (Csiszár, 1967), which measures the expectation of a convex function of the likelihood ratios induced by the two distributions. The ubiquitous Kullback-Leibler (KL) divergence is a special case, which we employ in our applications. Note that the plausibility judgment depends not only on the specific details of the research opportunity but also on the community's beliefs over the fixed parameters. As a result, whether or not the research is conducted or processed depends on past decisions.

We study the dynamics and long-run behavior of this assumption-based learning process. Our focus is on how the propensity to conduct research (via the imposition of an identifying assumption) changes as the research community's beliefs evolve over time. We show that this propensity cannot always increase over time. In other words, the research community cannot consistently lower its standards for accepting research as time goes by. It may, however, continually raise these standards, leading to a slowdown in the rate of research but an increase in its quality. Intuitively, as the researcher's belief exhibits greater confidence, she becomes more sensitive to the assumption's rough edges and therefore more reluctant to impose it.

We also provide a sufficient condition for a time-invariant propensity to conduct research. Essentially, the condition states that if an observable variable is correlated with context parameters given the other observables, then it must be independent of the fixed parameters of interest (and other fixed parameters that the identifying assumption renders observationally relevant). Thus, our condition demands separation between the observable effects of fixed parameters of interest and the context parameters. This result employs tools from the literature on graphical probabilistic models (e.g., Pearl (2009)).

We then ask what the researcher eventually comes to believe. We define a stable belief to be one that the updating process converges to with positive probability. Stable beliefs concentrate on parameter values for which observables' distribution conditional on the assumption is closest (in the sense of KL divergence) to their empirical distribution given the true value and the contexts in which research is conducted. In turn, these contexts are determined by the stable belief itself. This two-way relation makes stable beliefs an equilibrium object. For most models of interest, stable beliefs almost always assign probability one to a wrong value of the fixed parameters of interest. This result resonates with the Manski's above-mentioned incredible-certitude critique of scientific learning based on strong identifying assumptions.

We demonstrate the model's scope with stylized examples of familiar empirical methodologies. One example considers experimental research contaminated by interference (the identifying assumption rules out the interference). We show that the propensity to conduct the experiment decreases over time. Another example examines causal inference contaminated by confounding effects (the identifying assumption is that no such confounding exists). A variation on this example addresses instrumental-variable designs (the identifying assumption is that the instrument is independent of a latent confounder). Both variants satisfy our sufficient condition for time-invariant propensity to conduct research. All three examples satisfy the condition that leads to incredible certitude in the long-run.

Finally, we expand the notion of identifying assumptions to cover fixed parameters, and allow the researcher to choose from a set of candidate identifying assumptions. We do so through a stylized model of inference from selective samples, where the researcher wishes to learn the returns from an activity. She considers two alternative identifying assumptions: agents' selection into this activity is purely random, or selection is systematically related to observables that do not directly affect returns. The latter is a structural identifying assumption that captures in stylized form the method of Heckman selection (Heckman, 1979). Learning dynamics exhibit history-dependence: using an identification method changes beliefs in a way that can reinforce its use.

Related literature. Our paper continues a recent literature on Bayesian learning under misspecified prior beliefs (Esponda and Pouzo (2016), Fudenberg et al. (2017), Frick et al. (2020), Heidhues et al. (2018), Bohren and Hauser (2021), Esponda and Pouzo (2021)).[2] One strand in this literature (e.g., Cho and Kasa (2015), Ba (2024)) incorporates continual model selection and misspecification tests into the learning process. Our paper departs from the literature in several respects. First, the identification motive for adopting a misspecified model and the belief-based plausibility criterion that governs it are novel. Second, the rate of learning in our model is endogenous and can vary over time. In general, our economics-of-science angle is new to the misspecified-learning literature.

The econometrics literature contains methodological discussions of the role of identifying assumptions (Rothenberg (1971), Manski (2007), Lewbel (2019)). However, we are unaware of earlier discussions of how identification methods can be reconciled with the Bayesian approach. As we saw, Manski himself is a critic of using strong

[2]For a review of alternatives to Bayesian updating, see Ortoleva (2024).

identifying assumptions. Our paper can be viewed as a *descriptive* model of the phenomenon that Manski criticizes.

There have been recent attempts to model non-Bayesian researchers. Andrews and Shapiro (2021) show that conventional loss-minimizing estimators may be suboptimal when consumers of the researcher are Bayesian with heterogeneous priors. Banerjee et al. (2020) describe researchers as ambiguity averse max-minimizers. Spiess (2024) models strategic choice of model misspecification by researchers. In relation to this literature, our paper makes (to our knowledge) the first attempt to model the role of assumptions in how researchers interpret empirical observations.

2. A Model

A *representative researcher* is interested in a *question* whose answer is determined by *fixed parameters* $\omega \in \Omega \subset \mathbb{R}^n$. The question is a subset $Q \subseteq \{1, ..., n\}$, indicating which of the parameters the researcher wishes to learn. The researcher has a prior belief about the fixed parameters. We assume that Ω is compact and convex, that her beliefs admit a continuous probability density μ, and that $\mu(\omega) = \prod_{i=1}^{n} \mu_i(\omega_i) > 0$ for all $\omega \in \Omega$. We sometimes refer to μ as the researcher's beliefs.

Time is discrete. In every period $t = 1, 2, ...$, a real-valued vector $\theta^t \in \Theta$ of *context parameters* is realized. We refer to a realization of θ^t as a *context*. While ω represents a constant feature of the phenomenon of interest (e.g., returns to education), θ^t represents transient, circumstantial aspects of a particular dataset or experiment (e.g., whether assignment of students to educational treatments in a particular setting is random). We assume that Θ is compact and convex, and that θ^t admits a density p_θ.

In period t, the researcher observes θ^t and makes a decision $a^t \in \{0, 1\}$, indicating whether to conduct research. If the researcher chooses $a^t = 0$, then she passes over the opportunity to conduct research. If research is conducted in period t, then a vector of observed variables (referred to as *statistics*) $s^t \in S$ and a vector of *unobserved variables* $u^t \in U$ are generated; both S and U are subsets of Euclidean space. For expositional convenience, our definitions and general results proceed as if S and U are both finite; extension to the continuum case is straightforward and we make use of it in most of our examples. Throughout, we adopt the notational convention that for a vector x, $x_B = (x_i)_{i \in B}$ when B is a subset of the indices for x and $x_{-i} = (x_j)_{j \neq i}$.

The data-generating process p that governs the realization of (u, s) at every time period satisfies $p(u^t, s^t|\theta^t, \omega) = p_u(u^t)\, p(s^t|u^t, \theta^t, \omega)$. We assume that p is continuous and has full support at every (θ, ω). Then, the density

$$p\left(s^t, u^t, \theta^t, \omega\right) = \mu(\omega)\, p_\theta\left(\theta^t\right) p_u\left(u^t\right) p\left(s^t|u^t, \theta^t, \omega\right)$$

describes her prior beliefs.[3] The context parameters and unobserved variables are distributed independently and identically across periods.

An *assumption* is an element $\theta^* \in \Theta$. We say that an assumption θ^* is *identifying* for Q if for every $\omega, \psi \in \Omega$ such that $\omega_Q \neq \psi_Q$, there exists $s \in S$ such that $p(s \mid \theta^*, \omega) \neq p(s \mid \theta^*, \psi)$. If an identifying assumption holds, then repeated observation of s would eventually provide a definitive answer to the research question. We assume that there is a single feasible identifying assumption. In Section 6, we extend our model to allow for multiple feasible identifying assumptions (including assumptions about fixed parameters).

The researcher knows p. Entering period t, she observes the current context θ^t and has beliefs about the fixed parameters described by the density $\mu(\cdot|h^t)$, which depends on the history $h^t = (a^\tau, s^\tau, \theta^\tau)_{\tau<t}$; beliefs given the empty history equal the prior. There are a constant K and a divergence D measuring how much one probability measure on $S \times U$ deviates from another.[4] Let $p_{S,U}(\cdot|\theta^t, h^t)$ and $p_S(\cdot|\theta^t, h^t)$ denote the researcher's beliefs given (θ^t, h^t) over (s^t, u^t) and s^t, respectively. If

$$D\left(p_{S,U}\left(\cdot|\theta^t, h^t\right) || p_{S,U}\left(\cdot|\theta^*, h^t\right)\right) > K,$$

then the assumption is deemed implausible and the researcher chooses $a^t = 0$. Otherwise, the assumption is deemed sufficiently plausible and the researcher chooses $a^t = 1$.

If the researcher chooses $a^t = 0$, then she does not update her beliefs, and so the next research opportunity, arising at period $t + 1$, is evaluated according to the same

[3]The density is the Radon-Nikodyn derivative with respect to the Lebesgue measure on $\Theta \times \Omega$ multiplied by the uniform measure on $S \times U$.

[4]Recall that D is a divergence if $D(m||m')$ is positive, is continuous in (m, m'), and equals 0 if and only if $m = m'$ (a.s.).

belief as in period t. If the researcher chooses $a^t = 1$, then she updates her belief over Ω *as if* the assumption θ^* held exactly. That is, she now uses the density $\mu\left(\cdot|h^{t+1}\right)$ given by

$$\mu\left(\omega|h^t, s^t, a^t = 1, \theta^t\right) = \frac{p\left(s^t|\theta^*, \omega\right)\mu\left(\omega|h^t\right)}{p\left(s^t|\theta^*, h^t\right)} \tag{1}$$

for almost every ω. Denote by $\mu_t(\cdot)$ the Borel probability measure with density $\mu\left(\cdot|h^t\right)$, noting that μ_t is implicitly a function of h^t. Observe that the dependence of $p_{S,U}\left(\cdot|\theta^t, h^t\right)$ on h^t is restricted to its public part, namely $\left(s^\tau\right)_{\{\tau:a^\tau=1\}}$, because every time the researcher updates her belief, she does so as if the context is θ^*.

We assume that D belongs to the class of f-divergences (Rényi, 1961, Csiszár, 1967), i.e.,

$$D\left(m||m'\right) = \sum_{s,u} m'\left(s,u\right) f\left(\frac{m\left(s,u\right)}{m'\left(s,u\right)}\right)$$

for a strictly convex function $f\left(\cdot\right)$ satisfying $f(1) = 0$. In all of our examples, we take D to be KL divergence,

$$D\left(m||m'\right) = D_{KL}\left(m||m'\right) = \sum_{s,u} m\left(s,u\right)\ln\left(\frac{m\left(s,u\right)}{m'\left(s,u\right)}\right)$$

the special case where $f\left(x\right) = x\ln x$.

2.1. **An example: A contaminated experiment.** To illustrate our model, consider a researcher who wants to identify a behavioral effect from experimental data that is contaminated by "friction." When the researcher conducts the experiment at period t, she observes the statistic s^t given by $s^t = \omega_1 + \theta^t\omega_2 + \varepsilon^t$. She wants to learn the fixed parameter ω_1, i.e., $Q = \{1\}$. The fixed parameter ω_2 represents the friction's strength. The context $\theta \in [0,1]$ captures how well the experimental design manages to curb the friction, and $\varepsilon^t \sim N(0,1)$ represents noise and is independently drawn across periods. There are no latent variables. The researcher's prior holds that each $\omega_i \sim N\left(m_i, (\sigma_i)^2\right)$, independently of the other component of ω. For example, ω_1 is the degree of *intrinsic* altruism in a certain social setting, and ω_2 is how much the subjects want an outside observer to *perceive* them as altruistic. The only identifying assumption is $\theta^* = 0$: under any $\theta^t \neq 0$, both ω and $\left(\omega_1 - 1, \omega_2 + 1/\theta^t\right)$ generate the same distribution.

2.2. **Discussion.** The interpretation of the learning process is as follows. The researcher can only update her beliefs under an identifying assumption, but will do so only if she deems the assumption sufficiently plausible. We refer to the decision to process the data at a given period as if it is a decision whether to conduct the research at that period. This fits an interpretation that the plausibility judgment is made by the researcher herself. Alternatively, it could be viewed as a decision by the *research community* (embodied by seminar audiences and journal referees) whether to "take the research seriously" and incorporate it into its collective knowledge. Under both interpretations, the plausibility judgment at any given period is made *before* the research results are observed.

Plausibility is captured by how likely variable realizations are under the actual context θ^t relative to the assumed one θ^*. The likelihood judgment is based on the researcher's current beliefs. When relative likelihoods become farther away from 1, the assumption is deemed less plausible because it tends to produce unlikely predictions about the variables, according to the researcher's belief. Convexity of the function f in an f-divergence means that the penalty for a marginal shift of the likelihood ratio away from 1 is increasing.

The plausibility judgment has additional noteworthy features. First, it depends only on the current period's context and the current belief μ_t. Accordingly, the set of values of θ for which the researcher conducts research depends on her belief μ_t and is denoted by $\Theta^R(\mu_t)$. Second, since $p(\cdot|\theta,\omega)$ has full support, the data never definitively refute a wrong assumption. Consequently, every assumption has finite f-divergence from the truth. Third, the plausibility judgment takes into account the assumption's effect on the distribution of *both* observed (s) and latent (u) variables. This reflects our observation of real-life discussions of identification strategies in empirical economics. For instance, evaluation of an instrumental variable is based on a judgment of whether the (observed) instrument is correlated with (unobserved) confounding variables. Finally, the constant K captures the research community's tolerance to implausible assumptions. While this tolerance can reflect an underlying calculation of costs and benefit of doing research, we do not explicitly model this calculus. Since the research community knowingly chooses to distort its beliefs by making wrong assumptions, it is not obvious how one should model such a cost-benefit analysis.

An alternative interpretation of the plausibility constraint is that it reflects limitations on the ability to communicate the circumstances of a scientific study across the research community. Under this interpretation, a context parameter represents soft, high-dimensional information that is difficult to communicate. The assumption represents an exogenous communicability constraint. The researcher with the opportunity to conduct research knows the true context (or has a good signal about it) but cannot communicate it, yet she does not want to mislead the community. The plausibility constraint captures the idea that she only disseminates the research when she thinks it will not mislead the community too much by ignoring the context.

We conclude with a comment on the notion of identifying assumptions. Recall that in the Contaminated Experiment example of Section 2.1, we noted that the only feasible identifying assumption is $\theta^* = 0$. This claim rests on an implicit feature of our model: identification assessments are made for each time period in isolation. Suppose we observe the long-run distribution of s for two known values of θ. Then, we have two equations with two unknowns (ω_1 and ω_2), and we can therefore pin down both. It follows that if identification were assessed by combining multiple contexts (given by different values of θ), there would be no need for identifying assumptions. This "triangulating" identification strategy would work in most of the examples in this paper. However, it is natural to assume that researchers cannot perform this kind of triangulation, as they do not know future values of θ and past research opportunities cannot be repeated. Moreover, triangulation is inconsistent with the research practice we are familiar with, where the identification constraint is applied separately to each piece of research.

3. Research slowdown

In this section we begin our analysis of the learning process spelled out in the previous section. This section and the next are devoted to the question of how the rate of conducting research evolves over time.

3.1. **The Contaminated Experiment Continued.** Let us analyze learning dynamics for the example given by Section 2.1. Entering period t, the researcher believes that $\omega_i \sim N(m_i^t, (\sigma_i^t)^2)$ for each i. Whenever the researcher updates her beliefs, she does so as if $\theta = 0$, so her beliefs over ω_2 never evolve; accordingly, we remove the time index from the mean and variance of ω_2.

Entering period t, the researcher believes that the distribution of s^t conditional on θ^t is

$$N\left(m_1^t + \theta^t m_2, 1 + \left(\sigma_1^t\right)^2 + \left(\theta^t\right)^2 \sigma_2^2\right).$$

Using the standard formula for KL divergence between two scalar Gaussian variables,

$$D_{KL}\left(p_S\left(\cdot|h^t, \theta^t\right)||p_S\left(\cdot|h^t, \theta^*\right)\right) = \frac{1}{2}\left[\left(\theta^t\right)^2 \frac{\sigma_2^2 + m_2^2}{1 + (\sigma_1^t)^2} - \ln\left(1 + \frac{(\theta^t)^2 \sigma_2^2}{1 + (\sigma_1^t)^2}\right)\right].$$

Thus, the only time-varying elements that affect the propensity to experiment are σ_1^t and θ^t.

The KL divergence is continuous and increasing in θ^t, and vanishes when $\theta^t = 0$. Consequently, there exists a threshold $\bar{\theta}\left(\sigma_1^t\right) > 0$ such that the researcher conducts research if and only if $\theta^t \in \left[0, \bar{\theta}\left(\sigma_1^t\right)\right]$. Holding θ^t fixed, divergence decreases in σ_1^t, so the threshold for conducting research $\bar{\theta}\left(\cdot\right)$ increases in σ_1^t.

When the researcher conducts the experiment in period t, the variance of her belief about ω_1 decreases to $\sigma_1^{t+1} = \sigma_1^t\left((\sigma_1^t)^2 + 1\right)^{-\frac{1}{2}}$. That is, σ_1^t decreases monotonically over time. Consequently, the propensity to conduct research uniformly decreases over time. As the researcher becomes more certain of her belief over ω_1, she also becomes more sensitive to the noise and so more reluctant to assume it away. In other words, her standards for what passes as adequate research design increase over time. This slows down the rate of learning.

However, learning takes place with positive frequency in the long run. To see why, note that as $\sigma_1^t \to 0$, the divergence converges to

$$\frac{1}{2}\left[\left(\sigma_2^2 + m_2^2\right)\left(\theta^t\right)^2 - \ln\left(1 + \left(\theta^t\right)^2 \sigma_2^2\right)\right],$$

which is finite for every θ^t. This means that $\bar{\theta}\left(0\right) > 0$, and research takes place with positive probability, regardless of the researcher's current belief. This non-vanishing learning implies that $\sigma_1^t \to 0$ as $t \to \infty$. In the long-run, research is carried out whenever $\theta \in \left[0, \bar{\theta}(0)\right]$, and that the researcher's belief over ω_1 assigns probability one to

$$\omega_1 + \mathbb{E}\left(\theta | \theta < \bar{\theta}(0)\right) \omega_2.$$

Thus, the long-run estimate of the effect of interest is biased in proportion to the true value of the friction parameter ω_2. The magnitude of the bias also increases with σ_2^2 (the researcher's time-invariant uncertainty over the friction parameter) since $\bar{\theta}(0)$ increases with σ_2^2.

To summarize our findings in this example, the researcher's propensity to learn decreases over time but remains positive in the long-run. This in turn means that the long-run answer to the research question is biased. The bias is proportional to the true value of the fixed friction parameter, and increases (in absolute terms) with the researcher's uncertainty over it.

3.2. **General results.** In the above example, the set of contexts for which research takes place (weakly) contracts over time. This shows the possibility of a uniform (i.e., with probability one) research slowdown. Our first two results show that the opposite pattern, namely a uniform increase in the propensity to conduct research, cannot occur. Consequently, the rate of research decreases at least with some probability.

Proposition 1. *For any $\theta \in \Theta$ and history h^t, if $\mathbb{P}\left[\theta \in \Theta^R(\mu_{t+1}) \setminus \Theta^R(\mu_t) | h^t\right] > 0$, then there exists $t^* > t+1$ such that $\mathbb{P}\left[\theta \notin \mathit{\Theta}^R(\mu_{t^*}) | h^{t+1}\right] > 0$.*

This result states that any expansion in the set of parameters for which research is conducted reverses itself with positive probability. Consider a context for which research does not takes place in some period. Suppose that there is some piece of evidence that would lead to research being performed for that same context in the following period. The result shows that with positive probability, there is a point in the future at which the researcher would once again refrain from conducting research in that same context.

When a larger difference in context maps naturally to a bigger divergence, we can be more explicit about how the propensity to conduct research evolves.

Proposition 2. *Suppose that* $D\left(p_{S,U}\left(\cdot|\theta,h^{t}\right)||p_{S,U}\left(\cdot|\theta^{*},h^{t}\right)\right)$ *is quasi-convex in* θ *for every history* h^{t}. *If*

$$\mathbb{P}\left[\Theta^{R}\left(\mu_{t+1}\right)\setminus\Theta^{R}\left(\mu_{t}\right)\neq\emptyset|h^{t}\right]>0,$$

then

$$\mathbb{P}\left[\Theta^{R}\left(\mu_{t}\right)\setminus\Theta^{R}\left(\mu_{t+1}\right)\neq\emptyset|h^{t}\right]>0.$$

This result says that when there are contexts for which research takes place at period $t+1$ but not at t (i.e., $\Theta^{R}\left(\mu_{t+1}\right)\setminus\Theta^{R}\left(\mu_{t}\right)\neq\emptyset$), then with positive probability, there are contexts for which research takes place at t but not at $t+1$ (i.e., $\Theta^{R}\left(\mu_{t}\right)\setminus\Theta^{R}\left(\mu_{t+1}\right)\neq\emptyset$). That is, at any history, when the community conducts research in new contexts with positive probability, it also stops conducting research in others. When one value of the statistic increases the chance of conducting research, a different value decreases the chance.

The result relies on the assumption that the divergence is quasi-convex in θ. In particular, this holds when a larger Euclidean distance between θ and θ^{*} implies a larger divergence. In our examples, $\theta\in\mathbb{R}_{+}$, $\theta^{*}=0$, and the divergence strictly increases in θ. Consequently, Proposition 2 applies to all of our examples.

The proofs of Propositions 1 and 2 both rely on convexity of f-divergence. This implies that the divergence of $p_{S,U}(\cdot|\theta,h^{t})$ from $p_{S,U}(\cdot|\theta^{*},h^{t})$ goes up in expectation for every θ when h^{t} is concatenated by additional observations. If it decreases for some histories, then it must rise for others. Both proofs exploit this insight to show that expansions in Θ^{R} must be offset by contractions in it.

4. Constant Propensity to Research

In this section we continue to explore the rate of research in our model, focusing on cases in which this rate remains constant over time.

4.1. **An Example: Confounded causal inference.** Determining the causal effect of one variable on another is a central task for empirical researchers. To do so, researchers must account for (unobserved) confounding variables that affect both the (observed) cause and effect. We present a stylized example of causal inference from observational data in the presence of a potential confounder.

There are two observable variables, a potential cause s_1 and an outcome s_2. The researcher wants to learn the causal effect of the former on the latter. This effect is parameterized by $\omega_1 \in (-1, 1)$, i.e., $Q = \{1\}$. However, the observed correlation between the two variables is confounded by a latent variable u that affects both. The fixed parameter $\omega_2 \in (-1, 1)$ captures the strength of this confounding effect. The context parameter $\theta \in [0, 1]$ captures the extent to which it persists in a given dataset. A lower value of θ captures better "research design" in the sense of attenuating the confounding effect. If she conducts research in period t, then she observes

$$s_1^t = \theta^t \omega_2 u^t + \varepsilon_1^t$$
$$s_2^t = \omega_1 s_1^t + \omega_3 u^t + \varepsilon_2^t$$

where $u^t \sim N(0, 1)$ and $\varepsilon_i^t \sim N\left(0, \sigma_{i,t}^2\right)$ for $i = 1, 2$, independently of each other. There is no uncertainty regarding $\omega_3 > 0$. Set this parameter and the variances $\sigma_{1,t}^2$ and $\sigma_{2,t}^2$ such that $s_i|\omega, \theta^t \sim N(0, 1)$ for each $i = 1, 2$ and (θ, ω).[5] It follows that the only aspect of the long-run distribution of $s^t = (s_1^t, s_2^t)$ that sheds light on the fixed parameters is the pairwise correlation between the two statistics,

$$\rho_{12}\left(\theta^t, \omega\right) = (\theta^t)^2 \omega_2^2 \omega_2 + \theta^t \omega_2 \omega_3 + \omega_1.$$

The equation for $\rho_{12}(\theta^t, \omega)$ reveals that the only identifying assumption is $\theta^* = 0$. Under that assumption, the correlation between s_1 and s_2 equals $\rho_{12}(\theta^*, \omega) = \omega_1$ and so pins down the causal effect of interest. Note that, as in the previous example, this assumption prevents the researcher from learning anything about the other fixed parameter (ω_2).

We derive an expression for the KL divergence between the true and assumed distributions over (u, s). The joint density of the statistics conditional on the parameters factorizes as $p(u, s|\omega, \theta) = p(u)p(s_1|u, \omega_1, \theta)p(s_2|u, s_1, \omega_2)$. It follows that

[5]That is, $\sigma_{1,t}^2 = 1 - (\theta^t)^2 \omega_1^2$ and $\sigma_{2,t}^2 = 1 - \omega_2^2 - \omega_3^2 - 2\theta^t \omega_1 \omega_2 \omega_3$.

$$\begin{aligned}
&D_{KL}\left(p_{S,U}\left(\cdot|\theta^t,h^t\right)||p_{S,U}\left(\cdot|\theta^*,h^t\right)\right)\\
&=\int \ln\left(\frac{\int p(u)p(s_1|u,\omega_2,\theta^t)p(s_2|u,s_1,\omega_1)\mu\left(\omega_1,\omega_2|h^t\right)d\omega}{\int p(u)p(s_1|u,\omega_2,\theta^*)p(s_2|u,s_1,\omega_1)\mu\left(\omega_1,\omega_2|h^t\right)d\omega}\right)dp\left(s,u|\theta^t\right)\\
&=\int \ln\left(\frac{\int p\left(s_1|u,\omega_2,\theta^t\right)\mu\left(\omega_2|h^t\right)d\omega_2\int p\left(s_2|s_1,u,\omega_1\right)\mu\left(\omega_1|h^t\right)d\omega_1}{\int p\left(s_1|u,\omega_2,\theta^*\right)\mu\left(\omega_2|h^t\right)d\omega_2\int p\left(s_2|s_1,u,\omega_1\right)\mu\left(\omega_1|h^t\right)d\omega_1}\right)dp\left(s,u|\theta^t\right)\\
&=\int \ln\left(\frac{\int p\left(s_1|u,\omega_2,\theta^t\right)\mu\left(\omega_2\right)d\omega_2}{\int p\left(s_1|u,\omega_2,\theta^*\right)\mu\left(\omega_2\right)d\omega_2}\right)dp\left(s_1,u|\theta^t\right)
\end{aligned}$$

Note that only the researcher's belief about ω_2 enters the expression. Because this belief is stationary, the divergence for any given θ^t does not change over time. Therefore, the researcher's propensity to research is time-invariant: there is a fixed $\bar{\theta}$ such that she conducts research if and only if $\theta^t \in \left[0,\bar{\theta}\right]$. As $t \to \infty$, the researcher's belief is concentrated on

$$\widehat{\omega}_2 = \mathbb{E}\left[\rho_{12}(\theta,\omega)|\theta < \bar{\theta}\right].$$

Clearly, this long-run estimate is biased when $\omega_2 \neq 0$, i.e., when there is a confounding effect.

4.2. **A General Result.** In the above example, the set of contexts for which research takes place is history-independent. Our next result provides a general sufficient condition for this property. The condition concerns the probabilistic relationship between the context and the fixed parameters about which the researcher learns under the identifying assumption. We call these parameters *active*, and formally define the set of *active parameters* Q^* to be the smallest set of indices for which $p_S\left(\cdot|\theta^*,\omega\right) = p_S\left(\cdot|\theta^*,\omega'\right)$ whenever $\omega_{Q^*} = \omega'_{Q^*}$. Under the identifying assumption, no other fixed parameters affect the long-run distribution of s, and so repeated observation teaches the researcher nothing about them. The set of active parameters is defined with respect to θ^*, and by the definition of identifying assumptions, $Q \subseteq Q^*$. In both examples, the set of active parameters was $Q^* = Q = \{1\}$, though this need not be the case (see Section 4.3).

Our result makes use of the underlying recursive structure of the data-generating process as described by a directed acyclic graph (DAG), a graph whose edges correspond to an acyclic binary relation. Our analysis employs tools from the AI/Statistics literature on DAGs (see Pearl (2009) or Koller and Friedman (2009) for a general introduction, and Spiegler (2016, 2020) or Ellis and Thysen (2024) for earlier economic-theory applications). Denote $x^t = (s^t, u^t, \theta^t, \omega) \in \mathbb{R}^m$, $N = \{1, \ldots, m\}$, and $N^s \subset N$ the set of indices such that $x^t_{N^s} = s^t$.[6] Let $G = (N, R)$ be a DAG with nodes N and set of directed edges R such that no edge goes into any node in $N \setminus N^s$. We say that p has the *recursive structure* G if the density of $x^t = (s^t, u^t, \theta^t, \omega)$ can be written as

$$\mu(\omega)\, p_\theta\left(\theta^t\right) p_u\left(u^t\right) \prod_{i \in N^s} p\left(s^t_i | x^t_{R(i)}\right) \tag{2}$$

for every x^t. In this formula, $R(i)$ is the set of nodes that send directed edges into i, with customary abuse of notation.

Every p has a recursive structure (in particular, one where every statistic node is linked to every other node, in which case Equation (2) holds by a standard chain rule). However, sparser structures describe essential features of the data-generating process. Whenever p is described by a recursive system of regression equations, as in all the examples in this paper, this system defines a recursive structure for p: for every $i \in N^s$, $R(i)$ is the set of R.H.S. variables in the equation for s_i. For instance, the recursive structure of p in the Causal Inference example is

$$\begin{array}{ccccc} \theta^t & \rightarrow & s^t_1 & \leftarrow & \omega_2 \\ & \nearrow & \downarrow & & \\ u^t & \rightarrow & s^t_2 & \leftarrow & \omega_1 \end{array}.$$

A DAG G satisfies a conditional-independence property if every probability measure that is consistent with G satisfies this property. We say that θ and ω_{Q^*} are *G-separable* if for every i, G satisfies $s^t_i \perp \omega_{Q^*}$ whenever it satisfies $s^t_i \not\perp \theta^t | \left(s^t_{-i}, u^t\right)$. If θ and ω_{Q^*} are G-separable, then the active parameters do not influence any statistic whose distribution depends on the context (conditional on the other variables). Thus, G-separability describes a sense in which context and active parameters have distinct observable effects. Structural conditional-independence properties such as

[6] The components of x^t are ordered such that $s^t_i = x^t_i$.

G-separability have a graphical characterization known as "d-separation" (see Pearl (2009)), which makes the condition easy to check. The Contaminated Experiment example violates G-separability, since the context parameter and the fixed parameter of interest both send links to the statistic. In contrast, the Causal Inference example satisfies the property. The context parameter θ^t is d-separated from s_2^t given (s_1^t, u^t) since the latter block every path from θ^t to s_2^t, hence $s_i^t \perp \omega_{Q^*}|(s_1^t, u^t)$. The fixed parameter $\omega_1 = \omega_{Q^*}$ and s_1^t are independent because they have no common ancestor.

Proposition 3. *If p has a recursive structure G for which θ and ω_{Q^*} are G-separable, then $\Theta^R(\cdot)$ is constant.*

The conditional-independence property that underlies Proposition 3 is not imposed directly on the researcher's belief, but rather on its underlying structure G. Under G-separability, the set of contexts in which the researcher updates does not change over time. That is, when context and active parameters have separate observable effects, the propensity to conduct research is constant. The proof uses d-separation to factorize beliefs into conditional-probability terms. We show that every statistic whose distribution is sensitive to the identifying assumption must be conditionally independent of the active parameters. This in turn implies that every term involving ω_{Q^*} cancels out or integrates out of the expression for the f-divergence. Since the researcher's beliefs about the other fixed parameters do not change, the divergence for any given context remains fixed over time.

Proposition 3 relies on D being an expectation of some function of likelihood ratios (induced by the true and assumed values of θ). However, it does not rely on the convexity of f. In contrast, Propositions 1 and 2 rely on convexity of D, which follows from the convexity of f.

An identifying assumption makes G-separability non-trivial. If all fixed parameters are active (as would be the case for most values of θ), then G-separability requires that when a statistic depends on a context parameter, it must be independent of all fixed parameters, which is an extreme requirement that trivializes the problem. However, identifying assumptions typically imply the existence of inactive fixed parameters. Identification problems arise when the number of fixed parameters exceeds the number of "moments" that the researcher can extract from the statistic. This leads to a system of equations with too many unknowns. An identifying assumption effectively shuts down some of these unknowns. In our language, these shutdown unknowns are inactive parameters. When there are such inactive parameters, G-separability is not a trivially demanding condition.

4.3. **An Application: Instrumental Variables.** Proposition 3 allows a convenient analysis of whether the propensity to adopt identification strategies for causal inference changes over time. Consider a data-generating process described by the following recursive system of regression equations:

$$\begin{aligned} s_1^t &= \omega_1 \theta^t u^t + \varepsilon_1^t \\ s_2^t &= \omega_2 s_1^t + \omega_3 u^t + \varepsilon_2^t \\ s_3^t &= \omega_4 s_2^t + \omega_5 u^t + \varepsilon_3^t \end{aligned}$$

where u^t and the ε_i^t variables are all independent Gaussians. Their variances and the range of possible values of the parameters are set so that $s_i^t | \left(\theta^t, \omega\right) \sim N(0,1)$ for every $i = 1,2,3$.

The statistic s_2^t represents a potential cause of s_3^t, which represents an outcome. The statistic s_1^t is a potential instrument. The researcher wants to learn ω_4, the causal effect of s_2^t on s_3^t, i.e. $Q = \{4\}$. The latent variable u^t obfuscates the causal effect because it influences s_1^t, s_2^t, and s_3^t. Since the statistic variables are all standard normal, the only aspects of the long-run distribution of s^t that the researcher can use to learn ω are $E\left[s_1^t s_2^t\right]$, $E\left[s_2^t s_3^t\right]$ and $E\left[s_1^t s_3^t\right]$. This gives three equations with five unknowns, and so ω_4 cannot be identified. However, when we make the assumption $\theta^* = 0$, we get $\omega_4 = E\left[s_1^t s_3^t\right] / E\left[s_1^t s_2^t\right]$. This is the textbook 2SLS procedure that uses s_1^t as an instrument for s_2^t. The identifying assumption is that the instrument is independent of the confounding variable u^t.

We can apply Proposition 3 to this example. The process has a structure given by the DAG

$$\begin{array}{ccccccc}
\theta^t & & & & u^t & & \omega_5 \\
 & \searrow & & \swarrow & \downarrow & \searrow & \downarrow \\
\omega_1 & \rightarrow & s_1^t & \rightarrow & s_2^t & \rightarrow & s_3^t \\
 & & & \nearrow & \uparrow & & \uparrow \\
 & & \omega_2 & & \omega_3 & & \omega_4
\end{array}.$$

The active parameters are ω_2, ω_3, ω_4 and ω_5, i.e., $Q^* = \{2, 3, 4, 5\}$, and θ and ω_{Q^*} are G-separable. Only s_1^t is not independent of θ^t conditional on the other variables: the DAG includes a link $\theta^t \rightarrow s_1^t$, and using d-separation, one can show that the other two statistics, s_2^t and s_3^t, are both independent of θ^t given s_1^t and u^t. However, s_1^t is independent of ω_{Q^*} since they have no common ancestor. By Proposition 3, the researcher's propensity to employ the instrumental-variable identification strategy is time-invariant.[7]

5. Limiting Beliefs

We now turn to the question of what the research community believes about the fixed parameters of interest in the long run. We begin with a definition of stable beliefs. Convergence of beliefs is according to the weak* topology.

Definition 1. A Borel measure μ^* is *stable* for ω^* if $\mathbb{P}(\mu_t \rightarrow \mu^* | \omega^*) > 0$.

A belief is stable when the posterior beliefs generated by the learning process converge to it with positive probability in the long run. Similar definitions appear in, e.g., Fudenberg and Kreps (1993) and Esponda and Pouzo (2016). In all our prior examples, unique stable beliefs exist for all values of the fixed parameters, and the research process converges to that stable belief with probability one.

5.1. **A General Characterization.**

The following result characterizes stable beliefs.

[7] In a previous version of the paper (Ellis and Spiegler (2024)), we examined another causal-inference identification strategy, known as "front door identification" (see Pearl (2009)), and showed that it violates the condition for time-invariant propensity to learn.

Proposition 4. *For any $\omega^* \in \Omega$, if μ^* is stable for ω^* and $\Theta^R(\cdot)$ is continuous in a neighborhood of μ^*, then $\mu^*(O) = 1$ for any open $O \subset \Omega$ with*

$$O \supset \arg\min_{\omega' \in \Omega} D_{KL}\left(p_S\left(\cdot|\theta \in \Theta^R(\mu^*), \omega^*\right) || p_S(\cdot|\theta^*, \omega')\right). \tag{3}$$

Recall that observed statistics are affected both by the contexts in which research is conducted and the true value of the fixed parameters, ω^*. If the researcher consistently holds a belief close to μ^* for a long stretch of time, then the set of values of θ for which research takes place during that stretch is close to $\Theta^R(\mu^*)$ and the long-run frequency of the statistic is close to $p_S\left(s|\theta \in \Theta^R(\mu^*), \omega^*\right)$. However, the researcher updates her belief as if the context parameter is θ^*, and under that assumption, s occurs with probability $p_S(s|\theta^*, \omega)$ with parameter ω. Following Berk (1966) and Esponda and Pouzo (2016), the long-run belief that emerges from this misspecified Bayesian learning rules out parameter values that do not minimize the KL divergence from the true distribution to the subjective one. Therefore stable beliefs attach positive probability only to the parameters close to those that minimize this divergence.

We should not confuse the KL divergence in the result with the divergence governing the decision whether to conduct research. In the latter case, the divergence plays a similar role to a utility function that captures the researcher's preferences and dictates her actions at each period. In the former case, it is a statistical property of the long-run empirical frequency of the observable variables.

In what follows, we assume that every primitive function (e.g. $\mu(\cdot)$ and $f(\cdot)$) is smooth, that $\Theta = [0, 1]$, and that $\theta^* = 0$. Our next result studies whether the stable beliefs are correct. We show they are not for "most" models in the following class, where a data-generating process p and cutoff K constitute a model.

We say that a model (p, K) is *regular* if *(i)* $Q^* \neq \{1, \ldots, n\}$; *(ii)* the divergence

$$D_{KL}\left(\int p_{S,U}(\cdot|\theta, \omega)\, dm || \int p_{S,U}(\cdot|\theta^*, \omega)\, dm\right)$$

is strictly increasing in θ for every probability measure m on Ω; and *(iii)* for every ω and some $j \in Q$, $\frac{d}{d\omega_j} p_S(s|\omega, \theta^*) \neq 0$ for some s. All of our examples have regular models. The first requirement is that the assumption shuts down at least one fixed parameter. Therefore, the assumption disregards some parameters that affect the statistics. The second requires that the context parameter captures how far the assumed distribution is from the realized distribution when distance is captured by $D(\cdot)$. This implies that there exist a function $\bar{\theta}(\cdot) > 0$ so that $\Theta^R(\mu_t) = \left[0, \bar{\theta}(\mu_t)\right]$. The function $\bar{\theta}(\mu_t)$ is continuous, so we can apply Proposition 4. The third requires that the distribution of statistics is sensitive to some parameter $j \in Q$ under θ^*.

In all of our examples, the answer to the question at which the researcher eventually arrives is incorrect with probability one. For every ω^* with $\omega_2^* \neq 0$, the researcher almost surely comes to believe that ω_Q equals $\hat{\omega}_Q \neq \omega_Q^*$ with certainty. That is, researcher's use of an identifying assumption leads to wrong beliefs and what Manski terms "incredible certitude." We say that a model exhibits *incredible certitude* if for almost every $\omega^* \in \Omega$, every stable belief μ^* for ω^* attaches probability one to some $\hat{\omega}_Q \neq \omega_Q^*$.

Proposition 5. *Suppose that $D_{KL}\left(p_S\left(\cdot|\theta,\omega^*\right)||p_S\left(\cdot|\theta^*,\omega'\right)\right)$ is strictly convex in ω'_{Q^*} for any $\theta \in \Theta$ and any $\omega^* \in \Omega$. Then, every stable belief μ^* for any $\omega^* \in \Omega$ has a degenerate marginal on ω_Q when (p, K) is regular. Moreover, the models that exhibit incredible certitude are dense within the set of regular models.*

The result shows that two features of the examples lead to incredible certitude. First, the KL divergence between the empirical distribution of statistics in ω^* and the distribution of statistics in state ω' under the identifying assumption is strictly convex in ω'_{Q^*}. Since KL divergence is a convex function of probability measures, this holds when the assumed distribution resulting from $\alpha\omega'' + (1-\alpha)\omega'$ is in between the ones resulting from ω' and ω''. This ensures that every stable belief is degenerate. Second, the examples have regular models.

The proof of the proposition shows that for any Q^*, almost all small perturbations of a regular model (within a class that leaves Q^* and θ^* unchanged) exhibit incredible certitude.[8] We construct a necessary condition for a fixed parameter to lead to correct beliefs. Then, we apply the transversality theorem to a parameterization of the set of perturbations. This establishes that almost all of them do not satisfy it for a full measure of states.

[8]Specifically, the perturbations are changes in K or adding any polynomial (with small coefficients) to the data-generating process that does not change Q^*.

5.2. **An Example: Contaminated experiments, revisited.** There may be multiple stable beliefs. In this subsection, we demonstrate this possibility using a variant on the example from Section 3.1. Let $s^t \in \{0,1\}$, $\omega = (\omega_1, \omega_2) \in \Omega = [\varepsilon, 1-\varepsilon]^2$ for $\varepsilon \in \left(0, \frac{1}{2}\right)$, and $\theta \in [0,1]$. The data-generating process is

$$p\left(s^t = 1 \mid \theta^t, \omega\right) = \left(1 - \theta^t\right)\omega_1 + \theta^t \omega_2,$$

and the researcher wants to learn ω_1, i.e., $Q = \{1\}$. The only identifying assumption is $\theta^* = 0$, which prevents learning anything about ω_2. The divergence for the assumption is an increasing function of $\bar{\mu}_1^t \equiv \mathbb{E}[\omega_1 | h^t]$, and the researcher imposes the assumption whenever $\theta^t \leq \bar{\theta}(\bar{\mu}_1^t)$ for some threshold $\bar{\theta}(\cdot) > 0$.

Suppose that $\mathbb{E}_\mu[\omega_2] = \frac{1}{2}$, and that the distribution over θ is smooth with full support and mean $\frac{1}{2}$. Any candidate for a stable belief given the true ω^* assigns probability one to a ω_1 that satisfies

$$\omega_1 = E\left[\theta | \theta \leq \bar{\theta}(\omega_1)\right] \cdot (\omega_2^* - \omega_1^*) + \omega_1^* \tag{4}$$

by Proposition 4. The next result shows that Equation (4) has at least two solutions when K is small.

Proposition 6. *For any $K > 0$, $a_1 \in \left(\frac{1}{2}, 1-\varepsilon\right)$, and $\zeta > 0$, there exists $\delta > 0$ so that for every $\omega^* \in B_\delta((a_1, 1-a_1))$, there exists an attracting solution to Equation (4) in $\left(\frac{1-\zeta}{2}, \frac{1+\zeta}{2}\right)$. Moreover, if K and ζ are sufficiently small, this δ can be chosen so that there also exists an attracting solution to Equation (4) that is greater than $\frac{1+\zeta}{2}$.*

The result establishes the possibility of multiple stable beliefs. Moreover, even if the threshold K is small, there may be stable beliefs far from the truth. For any realized ω^*, including one with ω_1^* close to $1-\varepsilon$, there is a stable belief that attaches probability one to a ω_1 close to $\frac{1}{2}$, regardless of the size of K. Put another way, the researcher becomes certain that the question's answer is about $\frac{1}{2}$ when it instead is close to $1-\varepsilon$, the maximum.

Notice that when research is conducted in period t, $\bar{\mu}_1^{t+1} > \bar{\mu}_1^t$ whenever $s^t = 1$ and $\bar{\mu}_1^{t+1} < \bar{\mu}_1^t$ whenever $s^t = 0$. Since the threshold $\bar{\theta}$ is a strictly increasing function of $\bar{\mu}_1^t$, there will be phases of both accelerating and decelerating rates of research. This is in contrast to the uniform research slowdown that emerged when the statistic was Gaussian.

6. An Extension: Multiple and "Structural" Assumptions

Our model is restrictive in two respects. First, it assumes a single feasible identifying assumption rather than a set of identification strategies that the researcher can choose from. Second, it focuses on assumptions about context parameters rather than more "structural" assumptions about the fixed parameters. In this section we present an example that goes beyond these restrictions and offers a stylized representation of a familiar identification method in empirical economics. An earlier version of this paper (Ellis and Spiegler, 2024) also considered a researcher choosing between answering multiple questions.

We consider the classic problem of drawing causal inferences from a selective sample. An observational dataset has three variables. The first (s_1) indicates whether an agent enters some market ($s_1 = 1$ means entry). It also contains the agent's income conditional on entry (s_3) and an exogenous variable (s_2) that may affect both the entry decision and the income conditional on entry. Formally, there are three statistics, s_1, s_2 and s_3, where $s_1, s_2 \in \{0, 1\}$ and $s_3 \in \mathbb{R}$. Data about income is available only for agents who enter the market.

Our researcher has two feasible identification strategies to deal with selection. First, she can make a contextual assumption that market entry is purely random, thus assuming away selective entry. Second, she can make an assumption about the fixed parameters in the manner of "Heckman correction" (Heckman, 1979). We explore the trade-off between the two methods and how it affects research dynamics.

Formally, the true data-generating process is given by the following equations. First, s_2^t is uniformly and independently distributed over $\{0, 1\}$. Second,

$$s_1^t = \begin{cases} \mathbb{I}_+\left(s_2^t + u^t\right) & \text{with probability } \theta^t \\ \mathbb{I}_+\left(s_2^t + \varepsilon_1^t\right) & \text{with probability } 1 - \theta^t \end{cases}$$

where $\mathbb{I}_+(x) = 1$ if $x \geq 0$ and otherwise equals 0. Finally, given $s_1^t = 1$,

$$s_3^t = \omega_1 + \omega_2 s_2^t + \omega_3 \mathbb{E}\left[u^t | s_1^t = 1, s_2^t, \theta^t\right] + \varepsilon_2^t$$

where u^t, ε_1^t and ε_2^t are all independent normal variables with mean zero, and where the variances of u^t and ε_1^t are the same. The statistic s_3^t is not measured when $s_1^t = 0$. The context parameter $\theta^t \in [0,1]$ indicates the probability that an agent's assignment into the market is based on the agents' latent characteristics. Thus, $\theta^t = 0$ means purely random, non-selective assignment.

There are three fixed parameters in this specification, all of which enter the equation for s_3^t. These parameters represent the causal effects of three factors on agents' income: market entry itself (ω_1), the exogenous variable s_2^t (ω_2), and the latent variable u^t (ω_3). The researcher is interested in learning ω_1, i.e., $Q = \{1\}$. Long-run observation of s_3^t for each s_2^t provides two equations with three unknowns, hence ω_1 cannot be identified unless the researcher imposes an assumption. Parameterize beliefs μ so that $\omega_i \sim N(m_i, \sigma_i^2)$.

There are two feasible identifying assumptions. One assumption is $\theta^* = 0$, i.e., market entry is independent of u^t in the dataset. Under this assumption, $\mathbb{E}[u^t | s_1^t = 1, s_2^t] = 0$ for every s_2^t, and the long-run average of s_3^t given $s_1^t = 1$ and s_2^t is equal to $\omega_1 + \omega_2 s_2^t$. This gives two equations with two unknowns and enables the researcher to pin down ω_1.

Alternatively, the researcher could assume that the fixed parameter ω_2 equals $\omega_2^* = 0$. That is, the researcher assumes that in the dataset, the exogenous variable that may affect market entry does not have a direct causal effect on income conditional on entry. It is an "exclusion" restriction that turns s_2^t into a valid instrument for estimating ω_1, albeit with different parameterization than in the IV example we examined in Section 5.

The second identification method is based on Heckman's correction method (Heckman, 1979). For the sake of tractability, we simplified the model by admitting no fixed parameters into the distribution of s_1^t conditional on s_2^t. This enables us to treat $\mathbb{E}[u^t | s_1^t = 1, s_2^t]$ as a known quantity, whereas in practice it would be an estimated one. Our example thus trivializes the first stage of Heckman's procedure, and focuses on the second stage.

In any given period, the researcher selects the KL divergence minimizing assumption ($\theta^* = 0$ or $\omega_2^* = 0$), as long as this divergence does not exceed the constant K.

Under the assumption $\theta^* = 0$, the researcher thinks the statistic is independent of ω_3 and therefore does not update her beliefs about it. Similarly, under the assumption $\omega_2^* = 0$, she does not update her beliefs about ω_2.

The following result characterizes the researcher's selection strategy.

Proposition 7. *For almost every history h^t, there exist thresholds*

$$0 < \bar{\theta}^{RD}\left(\mu\left(h^t\right)\right) \leq \bar{\theta}^{S}\left(\mu\left(h^t\right)\right) \leq 1$$

such that the researcher assumes $\theta^* = 0$ when $\theta^t \in \left[0, \bar{\theta}^{RD}\left(\mu\left(h^t\right)\right)\right)$; assumes $\omega_2^* = 0$ when $\theta^t \in \left(\bar{\theta}^{S}\left(\mu\left(h^t\right)\right), 1\right)$; and passes when $\theta^t \in \left(\bar{\theta}^{RD}\left(\mu\left(h^t\right)\right), \bar{\theta}^{S}\left(\mu\left(h^t\right)\right)\right)$. The thresholds $\bar{\theta}^{RD}\left(\mu\left(h^t\right)\right)$ and $\bar{\theta}^{S}\left(\mu\left(h^t\right)\right)$ increase in $\left(\mathbb{E}_{\mu(h^t)}(\omega_2)\right)^2$ and $Var_{\mu(h^t)}(\omega_2)$, and decrease in $\left(\mathbb{E}_{\mu(h^t)}(\omega_3)\right)^2$. If K is large enough, then $\bar{\theta}^{RD}\left(\mu\left(h^t\right)\right) = \bar{\theta}^{S}\left(\mu\left(h^t\right)\right)$.

Thus, when market entry exhibits little selection (i.e., θ is small), the researcher employs the contextual assumption $\theta^* = 0$. In contrast, when entry is highly selective, the researcher passes or imposes the assumption $\omega_2^* = 0$. Her willingness to impose the latter assumption increases with its perceived accuracy (i.e., as $\mathbb{E}(\omega_2)$ gets closer to zero) and with her confidence of her estimate — i.e., as the variance of her belief over ω_2 goes down. Finally, the researcher is less likely to employ the contextual assumption when she believes that selective entry has a large effect on income (i.e., when $\mathbb{E}(\omega_3)$ is far from zero).

This setting has self-reinforcing learning dynamics. The researcher never updates her beliefs about ω_3 when she assumes $\theta^* = 0$. Likewise, she never updates her beliefs about ω_2 when she assumes $\omega_2^* = 0$. When she is confident that ω_2 is close to zero, she usually assumes $\omega_2^* = 0$ and rarely updates her belief over ω_2. Therefore, if this belief is inaccurate, it will take a long time to correct it. Moreover, when the researcher assumes $\omega_2^* = 0$, she misattributes part of the actual effect of ω_2 on income to ω_3. Depending on the true values of these parameters, this misattribution can make the researcher even less likely to employs the contextual assumption. Similarly, if the researcher is confident that ω_3 is low, she tends to assume $\theta^* = 0$. This leads her to misattribute part of the actual effect of ω_3 on income to ω_2, which may further strengthen her tendency to employ the contextual assumption. Thus, the researcher's predilection to stick to a particular identifying strategy for a long stretch of time is history-dependent.

7. Assumption-Based Learning vs. (Misspecified) Bayesian Learning

We conclude with a discussion of how assumption-based learning differs from Bayesian learning, both with and without misspecified prior beliefs. Bayesian learning corresponds to always taking $a^t = 1$ and updating according to Bayes rule, i.e. $\mu(\omega|h^t, s^t) = \frac{\mu(\omega|h^t) p(s^t|\omega,\theta^t)}{p(s^t|h^t,\theta^t)}$. Misspecified Bayesian learning is the limiting case of our model where $K = \infty$ and $\mu(\omega) p_u(u^t) p(s^t|\theta^*, u^t, \omega)$ is the misspecified prior. We draw a contrast between the two and assumption-based learning.

In the most comparable interpretation of misspecified Bayesian learning, a researcher makes an assumption (before the learning begins) that determines her prior beliefs about the data-generating process. She then imposes it at every research opportunity without reconsidering its appropriateness. Under assumption-base learning, the researcher re-evaluates the assumption's plausibility in every period against her current beliefs and the current context. The latter generates non-trivial predictions regarding the evolution of the rate of learning. Moreover, stable beliefs are an equilibrium object, as evident from the characterization in Proposition 4. Under misspecified Bayesian learning, experiments take place in every period, so all equilibrium effects vanish. Stable beliefs are still given by (3) with $\Theta^R(\mu^*) = \Theta$ for every μ^*, which delivers the original characterization by Berk (1966).

This highlights a difference in the assumptions' epistemic status between assumption-based and misspecified Bayesian learning. In the former, assumptions do not reflect the researcher's genuine beliefs. Instead, they are "working hypotheses" in service of her mission to infer the fixed parameters from a particular dataset. For instance, in the Contaminated Experiment example, the researcher does not assume $\theta^* = 0$ because she "truly" believes the interference does not exist. Rather, she makes it because it enables her to pin down the effect of interest from the data. If the researcher could somehow be informed of the true value of the interference fixed parameter ω_2 (say, by a "dry run" that calibrates the measurement instrument), she would not impose the assumption $\theta^* = 0$ because this would not be required to identify ω_1.

The models also differ in how predictably the researcher's subjective beliefs evolve. Correctly specified Bayesian learning satisfies the Martingale property: the expectation of $\mu_{t+1}(E)$ equals to $\mu_t(E)$ for all events E conditional on any history. Under misspecified Bayesian learning, this property generally fails. A Bayesian outsider with the same prior as the researcher who observes and correctly utilizes all information expects the researcher's beliefs to violate it. However, a misspecified Bayesian learner expects her own beliefs to satisfy the Martingale property because she does not question whether her model is misspecified. In contrast, an assumption-based learner would expect to violate the property. In other words, both the objective and subjective expectations over the researcher's posterior belief depart from her prior. If she were to form a belief at period t regarding her beliefs at period $t+1$, she would use the true context θ^t to calculate the distribution over $\mu_{t+1}(E)$. Yet, the posteriors themselves are arrived at by updating μ_t as if $\theta^* = 0$.

Appendix A. Proofs

For the proofs of Propositions 1, 2, and 4, we economize on notation by taking a history h^t and writing (h^t, s) for the history that concatenates h^t with the tuple $(s^t = s, a^t = 1, \theta^t)$ for arbitrary $\theta^t \in \Theta^R(\mu(h^t))$ (similarly for $(h^t, s, s', s'', ...)$). In addition, for any θ, h, s, let $q(\theta, h)$ and $q(\theta, h, s)$ serve as shorthand notation for the conditional distributions $p_{S,U}(\cdot|\theta, h)$ and $p_{S,U}(\cdot|\theta,(h,s))$, respectively.

A.1. **Proof of Proposition 1.** We begin the proof by drawing a simple implication from the fact that f-divergences are convex (see page 56 in Amari and Nagaoka (2000); it follows from arguments similar to Theorem 2.7.2 of Cover and Thomas (2006)).

The researcher always updates her belief as if the assumption θ^* is correct. Therefore, her belief over variables satisfies the Martingale property with respect to the distribution over statistics conditional on the assumption. That is, for every history h^τ, context θ^τ, and statistic realization s^τ,

$$\sum_{s^\tau \in S} p\left(s^{\tau+1}, u^{\tau+1}|\theta^{\tau+1}, (h^\tau, s^\tau)\right) p\left(s^\tau|\theta^*, h^\tau\right) = p\left(s^{\tau+1}, u^{\tau+1}|\theta^{\tau+1}, h^\tau\right). \tag{5}$$

For completeness, we derive this equation in greater detail. The researcher's belief over the variable realizations at $\tau+1$ given $\theta^{\tau+1}, (h^\tau, s^\tau)$ is

$$p\left(s^{\tau+1}, u^{\tau+1} | \theta^{\tau+1}, (h^\tau, s^\tau)\right) = \int_\omega p\left(s^{\tau+1}, u^{\tau+1} | \theta^{\tau+1}, \omega\right) \mu\left(\omega | h^\tau, s^\tau\right) d\omega.$$

Since the researcher always updates her beliefs as if $\theta = \theta^*$, as given by Eq (1),

$$p\left(s^{\tau+1}, u^{\tau+1} | \theta^{\tau+1}, (h^\tau, s^\tau)\right) = \int_\omega p\left(s^{\tau+1}, u^{\tau+1} | \theta^{\tau+1}, \omega\right) \frac{p\left(s^\tau | \theta^*, \omega\right)}{p\left(s^\tau | \theta^*, h^\tau\right)} \mu\left(\omega | h^\tau\right) d\omega.$$

Multiplying by $p\left(s^\tau | \theta^*, h^\tau\right)$ and summing across S, we obtain

$$\begin{aligned}
&\sum_{s^\tau \in S} p\left(s^{\tau+1}, u^{\tau+1} | \theta^{\tau+1}, (h^\tau, s^\tau)\right) p\left(s^\tau | \theta^*, h^\tau\right) \\
&= \sum_{s^\tau \in S} \int_\omega p\left(s^{\tau+1}, u^{\tau+1} | \theta^{\tau+1}, \omega\right) p\left(s^\tau | \theta^*, \omega\right) \mu\left(\omega | h^\tau\right) d\omega \\
&= \int_\omega p\left(s^{\tau+1}, u^{\tau+1} | \theta^{\tau+1}, \omega\right) \sum_{s^\tau \in S} p\left(s^\tau | \theta^*, \omega\right) \mu\left(\omega | h^\tau\right) d\omega \\
&= \int_\omega p\left(s^{\tau+1}, u^{\tau+1} | \theta^{\tau+1}, \omega\right) \mu\left(\omega | h^\tau\right) d\omega = p\left(s^{\tau+1}, u^{\tau+1} | \theta^{\tau+1}, h^\tau\right).
\end{aligned}$$

By Eq (5) and the convexity of f-divergence,

$$D\left(q(\theta^\tau, h^\tau) || q(\theta^*, h^\tau)\right) \leq \sum_{s^\tau \in S} D\left(q(\theta^\tau, h^\tau, s^\tau) || q(\theta^*, h^\tau, s^\tau)\right) p\left(s^\tau | \theta^*, h^\tau\right).$$

Our next step is to show that if the divergence drops after some statistic realization, this implies a lower bound on its increase under some alternative realization. For any history h and statistic realization s, let

$$\delta\left(\theta, h, s\right) \equiv D\left(q(\theta, h) || q(\theta^*, h)\right) - D\left(q(\theta, h, s) || q(\theta^*, h, s)\right)$$

be the change in the divergence given θ when the history h is concatenated by s.

Claim 1. If $\delta\left(\theta, h, s\right) > 0$, then there exists s' such that

$$D\left(q(\theta, h, s') || q(\theta^*, h, s')\right) \geq D\left(q(\theta, h) || q(\theta^*, h)\right) + p(s | \theta^*, h) \delta\left(\theta, h, s\right).$$

Proof. Denote $q(\theta, h, -s) = p_{S,U}\left(\cdot | \theta, \{(h, s')\}_{s' \neq s}\right)$. By convexity of f-divergence,

$$\begin{aligned}
&D\left(q(\theta, h) || q(\theta^*, h)\right) \\
\leq & p(s | \theta^*, h) D\left(q(\theta, h, s) || q(\theta^*, h, s)\right) + \left(1 - p(s | \theta^*, h)\right) D\left(q(\theta, h, -s) || q(\theta^*, h, -s)\right)
\end{aligned}$$

This inequality can be rewritten as

$$\begin{aligned} & D\left(q(\theta, h, -s)||q(\theta^*, h, -s)\right) - D\left(q(\theta, h)||q(\theta^*, h)\right) \\ \geq & p(s|\theta^*, h)\left[D\left(q(\theta, h, -s)||q(\theta^*, h, -s)\right) - D\left(q(\theta, h, s)||q(\theta^*, h, s)\right)\right]. \end{aligned}$$

Moreover, since $\delta(\theta, h, s) > 0$, the R.H.S. of this inequality is weakly above

$$p(s|\theta^*, h)\left[D\left(q(\theta, h)||q(\theta^*, h)\right) - D\left(q(\theta, h, s)||q(\theta^*, h, s)\right)\right] = p(s|\theta^*, h)\delta(\theta, h, s).$$

By convexity of f-divergence and the definition of $q(\theta, h, -s)$, there is $s' \neq s$ such that

$$D\left(q(\theta, h, s')||q(\theta^*, h, s')\right) \geq D\left(q(\theta, h, -s)||q(\theta^*, h, -s)\right),$$

which completes the proof. □

Fix a history h^{t+1} and θ^{t+1} so that $\theta^{t+1} \in \Theta^R\left(\mu\left(h^{t+1}\right)\right) \setminus \Theta^R\left(\mu\left(h^t\right)\right)$. Consider any h^t and s^* such that

$$D\left(q\left(\theta, h^t, s^*\right)||q\left(\theta^*, h^t, s^*\right)\right) \leq K < D\left(q\left(\theta, h^t\right)||q\left(\theta^*, h^t\right)\right).$$

Assume the proposition is false. Then, then for any $T \geq t+2$ and any continuation of statistic realizations $\left(s^{t+2}, s^{t+3}, \ldots, s^T\right)$, we have

$$D\left(q\left(|\theta, \left(h^t, s^*, s^{t+2}, s^{t+3}, \ldots, s^T\right)\right)||q\left(\theta^*, \left(h^t, s^*, s^{t+2}, s^{t+3}, \ldots, s^T\right)\right)\right) \leq K.$$

Since Bayesian posterior beliefs are invariant to the ordering of a sequence of conditionally independent signals, this inequality can be equivalently rewritten as

$$D\left(q\left(\theta, \left(h^t, s^{t+2}, s^{t+3}, \ldots, s^T, s^*\right)\right)||q\left(\theta^*, \left(h^t, s^{t+2}, s^{t+3}, \ldots, s^T, s^*\right)\right)\right) \leq K. \tag{6}$$

We inductively construct a sequence of histories $\bar{h}^T$ for $T \geq t$, starting with $\bar{h}^t = h^t$, whose divergence goes to infinity. Inductively assume that

$$D\left(q\left(\theta, \bar{h}^T\right)||q\left(\theta^*, \bar{h}^T\right)\right) \geq D\left(q\left(\theta, h^t\right)||q\left(\theta^*, h^t\right)\right) > K, \tag{7}$$

which is true by assumption for $T = t$. Then, by Equation (6),

$$D\left(q\left(\theta, \left(\bar{h}^T, s^*\right)\right)||q\left(\theta^*, \left(\bar{h}^T, s^*\right)\right)\right) \leq K$$

By the claim, we can find s^{T+1} such that

$$\begin{aligned}&D\left(q\left(\theta,\left(\bar{h}^T,s^{T+1}\right)\right)||q\left(\theta^*,\left(\bar{h}^T,s^{T+1}\right)\right)\right)\\ \geq &D\left(q\left(\theta,\bar{h}^T\right)||q\left(\theta^*,\bar{h}^T\right)\right)+p\left(s^*|\theta^*,\bar{h}^T\right)\delta\left(\theta,\bar{h}^T,s^*\right).\end{aligned}$$

Moreover, by Equation (7),

$$\delta\left(\bar{h}^T,s^*\right)\geq D\left(q\left(\theta,h^t\right)||q\left(\theta^*,h^t\right)\right)-K\equiv\delta^*>0 \tag{8}$$

and

$$p\left(s^*|\theta^*,\bar{h}^T\right)\geq\min_{\omega}p\left(s^*|\theta^*,\omega\right)\equiv p^*>0.$$

Set $\bar{h}^{T+1}=\left(\bar{h}^T,s^{T+1}\right)$. Applying the inductive step, we conclude that

$$D\left(q\left(\theta,\bar{h}^{T+1}\right)||q\left(\theta^*,\bar{h}^{T+1}\right)\right)\geq D\left(q\left(\theta,\bar{h}^T\right)||q\left(\theta^*,\bar{h}^T\right)\right)+p^*\delta^*>K.$$

This means that as $T\to\infty$,

$$D\left(p_{S,U}\left(\cdot|\theta,\bar{h}^T\right)||p_{S,U}\left(\cdot|\theta^*,\bar{h}^T\right)\right)\to\infty.$$

However, by convexity of D, the L.H.S. is bounded from above by

$$\max_{\omega\in\Omega}D\left(p_{S,U}\left(\cdot|\theta,\omega\right)||p_{S,U}\left(\cdot|\theta^*,\omega\right)\right)\leq\infty,$$

a contradiction.□

A.2. **Proof of Proposition 2.** Fix any h^t and s so that $\Theta^R\left(\mu\left(h^t,s\right)\right)\setminus\Theta^R\left(\mu\left(h^t\right)\right)\neq\emptyset$. Pick $\theta^1\in\Theta^R\left(\mu\left(h^t,s\right)\right)\setminus\Theta^R\left(\mu\left(h^t\right)\right)$ for which

$$D\left(q\left(\theta^1,h^t\right)||q\left(\theta^*,h^t\right)\right)>K>D\left(q\left(\theta^1,h^t,s\right)||q\left(\theta^*,h^t,s\right)\right)=\Delta.$$

By continuity, there exists $\theta=\beta\theta^1+(1-\beta)\theta^*$ so that

$$D\left(q\left(\theta,h^t\right)||q\left(\theta^*,h^t\right)\right)=K,$$

and so $\theta\in\Theta^R\left(\mu(h^t)\right)$. By quasi-convexity, $D\left(q\left(\theta,h^t,s\right)||q\left(\theta^*,h^t,s\right)\right)<K$. By convexity of f-divergence (page 56, Amari and Nagaoka (2000)),

$$
\begin{aligned}
K &= D\left(q\left(\theta, h^t\right)||q\left(\theta^*, h^t\right)\right) \\
&\leq \sum_{s'} D\left(q\left(\theta, h^t, s'\right)||q\left(\theta^*, h^t, s'\right)\right)p(s'|\theta^*, h^t) \\
&< \sum_{s'\neq s} D\left(q\left(\theta, h^t, s'\right)||q\left(\theta^*, h^t, s'\right)\right)p(s'|\theta^*, h^t) + \Delta p(s|\theta^*, h^t)
\end{aligned}
$$

Therefore,

$$
D\left(q\left(\theta, h^t, s''\right)||q\left(\theta^*, h^t, s''\right)\right) > K
$$

for some $s'' \in S\backslash\{s\}$. This also holds for all θ' sufficiently close to θ, including some of those for which $D\left(q\left(\theta', h^t\right)||q\left(\theta^*, h^t\right)\right) < K$. Therefore, $\Theta^R\left(\mu\left(h^t\right)\right)\backslash\Theta^R\left(\mu\left(h^{t+1}\right)\right) \neq \emptyset$ with a probability of at least $p\left(s'', \Theta^R\left(\mu\left(h^t\right)\right)|\omega\right)$, for any given value ω of the fixed parameters and any history h^t. $\square$

A.3. **Proof of Proposition 3.** Let the DAG $G = (N, R)$ be the structure of p. We introduce a few pieces of DAG-based notation. First, just as N^s is the set of nodes that represent statistics, N^ω is the set of nodes that represent fixed parameters. Define N^u and N^θ in the same manner. Recall that by definition, the nodes in N^ω, N^θ and N^u are ancestral. For convenience, we will sometimes abuse notation and identify ω_i, s_i, u_i, and θ_i with their corresponding nodes. The proof proceeds step-wise.

Step 1: The researcher never learns anything about ω_{-Q^*}.

By definition of Q^*, $p(s|\theta^*, \omega_{Q^*}, \omega_{-Q^*}) = p(s|\theta^*, \omega_{Q^*})$ for every s, so for almost every ω,

$$
\begin{aligned}
\mu\left(\omega_{-Q^*}|h^t, s^t, a^t = 1, \theta^t\right) &= \frac{\int p\left(s^t|\theta^*, \omega\right)\mu\left(\omega_{Q^*}|h^t\right)\mu\left(\omega_{-Q^*}|h^t\right)d\omega_{Q^*}}{\int p\left(s^t|\theta^*, \omega'\right)\mu\left(\omega'|h^t\right)d\omega'} \\
&= \frac{\int p\left(s^t|\theta^*, \omega\right)\mu\left(\omega_{Q^*}|h^t\right)d\omega_{Q^*}}{\int p\left(s^t|\theta^*, \omega'_Q\right)\mu\left(\omega'_{Q^*}|h^t\right)d\omega'_{Q^*}}\mu\left(\omega_{-Q^*}|h^t\right) = \mu\left(\omega_{-Q^*}|h^t\right)
\end{aligned}
$$

Therefore, beliefs about almost every ω_{-Q^*} are history-independent.

In preparation for the next step, define a subset $I \subseteq N^\omega$ consisting of all the parameters that are not independent of θ conditional on (s, u) in the following recursive manner. First, I_0 is the set of nodes $i \in N^\omega$ for which there exist $j \in N^\theta$ and $k \in N^s$ such that $i, j \in R(k)$. For every integer $n > 0$, I_n is the set of nodes $i \in N^\omega$ for which there exist $j \in I_{n-1}$ and $k \in N^s$ such that $i, j \in R(k)$. Define $I = \cup_{n\geq 0} I_n$.

Let N^I be the nodes in N^s with a parent in I. By construction, $j \in N^I$ implies that $R(j) \cap N^\omega \subset I$, whereas $j \notin N^I$ implies that $R(j) \cap I = \emptyset$.

Step 2: $I \cap Q^* = \emptyset$.

For contradiction, suppose that $\omega_i \in I \cap Q^*$. Then, there is a sequence $\omega_{i_1}, \ldots, \omega_{i_n}$ of structural-parameter nodes and a sequence $s_{i_1}, \ldots, s_{i_n}$ of statistics nodes, such that: $\omega_i = \omega_{i_1}$; every node s_{i_k} along the sequence ($k = 1, ..., n-1$) is a child of ω_{i_k} and $\omega_{i_{k+1}}$; and s_{i_n} is a child of a node in N^θ. The following diagram illustrates this configuration for $n = 3$.

$$\begin{array}{ccccccc} \omega_{i_1} & & \omega_{i_2} & & \omega_{i_3} & & \theta_j \\ \downarrow & \swarrow & \downarrow & \swarrow & \downarrow & \swarrow & \\ s_{i_1} & & s_{i_2} & & s_{i_3} & & \end{array}$$

We show that G does not satisfy the conditional-independence property $s_{i_1} \perp \theta | (s_{-i_1}, u)$. By a basic result in the Bayesian-network literature (e.g., Koller and Friedman (2009)), this property has a simple graphical characterization, known as d-separation. Perform the following two-step procedure.[9] First, take every triple of nodes i, j, k such that $i, j \in R(k)$ whereas i and j are not linked. Modify the DAG by connecting i and j. Second, remove the directionality of all links in the modified graph, such that it becomes a non-directed graph. In this so-called "moral graph" induced by G, check whether every path between s^{i_1} and a node in N^θ contains a node in $N^s \cup N^u$. This is not the case, by construction, as the moral graph contains a path from s_{i_1} to θ that goes through the nodes $\omega_{i_1}, \ldots, \omega_{i_n}$. For illustration, note that procedure generates the following moral graph from the DAG above:

$$\begin{array}{ccccccc} \omega_{i_1} & - & \omega_{i_2} & - & \omega_{i_3} & - & \theta_j \\ | & / & | & / & | & / & \\ s_{i_1} & & s_{i_2} & & s_{i_3} & & \end{array}$$

It follows that $s_{i_1} \not\perp \theta | (s_{-i_1}, u)$. By hypothesis, this implies $s_{i_1} \perp \omega_{Q^*}$, and hence $s_{i_1} \perp \omega_i$ (because ω_i is in Q^*). Since s_{i_1} is a child of ω_i, this property is violated, a contradiction. Therefore, we conclude that I and Q^* are disjoint.

[9] In general, there is a preliminary step, in which all nodes that appear below the nodes that represent ω_i, θ, s, u are removed. Since there are no such nodes in G, this step is vacuous.

Step 3: For every s, u, the likelihood ratio $p(s, u|\theta^t, h^t)/p(s, u|\theta^*, h^t)$ is history-independent.

For every $j \in N^s$, we use $(s, u, \omega, \theta)_{R(j)}$ to denote the values of the variables and parameters that are represented by the nodes in $R(j)$. Then, $p(s, u|\theta^t, h^t) = p(u)p(s|\theta^t, h^t, u)$ and we can write $p(s^t|\theta^t, h^t, u^t)$ equals

$$\int \prod_{j\in N^s} p\left(s_j^t|\left(s^t, u^t, \theta^t, \omega\right)_{R(j)}\right)\mu(\omega|h^t)d\omega$$
$$= \int\int \prod_{j\in N^I} p\left(s_j^t|\left(s^t, u^t, \theta^t, \omega\right)_{R(j)}\right)\prod_{j\notin N^I} p\left(s_j^t|\left(s^t, u^t, \theta^t, \omega\right)_{R(j)}\right)\mu(\omega_{-I}|h^t)\mu(\omega_I|h^t)d\omega_I d\omega_{-I}$$
$$= \int \prod_{j\in N^I} p\left(s_j^t|\left(s^t, u^t, \theta^t, \omega\right)_{R(j)}\right)\mu(\omega_I|h^t)d\omega_I \int \prod_{j\notin N^I} p\left(s_j^t|\left(s^t, u^t, \theta^t, \omega\right)_{R(j)}\right)\mu(\omega_{-I}|h^t)d\omega_{-I}$$

where the second equality follows from the relationship between N^I and I we articulated above.

By the definition of N^I, $p\left(s_j^t|\left(s^t, u^t, \theta^t, \omega\right)_{R(j)}\right)$ is constant in θ^t for every $j \notin N^I$. By Step 2, $Q^* \cap I = \emptyset$. By Step 1, $\mu(\omega_I|h^t)$ is constant in h^t. It follows that we can write the likelihood ratio as

$$\frac{p(s^t, u^t|\theta^t, h^t)}{p(s^t, u^t|\theta^*, h^t)} = \frac{\int \prod_{j\in N^I} p\left(s_j^t|\left(s^t, u^t, \theta^t, \omega\right)_{R(j)}\right)\mu(\omega_I)d\omega_I}{\int \prod_{j\in N^I} p\left(s_j^t|\left(s^t, u^t, \theta^*, \omega\right)_{R(j)}\right)\mu(\omega_I)d\omega_I}$$

because the other multiplicative terms in $p(s, u|\theta)$ cancel out. Therefore, the likelihood ratio is history-independent.

Step 4: Completing the proof

Let $R\left(N^I\right) = \bigcup_{j\in N^I} R(j)$. Suppose s_k is in N^I. As we saw in the proof of Step 2, s_k is not independent of θ conditional on (s_{-k}, u). By hypothesis, $s_k \perp \omega_{Q^*}$. This means that s_k cannot be a descendant of any node in ω_{Q^*} according to G. It follows that the parents of s_k also cannot be descendants of nodes in ω_{Q^*}. Therefore, for every s_j node in $N^I \cup R\left(N^I\right)$, $p\left(s_j^t|\left(s^t, u^t, \theta^*, \omega\right)_{R(j)}\right)$ is constant in ω_{Q^*}, and so by Step 1, $p\left(s_{N^I\cup R(N^I)}|h^t, \theta^*\right) = p\left(s_{N^I\cup R(N^I)}|\theta^*\right)$ for every history h^t.

Note that $D(p_{S,U}(\cdot|h^t,\theta^t)||p_{S,U}(\cdot|h^t,\theta^*))$ equals

$$\sum_{(s,u)} p(u^t)p(s^t|h^t,\theta^*,u)f\left(\frac{p(s^t,u^t|\theta^t,h^t)}{p(s^t,u^t|\theta^*,h^t)}\right)$$

$$=\sum_{(s,u)} p(u^t)p(s^t|h^t,\theta^*,u)f\left(\frac{\int\prod_{j\in N^I} p\left(s_j^t|\left(s^t,u^t,\theta^t,\omega\right)_{R(j)}\right)\mu(\omega_I)d\omega_I}{\int\prod_{j\in N^I} p\left(s_j^t|\left(s^t,u^t,\theta^*,\omega\right)_{R(j)}\right)\mu(\omega_I)d\omega_I}\right)$$

using the simplified expression for the likelihood ratio that we derived at the end of the proof of Step 3. The only s variables it involves are those represented by nodes in $N^I\cup R(N^I)$. Therefore, the likelihood ratio is independent of $s_{-(N^I\cup R(N^I))}$. It follows that for each u, when we sum over the values of $s_{-(N^I\cup R(N^I))}$, their contributions to the divergence are integrated out, and we can replace $p(s|h^t,\theta^*,u)$ with $p\left(s_{N^I\cup R(N^I)}|\theta^*,u\right)$ in the expression above. We have already observed that the likelihood ratio is history-independent for every s_{N^I}, as is the distribution of $s_{N^I\cup R(N^I)}$. Therefore, the divergence simplifies into the following history-independent expression

$$\sum_{(s,u)} p(u^t)p\left(s^t_{N^I\cup R(N^I)}|\theta^*,u\right)f\left(\frac{\int\prod_{j\in N^I} p\left(s_j^t|\left(s^t,u^t,\theta^t,\omega\right)_{R(j)}\right)\mu(\omega_I)d\omega_I}{\int\prod_{j\in N^I} p\left(s_j^t|\left(s^t,u^t,\theta^*,\omega\right)_{R(j)}\right)\mu(\omega_I)d\omega_I}\right),$$

completing the proof. □

A.4. **Proof of Proposition 4.** Consider any stable density μ. There is a positive probability set of histories H so that $\mu\left(\cdot|h^t\right)\rightarrow\mu$ for every $h\in H$. Denote by $\mu_{t+1}(\cdot)$ the Borel measure with density $\mu\left(\cdot|h^t\right)$ and μ^* the Borel measure with density μ; note μ_{t+1} is implicitly a function of h^t. By the Portmanteau theorem, μ_{t+1} weak* converges to μ^* for every $h\in H$. Take any open set $O\supset\arg\min_{\omega'\in\Omega} D_{KL}\left(p_S\left(\cdot|\theta\in\Theta^R\left(\mu^*\right),\omega^*\right)||p_S\left(\cdot|\theta^*,\omega'\right)\right)$ and let $C=\Omega\setminus O$, noting C is compact as a closed subset of Ω.

Pick any $w\in C$ and any $\hat{w}\in\arg\min_{\omega'\in\Omega} D_{KL}\left(p_S\left(\cdot|\theta\in\Theta^R\left(\mu^*\right),\omega^*\right)||p_S\left(\cdot|\theta^*,\omega'\right)\right)$. For any $s\in S$ and states $\omega,\omega'\in\Omega$, define

$$\Delta\left(s,\omega,\omega'\right)\equiv p\left(s|\theta^*,\omega\right)-p\left(s^t|\theta^*,\omega'\right).$$

We show there is an open $E\ni w$ with $\mu^*(E)=0$. By continuity and that S is finite, we can find open sets $\hat{E}\ni\hat{w}$ and $E\ni w$ so that $|\Delta\left(s,\omega,\hat{w}\right)|,|\Delta\left(s,\omega',w\right)|<\epsilon$ for all

$s \in S$, $\omega \in E$, and $\omega' \in \hat{E}$, and where $\epsilon > 0$ satisfies

$$\sum_{s \in S} p\left(s|\Theta^R\left(\mu^*\right), \omega^*\right)\left[\ln\left(\frac{p\left(s|\theta^*, \hat{w}\right) - \epsilon}{p\left(s|\theta^*, w\right) + \epsilon}\right)\right] > 0.$$

Such an ϵ exists since $\hat{w}$ minimizes divergence, because then

$$\begin{aligned} 0 >& D_{KL}\left(p\left(s|\Theta^R\left(\mu^*\right), \omega^*\right) || p(s|\theta^*, \hat{w})\right) - D_{KL}\left(p\left(s|\Theta^R\left(\mu^*\right), \omega^*\right) || p(s|\theta^*, w)\right) \\ =& \sum_{s \in S} p\left(s|\Theta^R\left(\mu^*\right), \omega^*\right)\left[\ln\left(\frac{p\left(s|\theta^*, w\right)}{p\left(s|\theta^*, \hat{w}\right)}\right)\right]. \end{aligned}$$

Now, for every history

$$\begin{aligned} \mu_{t+1}\left(\hat{E}\right) &= p\left(s^t|\theta^*, h^t\right)^{-1} \int_{\omega \in \hat{E}} p\left(s^t|\theta^*, \omega\right) d\mu_t \\ &= p\left(s^t|\theta^*, h^t\right)^{-1} \int_{\omega \in \hat{E}} \left\{p\left(s^t|\theta^*, \hat{w}\right) + \Delta\left(s^t, \omega, \hat{w}\right)\right\} d\mu_t \end{aligned}$$

and so

$$\mu_{t+1}\left(\hat{E}\right) p\left(s^t|\theta^*, h^t\right) \mu_t\left(\hat{E}\right)^{-1} - p\left(s^t|\theta^*, \hat{w}\right) \in \left(\inf_{s, \omega \in \hat{E}} \Delta\left(s, \omega, \hat{w}\right), \sup_{s, \omega \in \hat{E}} \Delta\left(s, \omega, \hat{w}\right)\right).$$

Similarly,

$$\mu_{t+1}\left(E\right) p\left(s^t|\theta^*, h^t\right) \mu_t\left(E\right)^{-1} - p\left(s^t|\theta^*, w\right) \in \left(\inf_{s, \omega \in E} \Delta\left(s, \omega, w\right), \sup_{s, \omega \in E} \Delta\left(s, \omega, w\right)\right).$$

Therefore, for every t there exist $\delta_t, \hat{\delta}_t \in (-\epsilon, \epsilon)$ so that

$$\frac{\mu_{t+1}(\hat{E})}{\mu_{t+1}(E)} = \frac{\mu_t(\hat{E})}{\mu_t(E)} \frac{p\left(s^t|\theta^*, \hat{w}\right) + \hat{\delta}_t}{p\left(s^t|\theta^*, w\right) + \delta_t}$$

when s^t occurs and $\theta^t \in \Theta^R\left(\mu\left(h^t\right)\right)$.

We claim that $\mu^*(E) = 0$. Suppose not, so $\mu^*(E) > 0$ and so $\mu_0(E) > 0$. In the history $h^t = (a^1, s^1, \theta^1; a^2, s^2, \theta^2; \ldots; a^{t-1}, s^{t-1}, \theta^{t-1})$ we have

$$\begin{aligned} \ln\frac{\mu_{t+1}(\hat{E})}{\mu_{t+1}(E)} &= \ln\frac{\mu_t(\hat{E})}{\mu_t(E)} + \mathbb{I}_{\Theta^R(\mu(h^t))}(\theta^t) \ln\frac{p\left(s^t|\theta^*, \hat{w}\right) + \hat{\delta}_t}{p\left(s^t|\theta^*, w\right) + \delta_t} \qquad (9) \\ &= \ln\frac{\mu_0(\hat{E})}{\mu_0(E)} + \sum_{\tau=1}^{t} \mathbb{I}_{\Theta^R(\mu(h^\tau))}(\theta^\tau) \ln\frac{p\left(s^t|\theta^*, \hat{w}\right) + \hat{\delta}_t}{p\left(s^t|\theta^*, w\right) + \delta_t} \\ &\geq \ln\frac{\mu_0(\hat{E})}{\mu_0(E)} + \sum_{\tau=1}^{t} \mathbb{I}_{\Theta^R(\mu(h^\tau))}(\theta^\tau) \ln\frac{p\left(s|\theta^*, \hat{w}\right) - \epsilon}{p\left(s|\theta^*, w\right) + \epsilon} \end{aligned}$$

Let

$$\bar{l}\left(\mu\right)=\int_{\Theta^{R}(\mu)}\left[\sum_{s'\in S}p\left(s'|\theta,\omega^{*}\right)\ln\frac{p\left(s|\theta^{*},\hat{w}\right)-\epsilon}{p\left(s|\theta^{*},w\right)+\epsilon}\right]dp(\theta)$$

Then,

$$\begin{aligned}\frac{1}{t}\ln\frac{\mu_{t+1}(\hat{E})}{\mu_{t+1}(E)}\geq&\frac{1}{t}\left[\ln\frac{\mu_{0}(\hat{E})}{\mu_{0}(E)}+\sum_{\tau=1}^{t}\bar{l}\left(\mu_{\tau}\right)\right]\\&+\frac{1}{t}\sum_{\tau=1}^{t}\left[\mathbb{I}_{\Theta^{R}(\mu_{\tau})}(\theta^{\tau})\ln\frac{p\left(s|\theta^{*},\hat{w}\right)-\epsilon}{p\left(s|\theta^{*},w\right)+\epsilon}-\bar{l}\left(\mu_{\tau}\right)\right].\end{aligned}$$

By arguments that are substantively identical to Claim B of Esponda and Pouzo (2016),

$$\frac{1}{t}\sum_{\tau=1}^{t}\left[\mathbb{I}_{\Theta^{R}(\mu_{\tau})}(\theta^{\tau})\ln\frac{p\left(s|\theta^{*},\hat{w}\right)-\epsilon}{p\left(s|\theta^{*},w\right)+\epsilon}-\bar{l}\left(\mu_{\tau}\right)\right]\to 0 \tag{10}$$

almost surely given ω^{*} and that $h^{\tau}\in H$. It follows from $\mathbb{P}\left(\mu\left(\cdot|h^{t}\right)\to\mu^{*}|H,\omega^{*}\right)=1$ and continuity of $\Theta^{R}(\cdot)$ at μ^{*} that

$$\mathbb{P}\left(\lim_{t}\left|\bar{l}\left(\mu_{t}\right)-\bar{l}\left(\mu^{*}\right)\right|=0|H,\omega^{*}\right)=1.$$

Therefore,

$$\mathbb{P}\left(\lim_{t}\frac{1}{t}\ln\frac{\mu_{t+1}(\hat{E})}{\mu_{t+1}(E)}\geq\bar{l}(\mu^{*})|H,\omega^{*}\right)=1$$

and since $\bar{l}(\mu^{*})>0$ by construction, we must have $\lim\frac{1}{t}\ln\frac{\mu_{t+1}(\hat{E})}{\mu_{t+1}(E)}>0$, which requires $\mu_{t+1}(E)\to 0$, contradicting that $\mu^{*}(E)>0$.

Now, for each $w\in C$, let E_{w} be the open set constructed above. $\{E_{w}:w\in C\}$ is an open cover of C and so has a finite sub-cover $\{E_{1},\ldots,E_{k}\}$. The set $E^{*}=\cup_{i=1}^{k}E_{i}\supset C$, and $\mu^{*}(E^{*})=0$. Conclude $\mu^{*}(O)=1$. □

A.5. **Proof of Proposition 5.** Because D is strictly increasing in θ, $\Theta^{R}(\mu')=[0,\bar{\theta}(\mu')]$ for some $\bar{\theta}:\Delta\Omega\to(0,1]$. Since D is continuous, $\bar{\theta}$ is continuous in μ by the Berge Maximum Theorem, and thus Θ^{R} is a continuous correspondence. Observe that by definition of Q^{*}, (3) does not depend on $\omega'_{-Q^{*}}$. Because the divergence for each θ is strictly convex in $\omega'_{Q^{*}}$, so is the divergence in (3) for a fixed μ^{*}. Therefore, if ω' and ω'' both minimize it, then $\omega'_{Q^{*}}=\omega''_{Q^{*}}$. By Proposition 4, any stable μ^{*} attaches probability one to that value of $\omega_{Q^{*}}$.

Notice also that by the same arguments as in Proposition 3, the marginal of μ_t on ω_{-Q^*} is equal to the marginal of the prior. Denote $\bar{\theta}(\omega_{Q^*})$ to be $\bar{\theta}(\mu^{\omega_{Q^*}})$, where $\mu^{\omega_{Q^*}}$ agrees with the prior over ω_{-Q^*} and attaches probability one to ω_{Q^*}. By the implicit function theorem and that all primitive functions are smooth, $\bar{\theta}(\omega_{Q^*})$ is a smooth function of ω_{Q^*}

For any model (q, K) and state ω^*, define $\bar{\theta}\left(\omega^*_{Q^*}, q, K\right)$ so that

$$D\left(\int q\left(\cdot|\bar{\theta}\left(\omega^*_{Q^*}, q, K\right), \omega\right) d\mu^{\omega^*_{Q^*}} || \int q\left(\cdot|\theta^*, \omega\right) d\mu^{\omega^*_{Q^*}}\right) = K,$$

i.e. $\bar{\theta}\left(\omega^*_{Q^*}, q, K\right)$ is the cutoff context when ω^* is true and the researcher's beliefs attach probability 1 to $\omega^*_{Q^*}$, given the model (q, K). For any model (q, K) and $j \in Q$, define

$$G_j(w, q, K) = \sum_s \int_0^{\bar{\theta}\left(\omega^*_{Q^*}, q, K\right)} q\left(s|\theta, w\right) \frac{\frac{d}{dw_j} q\left(s|\theta^*, w\right)}{q\left(s|\theta^*, w\right)} p_\theta(\theta) d\theta$$

for $w \in \Omega$. Notice that $G_j\left(\omega^*, q, K\right)$ is

$$\frac{d}{d\omega_j} D_{KL}\left(\left(\int_0^{\bar{\theta}\left(\omega^*_{Q^*}, q, K\right)} q\left(\cdot|\theta, \omega^*\right) p_\theta\left(\theta|\theta \leq \bar{\theta}\left(\omega^*_{Q^*}, q, K\right)\right) d\theta\right) || q\left(\cdot|\theta^*, \omega\right)\right) |_{\omega=\omega^*}$$

multiplied by the probability that θ^t is less than $\bar{\theta}\left(\omega^*_{Q^*}, q, K\right)$, which does not affect whether the expression is equal to zero. Therefore, if μ^* is a stable belief for ω^* attaching probability 1 to ω^*_Q, then $G_j\left(\omega^*, p_S, K\right) = 0$.

We want to show that $\{\omega \in \Omega : G_j\left(\omega, q, K\right) = 0\}$ has measure zero for "most" regular models (q, K). Since μ is a product measure that admits a density with support Ω and Ω is compact and convex, Ω is a manifold with boundary of dimension n, and we can work only with its interior since its boundary has dimension $n - 1$ (p. 59, Guillemin and Pollack, 1974) and so zero measure. Towards that end, pick any regular model (q, K) and its corresponding Q^*. Fix $k \notin Q^*$ and a $j \in Q$ so that for every ω, $\frac{d}{d\omega_j} q(s|\omega, \theta^*) \neq 0$ for some s. Pick any $s^* \in S$ and define

$$q\left(\zeta\right)\left(s'|\theta, w\right) = q\left(s'|\theta, w\right) + \sum_{s \neq s^*} \left(\mathbb{I}_{\{s\}}\left(s'\right) - \mathbb{I}_{\{s^*\}}\left(s'\right)\right) \omega_k \theta \zeta_s$$

for every $\zeta \in B_\varepsilon(0) = Z \subset \mathbb{R}^{\#S-1}$ for a sufficiently small $\varepsilon > 0$. When ε is small enough, $(q(\zeta), K')$ is a regular model with the same θ^* and Q^* as q for every $\zeta \in Z$ and $K' > 0$, and $q(\zeta)$ has a strictly convex divergence whenever q does.[10]

Note that $q(\zeta)(\cdot|\theta^*, w) = q(\cdot|\theta^*, w)$ for all w. Then, $\frac{dG_j(w,q(\zeta),K)}{d\zeta_s}$ equals

$$w_k \left(\frac{\frac{d}{dw_j} q(s|\theta^*, w)}{q(s|\theta^*, w)} - \frac{\frac{d}{dw_j} q(s^*|\theta^*, w)}{q(s^*|\theta^*, w)} \right) \int_0^{\bar{\theta}\left(w_{Q^*}, q(\zeta), K\right)} \theta p_\theta(\theta) d\theta$$

$$+ \frac{d\bar{\theta}(w_{Q^*}, q(\zeta), K)}{d\zeta_s} \sum_s q(\zeta) \left(s | \bar{\theta}(w_{Q^*}, q(\zeta), K), w \right) \frac{\frac{d}{dw_j} q(s|\theta^*, w)}{q(s|\theta^*, w)} p_\theta \left(\bar{\theta}(w_{Q^*}, q(\zeta), K) \right).$$

and $\frac{dG_j(w,q(\zeta),K)}{dK}$ equals

$$\frac{d\bar{\theta}(w_{Q^*}, q(\zeta), K)}{dK} \sum_s q(\zeta) \left(s | \bar{\theta}(w_{Q^*}, q(\zeta), K), w \right) \frac{\frac{d}{dw_j} q(s|\theta^*, w)}{q(s|\theta^*, w)} p_\theta \left(\bar{\theta}(w_{Q^*}, q(\zeta), K) \right).$$

By our definition of $j \in Q$, for every w there exist $s', s'' \in S$ so that

$$\frac{d}{dw_j} q(s'|\theta^*, w) > 0 > \frac{d}{dw_j} q(s''|\theta^*, w)$$

so for some $s^{**} \in \{s', s''\}$,

$$\frac{\frac{d}{dw_j} q(s^{**}|\theta^*, w)}{q(s|\theta^*, w)} \neq \frac{\frac{d}{dw_j} q(s^*|\theta^*, w)}{q(s^*|\theta^*, w)}.$$

When $w_k \neq 0$, either

$$\sum_s q(\zeta) \left(s | \bar{\theta}(w_{Q^*}), w \right) \frac{\frac{d}{dw_j} q(s|\theta^*, w)}{q(s|\theta^*, w)} p_\theta(\bar{\theta}(\omega^*_{Q^*})) = 0$$

and then $\frac{dG_j(w,q(\zeta),K)}{d\zeta_{s^{**}}} \neq 0$, or it does not equal zero, in which case $\frac{dG_j(w,q(\zeta),K)}{dK} \neq 0$ (since $\frac{d\bar{\theta}\left(w_{Q^*},q(\zeta),K\right)}{dK} \neq 0$). Therefore, G_j viewed as a function from $\Omega \setminus \{\omega : \omega_k = 0\} \times Z \times \mathbb{R}_{++}$ to $\mathbb{R}$ is transversal to $\{0\}$. By the Transversality Theorem (p 68, Guillemin and Pollack, 1974),

$$\{\omega \in \Omega : G_j(\omega, q(\zeta), K') = 0\}$$

is a dimension $n-1$ subset of Ω for almost all (ζ, K') in $Z \times \mathbb{R}_{++}$. That is, for any regular (q, K), there are many $q(\zeta)$ arbitrarily close to q and cutoffs K' arbitrarily

[10] We can extend Z to paramterize any family of polynomials with the same θ^*, Q^*, and maximum degree in a similar fashion without changing arguments.

close to K for which the set of fixed parameters ω^* that have ω^* as a minimizer of (3) has measure zero. By Proposition 4 and the first part of this proposition, no stable beliefs for any of those parameters attach positive probability to ω^*_Q.

A.6. **Proof of Proposition 6.** Let $\bar{\mu}^t_1$ denote the mean belief about ω_1 according to the belief μ_t, and $\bar{\mu}_2$ denote the time-invariant mean belief about ω_2. Denote $q^t = (1-\theta^t)\bar{\mu}^t_1 + \theta^t \bar{\mu}_2$. The KL divergence that determines whether research is conducted in period t is

$$D_{KL}\left(p_S(\cdot|\theta^t, h^t)||p_S(\cdot|\theta^*, h^t)\right) = q^t \ln \frac{q^t}{\bar{\mu}^t_1} + \left(1-q^t\right) \ln \frac{1-q^t}{1-\bar{\mu}^t_1},$$

a function of only θ^t and $\bar{\mu}^t_1$. Moreover, $\Theta^R(\mu_t)$ is an interval $\left[0, \bar{\theta}\left(\bar{\mu}^t_1\right)\right]$.

For $\bar{\theta} \in [0,1]$, define the expected frequency with which $s^t = 1$ given $\theta \leq \bar{\theta}$ and ω^* as

$$s\left(\bar{\theta}, \omega^*\right) = E\left[\theta|\theta \leq \bar{\theta}\right] \cdot (\omega^*_2 - \omega^*_1) + \omega^*_1.$$

Given $a_1 > \frac{1}{2}$ and $\zeta > 0$, we show that for every K, there exists $\delta > 0$ so that $\hat{w}_1 = \frac{1}{2}(\omega^*_2 + \omega^*_1) \in \left(\frac{1-\zeta}{2}, \frac{1+\zeta}{2}\right)$ is a solution to (4) when $\omega^* \in B_\delta(a_1, 1-a_1)$. To see why, note that when $\bar{\mu}^t_1 = \frac{1}{2}$, we have $q^t = \bar{\mu}^t_1$ for every θ^t. This means that $D_{KL}(p_S(\cdot|\theta^t, h^t)||p_S(\cdot|\theta^*, h^t)) = 0$ for all $\theta^t \in [0,1]$, hence $\bar{\theta}\left(\frac{1}{2}\right) = 1$. By continuity, there exists $\delta^* > 0$ so that $\bar{\theta}(x) = 1$ for all $x \in \left(\frac{1-\delta^*}{2}, \frac{1+\delta^*}{2}\right)$. Notice that when $\omega^* \in B_\delta(a_1, 1-a_1)$,

$$s(1, \omega^*) \in \left[\frac{1}{2} - \delta, \frac{1}{2} + \delta\right].$$

Take $\delta = \min\{\delta^*, \zeta\}$, and then $\hat{w}_1(\omega^*) = s(1, \omega^*) \in \left(\frac{1-\zeta}{2}, \frac{1+\zeta}{2}\right)$ satisfies Equation (4) for all $\omega^* \in E$.

For small enough K and ζ, a second steady state exists for the same $\omega^* \in B_\delta(a_1, a_2)$ as above. Notice that when $\zeta < \min\left\{\frac{1}{3}(2a_1 - 1), 1 - \varepsilon - a_1\right\}$,

$$1 - \varepsilon > a_1 + \zeta > a_1 - \zeta > \frac{1+\zeta}{2} > \frac{1}{2}(\omega^*_2 + \omega^*_1),$$

$\frac{d}{d\theta}s\left(\bar{\theta},\omega^*\right) < 0$, $s(0,\omega^*) = \omega_1^*$, and $s(1,\omega^*) = \frac{1}{2}\omega_2^* + \frac{1}{2}\omega_1^*$. Pick $\theta' > 0$ so that $s\left(\theta,\omega^*\right) > \frac{1+\zeta}{2}$ for all $\theta < \theta'$ and ω^* as above. For every small $\bar{\theta} \in (0,\theta')$, define

$$K(\bar{\theta}) \equiv \left[\left(1-\bar{\theta}\right)(a_1-\zeta)+\bar{\theta}\frac{1}{2}\right]\ln\frac{\left(1-\bar{\theta}\right)(a_1-\zeta)+\bar{\theta}\frac{1}{2}}{a_1+\zeta} + \left[\left(1-\bar{\theta}\right)(1-a_1+\zeta)+\bar{\theta}\frac{1}{2}\right]\ln\frac{\left(1-\bar{\theta}\right)(1-a_1-\zeta)+\bar{\theta}\frac{1}{2}}{1-a_1+\zeta}$$

Because $D_{KL}\left(p_S\left(\cdot|\bar{\theta},h_t\right)||p_S\left(\cdot|\theta^*,h_t\right)\right)$ equals

$$\left[\left(1-\bar{\theta}\right)\bar{\mu}_1^t+\bar{\theta}\frac{1}{2}\right]\ln\frac{\bar{\theta}\frac{1}{2}+\left(1-\bar{\theta}\right)\bar{\mu}_1}{\bar{\mu}_1}+\left[\left(1-\bar{\theta}\right)\left(1-\bar{\mu}_1^t\right)+\bar{\theta}\frac{1}{2}\right]\ln\frac{1-\left(\bar{\theta}\frac{1}{2}+\left(1-\bar{\theta}\right)\bar{\mu}_1\right)}{1-\bar{\mu}_1},$$

$K(\bar{\theta})$ is weakly below $D_{KL}\left(p_S\left(\cdot|\bar{\theta},h_t\right)||p_S\left(\cdot|\theta^*,h_t\right)\right)$ for all histories so that $\bar{\mu}_1^t \in [a_1-\zeta,a_1+\zeta]$. Therefore, when $K < K\left(\bar{\theta}\right)$ and $\bar{\mu}_1^t \in [a_1-\zeta,a_1+\zeta]$, $\bar{\theta}\left(\bar{\mu}_1^t\right) \leq \bar{\theta}$. This implies

$$a_1-\zeta \leq s(\bar{\theta},w^*) \leq E\left[\theta|\theta \leq \bar{\theta}\left(\bar{\mu}_1^t\right)\right]\cdot(w_2^*-w_1^*)+w_1^* \leq 1-\varepsilon$$

By the intermediate value theorem, for any $K \in \left(0,K\left(\bar{\theta}\right)\right)$, there exists a $\hat{\hat{w}}_1 \in [a_1-\zeta,1-\varepsilon]$ solving Equation (4). Therefore, both $\hat{w}_1$ and $\hat{\hat{w}}_1$ both solve Equation (4) for ω^*.

Moreover, the two solutions are attractors of the dynamic process. For a mean belief x that is sufficiently close to either solution, there are fewer (more) "success" realizations $s = 1$ than expected when x is above (below) the fixed point. For $\omega_1 = \hat{\omega}_1\left(\theta_2(K)\right)$, this follows from there being too few successes at $1-\varepsilon$ and too many successes at $\frac{1}{2}$. For $\omega_1 = \frac{1}{2}$, this follows because $\bar{\theta}\left(x\right) = 1$ for all x sufficiently close to $\frac{1}{2}$, and so the number of successes is locally constant in x. Consequently, when a belief is near one of these two fixed points, it tends to drift toward it. □

A.7. **Proof of Proposition 7.** For almost every history h^t, $\mu\left(h^t\right)$ is normally distributed with variables independent. Let η denote any such beliefs with η_i the marginal on the i-th dimension. Slightly abusing notation,[11]

$$S(\eta,\theta) = D_{KL}\left(p_{S,U}(\cdot|\eta,\theta)||p_{S,U}(\cdot|\eta_1,\eta_3,\omega_2^*=0,\theta)\right)$$

[11]Namely, by the "conditioning" on η. The meaning is that the distribution $p_{S,U}$ is induced by the distribution η over ω.

and

$$R(\eta,\theta) = D_{KL}\left(p_{S,U}(\cdot|\eta,\theta)||p_{S,U}(\cdot|\eta,\theta^*=0)\right).$$

Denote $g(x) = x - \ln x - 1$, noting that $g'(x) > 0$ when $x > 1$ and that $g(1) = 0$, and

$$h(x,y) = x\ln\left(\frac{x}{y}\right) + (1-x)\ln\left(\frac{1-x}{1-y}\right).$$

Then,

$$\begin{aligned} S(\eta,\theta) =& \frac{1}{4}\left[g\left(1+\frac{\sigma_2^2}{\sigma_1^2+\lambda_1^2\theta^2\sigma_3^2}\right)+\frac{m_2^2}{\sigma_1^2+\lambda_1^2\theta^2\sigma_3^2}\right] \\ R(\eta,\theta) =& \frac{1}{4}\left[g\left(1+\frac{\lambda_1^2\theta^2\sigma_3^2}{\sigma_2^2+\sigma_1^2}\right)+\frac{\lambda_1^2\theta^2m_3^2}{\sigma_2^2+\sigma_1^2}+g\left(1+\frac{\lambda_0^2\theta^2\sigma_3^2}{\sigma_1^2}\right)+\frac{\lambda_0^2\theta^2m_3^2}{\sigma_1^2}\right]+D_{S_1|S_2,U}(\theta) \end{aligned}$$

where

$$\lambda_i = \mathbb{E}\left[u|s_1=1,s_2=i\right] = \frac{\phi(-i)}{1-\Phi(-i)},$$

and $D_{S_1|S_2,U}(\theta)$ is

$$\int \frac{1}{2}\phi(u)\left(h\left(\theta\Phi(-1-u)+(1-\theta)\Phi(-1),\Phi(-1)\right)+h\left(\theta\Phi(-u)+(1-\theta)\frac{1}{2},\frac{1}{2}\right)\right)du,$$

namely the expected KL divergence of $p_{S_1}(\cdot|s_2,u,\theta)$ from $p_{S_1}(\cdot|s_2,u,\theta^*=0)$. This follows from the formula for KL divergence of two normal distributions, and from the observation that $D_{KL}\left(p_{S,U}(\cdot|\theta)||p_{S,U}(\cdot|\theta^*=0)\right)$ equals

$$\begin{aligned} \sum_{s_2} p(s_2)\Bigg[&\int D_{KL}\left(p_{S_3}(\cdot|\theta,s_2,u)||p_{S_3}(\cdot|s_2,\theta^*=0,u)\right)d\Phi(u) \\ &+\int D_{KL}\left(p_{S_1}(\cdot|\theta,s_2,u)||p_{S_1}(\cdot|s_2,\theta^*=0,u)\right)d\Phi(u)\Bigg]. \end{aligned}$$

Clearly, S decreases in θ, R increases in θ, $R(\eta,0)=0$, and $S(\eta,0)>0$. Therefore, there is an interval $[0,x]$ with $0<x$ such that $R(\eta,\theta)\geq S(\eta,\theta)$ if and only if $\theta\in[0,x]$. Similarly, there is an interval $[0,y]$ with $y>0$ such that $R(\eta,\theta)\leq K$ if and only if $\theta\in[0,y]$. Finally, there is an interval $(z,1]$ (with z possibly equal to 1) such that $S(\eta,\theta)<K$ if and only if $\theta\in(z,1]$. Then, $\left[0,\bar{\theta}^{RD}(\eta)\right]=[0,x]\cap[0,y]=[0,\min\{x,y\}]$, and $\left(\bar{\theta}^{S}(\eta),1\right]=(x,1]\cap(z,1]=(\max\{x,z\},1]$. In the former interval, $\theta^*=0$ induces a lower KL divergence than does $\omega_2^*=0$, and the divergence is below K. In the latter interval, $\omega_2^*=0$ induces a lower KL divergence than does $\theta^*=0$, and the

divergence is below K. If K is sufficiently large, then $z = 0$ and $y = 1$, so the two intervals are adjacent.

Notice that S strictly increases in m_2^2, while R is constant in it. Therefore, an increase in m_2^2 leads to an increase in $\bar{\theta}^{RD}(\eta)$ (weakly) and $\bar{\theta}^{S}(\eta)$ (strictly). Also, R strictly increases in m_3^2, while S is constant in it. Therefore, an increase in m_3^2 leads to a decrease in $\bar{\theta}^{RD}(\eta)$ (strictly) and $\bar{\theta}^{S}(\eta)$ (weakly). Finally, S strictly increases in σ_2^2, and R strictly decreases in it. Therefore, an increase in σ_2^2 leads to a (strict) decrease in both $\bar{\theta}^{RD}(\eta)$ and $\bar{\theta}^{S}(\eta)$. □

References

Shun-ichi Amari and Hiroshi Nagaoka. *Methods of Information Geometry*, volume 191 of *Translations of Mathematical Monographs*. American Mathematical Society, 2000.

Isaiah Andrews and Jesse M. Shapiro. A model of scientific communication. *Econometrica*, 89(5):2117–2142, 2021.

Cuimin Ba. Robust misspecified models and paradigm shift. *Working Paper*, 2024.

Abhijit V. Banerjee, Sylvain Chassang, Sergio Montero, and Erik Snowberg. A theory of experimenters: Robustness, randomization, and balance. *American Economic Review*, 110(4):1206–1230, April 2020.

Robert H. Berk. Limiting Behavior of Posterior Distributions when the Model is Incorrect. *The Annals of Mathematical Statistics*, 37(1):51 – 58, 1966. doi: 10.1214/aoms/1177699597.

J. Aislinn Bohren and Daniel N. Hauser. Learning with heterogeneous misspecified models: Characterization and robustness. *Econometrica*, 89(6):3025–3077, 2021.

In-Koo Cho and Kenneth Kasa. Learning and model validation. *Review of Economic Studies*, 82(1):45–82, 2015.

Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2nd edition, 2006. ISBN ISBN Number.

I. Csiszár. Information-type measures of difference of probability distributions and indirect observations. *Studia Scientiarum Mathematicarum Hungarica*, 2:299–318, 1967.

Bruno de Finetti. *Theory of Probability: A Critical Introductory Treatment, Volume 1&2*. Wiley, Chichester, 1974/5.

Andrew Ellis and Ran Spiegler. Identifying assumptions and research dynamics. *arXiv preprint arXiv:2402.18713*, 2024.

Andrew Ellis and Heidi C. Thysen. Subjective causality in choice. *Working Paper*, 2024.

Ignacio Esponda and Demian Pouzo. Berk-nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica*, 84(2):1093–1130, 2016.

Ignacio Esponda and Demian Pouzo. Equilibrium in misspecified markov decision processes. *Theoretical Economics*, 16(2):717–757, 2021.

Mira Frick, Ryota Iijima, and Yuhta Ishii. Misinterpreting others and the fragility of social learning. *Econometrica*, 88(6):pp. 2281–2328, 2020.

Drew Fudenberg and David M. Kreps. Learning mixed equilibria. *Games and Economic Behavior*, 5:320–367, 1993. doi: 10.1006/game.1993.1021. URL `https://www.sciencedirect.com/science/article/pii/S0899825605800211`.

Drew Fudenberg, Gleb Romanyuk, and Philipp Strack. Active learning with a misspecified prior. *Theoretical Economics*, 12(3):1155–1189, 2017.

Andrew Gelman and Cosma Rohilla Shalizi. Philosophy and the practice of bayesian statistics. *British Journal of Mathematical and Statistical Psychology*, 66(1):8–38, 2013.

Victor Guillemin and Alan Pollack. *Differential Topology*, volume 370 of *AMS Chelsea Publishing Series*. AMS Chelsea Publishing, 1974.

James J. Heckman. Sample selection bias as a specification error. *Econometrica*, 47 (1):153–161, 1979.

Paul Heidhues, Botond Koszegi, and Philipp Strack. Unrealistic expectations and misguided learning. *Econometrica*, 86(4):1159–1214, 2018.

Colin Howson and Peter Urbach. *Scientific Reasoning: The Bayesian Approach*. Open Court Publishing, Chicago, 3rd edition, 2006.

Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.

Arthur Lewbel. The identification zoo: Meanings of identification in econometrics. *Journal of Economic Literature*, 57(4):835–903, 2019.

Charles F. Manski. *Identification for Prediction and Decision.* Harvard University Press, 2007.

Charles F. Manski. The lure of incredible certitude. *Economics and Philosophy*, 36 (2):216–245, 2020.

Pietro Ortoleva. Alternatives to bayesian updating. *Annual Review of Economics*, 16:545–570, 2024.

Judea Pearl. *Causality: Models, Reasoning, and Inference.* Cambridge University Press, 2009.

Alfred Rényi. On measures of entropy and information. In *The 4th Berkeley Symposium on Mathematics, Statistics and Probability, 1960*, pages 547–561. University of California Press, 1961.

Thomas J. Rothenberg. Identification in parametric models. *Econometrica*, 39(3): 577–591, 1971.

Leonard Savage. *The Foundations of Statistics.* John Wiley and Sons, 1954.

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

Ran Spiegler. Behavioral implications of causal misperceptions. *Annual Review of Economics*, 12(1):81–106, 2020.

Jann Spiess. Optimal estimation when researcher and social preferences are misaligned. *Working Paper*, 2024.