IDENTIFYING ASSUMPTIONS AND RESEARCH DYNAMICS

ANDREW ELLIS AND RAN SPIEGLER

Abstract. A representative researcher pursuing a question has repeated opportunities for empirical research. To process findings, she must impose an "identifying assumption", which ensures that repeated observation would provide a definitive answer to her question. Research designs vary in quality and are implemented only when the assumption is plausible enough according to a KL-divergence-based criterion, and then beliefs are Bayes-updated as if the assumption were perfectly valid. We study the dynamics of this learning process and its induced long-run beliefs. The rate of research cannot uniformly accelerate over time. We characterize environments in which it is stationary. Long-run beliefs can exhibit history-dependence. We apply the model to stylized examples of empirical methodologies: experiments, causal-inference techniques, and (in an extension) "structural" identification methods such as "calibration" and "Heckman selection."

1. Introduction

When social scientists and their audiences interpret the findings of an empirical study, they regularly rely on *identifying assumptions.* These (often irrefutable) assumptions enable the research community to draw clear-cut conclusions from observations. For instance, assuming that the assignment of agents into treatments was random, and therefore treatable as a "natural experiment," legitimizes a causal interpretation of the difference between the treatments' outcomes. Sometimes the identifying assumption is unstated. For example, interpreting the behavior of experimental subjects relies on the implicit assumption that subjects do not change their behavior merely because the experimenter observes them. Such assumptions are rarely perfectly valid, and yet researchers interpret observations as if they are.

Date: 27 February 2024.
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 Tim Christensen, Martin Cripps, Ignacio Esponda, Kate Smith and audiences at a Stony Brook workshop on bounded rationality and learning and an LSE/UCL theory conference for helpful comments.

The role of assumptions in the research process is at odds with Bayesian learning, the normative benchmark for how scientific learning should take place. A Bayesian research community should hold a prior belief regarding a research question, accumulate evidence (in the form of controlled experiments or observational data), and update its beliefs in light of it via Bayes' rule. By repeating the process, the community's beliefs converge to a correct answer to the research question, provided that the prior belief is not misspecified and that the evidence is informative. Assumptions do not play any role in this process. Thus, there appears to be a gap between the reality of scientific learning (which admits assumptions) and the strict Bayesian prescription (which leaves no room for assumptions). As we discuss in Section 2, there are good reasons for this gap: strict, assumption-free Bayesianism may be an impractical ideal for empirical researchers, at least in the social sciences with which we are familiar.[1]

Motivated by everyday observation of our "applied" colleagues, we propose an account of the research process that incorporates identifying assumptions into otherwise-standard Bayesian learning. When a researcher engages in a piece of empirical research she examines the identifying assumption it employs. When the assumption fails to clear some plausibility threshold, she abandons the research and waits for the next learning opportunity (involving a new experiment or observational data set). When the assumption clears the plausibility bar, the researcher and her community accept the assumption and take it *as given* when processing and communicating the research findings.

This process departs from Bayesian updating in two ways. First, a Bayesian never misses an opportunity to learn from any available evidence. In contrast, our researcher foregoes opportunities to learn whenever the feasible identification strategies are deemed implausible, a judgment that itself relies on the researcher's current beliefs. Second, a Bayesian always processes all relevant uncertainties, including those pertaining to whatever tools one uses to draw inferences from data. For example, if there are doubts as to whether the assignment of

[1]While the econometrics literature contains methodological discussions of the role of identifying assumptions in econometric inference (e.g., Rothenberg (1971), Manski (2007), and Lewbel (2019)), we are not aware of earlier discussions of how identification methods should be reconciled with the normative model of Bayesian inference.

experimental subjects into treatments is fully random, then these doubts should be incorporated into the Bayesian inference regarding the causal effect that the experiment is designed to reveal. By contrast, our researcher updates her beliefs as if the identifying assumption holds *exactly*, leaving out the uncertainty about its validity. These deviations from Bayesian learning raise several natural questions. How does the propensity to conduct research evolve over time? Do certain identifying assumptions die out or become more common? What, if any, biases arise? Can the researcher's long-run beliefs exhibit history-dependence?

We address these questions with a simple model of dynamic learning by a representative researcher, a stand-in for the relevant research community. The researcher has a prior belief over a fixed, multi-dimensional state of Nature. We refer to each component of the state as a "structural parameter." The researcher is interested in determining the values of certain structural parameters (e.g., the returns to education). 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 agents into education levels in a data set is random). To interpret her findings, the researcher must make an *identifying assumption*, which ensures that repeated measurements of observed variables would produce a definitive answer to the research question (e.g., agents' assignment is perfectly random). Effectively, the assumption says that all sources of noise other than sampling error can be safely ignored, as far as answering the question is concerned.

The researcher's decision whether to impose the assumption is based on a judgment of its *plausibility*, given the quality of the research design at hand. We capture this plausibility judgment with the Kullback-Leibler (KL) divergence of the distribution of all variables (both observable and latent) given the actual context from the distribution conditional on the assumed value of the context parameters. If the assumption is deemed implausible, i.e., the KL divergence exceeds some given threshold, then the researcher passes over the opportunity to conduct research and waits for the next period. When the researcher deems the assumption plausible, she updates her belief as if it holds, i.e., as if the research design is *perfect*.

We study the dynamics and long-run behavior of this process. We focus on how the propensity to conduct research (via the imposition of the identifying assumption) changes as the research community's beliefs change over time. We provide a sufficient condition for this propensity to be time-invariant. The condition requires that if an observable variable is not independent of the context parameters conditional on the other observable variables, then it must be independent of all the structural parameters that affect observable variables under the identifying assumption (including those the researcher wishes to learn). We also show that the propensity to conduct research cannot uniformly increase over time. In other words, the research community cannot consistently lower its standards for accepting research as time goes by. It may, however, uniformly *raise* these standards over time, leading to a slowdown in the rate of research. The basic intuition in such cases is that as the researcher's belief becomes more precise (even if the precision is illusory because it is based on wrong assumptions), she becomes more sensitive to the assumptions' rough edges and therefore more reluctant to impose them.

We then turn to the beliefs that the learning process induces in the long run. We define a stable belief to be one that the updating process converges to with positive probability. We show that stable beliefs concentrate on states that minimize KL-divergence of the empirical distribution of observable variables, conditional on the contexts for which research is conducted, from the distribution conditional on the identifying assumption. In turn, the contexts in which research is conducted are determined by the stable belief. This two-way relation between stable beliefs and the contexts in which research takes place means that stable beliefs are an equilibrium object. Indeed, our concept of stable beliefs is subtly related to Berk-Nash equilibrium (Esponda and Pouzo, 2016), a basic notion of stable behavior when agents operate under misspecified models.

Our second main task is to demonstrate the model's scope with stylized examples of our colleagues' empirical methodologies. One example considers experimental research contaminated by "interference" (the identifying assumption rules out the interference). Another example examines causal inference contaminated by confounding effects (the identifying assumption is that no such confounding exists). A variant on this example addresses instrumental-variable design (the identifying assumption is that the instrument is independent of a latent confounder).

Later in the paper, we present two examples that expand the notion of identifying assumptions to include not only "research design" assumptions but also "structural" assumptions that involve the fixed state of Nature. These examples also shed light on situations in which the representative researcher chooses from a *set* of candidate identifying assumptions. First, we consider an example in which the researcher tries to identify two structural parameters but can only do so in piecemeal fashion, employing an identifying strategy that is reminiscent of the "calibration" method in quantitative macroeconomics. Second, we present a stylized model of inference from selective samples, where the researcher's goal is learn the returns to a certain activity. The researcher considers two identifying assumptions: (1) agents' selection into this activity is purely random, or (2) selection is systematically related to observable variables 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). Hopefully, these examples demonstrate that our modeling approach can shed light on the evolution of inference methods in empirical economics and neighboring disciplines.

Our paper is related to two main strands in the economics literature. First, it continues a recent literature on Bayesian learning under misspecified subjective prior beliefs (e.g., Esponda and Pouzo, 2016, Fudenberg et al., 2017, Heidhues et al., 2021, Bohren and Hauser, 2021, Esponda and Pouzo, 2021, Frick et al., 2020). One difference is in motivation, as our paper is an attempt to explore the dynamics of scientific research, rather than learning by boundedly rational agents. Another difference is that in our model, the subjective prior is an endogenous choice by the researcher.

A second related literature consists of recent models of 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 is (to our knowledge) the first descriptive model of researchers that incorporates the role of assumptions in how they interpret empirical observations.

2. A Model

A research community cares about a *question* whose answer is determined by a fixed but unknown collection of *structural parameters* $\omega \in \Omega \subset \mathbb{R}^n$. We occasionally refer to ω as the *state*. The research question is formalized as a subset $Q \subseteq \{1, ..., n\}$, indicating the structural parameters that the researcher wishes to learn.

Time is discrete. At every period $t = 1, 2, ...$, a real-valued vector $\theta^t \in \Theta$ of *context parameters* is realized. We will often refer to a realization of θ as a *context*. While ω represents structural constants of a phenomenon of interest (e.g., returns to education), θ represents transient, circumstantial aspects of a periodic data set (e.g., whether assignment of students to educational treatments in a particular setting is random). We assume that Θ is compact and convex. 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. We require that each of S, U, Θ is a subset of some Euclidean space. For expositional convenience, the key definitions in this section proceed as if S, U, Ω are all finite; extension to the continuum case is straightforward.

The data-generating process p that governs the realization of (u, s) at every time period satisfies

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

We assume that p is continuous in θ and that $p(\cdot|\theta, \omega)$ has full support for every θ, ω. The context parameters and unobserved variables are distributed independently 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 $\omega_i \neq \psi_i$ for some $i \in Q$, there exists $s \in S$ such that $p(s \mid \omega, \theta^*) \neq p(s \mid \psi, \theta^*)$. The interpretation is that if the assumption holds, repeated observation of s eventually provides a definitive answer to the research question. We assume that there is a single feasible identifying assumption, and denote it θ^* (we relax our assumption in Section 5).

We are now ready to describe the learning process. At the beginning of period 1, a *representative researcher* has a prior belief $\mu \in \Delta(\Omega)$. We assume that according to this belief, all n components of ω are statistically independent of each other. The researcher knows p, as well as the distribution from which θ is drawn (independently across periods).

At every period t, the researcher makes the binary decision $a^t \in \{0,1\}$, indicating whether to conduct research. Entering the period, she has beliefs described by $\mu(\cdot|h^t)$ that depend on the history $h^t = (a^\tau, s^\tau, \theta^\tau)_{\tau<t}$ and she observes the current context θ^t. If the researcher chooses $a^t = 0$, then she passes over the opportunity to conduct research. She does not update her beliefs, and so the next research opportunity, arising at period $t+1$, is evaluated according to the same belief as in period t. If the researcher chooses $a^t = 1$, then she conducts research and updates her beliefs so that

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

for every $\omega, \psi \in \Omega$. When she observes s^t and updates her belief over Ω, she does so *as if* the assumption θ^* held.

We denote by $p_{S,U}(\cdot|\theta^t, h^t)$ the researcher's marginal probability over (s^t, u^t) at period t, given (θ^t, h^t). Because every time the researcher updates, she does so as if the context is θ^*, $p_{S,U}(\cdot|\theta^t, h^t)$ only depends on the public part of h^t, namely $(s^\tau)_{\{\tau:a^\tau=1\}}$, and the context θ^t. Since the distribution of u^t is known and independent of (θ^t, h^t), the only non-trivial aspect of $p_{S,U}$ is the conditional distribution of s^t.

We now describe the researcher's choice of a^t. The KL divergence of the variables' distribution given θ^t, h^t from the distribution given θ^*, h^t is

$$D_{KL}\left(p_{S,U}\left(\cdot|\theta^t, h^t\right) || \, p_{S,U}\left(\cdot|\theta^*, h^t\right)\right) = \sum_{s,u} p\left(s, u|\theta^t, h^t\right) \ln\left(\frac{p(s, u|\theta^t, h^t)}{p(s, u|\theta^*, h^t)}\right).$$

If this quantity exceeds a constant $K > 0$, then the assumption is deemed implausible and the researcher chooses $a^t = 0$. Otherwise, she chooses $a^t = 1$ and conducts research.

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. Plausibility is captured by how likely on average the variable realizations are under the actual context θ^t relative to the assumed one θ^*. KL divergence is a standard measure of this likelihood-based notion of plausibility. The likelihood judgment is based on the researcher's current beliefs. 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.

The plausibility judgment has a few noteworthy features. First, it depends only on the current period's context and the current belief $\mu(\cdot|h^t)$. Accordingly, the set of values of θ for which the researcher conducts research given the belief μ is denoted $\Theta^R(\mu)$. Second, since $p(\cdot|\theta,\omega)$ has full support, the KL divergence is always finite, and a wrong assumption can never be categorically refuted by data. Third, the plausibility judgment takes into account the assumption's effect on the distribution of *both* observed (s) and latent (u) variables. This aspect of our model reflects our observation of real-life discussions of identification strategies in empirical economics. To give a concrete example, evaluation of the plausibility 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.

In our model, an assumption that underlies a particular study is subjected to a binary, "up or out" plausibility judgment. When the outcome of this evaluation is affirmative, beliefs regarding the research question are updated as if the assumption were perfectly sound. Once an assumption is accepted in a certain context, subsequent research never put its contextual plausibility in doubt again. This feature seems to be consistent with our casual observation that debates over the adequacy of an identification strategy for a particular study play an important role in the research community's decision whether to admit the study (amplifying its exposure in seminars and conferences, accepting it for publication in prestigious journals, etc.), yet subsequent references to the published study rarely re-litigate the identification strategy's appropriateness for that particular study.

Throughout the paper, we take the assumption-based, semi-Bayesian learning process as given, without trying to derive it from some explicit optimization problem. Informally, however, we can think of two broad motivations behind the reliance on identifying assumptions. First, if repeated observations did not produce a definitive answer to the research question, long-run beliefs about it would remain sensitive to subjective prior beliefs, thus defeating one purpose of the scientific enterprise, which is to produce consensus answers. Second, the strict Bayesian model requires the research community to hold, process and communicate multi-dimensional uncertainty. When researchers interpret empirical evidence, they need to take into account various sources of noise that interfere with the mapping from the underlying object of study to empirical evidence. When researchers are uncertain about the magnitude and direction of such interferences, Bayesian learning requires them to carry this "secondary" uncertainty throughout the updating process in addition to the uncertainty regarding the research question. This multi-dimensional updating is inherently difficult to conduct and to communicate to other members of the research community. Identifying assumptions reduce this complexity by removing secondary uncertainties.

We adopt the following standard notational conventions. For a vector x and a subset of its indices E, x_E is the vector $(x_i)_{i\in E}$ and x_{-E} is the vector $(x_i)_{i\notin E}$. Similarly for an index i, x_{-i} denotes the vector $(x_j)_{j\neq i}$. For two vectors x and y with disjoint indices, (x, y) denotes their concatenation. Finally, we denote by $\mathbb{P}(\cdot)$ the equilibrium probability distribution over all variables.

3. Examples

In this section we illustrate the model with two examples. Our aim is to showcase the model's expressive scope, as well as give a taste for the kind of learning dynamics that it can give rise to. Throughout, we refer to the first as the "Contaminated Experiment" and the second as the "Causal Inference" example

3.1. **Contaminated Experiment.** Suppose that our representative researcher wants to identify a behavioral effect, but she is concerned that observations of this effect are contaminated by "friction." Specifically, there is a single observable variable, given by

$$s^t = \omega_1 + \theta^t \omega_2 + \varepsilon^t$$

where ω_1 is the structural parameter the researcher wants to learn, i.e., $Q = \{1\}$. The parameter ω_2 represents the friction's strength, the context parameter θ captures how well experimental design manages to curb the friction, and $\varepsilon \sim N(0,1)$ is independently drawn each period. There are no latent variables. To illustrate this specification, think of ω_1 as the degree of *intrinsic* altruism in a certain social setting, while ω_2 represents how much the subjects want an outside observer to *perceive* them as altruistic.

The researcher's prior belief over ω_i at the beginning of period t is $N\left(m_1^t, (\sigma_i^t)^2\right)$, independently of the other component of ω. The distribution of s^t conditional on ω and θ^t is thus $N\left(\omega_1 + \theta^t \omega_2, 1\right)$. It follows that the only feasible identifying assumption is $\theta^* = 0$, since under any $\theta \neq 0$, ω and $\left(\omega_1 - k, \omega_2 + \frac{k}{\theta}\right)$ generate the same distribution of s. Furthermore, when the researcher assumes $\theta^* = 0$, she can learn nothing about ω_2 from observations of s. It follows that whenever the researcher updates her beliefs, she does so as if $\theta = 0$, and her beliefs over ω_2 never evolve (accordingly, we will remove the time index from the mean and variance of ω_2). The distribution of s conditional on θ is

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

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

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

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

The 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.

Recall that when she does so, she updates her belief as if $\theta^t = 0$. Using the standard formula for updating a normal distribution, $\sigma_1^{t+1} = \sigma_1^t \left(\left(\sigma_1^t\right)^2 + 1\right)^{-\frac{1}{2}}$. That is, σ_1^t decreases monotonically over time. Therefore, 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

$$\bar{\theta}\left(0\right) = \frac{1}{2}\left[\left(\sigma_2^2 + m_2^2\right)\left(\theta^t\right)^2 - \ln\left(1 + \left(\theta^t\right)^2 \sigma_2^2\right)\right] < \infty$$

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 this long-run limit, research is carried out when $\theta \in \left[0, \bar{\theta}\left(0\right)\right]$. This means that the researcher's long-run belief over ω_1 assigns probability one to

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

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 $\bar{\theta}\left(0\right)$ 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.

Comment on feasible identification strategies. Our claim that the only feasible identifying assumption in this example is $\theta^* = 0$ rests on our assumption that this judgment is 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 the identification judgment could be made by combining multiple contexts (given by different values of θ), there would be no need to make wrong identifying assumptions. This "triangulating" identification strategy would work in most of the examples in this paper. However, it is inconsistent with the research practice we are familiar with, where the identification constraint is applied to each research *in isolation*.

3.2. **Confounded causal inference.** Determining the causal effect of one variable on another is a central task for empirical researchers. A key difficulty here is that the effect is often masked by an unobserved confounding variable that affects both observable variables. We now present a stylized example of causal inference from observational data in the presence of a potential confounder.

There are two observable variables, s_1 and s_2. The researcher wants to learn the causal effect of the former on the latter. This effect is parameterized by $\omega_2 \in (-1, 1)$, i.e., $Q = \{2\}$. However, the observed correlation between the two variables is confounded by a latent variable u that affects both. The fixed parameter $\omega_1 \in (-1, 1)$ captures the strength of this confounding effect. The context parameter $\theta \in [0, 1]$ captures the extent to which a given data set manages to shut down this confounding channel. More explicitly,

$$s_1 = \theta\omega_1 u + \varepsilon_1$$

$$s_2 = \omega_2 s_1 + \omega_3 u + \varepsilon_2$$

where $u \sim N(0, 1)$ and $\varepsilon_i \sim N(0, \sigma_i^2)$ for $i = 1, 2$, independently of each other. There is no uncertainty regarding $\omega_3 > 0$. Set this parameter and the variances σ_1^2 and σ_2^2 such that $s_i|\omega, \theta \sim N(0, 1)$ for each $i = 1, 2$ θ, and ω.[2] It follows that the only aspect of the long-run distribution of (s_1, s_2) that could potentially shed light on the state is the pairwise correlation between the two statistics,

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

[2]That is, $\sigma_1^2 = 1 - \theta^2\omega_1^2$ and $\sigma_2^2 = 1 - \omega_2^2 - \omega_3^2 - 2\theta\omega_1\omega_2\omega_3$.

It is evident from the equation for $\rho_{12}(\theta,\omega)$ that the only feasible identifying assumption is $\theta^* = 0$. As in the previous example, this assumption prevents any learning about the other structural parameter (ω_1). Observe that $\rho_{12}(\theta^* = 0, \omega) = \omega_2$. Hence, under the identifying assumption, the observed long-run correlation between s_1 and s_2 pins down the causal effect of interest.

We now derive an expression for the KL divergence between the true and assumed distributions over (u, s). Observe that the joint density of the variables conditional on the parameters can be factorized as

$$p(u, s|\omega, \theta) = p(u)p(s_1|u, \omega_1, \theta)p(s_2|u, s_1, \omega_2).$$

Thus, $D_{KL}\left(p_S\left(\cdot|h^t, \theta^t\right)||p_S\left(\cdot|h^t, \theta^*\right)\right)$ is equal to

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

Note that the researcher's belief over ω_2 (which evolves over time) does not appear in the final expression we have arrived at. The only aspect of μ that enters the divergence is the belief over ω_1. Because this belief is stationary, it follows that the expression for the divergence (for any given θ^t) does not change over time. This means that the researcher's propensity to research is time-invariant: there is $\bar{\theta}$ such that the researcher will update her beliefs over ω_2 if and only if $\theta^t \in [0, \bar{\theta}]$. 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 $\omega_1 \neq 0$, i.e., when there is a confounding effect.

4. General Analysis

This section presents results that describe properties of the learning process in general. We begin with results about how the propensity to learn changes over time, and illustrate these results with additional examples. We then turn to the long-run beliefs that the learning process induces. Throughout this section, we assume that S, U and Ω are finite for expositional simplicity.

4.1. **Evolution of the propensity to conduct research.** In the examples from Section 3, the set of contexts for which research takes place (weakly) contracts over time. For instance, the Contaminated Experiment example demonstrated the possibility of a uniformly decreasing rate at which research takes place. Our first two results show that the opposite pattern, namely a uniformly increasing 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 $\mathbb{P}\left(\theta \in \Theta^R\left(\mu\left(h^{t+1}\right)\right) \setminus \Theta^R\left(\mu\left(h^t\right)\right)|h^t\right) > 0$, then there exists $t^* > t+1$ such that $\mathbb{P}\left(\theta \notin \Theta^R\left(\mu\left(h^{t^*}\right)\right)|h^{t+1}\right) > 0$.*

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 at 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 research would once again not be conducted in that same context.

When the contexts map naturally to the KL divergence, we can be more explicit about how the propensity to research evolves.

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

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

then

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

This result says that when there are contexts for which research takes place at period $t+1$ but not at t ($\Theta^R\left(\mu\left(h^{t+1}\right)\right) \setminus \Theta^R\left(\mu\left(h^t\right)\right) \neq \emptyset$), then with positive probability, there are contexts for which research takes place at t but not at $t+1$ ($\Theta^R\left(\mu\left(h^t\right)\right) \setminus \Theta^R\left(\mu\left(h^{t+1}\right)\right) \neq \emptyset$).

That is, when the community conducts research in new contexts with positive probability, it also stops conducting research in others.

The result relies on the assumption that the KL 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 rely on convexity of relative entropy. This implies that relative entropy increases on average. Therefore, the divergence between $p_{S,U}(\cdot|\theta, h^t)$ and $p_{S,U}(\cdot|\theta^*, h^t)$ rises in expectation for every θ. 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.

In the Causal Inference example, the set of contexts for which research takes place is history-independent. Our next result provides a general sufficient condition for this property. We state the sufficient condition using language from the literature on graphical probabilistic models. (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.) A directed acyclic graph (DAG) consists of a set of nodes N representing variables and a set R of directed links between nodes, such that the graph contains no cycle of directed links.

We say that a data-generating process is *recursive* if it is described by a recursive system of structural equations, where the equations for the parameters and latent variables are degenerate (i.e., their R.H.S. includes no variable or parameter). A recursive data-generating process corresponds to an *underlying DAG*, where all the parameters and unobserved variables are represented by ancestral nodes, and there is an edge into s_i from each parameter or variable in the R.H.S. of the equation that defines s_i. All of our examples assume a recursive data-generating process. For instance, in the Causal Inference example, the underlying DAG is

$$\begin{array}{ccccc} \theta & \rightarrow & s_1 & \leftarrow & \omega_1 \\ & \nearrow & \downarrow & & \\ u & \rightarrow & s_2 & \leftarrow & \omega_2 \end{array} .$$

Following Spiegler (2016), let $R(i)$ denote the set of node i's "parents," i.e., the set of nodes that send directed links into i. Say that a joint distribution p with full support over a product set $X = \times_{i \in N} X_i$ is consistent with the DAG (N, R) if

$$p(x) = \prod_{i \in N} p\left(x_i | x_{R(i)}\right)$$

for every $x \in X$. A DAG G satisfies a conditional-independence property if every distribution that is consistent with G satisfies this property. Any such conditional-independence property has a graphical characterization known as "d-separation" (see Pearl (2009)).

We define the set of *active parameters* A to be the smallest set of indices for which $p_{S,U}\left(\cdot|\theta^*, \omega\right) = p_{S,U}\left(\cdot|\theta^*, \omega'\right)$ whenever $\omega_A = \omega'_A$. This means that under the identifying assumption, all other structural parameters do not affect the long-run distribution of s, and therefore repeated observation can teach the researcher nothing about them. In the Causal Inference example, the set of active parameters was $\{2\}$. The set of active parameters is defined with respect to θ^*. It is not purely determined by the DAG structure underlying p, because it depends on the value of θ^*. Note also that by the definition of identifying assumptions, $Q \subseteq A$.

Say that θ and ω_A are *G-separable* if for every i, G satisfies $s_i \perp \omega_A$ whenever it satisfies $s_i \not\perp \theta | (s_{-i}, u)$. If θ and ω_A are G-separable, then any statistic that is affected by the context (conditional on the other statistics and the latent variables) is not affected by the answer to the question (nor by the value of other active parameters). As we show in the proof of the next results, this property implies that $\omega_Q \perp \theta | (s, u)$, i.e., the structural parameters of interest and the context parameters are independent conditional on the variables. This in turn implies that the parameters of interest and the context directly affect different sets of statistics. The Contaminated Experiment example violates this condition, since the context parameter and the structural parameter of interest directly cause the only statistic. In contrast, the Causal Inference example satisfies the condition: the context parameter (which determines the strength of the confounding effect) has a direct effect only s_1, while the structural parameter of interest (which measures the causal effect of s_1 on s_2) has a direct effect only on s_2.

Proposition 3. *Suppose that the data-generating process is recursive with an underlying DAG G. If θ and ω_A are G-separable, then $\Theta^R(\cdot)$ is constant.*

Under G-separability, the set of contexts for which the researcher conducts research does not change over time, i.e., there is a constant propensity to research. The proof uses the DAG tool of d-separation to factorize belief into conditional-probability terms. Using this factorization, we show that every statistic whose distribution is changed by the identifying assumption must be conditionally independent of the active parameters. This in turn implies that every term involving ω_A cancels out or gets integrated out in the expression for the KL divergence. Since the researcher only learns about ω_A under the identifying assumption, the KL divergence for any given context remains fixed over time.

Note that the conditional-independence property that underlies Proposition 3 is not imposed directly on the researcher's belief. Instead, it holds for the system of recursive equations that generate the belief. It is thus "robust" in the sense that it does not depend on the specific distributions of the underlying variables, but only on their underlying qualitative relationships.

4.1.1. *An Example: Instrumental-variable causal identification.* The DAG language 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 system of recursive equations:

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

where u and the ε variables are all independent Gaussians. Set their variances and the range of possible values of the parameters such that $s_i \sim N(0,1)$ for every $i = 1,2,3$. Suppose that the researcher wants to learn ω_4, the causal effect of s_2 on s_3. Formally, $Q = \{4\}$. This relationship is obfuscated by the unknown effect of u on s_1, s_2, and s_3. Since the statistic variables are all standard normal, the only aspects of the long-run distribution of s that the researcher can use to learn ω are $E(s_1 s_2)$, $E(s_2 s_3)$ and $E(s_1 s_3)$. This gives three equations with five unknowns, and therefore ω_4 cannot be identified. However, when we make the assumption $\theta^* = 0$, we get $\omega_4 = E(s_1 s_3)/E(s_1 s_2)$, which is the textbook 2SLS procedure. The identification strategy uses s_1 as an instrument for s_2. The identifying assumption is that the instrument is independent of the confounding variable u.

Let us apply Proposition 3 to this example. The DAG structure of the system is

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

The active parameters are ω_2, ω_3, ω_4 and ω_5, i.e., $A = \{2, 3, 4, 5\}$. Observe that s_1 is not independent of θ conditional on the other variables (because there is a direct link between the two nodes). However, s_1 is independent of ω_A since they have no common ancestor. Using d-separation, we can show that the other two statistics, s_2 and s_3, are both independent of θ given s_1 and u. By Proposition 3, the researcher's propensity to employ the IV identification strategy is time-invariant.

In Appendix B, we examine another causal-inference identification strategy, known as "front door identification" (see Pearl (2009)), and show that it violates the condition for time-invariant propensity to learn.

4.2. **Long-run beliefs.** Finally, we turn to the question of what the community believes about the state. We begin with a definition of stable beliefs.

Definition 1. A belief μ^* is stable for ω^* if $\mathbb{P}(\lim_{t\to\infty} ||\mu(h^t) - \mu^*|| = 0|\omega^*) > 0$.

A belief is stable when the posterior beliefs generated by the learning process converge to it with positive probability in the long run. The following result characterizes stable beliefs.

Proposition 4. *For a parameter ω^*, if μ^* is stable for ω^* and $\Theta^R(\cdot)$ is continuous in a neighborhood of μ^*, then*

$$\mu^* \left(\arg\min_{\omega \in \Omega} D_{KL} \left(p_S \left(\cdot | \theta \in \Theta^R(\mu^*), \omega^* \right) || p_S \left(\cdot | \theta^*, \omega \right) \right) \right) = 1.$$

To understand this result, recall that observed statistics are affected by the contexts in which research is conducted and by the actual state ω^*. If the researcher consistently holds the belief μ^* for a long stretch of time, this means that the set of values of θ for which research takes place during that stretch is $\Theta^R(\mu^*)$. In this case, the long-run frequency of the statistic is $p \left(s | \theta \in \Theta^R(\mu^*), \omega^* \right)$. However, the researcher updates his belief according to the identifying assumption that the context parameter is θ^*. Under that assumption, the statistic s is realized with probability $p\left(s|\theta^*, \omega\right)$ in state ω. Following Berk (1966) and Esponda and Pouzo (2016), the long-run belief that emerges from this misspecified Bayesian learning assigns probability one to states that minimize the KL divergence of the true distribution from the subjective one. A belief μ^* is stable if it only attaches positive probability to the states that minimize this divergence.

We should not confuse the KL divergence in the result with the role of the divergence in researchers' 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 underlies the characterization of long-run beliefs.

Of course, Proposition 4 does not establish whether a stable belief exists, whether it is unique, and whether the process does indeed converge when there is a unique stable belief. As is often the case in the literature on misspecified learning, these are difficult questions, which we do not address here. The following example illustrates the possibility of multiple stable beliefs.

4.2.1. *An Example: Contaminated experiments, revisited.* This is a variant on the example from Section 3.1. The main difference is that the statistic s is now a binary variable that gets values in $\{0, 1\}$. As before, there are two structural parameters, ω_1 and ω_2, and a single context parameter θ. Both structural parameters take values in $[\varepsilon, 1-\varepsilon] \subset (0, 1)$. There are no latent variables. The conditional distribution of s is given by

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

The researcher wants to learn ω_1, i.e., $Q = \{1\}$. As in the original example, the only feasible identifying assumption is $\theta^* = 0$. As before, this assumption prevents learning anything about ω_2.

Let $\bar{\mu}_1^t$ denote the mean of ω_1 according to the belief μ_1^t. Let $\bar{\mu}_2$ denote the mean of ω_2 according to the time-invariant belief μ_2. Denote $q^t = (1 - \theta^t)\bar{\mu}_1^t + \theta^t\bar{\mu}_2$. The KL divergence that determines whether research is conducted at period t is:

$$D_{KL}\left(p_S(\cdot|\theta^t, h^t)||p_S(\cdot|\theta^*, h^t)\right) = q^t \ln \frac{q^t}{\bar{\mu}_1^t} + \left(1 - q^t\right) \ln \frac{1 - q^t}{1 - \bar{\mu}_1^t}$$

The derivative of this expression with respect to θ^t is negative. Therefore, $\Theta^R(\mu^t)$ is an interval $[0, \bar{\theta}(\mu^t)]$. In previous examples, our focus was on how $\bar{\theta}(\mu^t)$, and therefore the propensity to conduct research, evolve over time. Instead, this example focuses on long-run beliefs.

The divergence $D_{KL}(\cdot)$ is not constant in $\bar{\mu}_1^t$. Moreover, since $\bar{\mu}_1^t$ can move back and forth in the range $[\varepsilon, 1 - \varepsilon]$, there will be phases of both accelerating and decelerating rates of research. This is in contrast to the uniform research slowdown that emerged in the original contaminated-experiment example, where the statistic was Gaussian.

By definition, both $\bar{\mu}_1^t$ and $\bar{\mu}_2$ are restricted to $[\varepsilon, 1 - \varepsilon]$. Therefore, whatever the researcher's beliefs, the KL divergence is finite, such that $\bar{\theta}(\mu^t)$ is bounded away from zero. As a result, the probability that research is carried out is positive after every history. This means that the researcher will obtain infinitely many observations of s. Under the identifying assumption, $s^t = 1$ with independent probability ω_1 at every t. Therefore, the researcher identifies the long-run frequency of $s = 1$ with ω. Any candidate for a stable belief given the true ω is a degenerate distribution that assigns probability one to some ω_1^* (abusing notation, use ω_1^* to represent this belief), which satisfies the equation

$$\omega_1^* = E[\theta|\theta < \bar{\theta}(\omega_1^*)] \cdot (\omega_2 - \omega_1) + \omega_1. \tag{2}$$

We will now illustrate the possibility that this equation has multiple solutions, namely multiple candidates for stable beliefs. Suppose $\omega_1 + \omega_2 = 1$; $\bar{\mu}_2 = \frac{1}{2}$; and the distribution over θ is smooth with full support on $[0,1]$ and mean $\frac{1}{2}$. Under this specification, $\omega_1^* = \frac{1}{2}$ is a solution to (2). To see why, note that when $\omega_1^* = \frac{1}{2}$, we have $q = \bar{\mu}_1$ for any θ. This means that $D_{KL}(\cdot) = 0$ for all θ, hence $\bar{\theta}\left(\frac{1}{2}\right) = 1$. The R.H.S. of (2) thus becomes

$$E(\theta)(\omega_2 - \omega_1) + \omega_1 = \frac{1}{2}(1 - 2\omega_1) + \omega_1 = \frac{1}{2}$$

which coincides with the equation's L.H.S.

We now show that we can find ω_1 and $\omega_2 = 1 - \omega_1$ such that the equation has a second solution that lies strictly above $\frac{1}{2}$. Since $\omega_1, \omega_2 \in [\varepsilon, 1-\varepsilon]$ and $E(\theta) = \frac{1}{2}$, we have

$$1 - \varepsilon > E[\theta | \theta < \bar{\theta}(1-\varepsilon)] \cdot (\omega_2 - \omega_1) + \omega_1$$

That is, the R.H.S. of (2) lies below the L.H.S. at $\omega_1^* = 1 - \varepsilon$. Now consider $\omega_1^{**} \in (\frac{1}{2}, 1-\varepsilon)$. Let $\theta_2 \in (0,1)$ satisfy

$$\omega_1^{**} < E\left[\theta | \theta < \theta_2\right] \cdot (2\varepsilon - 1) + 1 - \varepsilon.$$

We can always find such θ_2, since $E(\theta | \theta < \theta_2)$ continuously decreases with θ_2 and converges to zero as $\theta_2 \rightarrow 0$. Now select K to satisfy

$$K = \omega_1^{**} \ln \frac{\frac{1}{2} + (1-\theta_2)\left(\omega_1^{**} - \frac{1}{2}\right)}{\omega_1^{**}} + [1 - \omega_1^{**}] \ln \frac{\frac{1}{2} - (1-\theta_2)\left(\omega_1^{**} - \frac{1}{2}\right)}{1 - \omega_1^{**}},$$

The R.H.S. of this equation is the value of the KL divergence given θ_2 and ω_1^{**}. It follows that $\bar{\theta}(\omega_1^{**}) = \theta_2$. We have thus established that when $\omega_1 = 1 - \varepsilon$ and $\omega_2 = \varepsilon$, the R.H.S. of (2) lies above the L.H.S. at ω_1^{**}. Since we established the opposite ranking at $\omega_1^* = 1 - \varepsilon$ as well as an intersection at $\omega_1^* = \frac{1}{2}$, there must be an additional intersection at some $\omega_1^* \in (\frac{1}{2}, 1-\varepsilon)$, by continuity of the expression for the KL divergence.

Of course, the fact that the formula that characterizes stable beliefs has multiple solutions does not imply by itself that all these solutions represent stable beliefs. However, it does establish that the question of whether our learning process converges to a stable belief is non-trivial, as is the case in other models of misspecified learning.

5. An Extension: Multiple/"Structural" Assumptions

Our model is restrictive in several respects. First, it assumes a single feasible identifying assumption, rather than a set of identification strategies the researcher could choose from. Second, it assumes the researcher has a single question, rather than a set of nested questions (such that she can choose between answering an ambitious question using a strong assumption or a modest question using a weak assumption).[3] Third, it focuses entirely on "research design" assumptions that pertain to the context parameters, rather than "structural" assumptions that pertain to the fixed state of Nature. In this section we present two examples that go beyond these restrictions and offer stylized representations of familiar identification methods in empirical economics.

5.1. **Learning by "Calibration".** Suppose that the statistic follows the process

$$s^t = (\omega_1 + \omega_2) + \varepsilon^t$$

where $\varepsilon^t \sim N(0,1)$. The researcher wants to learn both ω_1 and ω_2 , i.e., $Q = \{\omega_1, \omega_2\}$. There are no context parameters in this specification, hence our notion of identifying assumptions in the basic model is moot. Clearly, the researcher cannot identify both structural parameters from observations of s. However, the researcher can settle for identification of one of the structural parameters, by imposing an assumption about the value of the other structural parameter. This is an example of a structural identification strategy which does not aim at a complete answer to the research question and settles for a partial answer instead.

Formally, assume that at every period, the researcher can assume $\omega_2 = \omega_2^*$ or $\omega_1 = \omega_1^*$, where ω_2^* and ω_1^* can take any value. When the researcher assumes $\omega_i = \omega_i^*$, she interprets all the variation in s^t as a consequence of ω_{-i} and the sampling error ε^t. When the researcher assumes $\omega_i = \omega_i^*$, she updates only her belief about ω_{-i}. The researcher selects the assumption that minimizes the KL divergence (relative to the true data-generating process, given her current beliefs), and performs the research only if this minimal divergence does not exceed $K > 0$.

[3]For a systematic discussion of the similar dilemma between "point" and "partial" identification, see Manski (2007).

This learning process is a metaphor for the "calibration" method employed by quantitative macroeconomists. In this field, it is customary to confront a multi-parameter model with observational data lacking the richness that enables full identification of the model's parameters. Macroeconomists then proceed by assigning values to some of the parameters in order to identify the remaining parameters from the data.

We examine the learning dynamics that this procedure induces. Suppose that the researcher's belief at the beginning of period t is that $\omega_i \sim N\left(m_i^t, (\sigma_i^t)^2\right)$, independently across the components of ω. Then,

$$2D_{KL}\left(p_S(\cdot|h^t)||p_S(\cdot|h^t, \omega_i = \omega_i^*)\right) = \frac{((\sigma_1^t)^2 + (\sigma_2^t)^2) + (m_i^t - \omega_i^*)^2}{(\sigma_{-i}^t)^2} - \ln\left(\frac{(\sigma_1^t)^2 + (\sigma_2^t)^2}{(\sigma_{-i}^t)^2}\right) - 1.$$

The divergence minimizing value of ω_i^* is $\omega_i^* = m_i^t$, and then

$$2D_{KL}\left(p_S(\cdot|h_t)||p_S(\cdot|h_t, \omega_i = m_i^t)\right) = \frac{(\sigma_i^t)^2}{(\sigma_{-i}^t)^2} - \ln\left(1 + \frac{(\sigma_i^t)^2}{(\sigma_{-i}^t)^2}\right).$$

The researcher effectively chooses between setting $\omega_1 = m_1^t$ and setting $\omega_2 = m_2^t$. The former induces a lower divergence than the latter if and only if $\sigma_1^t < \sigma_2^t$. Therefore, the researcher will assume there is no uncertainty about the parameter she is more certain about. This again brings to mind the "calibration" methodology: The researcher "calibrates" the parameter she is more confident about, using her best estimate of this parameter.

We assume that K is large enough such that learning always take place, and that, w.l.o.g, $\sigma_1^1 \geq \sigma_2^1$. At $t = 1$, the researcher assumes $\omega_2 = m_2^1$, and then updates her belief about ω_1. Because the researcher's belief about each parameter is given by an independent Gaussian distribution, each update about ω_i shrinks σ_i by a deterministic percentage. After updating about ω_1 for some number k of periods, the variance of σ_1^{t+k} will fall below that of σ_2^{t+k}. At that point, the researcher switches to the other assumption, namely $\omega_1 = m_1^{t+k}$ and proceeds to update about ω_2. She then repeats, alternating between updating about ω_1 and ω_2.

The next result addresses long-run beliefs. For convenience, we assume that initial variances are identical. Then, w.l.o.g, in odd periods, the researcher will set $\omega_1 = m_1^t$ and update her beliefs about ω_2, and in even periods, she will set $\omega_2 = m_2^t$ and update her beliefs about ω_1.

Proposition 5. *As $t \to \infty$, $\sigma_i^t \to 0$ almost surely for each i. Conditional on the realized value of (ω_1, ω_2), $m_1^t + m_2^t \to \omega_1 + \omega_2$ with probability one, and there exists $v > 0$ such that m_i^t is normally distributed with variance greater than v for all t.*

In the long run, the researcher correctly learns the sum of the two structural parameters. She also becomes perfectly confident of her estimates of the individual parameters. However, these estimates are in fact noisy, and incorrect with probability one. The learning process also exhibits order effects. Early observations effectively get more weight than late ones, and they have a non-vanishing contribution to the limit belief.

5.2. **A "Heckman" Selection Model.** In this example, there are three statistic variables, s_1, s_2 and s_3, where $s_1, s_2 \in \{0,1\}$ and $s_3 \in \mathbb{R}$. The statistic s_1 indicates whether an agent enters some market ($s_1 = 1$ means entry). The statistic s_3 represents the agent's income. The statistic s_2 is an exogenous variable that may affect both the entry decision and the income conditional on entry. Data about income is available only for agents who enter the market.

This is a classic problem of drawing causal inferences from a selective sample. To deal with it, our researcher has two feasible identification strategies. First, she can assume that market entry is purely random, thus assuming away selective entry. Second, she can make a structural assumption 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 is uniformly distributed over $\{0,1\}$. Second,

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

Finally, given $s_1 = 1$ and each $s_2 = 0, 1$,

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

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

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

There are two feasible identifying assumptions. One assumption is $\theta^* = 0$, i.e., market entry is independent of u. Under this "research design" assumption, $\mathbb{E}\left[u|s_1 = 1, s_2\right] = 0$ for every s_2, such that the long-run average of s_3 given $s_1 = 1$ and s_2 is equal to $\omega_1 + \omega_2 s_2$. This gives two equations with two unknowns, which enables the researcher to pin down ω_1. An alternative assumption is $\omega_2^* = 0$. This is a "structural" assumption because it pins down a value of one of the fixed parameters. The assumption means that the exogenous variables that may affect market entry do not have a direct causal effect on income conditional on entry. It is an "exclusion" restriction that turns s_2 into a valid instrument for estimating ω_1, albeit with different parameterization than in the IV example we examined in Section 4.

The structural identification method is based on Heckman's correction method (Heckman, 1979). For the sake of tractability, we simplified the model by admitting no structural parameters into the distribution of s_1 conditional on s_2. This enables us to treat $\mathbb{E}\left[u|s_1 = 1, s_2\right]$ 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.

At 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. The following result characterizes the researcher's selection strategy.

Proposition 6. *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 research-design assumption $\theta^* = 0$. In contrast, when entry is highly selective, the researcher passes or imposes the structural assumption $\omega_2^* = 0$. Her willingness to impose the structural 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 research-design 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 characterization can lead to 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 research-design 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 research-design assumption. Thus, the researcher's predilection to stick to a particular identifying strategy for a long stretch of time is history-dependent.

6. Conclusion

The ethos of scientific inquiry involves the pursuit of evidence-based consensus answers to research questions. However, empirical observations are often open to multiple interpretations. Research communities employ assumptions in order to extract an unequivocal interpretation of data, such that repeated observations will lead to a consensus among researchers, regardless of their subjective prior beliefs. Assumptions are rarely undisputed.

However, researchers are willing to make them if they find them plausible enough. This paper articulated this process of assumption-based learning and explored how it affects the rate of learning and the long-run beliefs it may induce.

A key feature of the model is that the plausibility of assumptions is evaluated on a case-by-case basis, one research opportunity at a time. We believe this feature approximates actual practice in the empirical research fields we are familiar with as economists. One lesson from our findings is increased support for meta-studies. When the identification power of an assumption and its plausibility are examined in the context of an isolated research design, this leads researchers to impose strong assumptions in order to get identification. However, as we pointed out in the context of our Contaminated Experiment example, when researchers consider multiple research designs with different context parameters, they can get identification without being forced to make strong assumptions. This suggests that our practice of thinking about pieces of research in isolation leads to biases in the process of scientific learning.

The pursuit of clear-cut answers to questions is not particular to scientific researchers: Ordinary people seek them in their everyday decisions. From this point of view, our learning model also sheds light on how individual decision-makers learn from observations. It departs from the model of Bayesian learning under misspecified prior beliefs in two respects. First, it assumes that agents learn only when they can draw clear-cut conclusions from the data. Second, the misspecified beliefs are endogenous, resulting from assumptions that agents make in order to be make clear-cut inferences.

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 $(h^t, s, s', s'', ...)$).

A.1. **Proof of Proposition 1.** Fix a history h^{t+1} and θ so that $\theta \in \Theta^R(\mu(h^{t+1})) \setminus \Theta^R(\mu(h^t))$. Adopt the *sup* metric throughout.

For any history h^τ, it holds that

$$D_{KL}\left(p_{S,U}\left(\cdot|\theta,h^\tau\right)||p_{S,U}\left(\cdot|\theta^*,h^\tau\right)\right) \leq \sum_s D_{KL}\left(p_{S,U}\left(\cdot|\theta,h^\tau,s\right)||p_{S,U}\left(\cdot|\theta^*,h^\tau,s\right)\right)p_S\left(s|\theta^*,h^\tau\right).$$

by convexity of relative entropy (Theorem 2.7.2 of Cover and Thomas (2006)). In particular, if

$$\delta \leq D_{KL}\left(p_{S,U}\left(\cdot|\theta,\tilde{h}^\tau\right)||p_{S,U}\left(\cdot|\theta^*,\tilde{h}^\tau\right)\right)$$

for some δ and history $\tilde{h}^\tau$, then there exists $s' \in S$ such that

$$\delta \leq D_{KL}\left(p_{S,U}\left(\cdot|\theta,\left(\tilde{h}^\tau,s'\right)\right)||p_{S,U}\left(\cdot|\theta^*,\left(\tilde{h}^\tau,s'\right)\right)\right).$$

By assumption,

$$D_{KL}\left(p\left(\cdot|\theta,\left(h^t,s\right)\right)||p\left(\cdot|\theta^*,\left(h^t,s\right)\right)\right) \leq K+\epsilon < D_{KL}\left(p\left(\cdot|\theta,h^t\right)||p\left(\cdot|\theta^*,h^t\right)\right)$$

for some $\epsilon > 0$, so there exists $s^{t+1}, s^{t+2}, \ldots$ so that

$$K+\epsilon < D_{KL}\left(p\left(\cdot|\theta,\left(h^t,s^{t+1},s^{t+2},\ldots,s^{t+m}\right)\right)||p\left(\cdot|\theta^*,\left(h^t,s^{t+1},s^{t+2},\ldots,s^{t+m}\right)\right)\right)$$

for every $m > 1$.

For large m, $p\left(\cdot|\theta,\left(h^t,s^{t+1},s^{t+2},\ldots,s^{t+m},s\right)\right)$ and $p\left(\cdot|\theta^*,\left(h^t,s^{t+1},s^{t+2},\ldots,s^{t+m},s\right)\right)$ are arbitrarily close to $p\left(\cdot|\theta,\left(h^t,s^{t+1},s^{t+2},\ldots,s^{t+m}\right)\right)$ and $p\left(\cdot|\theta^*,\left(h^t,s^{t+1},s^{t+2},\ldots,s^{t+m}\right)\right)$. Note that $D_{KL}(p||q)$ is continuous in both p and q. Moreover, both are invariant to permutations of s^{t+i}. Therefore, for sufficiently large m,

$$D_{KL}\left(p\left(\cdot|\theta,\left(h^t,s,s^{t+1},s^{t+2},\ldots,s^{t+m}\right)\right)||p\left(\cdot|\theta^*,\left(h^t,s,s^{t+1},s^{t+2},\ldots,s^{t+m}\right)\right)\right) > K,$$

that is, $\theta \notin \mathit{\Theta}^R\left(\mu\left(h^t,s,s^{t+1},s^{t+2},\ldots,s^{t+m}\right)\right)$. We conclude that

$$\mathbb{P}\left(\theta \notin \mathit{\Theta}^R\left(\mu\left(h^{t+m+1}\right)\right)|h^t\right) > \mathbb{P}\left(\left(h^t,s,s^{t+1},s^{t+2},\ldots,s^{t+m}\right)|h^t\right) > 0.$$

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 $\theta^1 \in \Theta^R\left(\mu\left(h^t,s\right)\right) \setminus \Theta^R\left(\mu\left(h^t\right)\right)$ so

$$D_{KL}\left(p\left(\cdot|\theta^1,h^t\right)||p\left(\cdot|\theta^*,h^t\right)\right) > K > D_{KL}\left(p\left(\cdot|\theta^1,\left(h^t,s\right)\right)||p\left(\cdot|\theta^*,\left(h^t,s\right)\right)\right) = \Delta.$$

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

$$D_{KL}\left(p\left(\cdot|\theta, h^t\right)||p\left(\cdot|\theta^*, h^t\right)\right) = K,$$

and so $\theta \in \Theta^R(\mu(h^t))$. By quasi-convexity, $D_{KL}\left(p\left(\cdot|\theta,(h^t,s)\right)||p\left(\cdot|\theta^*,(h^t,s)\right)\right) < K$. By convexity of relative entropy (Theorem 2.7.2 of Cover and Thomas (2006)),

$$\begin{aligned} K &= D_{KL}\left(p\left(\cdot|\theta, h^t\right)||p\left(\cdot|\theta^*, h^t\right)\right) \\ &\leq \sum_S D_{KL}\left(p\left(\cdot|\theta, \left(h^t, \tilde{s}\right)\right)||p\left(\cdot|\theta^*, \left(h^t, \tilde{s}\right)\right)\right) p_S(\tilde{s}|\theta^*, h^t) \\ &< \sum_{\tilde{s}\neq s} D_{KL}\left(p\left(\cdot|\theta, \left(h^t, \tilde{s}\right)\right)||p\left(\cdot|\theta^*, \left(h^t, \tilde{s}\right)\right)\right) p_S(\tilde{s}|\theta^*, h^t) + \Delta p_S(s|\theta^*, h^t) \end{aligned}$$

Therefore, $D_{KL}\left(p\left(\cdot|\theta,(h^t,s')\right)||p\left(\cdot|\theta^*,(h^t,s')\right)\right) > K$ for some $s' \in S\backslash\{s\}$. This also holds for all θ' sufficiently close to θ, including some of those for which $D_{KL}\left(p\left(\cdot|\theta', h^t\right)||p\left(\cdot|\theta^*, h^t\right)\right) < K$. Therefore, $\Theta^R\left(\mu\left(h^t\right)\right)\backslash\Theta^R\left(\mu\left(h^{t+1}\right)\right) \neq \emptyset$ with probability of at least $p\left(s', \Theta^R\left(\mu\left(h^t\right)\right)|\omega\right)$ given state ω and history h^t.

A.3. **Proof of Proposition 3.** Let p be consistent with a DAG G as in the statement of the Theorem. We introduce a few pieces of DAG-based notation. First, let N^ω be the set of nodes that represent the structural parameters. In the same manner, define N^u, N^s and N^θ. Second, for any subset of graph nodes $M \subseteq N$, we use the shorthand notation ω_M for $\omega_{M\cap N^s}$. In the same manner, define u_A, s_M and θ_M. The proof proceeds step-wise.

Step 1: The researcher never learns anything about ω_{-A}.

Proof. By definition of A,

$$p(s|\theta^*, \omega_A, \omega_{-A}) = p(s|\theta^*, \omega_A)$$

for almost every s. Because $\omega_A \perp \omega_{-A}$,

$$\begin{aligned} p(\omega_{-A}|s,\theta^*) &= \frac{p(\omega_{-A})\int_{\omega_Q} p(s|\theta^*, \omega_A, \omega_{-A})p(\omega_A)d\omega_A}{\int_\omega p(s|\theta^*, \omega_A, \omega_{-A})p(\omega_A)p(\omega_{-A})d\omega} \\ &= \frac{p(\omega_{-A})\int_{\omega_J} p(s|\theta^*, \omega_A)p(\omega_A)d\omega_A}{\int_\omega p(s|\theta^*, \omega_A)p(\omega_A)p(\omega_{-A})d\omega} = p(\omega_{-A}), \end{aligned}$$

for almost every s. Therefore, beliefs about ω_{-A} are almost always 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: $I \cap A = \emptyset$.

Proof. For contradiction, suppose that $\omega_i \in I \cap A$. 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 $n = 3$.

$$\begin{array}{ccccccc} \omega^{i_1} & & \omega^{i_2} & & \omega^{i_3} & & \theta \\ \downarrow & \swarrow & \downarrow & \swarrow & \downarrow & \swarrow & \\ s^{i_1} & & s^{i_2} & & s^{i_3} & & \end{array}$$

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.[4] First, take every triple of nodes i,j,k such that $i,j \in R(k)$ whereas i and j are unlinked. 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 is blocked by 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

[4]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.

illustration, note that procedure generates the following moral graph from the DAG above:

$$
\begin{array}{ccccccc}
\omega^{i_1} & - & \omega^{i_2} & - & \omega^{i_3} & - & \theta \\
| & / & | & / & | & / & \\
s^{i_1} & & s^{i_2} & & s^{i_3} & &
\end{array}
$$

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

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.

Proof. 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, $p(s, u|\theta^t, h^t) = p(u)p(s|\theta^t, h^t)$ and we can write $p(s|\theta^t, h^t)$ equals

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

where the third inequality follows from the relationship between N^I and I we articulated above.

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

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

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

Step 4: Completing the proof

Proof. Let $R\left(N^I\right) = \bigcup_{j\in N^I} R(j)$. Suppose s_k is represented by a node 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, $s_k \perp \omega_A$. This means that s_k cannot be a descendant of any node in ω_A according to G. It follows that the parents of s_k also cannot be descendants of nodes in ω_A. Therefore, for every s_j node in $N^I \cup R\left(N^I\right)$, $p\left(s_j|\left(s,u,\omega,\theta^t\right)_{R(j)}\right)$ is constant in ω_A, and so by Step 1, $p\left(s_{N^I\cup R(N^I)}|h^t,\theta^t\right) = p\left(s_{N^I\cup R(N^I)}|\theta^t\right)$ for every history h^t.

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

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

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 D_{KL} are integrated out, and we can replace $p(s|h^t,\theta^t)$ with $p\left(s_{N^I\cup R(N^I)}|\theta^t\right)$ 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 KL divergence simplifies into the following history-independent expression

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

completing the proof. □

A.4. **Proof of Proposition 4.** Suppose not, so $\mathbb{P}\left(\lim_t|\mu_t-\mu^*|=0|\omega^*\right)>0$ and $\mu^*(w)>0$ for some w that does not minimize divergence. Pick any $\hat{w}$ that does. Let H be the set of histories for which $\lim_t|\mu_t-\mu^*|=0$. Now,

$$\frac{\mu_{t+1}(\hat{w})}{\mu_{t+1}(w)} = \frac{\mu_t(\hat{w})}{\mu_t(w)}\frac{p(s^t|\hat{w},\theta^*)}{p(s^t|w,\theta^*)}$$

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

Therefore in the history $h^T = (a^1, s^1; a^2, s^2; \dots)$ we have

$$\begin{aligned}\ln \frac{\mu_{t+1}(\hat{w})}{\mu_{t+1}(w)} &= \ln \frac{\mu_t(\hat{w})}{\mu_t(w)} + \mathbb{I}_{\Theta^R(\mu(h^t))}(\theta^t) \ln \frac{p(s^t|\hat{w},\theta^*)}{p(s^t|w,\theta^*)} \\ &= \ln \frac{\mu_0(\hat{w})}{\mu_0(w)} + \sum_{\tau=1}^{t} \mathbb{I}_{\Theta^R(\mu(h^\tau))}(\theta^\tau) \ln \frac{p(s^\tau|\hat{w},\theta^*)}{p(s^\tau|w,\theta^*)}.\end{aligned} \tag{3}$$

Let

$$\bar{l}\left(\mu\left(h^t\right)\right) = E\left[\ln \frac{p(s^t|\hat{w},\theta^*)}{p(s^t|w,\theta^*)} \mathbb{I}_{\Theta^R(\mu(h^t))}(\theta^t)|\omega^*\right] = \int_{\Theta^R(\mu(h^t))} \left[\sum_{\sigma\in S} \ln \frac{p(\sigma|\hat{w},\theta^*)}{p(\sigma|w,\theta^*)} p\left(\sigma|\theta,\omega^*\right)\right] dp(\theta)$$

Then,

$$\begin{aligned}\frac{1}{t}\ln \frac{\mu_{t+1}(\hat{w})}{\mu_{t+1}(w)} =& \frac{1}{t}\left[\frac{\mu_0(\hat{w})}{\mu_0(w)} + \sum_{\tau=1}^{t} \bar{l}\left(\mu\left(h^\tau\right)\right)\right] \\ &+ \frac{1}{t}\sum_{\tau=1}^{t}\left[\mathbb{I}_{\Theta^R(\mu(h^\tau))}(\theta^\tau) \ln \frac{p(s^\tau|\hat{w},\theta^*)}{p(s^\tau|w,\theta^*)} - \bar{l}\left(\mu\left(h^\tau\right)\right)\right].\end{aligned}$$

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

$$\frac{1}{t}\sum_{\tau=1}^{t}\left[\mathbb{I}_{\Theta^R(\mu(h^\tau))}(\theta^\tau) \ln \frac{p(s^\tau|\hat{w},\theta^*)}{p(s^\tau|w,\theta^*)} - \bar{l}\left(\mu\left(h^\tau\right)\right)\right] \to 0 \tag{4}$$

almost surely given ω^* and that $h^\tau \in H$. It follows from $\mathbb{P}\left(\lim_t |\mu_t - \mu^*| = 0|H\right) = 1$ and continuity of $\Theta^R(\cdot)$ at μ^* that

$$\mathbb{P}\left(\lim_\tau \left|\bar{l}\left(\mu\left(h^\tau\right)\right) - \bar{l}\left(\mu^*\right)\right| = 0|H\right) = 1.$$

Since $\hat{w}$ minimizes divergence,

$$D_{KL}\left(p(s|\Theta^R\left(\mu^*\right), \omega^*)||p(s|\theta^*, \hat{w})\right) < D_{KL}\left(p(s|\Theta^R\left(\mu^*\right), \omega^*)||p(s|\theta^*, w)\right),$$

and so $\bar{l}(\mu^*) > 0$. Therefore, $\mathbb{P}\left(\lim_t \left|\ln \frac{\mu_{t+1}(\hat{w})}{\mu_{t+1}(w)} - t\bar{l}(\mu^*)\right| = 0|H\right) = 1$, contradicting that $\mu^*(w) > 0$.

A.5. **Proof of Proposition 5.** Under the identifying assumption that $\omega_{-i} = m^t_{-i}$, beliefs evolve so that $m^{t+1}_{-i} = m^t_{-i}$ and

$$\begin{aligned} m^{t+1}_i &= m^t_i + \frac{(\sigma^t_i)^2}{(\sigma^t_i)^2+1}\left(s^t - m^t_1 - m^t_2\right) \\ &= \frac{1}{(\sigma^t_i)^2+1} m^t_i + \frac{(\sigma^t_i)^2}{(\sigma^t_i)^2+1}\left(s^t - m^t_{-i}\right). \end{aligned}$$

Suppose that $\left(\sigma^0_i\right)^2 = v$ for $i = 1, 2$, and that K is large enough that research is conducted at $t = 1$. W.l.o.g, the researcher updates her beliefs over ω_1 (ω_2) in odd (even) periods.

Break the time periods into blocks: block 1 corresponds to $t = 1, 2$; block 2 corresponds to $t = 3, 4$; etc. Let $s(\tau, k)$ denote the s realization in part k of block τ. Then, the variance after block τ is

$$\sigma^2_1(\tau) = \sigma^2_2(\tau) = \frac{v}{1+\tau v}$$

Denote

$$\alpha_\tau = \frac{1+\tau v}{1+(1+\tau)v}$$

The updated means $m_1(\tau+1)$ and $m_2(\tau+1)$ at the end of block $\tau+1$ are given by

$$m_1(\tau+1) = \alpha_\tau m_1(\tau) + (1-\alpha_\tau)(s(\tau+1,1) - m_2(\tau)) \tag{5}$$

and

$$\begin{aligned} m_2(\tau+1) &= \alpha_\tau m_2(\tau) + (1-\alpha_\tau)(s(\tau+1,2) - m_1(\tau+1)) \\ &= \alpha_\tau m_2(\tau) + (1-\alpha_\tau)(s(\tau+1,2) - m_1(\tau)) - (1-\alpha_\tau)(s(\tau+1,1) - m_1(\tau)) - m_2(\tau)) \\ &= (\alpha_\tau - (1-\alpha_\tau)^2) m_2(\tau) + (1-\alpha_\tau)s(\tau+1,2) - (1-\alpha_\tau)^2 s(\tau+1,1) - (1-\alpha_\tau)\alpha_\tau m_1(\tau). \end{aligned}$$

Add up the two equations for $m_i(\tau+1)$ and denote

$$\begin{aligned} x(\tau+1) &= m_1(\tau+1) + m_2(\tau+1) \\ &= \alpha^2_\tau x(\tau) + (1-\alpha^2_\tau)\left[\frac{\alpha_\tau}{1+\alpha_\tau} s(\tau+1,1) + \frac{1}{1+\alpha_\tau} s(\tau+1,2)\right]. \end{aligned}$$

We first consider the distribution of $x(\tau+1)$, then that of $m_i(\tau)$.

Since $x(0)$ is a given constant, we can write

$$x(1) = \beta_0^1 x(0) + \beta_1^1 s_1 + \beta_2^1 s_2$$

with $\beta_1^2, \beta_2^2 \leq 1 - \alpha_0$. For $\tau \geq 1$, suppose that

$$x(\tau) = \beta_0^\tau x(0) + \beta_1^\tau s_1 + \cdots + \beta_{2\tau}^\tau s_{2\tau}$$

with $\beta_j^\tau \leq 1 - \alpha_{\tau-1}$ for each $j > 0$. Then,

$$x(\tau+1) = \alpha_\tau^2(\beta_0^\tau x(0) + \beta_1^\tau s_1 + \cdots + \beta_{2\tau}^\tau s_{2\tau}) + (1-\alpha_\tau)s_{2\tau+1} + (1-\alpha_\tau)\alpha_\tau s_{2\tau+2}.$$

For all $0 < j \leq 2\tau$, when we let

$$\beta_j^{\tau+1} \equiv \alpha_\tau^2 \beta_j^\tau \leq \alpha_\tau \beta_j^\tau \leq \alpha_\tau \left(1 - \alpha_{\tau-1}\right) = \frac{1+\tau v}{1+(1+\tau)v} \cdot \frac{1}{1+\tau v} = 1 - \alpha_\tau,$$

it follows that

$$x(\tau+1) = \beta_0^{\tau+1} x(0) + \beta_1^{\tau+1} s_1 + \cdots + \beta_{2\tau+2}^{\tau+1} s_{2\tau+2} \tag{6}$$

with $\beta_j^{\tau+1} \leq 1 - \alpha_\tau$ for all $j > 0$.

By the above, $x(\tau)|\omega \sim N\left(m_\tau, v_\tau\right)$ with

$$v_\tau \leq \sum_{j=1}^{2\tau} \left(\beta_j^\tau\right)^2 \leq 2\tau \left[1 - \alpha_{\tau-1}\right]^2 = \frac{2\tau}{\left(1+(1+\tau)v\right)^2}$$

for all $\tau > 1$. This upper bound tends to zero as $\tau \to \infty$. Finally, notice that

$$\beta_0^{\tau+1} = \prod_{j=1}^{\tau+1} \alpha_\tau^2 = \prod_{j=1}^{\tau+1} \frac{(1+jv)^2}{\left(1+(1+j)v\right)^2} = \left(\frac{v+1}{\left(\tau+2\right)v+1}\right)^2 \to 0$$

as $\tau \to \infty$, and that $\beta_0^{\tau+1} + \beta_1^{\tau+1} + \cdots + \beta_{2\tau+2}^{\tau+1} = 1$. Therefore, in the $\tau \to \infty$ limit, $x(\tau+1)$ in (6) is a convex combination of s realizations. Hence, $x(\tau+1) \to \mathbb{E}[s_i|\omega] = \omega_1 + \omega_2$.

We now turn to beliefs about ω_i. Using recursive substitutions of Equation (5), we show by induction that

$$m_i(\tau) = k_0^{i,\tau} + (-1)^i \sum_{j=1}^{\tau} k_{j,2}^{i,\tau} s(j,2) + (-1)^{i+1} \sum_{j=1}^{\tau} k_{j,1}^{i,\tau} s(j,1) \tag{7}$$

for some $k^{i,\tau}_{j,h} \in [(1-\alpha_j)\alpha_j, 1-\alpha_j]$ for $j < \tau$, $k^{1,\tau}_{\tau,2} = 0$, and $k^{1,\tau}_{\tau,1} = k^{2,\tau}_{\tau,2} = 1-\alpha_\tau$. In particular, $m_1(\tau)$ is increasing in odd signals and decreasing in even signals, and vice versa for $m_2(\tau)$. If true, then non-vanishing weight gets placed on every signal.

Equation (7) holds with weights in appropriate bounds for $m_1(1)$ since

$$m_1(1) = (1-\alpha_1)s_1 + k^{1,1}_0$$

with $k^{1,1}_0 = \alpha_1 m_1(0)$, $k^{1,1}_{1,1} = (1-\alpha_1)$ and $k^{1,1}_{1,2} = 0$. Also for $m_2(1)$ since

$$m_2(1) = (1-\alpha_1)s_2 - \alpha_1(1-\alpha_1)s_1 + k^{2,1}_0$$

with $k^{2,1}_0 = \alpha_1 m_2(0)$, $k^{2,1}_{1,1} = \alpha_1(1-\alpha_1)$ and $k^{2,1}_{1,2} = (1-\alpha_1)$.

Assume that there exist weights $k^{i,\tau}_{j,h}$ as claimed so that equation (7) holds for τ and $i = 1, 2$. Substituting the inductive hypothesis into equation (5),

$$\begin{aligned} m_1(\tau+1) =& \alpha_\tau m_1(\tau) + (1-\alpha_\tau)s(\tau+1,1) - (1-\alpha_\tau)m_2(\tau) \\ =& \sum_{j=1}^{\tau}[\alpha_\tau k^{1,\tau}_{j,1} + (1-\alpha_\tau)k^{2,\tau}_{j,1}]s(j,1) + (1-\alpha_\tau)s(\tau+1,1) \\ & - \sum_{j=1}^{\tau}[\alpha_\tau k^{1,\tau}_{j,2} + (1-\alpha_{\tau+1})k^{2,\tau}_{j,2}]s(j,2) + [\alpha_\tau k^{1,\tau}_0 - (1-\alpha_\tau)k^{2,\tau}_0]. \end{aligned}$$

Equation (7) holds for $\tau+1$ and $i = 1$ when we let $k^{1,\tau+1}_0 = \alpha_\tau k^{1,\tau}_0 - (1-\alpha_\tau)k^{2,\tau}_0$, $k^{1,\tau+1}_{\tau+1,1} = (1-\alpha_\tau)$, $k^{1,\tau+1}_{\tau+1,2} = 0$, and $k^{1,\tau+1}_{j,h} = \alpha_\tau k^{1,\tau}_{j,h} + (1-\alpha_\tau)k^{2,\tau}_{j,h}$ for $h = 1, 2$ and $j \leq \tau$. These are clearly within the bounds. Similarly,

$$\begin{aligned} m_2(\tau+1) =& \sum_{j=1}^{\tau}[\alpha_\tau k^{2,\tau}_{j,2} + (1-\alpha_\tau)k^{1,\tau}_{j,2}]s(j,2) + (1-\alpha_\tau)\, s(\tau+1,2) - \alpha_\tau(1-\alpha_\tau)s(\tau+1,1) \\ & - \sum_{j=1}^{\tau}[\alpha_\tau k^{2,\tau}_{j,1} + (1-\alpha_{\tau+1})k^{1,\tau}_{j,1}]s(j,1) + [\alpha_\tau k^{2,\tau}_0 + (1-\alpha_\tau)k^{1,\tau}_0] \end{aligned}$$

so $k^{2,\tau+1}_{j,h}$ can be defined in a similar way so that equation (7) holds for $\tau+1$ and $i = 2$. Inductive arguments extend the formula to all τ.

Now, observe that $m_i(\tau)$ is a normally distributed random variable. Conditional on $\omega_1+\omega_2$, its variance is bounded from below by, say, $(k^{i,\tau+1}_{1,1})^2 \geq ((1-\alpha_1)\alpha_1)^2 > 0$. It is bounded from above by

$$\sum_{j=1}^{\tau-1}[(k_{j,1}^{i,\tau})^2+(k_{j,2}^{i,\tau})^2]\le 2\sum_{j=1}^{\infty}(1-\alpha_j)^2=2\sum_{j=1}^{\infty}\left(\frac{v}{1+jv}\right)^2.$$

This sum converges by the integral rule.

A.6. **Proof of Proposition 6.** 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 ith dimension. Slightly abusing notation,[5]

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

and

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

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,

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

where

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

and

$$D_{S_1|S_2,U}(\theta)=\int\frac{1}{2}\phi(u)\left(h\left(\theta\Phi(-1-u)+(1-\theta)\Phi(-1),\Phi(-1)\right)+h\left(\theta\Phi(-u)+(1-\theta)\frac{1}{2},\frac{1}{2}\right)\right)du$$

is 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

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

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

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

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 $y=1$, such that 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 $\bar{\theta}^{RD}(\eta)$ and $\bar{\theta}^{S}(\eta)$. 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 $\bar{\theta}^{RD}(\eta)$ and $\bar{\theta}^{S}(\eta)$. Finally, S strictly increases in σ_2^2, and R strictly decreases in it. Therefore, an increase in σ_2^2 leads to a decrease in both $\bar{\theta}^{RD}(\eta)$ and $\bar{\theta}^{S}(\eta)$.

Appendix B. An Example: Front-Door Causal Identification

Consider a data-generating process described by the following system of recursive equations:

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

As in previous examples, set the primitives such that every statistic is standard normal. Let $Q=\{2,4\}$, i.e., the researcher wants to learn ω_2 and ω_4 (which determine the causal effect of s_1 on $s_2$and the causal effect of s_2 on s_3).

The DAG underlying this process is

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

Since the s variables are all standard normal, the only aspects of the long-run distribution of s that the researcher can use to infer ω_2 are $E(s_1 s_2)$, $E(s_2 s_3)$ and $E(s_1 s_3)$. This gives us three equations with five unknowns, and ω_Q cannot be identified. However, when we make the assumption $\theta^* = 0$, we get

$$
\begin{aligned}
E(s_1 s_2) &= \omega_2 \\
E(s_3 s_2) &= \omega_4 + \omega_1 \omega_2 \omega_5 \\
E(s_1 s_3) &= \omega_2 \omega_4 + \omega_1 \omega_5
\end{aligned}
$$

such that

$$
\begin{aligned}
\omega_2 &= E(s_1 s_2) \\
\omega_4 &= \frac{E(s_3 s_2) - E(s_1 s_2) E(s_1 s_3)}{1 + (E(s_1 s_2))^2}
\end{aligned}
$$

This identification strategy, known as the "front door" method (see Pearl (2009)), is based on assuming away the causal effect of u on s_2.

In the DAG, both θ and ω_2 send direct links into s_2. This means that s_2 is not independent of θ conditional on the other variables, but at the same time s_2 is not independent of ω_2. Therefore, the condition for time-invariant propensity to learn is violated. Although we did not establish the tightness of this condition, can be shown directly (albeit tediously) from the formula for the KL divergence that in the present example, the divergence is not invariant to the researcher's belief over ω_Q, which evolves over time. Thus, unlike the causal identification strategies explored in the main text, the propensity to adopt front-door identification is history-dependent.

References

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

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 Hauser. Bounded rationality and learning: A framework and a robustness result. *Econometrica*, 89(1):345–376, 2021.

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

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

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*, 89(4):1631–1672, 2021.

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.

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

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

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.