Learning (to Disagree?) in Large Worlds*

Itzhak Gilboa
HEC Paris and
Tel Aviv University
TzachiGilboa@gmail.com

Larry Samuelson
Department of Economics
Yale University
Larry.Samuelson@yale.edu

David Schmeidler
Tel Aviv University
Tel Aviv 6997801, Israel
DavidSchmeidler@gmail.com

June 30, 2019

Abstract. Beginning with Robert Aumann's 1976 "Agreeing to Disagree" result, a collection of papers have established conditions under which it is impossible for rational agents to disagree, or bet against each other, or speculate in markets. The subsequent literature has provided many explanations for disagreement and trade, typically exploiting differences in prior beliefs or information processing. We view such differences as arising most naturally in a "large worlds" setting, where there is no commonly-accepted understanding of the underlying uncertainty. This paper develops a large-worlds model of reasoning and examines how agents learn in such a setting, with particular interest in whether accumulated experience will lead them to common beliefs (and hence to agree, and to cease trading). No learning process invariably ensures learning, leaving ample room for persistent disagreement and trade. However, there are intuitive processes that lead people with different models of the underlying uncertainty to a common view of the world *if* the data generating process is sufficiently structured.

*We thank Ken Binmore, Michael Greinecker and Jean-Marc Tallon for helpful discussions and insights. Itzhak Gilboa gratefully acknowledges ISF Grant 1077/17, the Investissements d'Avenir ANR-11-IDEX-0003/Labex ECODEC No. ANR-11-LABX-0047, the AXA Chair for Decision Sciences at HEC and the Foerder Institute at Tel-Aviv University. Larry Samuelson thanks the National Science Foundation (SES-1459158) for financial support.

Learning (to Disagree?) in Large Worlds

Contents

Learning (to Disagree?) in Large Worlds

1 Introduction

1.1 Motivation

The point of departure for a collection of "no-trade" results is Aumann's [3] agreeing-to-disagree theorem: if Bayesian agents share a common prior and their posteriors are common knowledge, then these posteriors must be common. Milgrom and Stokey [58] extend this type of reasoning to obtain the canonical no-trade theorem: if a state-contingent allocation is efficient, among (weakly) risk averse agents with concordant priors, and the arrival of new information makes it common knowledge that there exists a (weakly) Pareto improving trade, then it is impossible for that trade to make any agent strictly better off, and hence the agents may as well not trade.

No-trade results challenge our intuition that people readily disagree and trade.[1] A large literature, briefly discussed in Section 1.4, has examined escapes from no-trade results, typically revolving around the idea that people have different beliefs or process information differently.

We view differences in prior beliefs or information processing as arising most naturally in a large-worlds setting, and view agreeing-to-disagree and no-trade results as a clarion call for a large-worlds model of learning. This paper develops and examines such a model.

1.2 Large Worlds

What are large worlds, and why are they relevant?

Savage [68, p. 9] describes a state as "a description of the world, leaving no relevant aspect undescribed", while Arrow [2, p. 45] describes a state as "a description of the world so complete that, if true and known, the consequences of every action would be known." A common paraphrase of these descriptions is that a state specifies the resolution of all *relevant* uncertainty.

In some problems—such as betting on the draw of a ball from an urn or deciding whether to carry an umbrella—the relevant uncertainty and hence the appropriate states are reasonably obviously defined, and one can harness the laws of probability and statistical analysis to form beliefs. The cardinality of the state space may be prodigious, as in the case of a interval describing the amount of rainfall, but neither conceptual nor practical problems preclude imagining the state space and a probability over it. Savage [68] refers to such problems as "small worlds" problems.

[1] For example, Ross [65, p. 94] argues that "it is difficult to imagine that the volume of trade in securities markets has very much at all to do with the modest amount of trading required to accomplish the continual and gradual portfolio rebalancing inherent in our current intertemporal models." French [28, p. 1552] argues that speculative trading is a negative-sum game (after transactions costs), and that "it is hard to understand why equity investors pay to turn their aggregate portfolio over more than two times in 2007."

Other problems call for more intricate states. As Gilboa, Minardi, Samuelson, and Schmeidler [33] explain, the specification of a "state" has inexorably been expanded as the basic decision-theoretic model has been applied to new problems, including dynamic decisions, incomplete information, causal relationships, interfering measurements, and interactive beliefs. These expansions pose no difficulties for Savage's [68, p. 8] *idealized* conception of a state, which in the extreme specifies "[t]he exact and entire past, present, and future history of the universe, understood in any sense, however wide." But we can no longer be confident that different agents will adopt the same view of the "relevant" uncertainty. In addition, there are no commonly-accepted procedures for forming beliefs over such a state space. Indeed, the observations of the statistical procedures by which agents might estimate probabilities may themselves be part of the specification of the state, precluding an attempt to formulate beliefs via statistical analysis. In Savage's [68] terms, we have a "large worlds" problem.

Where is the boundary between small worlds and large worlds? Savage [68, p. 16] offers some examples, noting that playing a chess game and planning a picnic are undertakings sufficiently complicated as to preclude a small worlds model. Our criterion, discussed further in Section 2.4, is that we have a small worlds problem if (but only if) (i) there is a specification of a state space sufficiently obvious that we can reliably expect different agents to adopt this state space, and (ii) we can reliably expect agents to formulate equivalent probability measures over this state space, and indeed to agree on their probability measures if given identical, sufficiently rich data.[2] The distinction between small and large worlds has its grey zones, and it can shift according to the state of knowledge in statistics, machine learning, and science. Decision problems in which many agents make independent choices, such as many consumption decisions, are reasonably viewed as small worlds problems. The occurrence of financial crises and wars are large worlds problems.[3]

With agreement (on the state space) playing a central role in the criterion for a small world, it becomes nearly tautological that rational agents may disagree in large worlds, echoing the well-known result (cf. Section 1.4) that information can lead to trade when agents have different beliefs or process information differently. In contrast to this literature, we focus on the question of when such disagreements are expected to be transient, and what characterizes the "large worlds" in which disagreements may persist even in light of arbitrarily abundant data. To do so, we sweep away other sources of disagreement, such as cognitive limitations, biases in information processing, and insufficient data.

[2] Two probability measures are equivalent if they are mutually absolutely continuous (Billingsley [4, p. 448]). Probability measures that differ to the extent that one attaches positive probability to a set of states S to which others attach zero probability (i.e., that are not equivalent) are effectively a way of smuggling different state spaces into the analysis. Agents' (equivalent) probability measures may differ in a small-worlds setting, perhaps reflecting differences in their information.

[3] Interestingly, the classification of climate change is controversial. Many believe that scientific knowledge, coupled with vast amounts of data, allows one to recover probabilities from simulations. Others, pointing to the divergence of opinion among scientific teams, argue for a more agnostic approach (see Heal and Millner [43]).

1.3 Learning in Large Worlds

Section 2 presents our model of individual reasoning in a large world. In keeping with our interpretation of large worlds, we think of problems such as speculating in financial markets or betting on the identity of the next US president or predicting the effects of global warming, rather than flipping a coin or buying a state lottery ticket or even pricing an automobile insurance policy. Our model captures Bayesian learning as a special case, but allows for a variety of other modes of reasoning.

Section 3 applies the model of reasoning to a simple model of an economy. Section 4 confirms the expected result that in a large-worlds setting, in which different agents in the economy may model the economy differently, the arrival of information can readily give rise to disagreement and trade.

Section 5 turns to a dynamic analysis in which we can address our central question of whether disagreements are likely to be transient or persistent. Section 5.1 considers a small-worlds setting, in which agents are reasonably modeled as Bayesians with a common state space. We invoke standard merging arguments to note that if the agents' beliefs are different but compatible with each other, then according to each of these beliefs, trade will dissipate over time.[4]

Section 5.2 examines learning in a large-worlds setting. There is no analogous merging result for nonBayesian beliefs, even with common support. Indeed, *no* learning process invariably ensures learning, leaving ample room for persistent disagreement and trade. However, Section 5.3 shows that there are intuitive processes that lead people with different models to a common view of the world (and hence to agree, and to cease trading) *if* the data generating process is sufficiently structured, even though different agents employ various different modes of reasoning and potentially shift between modes of reasoning as they learn which is the most appropriate. We examine a setting in which agents attach particular salience to the possibility that they face either a deterministic or an *iid* process. Each agent considers a family of inference rules that may select one of the countably-many possible deterministic theories, or that may approximate the distribution of an *iid* process. The agents may vary in the a priori weights they assign to the various deterministic theories and to the parameters of the *iid* process, *as well as* to the two modes of reasoning. If the agents' beliefs are compatible with each other *and with reality*, learning will eventually take place, and trade will taper off. This conclusion can fail in either of two cases: (i) the agents' beliefs are compatible with each other, but not with reality; that is, the true data generating process is neither deterministic nor *iid*, (ii) the agents are sufficiently open-minded to consider the possibility that the true process is neither deterministic nor *iid*.

A key element of the learning process we describe is an individual's "style of explanation", reflecting their propensity to see either patterns or randomness in the data. We suggest that people differ in their styles of explanation, just as

[4]Standard though often less-emphasized arguments show that trade indeed *will* dissipate over time if the agents' beliefs are also compatible with the "truth", that is, with the underlying process used to analyze the problem.

they do in their tolerance for risk, patience, and ambiguity attitude, and that differences in styles of explanation play an important role in determining how people interact in large worlds. We believe the nature of learning in large worlds is important beyond the more narrow question of what motivates disagreement and trade.

1.4 Literature

This section provides a selected guide to the extensive literature that has responded to the agreeing-to-disagree and no-trade theorems.

To set the stage, suppose that each of a finite number of agents has a (possibly different) partition on a space $\Theta \times T$, where Θ is interpreted as a set of payoff-relevant states and T as a set of signals. A trade is a θ-trade if it can be expressed as a function only on Θ. Suppose we have an allocation that cannot be Pareto improved by any θ-trade. Now suppose that a state is drawn and that each agent is informed of her partition element containing that state. Milgrom and Stokey's [58] no-trade theorem states that if the agents are at least weakly risk averse (so that it is not utility-enhancing to bet on irrelevant signals) and have concordant priors (i.e., agree on the marginal distribution of the signal t, conditional on θ), then this information cannot make it common knowledge that there is a Pareto-improving θ-trade—information cannot give rise to trade.[5]

A first set of explanations for such trade appeals to elements missing from the model. One possibility is that markets are incomplete, so that agents cannot write contracts conditioning consumption on the realizations of the states in Θ, and hence trade occurs continually as information about the state unfolds.[6] Alternatively, there may be "noise traders" in the market who trade for reasons exogenous to the model (cf. Kyle [51]), and who in turn create opportunities for rational agents to profitably trade. Yet again, speculation may be fueled by sources of gains missing from the model, such as implicit bail-out guarantees.[7]

Perhaps the most obvious response to agreeing-to-disagree and no-trade results, developed in detail by Morris [60], is that people may have priors that disagree about the relationship between the payoff-relevant states θ and the signals t. Dow, Madrigal and Ribeiro da Costa Werlang [20] offer a similar argument, but in which beliefs are described by capacities rather than probability distributions, taking us beyond expected utility maximization. Dow, Madrigal and Ribeiro da Costa Werlang [20], Ma [54, Theorem 1] and Halevy [39] examine conditions on (not necessarily expected utility) preferences that lead to no-trade results.

[5] The motivation for the common prior that appears in agreeing-to-disagree theorems is typically traced to Harsanyi [40, Part I (Section 5) and Part III], [41, Section 4.4]. Geanakoplos and Polemarchakis [30] examine a protocol for sharing information that leads to common-knowledge posteriors. Sebenius and Geanakoplos [71] examine a protocol by which it can become common knowledge that an improving trade exists.

[6] Magill and Quinzii [55] examine general equilibrium models with incomplete markets, with Chapter 2 (Section 11) noting that equilibria in such markets are typically inefficient.

[7] See Ritholtz [64] for a popular account.

An alternative to different prior beliefs is a model in which some imperfection in information processing leads people with common beliefs to draw different conclusions from the signals they receive. Geanakoplos [29] examines a model in which an agent's information is described by a possibility correspondence rather than a partition. Brandenburger, Dekel and Geanakoplos [10] pursue this analysis, in the process showing that assuming agents have nonpartitional information structures is equivalent to assuming they have partitions but different prior beliefs. Similar results appear in Samet [66].

Information processing may be imperfect for many reasons. Zimper [78] argues that trade can arise in environments in which agents hold beliefs described by capacities reflecting differing degrees of optimism and pessimism. Eyster, Rabin and Vayanos [24] examine agents subject to correlation neglect, a variation on different prior beliefs in which agents do not perfectly understand the correlation between others' types and their actions.[8] Similar ideas appear in Jehiel and Koessler's [48, Section 7] demonstration that an analogy-based expectations equilibrium can give rise to trade. Brocas, Carrillo, Wang and Camerer [11] argue (and illustrate experimentally) that level-k thinking can give rise to trade under circumstances covered by the no-trade theorem.

The preceding literature typically retains the small-worlds hallmark of assuming the agents work with a common state space $\Theta \times T$, even while allowing prior beliefs or information processing modes to differ. Our contrasting focus on large worlds has antecedents. Blume, Easely and Halpern argue that [9, p. 14] "[m]ost interesting decision problems do not come with a state space ...", and they accordingly examine a setting in which a state space emerges as part of a representation result that begins with observations of behavior. More directly, Binmore [6, 7] argues for a large worlds approach, with motivation very similar to (and influential in developing) that which lies behind our analysis. In contrast to our analysis, Binmore [7] retains as much of the structure of expected utility theory as possible. For example, he retains the sure-thing principle, arguing that it is indispensable in a theory of rationality. We note in Section 2.3.2 that even "Bayesians" in our setting can violate the sure-thing principle.

2 A Model of the Individual

2.1 Structure

We extend the model of Gilboa, Samuelson and Schmeidler [35] by examining quantitative rather than merely qualitative beliefs and adding decisions to the model. We present the components of the model with a minimum of interpretation here and discuss their motivation and foundations in Appendices A.1–A.3.

Time is discrete and indexed by $0, \ldots, T$. We will sometimes consider single-period problems, but will present the model for our primary case of in-

[8] Eyster, Rabin and Vayanos [24] build correlation neglect into their model by examining a cursed equilibrium (Eyster and Rabin [23]). Jehiel [46] introduces the related concept of an analogy-based expectations equilibrium and uses it to examine optimism in Jehiel [47].

terest, in which $T = \infty$.

In each period t a signal x_t is drawn from the finite set X and an output y_t is drawn from the finite set Y. The analyst's state space is thus given by $\Omega = \{X \times Y\}^\infty$. Let $\omega(t) = (\omega_X(t), \omega_Y(t)) \in X \times Y$ identify the period-t values (x_t, y_t) corresponding to state ω. We defer assumptions about the process that determines the state, noting at this point only that the model does not capture endogenous processes, in which the agent's choice at time t affects the realizations at times $t' > t$. As a result, the agent faces no incentive to experiment, i.e., to sacrifice her current payoff in order to affect (or collect information about) the data generating process.

Let h_t denote a period-t history, i.e., a sequence $h_t = (x_0, y_0, \ldots, x_{t-1}, y_{t-1}, x_t)$, and let H_t be the set of such histories. Let

$$[h_t] = \{\omega \in \Omega | (\omega(0), \ldots, \omega(t-1), \omega_X(t)) = h_t\}$$

and, for $Y' \subseteq Y$,

$$[h_t, Y'] = \{\omega \in [h_t] | \omega_Y(t) \in Y'\}.$$

We refer to sets of the latter type as cylinder sets. Notice that $[h_t] = [h_t, Y]$, so that the former sets are special cases of cylinder sets.

In each period t, the following sequence of events transpires:

1. The agent observes a history $h_t = (x_0, y_0, \ldots, x_{t-1}, y_{t-1}, x_t)$.
2. The agent updates her beliefs about the data generating process.
3. The agent forms beliefs about y_t.
4. The agent decides.

2.2 Components

This section presents, in turn, the agent's updating, belief formation, and decision criterion. One might attempt to mimic Savage [68] in jointly characterizing beliefs and utilities, while viewing these components of the model as having (despite their name) no interpretations other than as a mathematically convenient representation of preferences. In contrast, we think of beliefs and the utilities that lie behind the agent's decisions as independently meaningful objects.[9]

We shall see that our model generalizes Bayesian expected utility theory, which is sufficiently widely used as a positive tool that we have no qualms about viewing our exercise as a positive analysis. Indeed, in defending this as a useful model, our concern will not be whether it is sufficiently permissive in the behavior it captures, but whether it captures too much.

[9]Viewing these objects as *independently* meaningful is consistent with Binmore's [6, p. 5] advocacy of *Aesop's principle*, which is that the decision maker's "preferences, her beliefs, and her assessments of what is feasible should all be independent of one another."

2.2.1 Updating Beliefs about the Data Generating Process

A conjecture is a subset $A \subseteq \Omega$, interpreted as an assertion about the world. A conjecture may contain only a single state (say) $\{\omega\}$, in which case it is an assertion that ω is the realized state. A conjecture may also contain multiple states. For example, a conjecture may be the set of states in which a trade war breaks out in the next period, with the conjecture's silence on other matters captured by the fact that the conjecture contains many states. Any conjecture can be described by enumerating the states it contains, but it is often more intuitive to describe conjectures in terms of a common characteristic of those states. For example a conjecture may be a set of states that satisfy some rule, such as that "members of a free-trade association do not impose tariffs on one another". Alternatively, a conjecture may be a set of states for which the current period's output is similar to that of the previous period, or for which the current period's output is similar to that from other periods with similar signals.

We interpret a conjecture as reflecting an argument (or process of reasoning) leading to the assertion embodied in the set of states defining the conjecture, and we interpret a restriction to certain types of conjectures as reflecting a corresponding type of reasoning. For example, we note in Section 2.3 that an agent who invokes only singleton conjectures can be viewed as a Bayesian, while a person who invokes only conjectures reflecting similarities between current and (selected) past outputs can be viewed as a case-based reasoner. However, there is no reason to expect people to restrict themselves to one type of reasoning.

In principle, the set of conjectures is the set 2^Ω of all subsets of Ω. Because we are interested in beliefs about y_t, it sacrifices no generality to take the set of conjectures to be the σ-algebra $\mathcal{A} \subset 2^\Omega$ generated by the cylinder sets.

A credence function ϕ is a probability measure on the measure space $(\mathcal{A}, \mathcal{E})$, where $\mathcal{E} \subseteq 2^{\mathcal{A}}$ is a σ-algebra on the set $\mathcal{A}$. We discuss the σ-algebra $\mathcal{E}$ in Section 2.2.2. We view the credence function ϕ as the agent's model of the situation she faces, or equivalently as identifying the agent's state space.

If we took the credence function to be a probability measure defined on $(\Omega, \mathcal{A})$, then we would have a standard model of Bayesian reasoning, with the credence function ϕ serving as the prior belief. Instead, our interpretation of the probability measure ϕ defined on $(\mathcal{A}, \mathcal{E})$ is that, when asked to assess an event, the agent will bring multiple arguments of various types to bear. One conjecture may reflect the agent's Bayesian belief about the event. Another may embody rules that the agent invokes to assess the event. Yet another may involve comparisons to similar events. Each of these arguments is captured by a conjecture, and the assessment of the event will depend upon the *collection* of such conjectures. This collection a *subset* of $\mathcal{A}$, and hence the credence function is defined on $(\mathcal{A}, \mathcal{E})$.

Given a credence function defined on $(\mathcal{A}, \mathcal{E})$, the agent's updating is straightforward. Having observed a history h_t, the agent discards all conjectures that have been falsified, i.e., the agent discards any conjecture A for which $A \cap [h_t] = \emptyset$. Once a conjecture is falsified, it never again enters into consideration.

2.2.2 Forming Beliefs about y_t

The agent uses the measure ϕ and the history h_t to induce a capacity $\tilde{\phi}_{h_t}$ over the set Y.[10] For each set $Y' \subseteq Y$, the agent identifies the conjectures that support Y', namely

$$\mathcal{A}(h_t, Y') = \{A \in \mathcal{A} | \emptyset \neq A \cap [h_t] \subseteq [h_t, Y']\}.$$

For any conjecture A, the relevant part of that conjecture given history h_t is taken to be $A \cap [h_t]$. This conjecture contributes to the support of the set Y' if and only if is contained in the cylinder set associated with Y', i.e., $A \cap [h_t] \subseteq [h_t, Y']$. The agent attaches to set Y' the (normalized) credence of the set of conjectures supporting Y':

$$\tilde{\phi}_{h_t}(Y') \ := \ \frac{\phi(\mathcal{A}(h_t, Y'))}{\phi(\mathcal{A}(h_t, Y))} \ = \ \frac{\phi(\{A \in \mathcal{A} | \emptyset \neq A \cap [h_t] \subseteq [h_t, Y']\})}{\phi(\{A \in \mathcal{A} | \emptyset \neq A \cap [h_t] \subseteq [h_t, Y]\})}. \tag{1}$$

Our minimum requirement on the set $\mathcal{E}$ is that it contain the σ-algebra generated by the sets

$$\mathcal{A}(h_t, Y'), \quad t \geq 0, \ h_t \in H_t, \ Y' \subseteq Y.$$

These are the sets that we need to be measurable in order for the agent's belief-formation in (1) to be well defined. We will feel free to assume that other sets are contained in $\mathcal{E}$ as we proceed.

2.2.3 Decisions

We let Z be a set of consequences and let Π be the set of finite-support probability distributions over Z. We let $\mathcal{L}$ be the set of functions from Y to Π, associating a probability distribution over Z with each output Y. Given a history h_t, the agent must choose a consumption bundle $\ell \in \mathcal{L}$ (subject to feasibility constraints). To make this decision, the agent appeals to a utility function $\tilde{u} : \mathcal{L} \to \mathbb{R}$ and chooses over $\mathcal{L}$ according to preferences defined by

$$\ell \succsim \ell' \iff \int_Y \tilde{u}(\ell(y)) d\tilde{\phi}_{h_t}(y) \geq \int_Y \tilde{u}(\ell'(y)) d\tilde{\phi}_{h_t}(y), \tag{2}$$

where the integral is a Choquet integral[11] and $\tilde{u}$ is an affine function $\tilde{u} : \mathcal{L} \to \mathbb{R}$. Because $\tilde{u}$ is affine, it is straightforward that there exists a function $u : Z \to \mathbb{R}$ such that

$$\tilde{u}(\ell(y)) = \sum_{z \in Z} u(z) \ell_{(y)}(z), \tag{3}$$

where $\ell_{(y)}(z)$ is the probability attached to consequence z by the lottery $\ell(y)$ associated with output y, and the function $u : Z \to \mathbb{R}$ has the interpretation that $u(z)$ is the utility attached by $\tilde{u}$ to a distribution that puts probability one on consequence z.

---

[10] Beliefs about y_t are thus given by a function $\tilde{\phi}_{h_t} : 2^Y \to \mathbb{R}_+$ that satisfies $\tilde{\phi}_{h_t}(\emptyset) = 0$, $\tilde{\phi}_{h_t}(Y) = 1$, and that is monotone with respect to set inclusion.

[11] Hence, $\int f d\nu = \int_{-\infty}^{0} (\nu(\{s | f(s) \geq x\}) - 1)\, dx + \int_0^{\infty} \nu(\{s | f(s) \geq x\})\, dx$.

2.3 Special Cases

This section illustrates some of the types of reasoning that can arise as special cases of our model.

2.3.1 Bayesians

Define the set of *Bayesian conjectures* to be

$$\mathcal{B} = \{\{\omega\} \mid \omega \in \Omega\} \subset \mathcal{A}.$$

We assume that $\mathcal{B}$ and, for every history h_t the set of unfalsified Bayesian conjectures $\mathcal{B}(h_t) = \{\{\omega\} \mid \omega \in [h_t]\}$, are elements of $\mathcal{E}$. A credence function ϕ is Bayesian if only Bayesian hypotheses matter in determining the weights of credence attached to a set of conjectures, i.e., if for any set $E \in \mathcal{E}$, we have

$$\phi(E) = \phi(E \cap \mathcal{B}).$$

The standard model of Bayesian reasoning captures beliefs via a probability measure φ on $(\Omega, \mathcal{A})$. Gilboa, Samuelson and Schmeidler [35] prove the following straightforward result, which is also a special case of Proposition 7 in Appendix A.1, establishing that a special case of our model is equivalent to the standard model of Bayesian reasoning:

Proposition 1 *Let φ be a probability measure on $(\Omega, \mathcal{A})$. There exists a Bayesian credence function ϕ such that for every history h_t with $\varphi([h_t]) > 0$, there is a constant $\lambda > 0$ for which, for every $Y' \subseteq Y$,*

$$\phi(\mathcal{A}(h_t, Y')) = \lambda\varphi([h_t, Y']|h_t).$$

For any Bayesian belief φ, we can thus find a credence function that attaches identical beliefs to every history to which φ attaches beliefs.

2.3.2 Savage Bayesians

Savage [68] has in mind that people work with states that resolve *relevant* uncertainty. Hence, people neglect some distinctions between Savage's entire-history-of-the-universe states, lumping states together into sets that become the states in a Bayesian model.[12] Toward this end, let $\mathcal{S}$ be a partition of Ω, whose elements we refer to as *Savage Bayesian conjectures.* Some Savage Bayesian conjectures may be singletons, but in general they need not be. We assume that $\mathcal{S} \in \mathcal{E}$, as are sets of the form $\{A \in \mathcal{S} | A \cap [h_t] \neq \emptyset\}$.

[12] For example, when modeling a duopoly with uncertainty about costs, the states are typically taken to identify (only) the cost parameters of the firms. This is not an assertion that the only uncertainty in the world concerns these firms' costs, but rather that all of the states corresponding to a particular cost vector can be lumped together into a single "Savage state", and then the agents are Bayesian over the set of such Savage states.

A credence function ϕ is Savage Bayesian (or, equivalently, beliefs are Bayesian) if only Savage Bayesian conjectures matter in determining the weights of credence attached to a set of conjectures, i.e., if for any set $E \in \mathcal{E}$, we have

$$\phi(E) = \phi(\{\tilde{E} \in \mathcal{S} | \tilde{E} \subseteq E\}).$$

The following straightforward analogue of Proposition 1 (requiring only notational adjustments in the proof) confirms that our model captures Bayesian reasoning, restricting attention to sets in $\sigma(\mathcal{S})$:

Proposition 2 *Let φ be a probability measure on $(\Omega, \mathcal{A})$ that is measurable with respect to $\sigma(\mathcal{S})$. There exists a Savage Bayesian credence function such that for every history h_t with $\varphi([h_t]) > 0$ and $[h_t, Y'] \in \sigma(\mathcal{S})$,*

$$\phi(\mathcal{A}(h_t, Y')) = \lambda\varphi([h_t, Y'] | h_t).$$

Savage Bayesians will look just like Bayesians, as long as they never have to assess events that are not measurable with respect to the sub-σ-algebra of $\mathcal{E}$ corresponding to their partition $\mathcal{S}$ of Ω. However, should an event appear that is not contained in $\sigma(\mathcal{S})$, Savage Bayesians will look decidedly nonBayesian. It is then straightforward to construct examples in which de Finetti's [17] cancellation axiom and Savage's [68] sure-thing principle fail,[13] and in which Savage Bayesians give rise to the Ellsberg paradox (cf. Example 4 in Section 5.2).

2.3.3 Learning from Experts

Cesa-Bianchi and Lugosi [13] provide an introduction to the literature on learning from experts, designed to address relatively unstructured prediction problems in which the agent has no information about the data generating process. The experts appearing in this literature are reminiscent of sets of conjectures, and our model readily captures the various learning procedures.[14] We offer two illustrations.

First, for the simplest example, let X be a singleton and let Y contain two elements. The agent has available a finite number N of experts, with each expert $i \in \{1, \dots, N\}$ predicting an element of Y in each period, as a function of the history h_t. It is known that one of the experts predicts correctly, though the agent has no idea which expert this is. The agent similarly predicts an element of Y in each period, with the objective of minimizing the number of incorrect predictions. A straightforward calculation is that if the agent makes the prediction made by the majority of surviving experts in each period (with an arbitrary

[13] Halevy [38] shows that dynamic consistency and consequentialism imply the sure-thing principle, while preferences that satisfy the sure-thing principle induce conditional preferences satisfying dynamic consistency and consequentialism. The latter conditions play key roles in his no-trade result. It is thus natural that models explaining speculative trade might exhibit violations of the sure-thing principle. We share with Binmore [6] a large-worlds motivations for our work, but we differ in that we do not insist on the sure-thing principle.

[14] It is immediate from Proposition 7 in Appendix A.1 that we can duplicate the resulting sequence of beliefs. The salient point here is that we can capture much of the contending-experts reasoning process.

tie-breaking rule), while discarding any expert who ever makes an incorrect prediction, then the agent can be assured of making at most $\log_2 N$ mistakes. To capture this process, we model each expert i as a singleton conjecture $\{\omega^i\}$.[15] The credence function attaches equal weight to each such conjecture. In each period, the agent then predicts that element of Y predicted by the majority of the surviving conjectures, automatically dismissing falsified conjectures until settling on the true conjecture, and making at most $\log_2 N$ mistakes.

Second, and somewhat more generally, we let each expert $i \in \{1, \dots, N\}$ make a prediction $p_t^i(h_t) \in \Delta Y$ in each period t, as a function of the history h_t. The agent's belief in period t is a weighted average of the predictions. The literature considers a variety of rules (and attendant objectives from which they are derived) for how the weights are adjusted as this history unfolds. To capture such a process, let the set of outputs be given by $W \times Y$, where each element of the finite set W corresponds to an expert. For any history h_t, the rule for adjusting weights over experts induces a probability distribution over W, and the corresponding predictions $p_t^i(h_t) \in \Delta Y$ extend this to a distribution over $W \times Y$. We can then apply Proposition 7 in Appendix A.1 to construct a credence function that gives the appropriate history-dependent probability distributions. As a history unfolds, this credence function will then mimic the learning-from-experts protocol for adjusting the relative influence of the various experts.

2.4 Small or Large Worlds?

When does the credence function ϕ correspond to a small-worlds model? The small-worlds criterion presented in Section 1.2 has two components.

The first small-worlds component concerns our model of individual reasoning. We require that there is a partition of the state space Ω, with the agent's credence function given by a probability measure on the σ-algebra defined by this partition. As Section 2.3.2 explains, such agents are Bayesian as envisioned by Savage (or "Savage Bayesians")—they partition the state space to achieve a manageably small world, the elements of this partition become the states in their model, and their beliefs are given by probability measures on the resulting set of states.

Why do we require Bayesian beliefs? As Appendix A explains in more detail, the weight attached to a conjecture can be interpreted as (i) the support given to any event that is a superset of that conjecture by (ii) arguments that are captured by the conjecture but that cannot be decomposed into any finer elements. Suppose an agent partitions the set Ω into a collection $\mathcal{S}$ of candidate states and forms a credence function ϕ on the σ-algebra $\sigma(\mathcal{S})$ generated by $\mathcal{S}$, but is not a Savage Bayesian, i.e., there exists an element E of $\sigma(\mathcal{S})$ whose weight does not equal the sum of the weights of an exhaustive collection of disjoint subsets of E drawn from $\sigma(\mathcal{S})$. We view this as an indication that the

[15] This seemingly robs experts of the ability to make predictions conditional on history. But since a single incorrect prediction suffices to dismiss an expert from consideration, this restriction sacrifices no generality.

agent lacks some understanding of her environment, in the sense that she has not been able to arrange her partition $\mathcal{S}$ so as to collect together unimportant distinctions and separate important distinctions. In essence, the agent is unable to form a consistent view of "relevant uncertainty", putting us in Savage's large-worlds category.[16]

For example, suppose that an agent can conceive of any state that can be described using the English language, and indeed leaves nothing out of such a description, but knows absolutely nothing, and so her conjecture puts positive weight only on the set Ω. We view this agent as reasoning in a large world, in the sense that she is not able even to get started on Savage's task of identifying the relevant uncertainty and grouping together states that differ only in irrelevancies.

Suppose, instead, that the agent is aggressive in lumping elements of Ω together, forming a model that contains only three states, say "rain", "hail" and "sun". Suppose, however, she has beliefs specified by a capacity that puts probability .2 on each of $\{rain\}$, $\{hail\}$, and $\{sun\}$, and probability .6 on the set $\{hail, sun\}$. This person may have partitioned the state space in such a way as to capture what matters for her payoffs, but cannot express her understanding of the environment in terms of her state space—there are some arguments supporting the states $\{hail, sun\}$ that cannot be further decomposed and hence cannot be associated with elements of her state space. We again view this person as reasoning in a large world.

The second small-worlds component comes into play when we consider multiple agents. We require the agents to formulate the *same* state space (i.e., the same set $\mathcal{S}$) in the course of formulating their models. Given that beliefs are given by probability measures, we require the beliefs to be equivalent—nonequivalent measures allow one agent to attach positive probability to states that another regards as impossible, effectively giving us different state spaces.

We can generally think of people as facing small-worlds problems if they have seen many independent iterations of the problem. We can then expect frequentist arguments to guide agents to (at least approximately) the same understanding of the relevant sources of uncertainty and to (again, at least approximately) Bayesian beliefs about this uncertainty. A professional bidder who has seemingly forever faced the same counterparts in an auction may reasonably be modeled as reasoning in a small world.

At the opposite extreme, we might appropriately model people as facing a large-worlds problem when the entire history of their experience is viewed as a single idiosyncratic occurrence, allowing for a variety of causal (and statistical) dependencies and precluding the collection of multiple, independent observations. For example, a macroeconomic policy intervention (such as lowering interest rates after a crisis) typically occurs in the face of a collection of idiosyncratic details that can only be observed once. In such models data do not identify probabilities, and hence agents might be loathe to commit to a proba-

[16] Mukerji and Shin [62, p. 3], building on Ghirardato [31] and Mukerji [61], explain how beliefs that do not correspond to probability measures are to be interpreted as an indication that "decision makers believe that they know the relevant state space only incompletely".

bility distribution, leading to nonBayesian beliefs or even different conceptions of the relevant states.[17]

The agent's choices ultimately lead to randomizations over the set of consequences Z. Why not simply take the state space to be ΔZ, plausibly putting us back into a small worlds setting? We return to the view that our model is not simply a revealed preference analysis. Instead, we view people as thinking of the world in terms of state spaces and beliefs that the agents regard as objects distinct from probability distributions over consequences, presumably because they believe this is more useful than simply thinking in terms of ΔZ.[18]

2.5 Learning about Credence Functions?

Appendix A elaborates on the foundations of this model and on the reasoning behind our various modeling choices. At this point, we simply draw attention to Proposition 7 in Appendix A.1, establishing that our protocol for updating credence functions imposes virtually no restrictions—for any coherent assignment of capacities over Y to histories, there is a credence function generating that assignment.

Proposition 7 is important in considering the following question. Instead of simply updating their credence functions, could the agents do better by updating at a higher level, sometimes altering the very structure of their credence functions? Is there scope for one agent to say, "Look, my credence function has done much better than yours, so you should switch to mine."?

Proposition 7 ensures that for any model in which such conversions occur, there is an equivalent model exhibiting no such conversions. More precisely, suppose there exist credence functions ϕ and ϕ' and a history h_t such that upon observing h_t, the agent changes her credence function from ϕ to ϕ'. Then there exists an alternative credence function ϕ'' that exhibits behavior identical to that produced by the combination of ϕ and the switch to ϕ' after history h_t. It thus sacrifices no generality to model agents as updating their credence function in light of their experience, but as never altering the structure of their credence functions, even in light of evidence that other credence functions are "performing better".

[17]Our distinction between small and large worlds inevitably leads to some discontinuities. An agent may have nonBayesian beliefs that are arbitrarily close to Bayesian, differing perhaps only in a minuscule weight attached to the "I'm not sure of anything" set Ω, and yet land in the large-worlds category. There is no escaping this if we are to use qualitative terms to describe quantitative objects. We take comfort in the fact that the same is true of some of the most useful concepts in economics.

[18]Similar considerations appear in Blume, Easely and Halpern [9]. In the simplest version of their analysis, the state space is the set of outcomes of the tests the agent can run. When there are multiple agents, one then wonders how the agents decide which tests to run, and whether they will run the same tests, giving rise to large-worlds considerations.

3 Allocations

We now consider an economy in which there are I agents, indexed by $i = 1, \ldots, I$. There is a single good, referred to as consumption or wealth.[19] The set of consequences Z_i for each agent i is $\mathbb{R}_+$. Each agent is characterized by a strictly increasing, concave utility function $u : Z_i \to \mathbb{R}_+$.[20]

In each period, the economy is endowed with a total e_y of the single consumption good in each output y. One interesting special case is that in which e_y is constant in y, but we do not insist on this. Alternatively, we could consider more general specifications that allow the endowment to vary across time and history as well, without affecting the reasoning.

In each period t, each of the I agents first observes the history $h_t = (x_0, y_0, \ldots, x_{t-1}, y_{t-1}, x_t)$, updates her beliefs about the data generating process, forms beliefs about y_t, and then uses the resulting preferences to evaluate consumption bundles. Notice that our agents observe the value of x precisely. This appears to be at odds with a literature that models agents as receiving noisy signals. More importantly, we give the agents a common history, and hence common knowledge. This again appears to be at odds with a literature that asks whether potentially *private* information can give rise to trade. However, the essence of the question in this literature—whether private information can give rise to the common knowledge of trade among agents whose concordant priors give them identical beliefs about the conditional distribution of signals (i.e., about the distribution of x conditional on y in our context)—is most transparently captured by asking whether common information can give rise to (necessarily the common knowledge of) trade. More importantly, our finding in Section 4 that information can give rise to trade in large worlds becomes all the stronger when this information is unambiguous and common.[21] In contrast, the learning results of Section 5 do not require this commonality.[22]

An allocation specifies, for each history h_t and each agent i, a consumption bundle ℓ associating an element of Π (the set of finite-support distributions

[19] Our interest lies in how, given an efficient allocation, differences in information and beliefs give rise for incentives for people to trade. Efficiency immediately ensures that gains from trade *conditional on an outcome* y will be exhausted, with interest centering on how agents trade across outcomes. An examination of the former gains from trade is most naturally conducted in a model with multiple goods, while the latter are most effectively examined in a single-good model.

[20] The requirement that u be (at least weakly) concave is familiar from the no-trade theorem, ensuring that information does not give rise to trade simply because it gives risk-seeking agents more contingencies on which to bet.

[21] Indeed, one technique for generating information-induced trade in conventional models is to make the agents imperfect information processors, breaking the link between the common knowledge of willingness to trade and the common knowledge of beliefs. Our formulation ensures that we have not simply repackaged this possibility. Alternatively, it is clear that with nonconcordant priors, agents who know the realization of y may still exploit differing beliefs about x to construct mutually beneficial trades conditional on x, and our formulation again ensures that we have not simply repackaged this result.

[22] Zimper [78] draws a contrast between cases in which agents have identical partitions and cases in which their partitions differ. In his model, however, the partitions affect the nature of the optimism or pessimism built into agents' updating rules.

over Z_i) with each output y. However, given that u is concave, there is no loss of generality in replacing each such lottery by its expected value. Hence, we can view an allocation as specifying, for each history h_t (the notation for which we suppress), a vector $(z_{iy})_{i=1,\ldots I, y\in Y}$ with z_{iy} identifying the quantity consumed by each agent i given output y in period t, given history h_t. The utility of such an allocation is given by the Choquet integral (cf. (2)–(3))

$$\int_Y u_i(z_{iy}) d\tilde{\phi}_{h_t}. \tag{4}$$

An allocation is feasible if, for every history h_t and output y, we have

$$\sum_{i=1}^{I} z_{iy} = e_y.$$

As usual, an allocation z is Pareto efficient if there is no alternative feasible allocation $\tilde{z}$ that yields a higher value of the expected utility for at least one agent, and a lower value for no agent. We make our efficiency comparisons operational as follows. We initially suppose that for each period t we have a history-contingent allocation $(z_{iy})_{i=1,\ldots I, y\in Y}$ that depends only on the history $(y_0, \ldots, y_{t-1})$, and that is efficient in the space of such allocations. Then, in each period t, having observed the history $h_t = (x_0, y_0, \ldots, x_{t-1}, y_{t-1}, x_t)$, we ask whether there exists a (necessarily common knowledge, given the common information structure) Pareto improvement. If so, we say that information leads to trade.

Remark 1 (Interactions between Periods?) The only interaction between periods arises out of beliefs. The agents cannot transfer goods or consumption across periods, cannot trade assets that have payoff implications across periods, and cannot borrow or save. This is precisely what we need in order to focus attention on the relationship between the information structure and trade. This appears to exclude many important intertemporal questions. However, if interested in such questions, we would interpret the elements of X and Y so as to build the relevant structure into each period. ■

4 Information and Trade in Static Models

In this section we assume that there is only a single period, so that the state space Ω is given by the product $X \times Y$.

The endowment is now an allocation vector $(z_{iy})_{i=1,\ldots I, y\in Y}$ specifying the amount consumed by each agent in each output. We assume the endowment is efficient, in the space of such allocations. Nature then draws a state, and reveals the signal x to the agents. Our central question is: does there now exist a Pareto-improving trade? If so, we say that information leads to trade.

Set	ϕ^m	$\tilde{\phi}$	ψ^m	$\tilde{\psi}$
$\{y'\}$	1/3	1/3	1/3	1/3
$\{y''\}$	0	0	0	0
$\{y'''\}$	0	0	2/3	2/3
$\{y', y''\}$	0	1/3	0	1/3
$\{y', y'''\}$	0	1/3	0	1
$\{y'', y'''\}$	2/3	2/3	0	2/3
$\{y', y'', y'''\}$	0	1	0	1

Figure 1: Beliefs for the example of Section 4.1.1. In the absence of information, beliefs are given by the capacity $\tilde{\phi}$ (with Möbius transform ϕ^m). Upon observing signal x', beliefs are given by the capacity $\tilde{\psi}$ (with Möbius transform ψ^m).

4.1 Common Credence Functions

We begin by investigating the possibility of information-induced trade when agents have common credence functions. We first present an example, showing that if beliefs are given by capacities rather than probability distributions, then information can indeed give rise to trade, even when agents have identical capacities.[23] We then examine the forces behind this example.

4.1.1 An Example

Consider an exchange economy with two agents, two signals, $X = \{x', x''\}$, and three outputs, $Y = \{y', y'', y'''\}$.

The agents have common utility function $u(z) = \ln z$. The endowment, giving consumption for each agent as a function of the output, is given by:

	Output		
Agent	y'	y''	y'''
1	1	1	1
2	1	1	3

The agents share a common credence function, determined by

$$\phi(\{x'y'\}) = \frac{1}{3}, \qquad \phi(\{x'y''', x''y''\}) = \frac{2}{3}.$$

These agents are Savage Bayesians. Before x is observed, these credence func-

[23]Dow, Madrigal and Ribeiro da Costa Werlang [20] offer an example involving capacities in which information gives rise to trade, but in which the agents have different capacities, which are not required to be totally monotone (cf. Appendix A.2). Zimper [78] also gives an example in which different capacities can lead to trade.

tions induce the capacity $\tilde{\phi}$, presented with its accompanying Möbius transform ϕ^m in Figure 1.[24] The endowment is efficient.

Now suppose the agents observe the signal x'. We then have the capacity $\tilde{\psi}$, presented with its accompanying Möbius transform ψ^m in Figure 1. Now there will exist terms on which the agents are willing to trade. Both agents now regard y' and y''' as the only possible outputs, with y''' twice as likely. Agent 2's marginal utility of income in output y''' is a third of that in output y', while agent 1 equalizes the marginal utility of income across signals. It is accordingly Pareto improving for agent 2 to shift consumption to agent 1 in output y''', in return for receiving an equal amount of consumption in output y'.

The forces behind this example are straightforward. Agent 2 is relatively well endowed with consumption in output y'''. One would then expect that there are mutual gains from transferring consumption from agent 2 to agent 1 in output y''', in return for a reverse transfer in output y'. The difficulty is that the agents' priors, described by the capacity $\tilde{\phi}$, give agent 1 no credit for increasing consumption in output y'''. Intuitively, any argument the agents can adduce in support of output y''' also motivates output y'', and there are then no gains to be realized from increasing y''' without also increasing y''. By excluding output y'', the arrival of information makes such gains possible.

4.1.2 No Trade without Aggregate Uncertainty

We explore the result in Section 4.1, coming to the conclusion that it is rather special, by identifying two circumstances that preclude trade. First, there is no scope for information to generate trade if the aggregate endowment is constant across outputs.

Proposition 3 *Suppose that the aggregate endowment is constant across outputs and agents are strictly risk averse. Then information cannot give rise to trade.*

This result directly follows from the literature. First, as we explain in footnote 46 in Appendix A.3, Gilboa and Schmeidler [36] show that the preferences represented by maximizing the Choquet integral of a utility function u against a (totally monotone and hence convex) capacity $\tilde{\phi}$ are equivalent to those represented by the maxmin of expected utility against the set of multiple priors given by the core of $\tilde{\phi}$. Let us accordingly think of our agents as maxmin expected utility maximizers. Since they share a common credence function, the sets of multiple priors with respect to which they are maximizing are identical, and hence certainly contain a nonempty intersection. We can then apply Theorem 1 of Billot, Chateauneuf, Gilboa and Tallon [5, p. 688] to conclude that every efficient allocation exhibits full insurance, and every full insurance allocation is efficient, no matter what the agents' beliefs. Indeed, we see here

[24]The Möbius transform of a capacity on Y is a probability distribution over the subsets of Y. The weight attached to a set Y' by the Möbius transform can be interpreted as the marginal contribution made by the set Y' to the total weight or credence attached to any set Y'' containing Y'. See Appendix A.2 for a definition and details.

that this result requires considerably less than that there are identical credence functions, and requires only a nonempty intersection of the cores of their credence functions. It is then not a coincidence that the aggregate endowment in Section 4.1.1 varies across outputs.

4.1.3 Balanced Allocations

Second, we show that the payoff profile induced by an efficient allocation can be produced by a "balanced" allocation, and that information cannot induce agents characterized by a balanced allocation to trade. We simplify the exposition by assuming utility functions are differentiable and strictly concave.

We say that an allocation is strictly comonotonic, analogously to the definition of comonotonic acts (cf. footnote 45 in Appendix A.3), if for any two agents i and j and outputs $y_k \in Y$ and $y_\ell \in Y$, we have

$$z_{ik} > z_{i\ell} \implies z_{jk} > z_{j\ell}.$$

We say that an allocation is balanced if, for any two agents i and j and outputs y_k and y_ℓ, we have

$$\frac{u_i'(z_{ik})}{u_i'(z_{i\ell})} = \frac{u_j'(z_{jk})}{u_j'(z_{j\ell})}.$$

The strict inequalities defining strict comonotonicity allow some leeway, and so it is no surprise that not every strictly comonotonic allocation is balanced. The straightforward proof of the following is given in Appendix A.4.

Lemma 1 *Let utility functions be differentiable and strictly concave. Then:*

[1.1] All balanced allocations are strictly comonotonic.

[1.2] For any common capacity $\tilde{\phi}$ on Y, every balanced allocation is efficient.

[1.3] For every efficient allocation, there exists a balanced allocation yielding identical payoffs.

The allocation in Section 4.1.1 is not balanced. The lack of balancedness is important to the result. In particular, notice that whether an allocation is balanced is independent of the specification of the capacity $\tilde{\phi}$. Indeed, the argument behind Lemma 1.2 builds on Lemma 1.1 to show that a balanced allocation is efficient for *any* capacity. Hence, given an initial capacity $\tilde{\phi}$ and a balanced allocation, the arrival of information giving rise to any alternative capacity $\tilde{\psi}$ cannot vitiate the efficiency of the allocation and hence cannot give rise to trade:

Corollary 1 *Given a balanced allocation, the arrival of information cannot give rise to trade.*

We thus have a class of allocations, namely the balanced allocations, that suffice to generate the efficient frontier, and that have the property that information cannot lead to trade. However, as we have seen in Section 4.1.1, there are other allocations generating efficient payoffs. These allocations are not balanced. The key to the construction of efficient allocations that are not balanced is that the Möbius transform does not attach positive probability to all subsets of Y, and the failure of balancedness affects payoffs attached to subsets of Y to which the Möbius transform attaches zero probability. Revisions in beliefs can attach positive probability to these subsets, generating trade.

These remarks point to a characteristic of unbalanced efficient allocations. Let us say that an allocation has slack if there are two agents, i and j and an output k such that i can transfer income to j in output k (only) without affecting the payoffs. The allocation in Section 4.1.1 has slack, since agent 2 could transfer income to agent 1 in output y''' without affecting payoffs. Section A.5 proves:

Lemma 2 *An efficient allocation that is not balanced exhibits slack.*

Corollary 2 *Given an efficient allocation, the arrival of information can lead to trade only if the allocation exhibits slack.*

4.1.4 The Example Redux

We return to the example of Section 4.1.1 to illustrate the previous results. This example begins with the allocation

Agent	y'	y''	y'''
1	1	1	1
2	1	1	3

,

with beliefs specified by $\phi^m(\{y'\}) = 1/3$, $\phi^m(\{y'', y'''\}) = 2/3$, and all other subsets of Y receiving probability 0 in the Möbius transform. This allocation, which gives payoffs

$$u_1 = 0, \qquad u_2 = 0,$$

is neither strictly comonotonic nor balanced, and exhibits slack—income can be transferred from agent 1 to agent 2 in output y''' without affecting utilities. The payoffs $(u_1, u_2) = (0, 0)$ appear on the efficient frontier. The balanced payoff allocation giving these payoffs is

Agent	y'	y''	y'''
1	1	1	2
2	1	1	2

.

Balanced allocations sweep out the efficient frontier, which gives agent 2 a payoff of $\ln 2 \approx .69$ in the limit as 1's payoff approaches $-\infty$, and gives agent 1 a payoff of $\ln 2 \approx .69$ in the limit as 2's payoff approaches $-\infty$.

We now suppose that the arrival of information gives us a Möbius transform with $\psi^m(\{y'\}) = 1/3$, $\psi^m(\{y'''\}) = 2/3$, with all other subsets of Y receiving 0. The utilities in the balanced allocation increase, to

$$u_1 = \frac{2}{3}\ln 2 \approx .46, \qquad u_2 = \frac{2}{3}\ln 2 \approx .46 \quad .$$

This allocation lies on the new, higher Pareto frontier and allows no scope for trade. In contrast, the previously efficient unbalanced allocation gives payoffs

$$u_1 = 0, \qquad u_2 = \frac{2}{3}\ln 3 \approx .73, \tag{5}$$

which is inefficient. In comparison, the efficient frontier is given by balanced allocations of the form.

Agent	y'	y''	y'''
1	z	z	$2z$
2	$2-z$	$2-z$	$4-2z$

.

The new efficient frontier has asymptotes $\frac{1}{3}\ln 2 + \frac{2}{3}ln4 \approx 1.15$, and in addition to (approximately)$(.46, .46)$, includes the payoffs (obtained by fixing one agent at the payoff given in (5) and pushing the other agent to the efficient frontier)

$$\begin{aligned} u_1 &= 0, & u_2 &\approx .78 \\ u_1 &\approx .09, & u_2 &\approx .73. \end{aligned}$$

There is thus scope for trade to take us from the payoffs given in (5) to any of the payoffs on the efficient frontier bounded by these two values.

4.2 Different Credence Functions

Once we allow agents to hold different capacities, examples in which information gives rise to trade abound. We offer one such example, involving an economy with two risk averse agents and three outputs, with initial allocation:

	State		
Agent	y'	y''	y'''
1	1	1	1
2	1	1	1

.

Suppose the agents initially have credence functions inducing the capacities and Möbius transforms given by

Set	ϕ_1^m	$\tilde{\phi}_1$	ϕ_2^m	$\tilde{\phi}_2$
$\{y'\}$	1/3	1/3	1/3	1/3
$\{y''\}$	1/3	1/3	0	0
$\{y'''\}$	1/3	1/3	0	0
$\{y', y''\}$	0	2/3	0	1/3
$\{y', y'''\}$	0	2/3	0	1/3
$\{y'', y'''\}$	0	2/3	2/3	2/3
$\{y', y'', y'''\}$	0	1	0	1

.

Hence, both agents are Savage Bayesians, but person 2 does not distinguish outputs 2 and 3. Suppose both agents are risk averse. The initial allocation is an efficient allocation.

Now suppose the agents draw a value of x that eliminates any conjecture comprised entirely of histories involving only the value y'', and leaves all other conjectures intact. Hence, their capacities and Möbius transforms are now

Set	ϕ_1^m	$\tilde{\phi}_1$	ϕ_2^m	$\tilde{\phi}_2$
$\{y'\}$	1/2	1/2	1/3	1/3
$\{y''\}$	0	0	0	0
$\{y'''\}$	1/2	1/2	0	0
$\{y', y''\}$	0	1/2	0	1/3
$\{y', y'''\}$	0	1	0	1/3
$\{y'', y'''\}$	0	1/2	2/3	2/3
$\{y', y'', y'''\}$	0	1	0	1

.

Now there will exist terms on which the agents are willing to trade. Intuitively, agent 1, knowing that output y'' is impossible, is willing to sell consumption in outputs y'' and y''' to agent 2, in return for consumption in output y', on terms that both will find advantageous.

5 The Dynamics of Reasoning

Section 4 joins the refrain common in the literature—whether beliefs are Bayesian or otherwise, information does not give rise to trade among agents with common prior beliefs (barring exceptional circumstances in the nonBayesian case), while readily prompting trade between agents with different prior beliefs. Were the analysis to stop here, we could claim only to have repackaged a familiar result.

We see no reason to expect beliefs to agree, and no reason to be surprised by trade, in novel situations.[25] For example, it is no surprise that trading was brisk when financial markets reopened after the attacks of December 7, 1941 or September 11, 2001. However, in other circumstances we expect agents to acquire ample experience with the economy, and to learn from this experience. Should we expect agents' conditional beliefs to (at least approximately) converge, and hence trade to dissipate, as the agents accumulate experience with the economy? If so, no-trade theorems become all the more challenging, in that conventional responses allow only temporary respite.

[25] Harsanyi [40, Part I, Section 5; Part III], [41, Section 4.4] might well argue that they *must* agree. In contrast, Luce and Raiffa [53], viewing priors as arising out of a process of "jockeying—making snap judgments, checking on their consistency, modifying them, again checking on consistency, etc. ..." ([53, p. 302]) would presumably regard coinciding priors as most unlikely.

Section 5.1 begins with a small worlds special case. We set $T = \infty$ to allow experience to accumulate. The Bayesian agents in the small-worlds model have prior beliefs that are equivalent. Standard merging arguments ensure that, over the course of repeated interactions, on a set of states to which both agents attach unitary probability, the agents' beliefs coincide. If the data generating process is absolutely continuous with respect to the agents' beliefs, then these expectations will be realized, and trade will indeed evaporate. In this setting, differences in prior beliefs provide only a transient foundation for trade.

Section 5.2 shows that similar results do not hold once we expand credence functions beyond Bayesian conjectures. This gives us two possible approaches to a model of persistent trade, one based on nonBayesian credence functions and one on Bayesians facing an "unanticipated" data generating processes. However, we regard either approach in isolation as unconvincing. Section 5.2 ensures that standard merging arguments do not force nonBayesian beliefs together, but might not other arguments or learning rules suffice? In essence might there not be a role for induction, as well as the deductive reasoning of updating? After all, Hume's [45, Section IV] classic skepticism about induction to the contrary, most people have no difficulty understanding the behavior of the sun. We turn to these questions in Section 5.3.

5.1 Small Worlds

Suppose that the state is drawn from $(X \times Y)^\infty$ and that the agents have equivalent Bayesian credence functions.

Before the process begins, the agents can negotiate trades in each future period t conditional on the realized values of $(y_0, \ldots, y_t)$, but not conditional on the draws of x_t. This gives us an efficient allocation z. We then ask whether, given this allocation, the additional information contained in $(x_0, \ldots, x_t)$ can give rise to the existence of Pareto superior allocation.

As we have seen in Proposition 1, a Bayesian credence function is equivalent to a measure φ on $(\Omega, \mathcal{A})$. Let each agent $i = 1, \ldots, I$ be characterized by such a measure φ_i. Let $\mathcal{A}_t$ be the sub-σ-algebra of $\mathcal{A}$ induced by the history h_t.

Proposition 4 *Let $\varphi_1, \ldots, \varphi_I$ be equivalent. Then for any pair of agents i and j,*

$$\varphi_i\left(\left\{\omega : \lim_{t\to\infty} \sup_{A\in\mathcal{A}} |\varphi_j(A|\mathcal{A}_t(\omega)) - \varphi_i(A|\mathcal{A}_t(\omega))| = 0\right\}\right) = 1.$$

This is Theorem 2.2 of Sorin [73], who attributes the result to Blackwell and Dubins [8]. This proposition indicates that each agent i attaches probability one to the event that her posterior beliefs will merge with the posterior beliefs of the other agents. Of course states may arise in which beliefs do not merge, but all agents attach probability zero to such states. If the state is generated by a process that is absolutely continuous with respect to the agents' beliefs, then their beliefs will indeed merge.[26] Merging in turn implies that the difference in the probabilities that any two agents attach to an event converges to zero, uniformly across events and agents. If the agents are risk neutral then even minuscule differences in beliefs can give rise to the desire to engage in unboundedly large trades. But if the agents are strictly risk averse, then agents will eventually be willing to engage in only arbitrarily small trades. Hence, in the limit, information will not give rise to trade.

5.2 Large Worlds

Once we move beyond Bayesian credence functions, unlimited data and merging do not suffice to drive beliefs together. Our first example illustrates.

Example 1 Let $X = \{x', x''\}$ and $Y = \{y', y''\}$. There are two agents, whose credence functions are such that in any period t and after any realizations in periods $0, \dots, t-1$, beliefs about the values of (x_t, y_t) are given by the following capacities (with their Möbius transforms):

Conjecture	ϕ_1^m	$\tilde{\phi}_1$	ϕ_2^m	$\tilde{\phi}_2$
$\{x'y'\}$	2/12	2/12	3/12	3/12
$\{x''y''\}$	4/12	4/12	6/12	6/12
$\{x'y', x'y'', x''y', x''y''\}$	6/12	1	3/12	1

.

In particular, let the agent 1's credence function attach probability 1/2 to Ω, and let 2's credence function attach probability 1/4 to Ω. Let the remaining probability in each case be attached to the set of Bayesian conjectures, arranged to correspond to a process in which *iid* draws from $\{x'y', x''y''\}$ are taken in each period, with the latter twice as likely. Importantly, notice that these credence

[26]If the latter condition fails, the agent's beliefs may neither merge with the data generating process nor with those of other agents. For example, let X be a singleton and $Y = \{y', y''\}$. Let there be two measures, φ' and φ''. The former states that, with probability 1, we'll see only y' from some period on, while the latter states that we will see only y'' from some period on. In particular, let

$$\begin{aligned}
\varphi'(y', y', y', ...) &= 1/4 \\
\varphi'(y'', y', y', ...) &= 1/4 \\
\varphi'(\cdot, y'', y', y', ...) &= 1/4 \\
\varphi'(\cdot, \cdot, y'', y', y' ...) &= 1/8 \\
\varphi'(\cdot, \cdot, \cdot, y'', y', y' ...) &= 1/16 \\
\varphi'(\cdot, \cdot, \cdot, \cdot, y'', y', y' ...) &= 1/32 \\
&\vdots
\end{aligned}$$

In general, $\varphi'(\tilde{\boldsymbol{y}}, y', y', y', y', y', \dots) = 2^{-t}2^{-(t+1)}$ if the last y'' occurs in period t for any $t \geq 1$, for any t-length vector $\tilde{\boldsymbol{y}}$ that precedes it. The measure φ'' is defined symmetrically.

Agent 1 has prior beliefs $0.9\varphi' + 0.1\varphi''$ and agent 2 has beliefs $0.1\varphi' + 0.9\varphi''$. Their beliefs are therefore mutually absolutely continuous, and they both believe that they will converge to the same belief. However, if the true data generating process is sequence of independent random tosses of a fair coin, then with probability one the agents will disagree forever.

functions are equivalent—but for the fact that the agents are not Bayesian, the hypothesis of Proposition 4 are satisfied.

The agents thus agree on a leading candidate for their explanation of the world, namely that values of y' and y'' are independently drawn in each period with y'' twice as likely, but the agents also attach some credence to the event that they have no idea how values are determined. They differ only in how cautious they are in thinking that they understand the process, with agent 2 more confident of her understanding.

Suppose that in the absence of information, we have an allocation in which each agent is endowed with one unit of consumption in each output. This allocation is efficient.

In each period, the observation of x prompts the agents to revise the Bayesian component of the agents' beliefs. If the Bayesian conjectures were the only ones in the agents' credence functions, then the information could not give rise to trade. As it stands, however, this information invariably gives rise to trade. Upon observing x', the agents' beliefs are described by:

Conjecture	ϕ_1^m	$\tilde{\phi}_1$	ϕ_2^m	$\tilde{\phi}_2$
$\{x'y'\}$	1/4	1/4	1/2	1/2
$\{x'y', x'y''\}$	3/4	1	1/2	1

,

and agent 1 will trade consumption to agent 2 in output y' in return for higher consumption in output y'', reflecting the relatively larger weight that agent 2 now puts on output y'. An observation of x'' gives rise to trade in the opposite direction. ■

This example illustrates the general principle that if Bayesian and non-Bayesian conjectures survive in agents' reasoning, and survive in different proportions, then circumstances can arise in which information can give rise to trade. Of course, one possibility is that *only* Bayesian conjectures survive, so that eventually standard merging arguments can be brought to bear. This simply adds another "eventually" to Proposition 4, indicating that if we are *eventually* in a small world, then trade eventually ceases. Example 1 shows that if some elements of a large world persist, then information can persistently give rise to trade.[27]

The agents in Example 1 differ in the initial weights they place on their (identical) Bayesian beliefs, with the learning process doing nothing to eliminate these differences. The following example shows that such differences can arise endogenously.

[27] One might argue that more sensible credence functions would allow the importance given to the set $\{x'y', x'y'', x''y', x''y''\}$ to depend on the history h_t. A history in which y'' has appeared about 2/3 of the time matches the Bayesian explanation quite well, and hence should make the agents relatively confident of their Bayesian explanation, relative to a history in which y'' has appeared only 1/3 of the time. As long as the agents still differ in their propensity to think they understand the process, and hence differ in the weights their capacities attach to the set $\{x'y', x'y'', x''y', x''y''\}$, the basic nature of the result remains.

Example 2 Let $Y = \{y', y''\}$, let there be two agents, and (temporarily) let X be a singleton. Let the agents' beliefs again have a Bayesian component, in which the values of y' and y'' are independently drawn in each period, either with y' twice as likely in each period or y'' twice as likely. Agent 1 puts a higher prior probability on the first possibility, while agent 2 puts higher prior probability on the second possibility. Let the data generating process equiprobably choose one of y' or y'' to be twice as likely as the other, and thereafter take independent draws from the chosen distribution.

Each agent's credence function also attaches some weight to the set Ω, and hence after each history each agent's capacity also attaches some weight to the set $\{y', y''\}$. The agents initially place the same probability on the set Ω, and then follow identical rules that specify the surviving weight attached to Ω as a strictly decreasing function of how likely the observed sequence is under their Bayesian beliefs.

We thus have a counterpart of Example 1, but now with some differences in the Bayesian components of the agents' beliefs and with identical initial weight attached to the nonBayesian component of their beliefs, as well as identical rules for how this weight evolves. A familiar merging argument ensures that the Bayesian component of their beliefs will converge. However, the realized sequence will have a higher likelihood for one agent than the other, and the former agent's capacity will attach lower weight to the set $\{y', y''\}$.

One possibility is that the weight attached to $\{y', y''\}$ disappears for both agents, in which case their beliefs will indeed merge. Suppose, however, that at least one of them retains some doubts about their understanding of the world, but that the weight attached to $\{y', y''\}$ does not converge to 1 for both agents. Then the agents will consistently attach different weights to the nonBayesian components of their credence functions. As in Example 1, this opens the door (now letting X be nonempty to as to convey information) to persistent information-induced trade. ■

Trade *may* dissipate in the presence of different NonBayesian beliefs:

Example 3 Suppose $Y = \{y', y'', y'''\}$ and X is a singleton. Consider a finite collection $\{\tilde{\phi}_n\}$ of capacities, each of which induces a Möbius transform that puts weight only on the sets $\{y', y''\}$, $\{y', y'''\}$ and $\{y'', y'''\}$. Each agent has a credence function corresponding to a full support probability distribution over the set $\{\tilde{\phi}_n\}$, though these distributions differ across agents, with subsequent draws then taken independently from the selected capacity.[28] The

[28] We say that a credence function ϕ treats the draws of y_t as independent draws from a capacity $\tilde{\phi}$ if, after any history h_t, the agent's beliefs about the next draw of y_t are described by $\tilde{\phi}$, i.e., if $\tilde{\phi}_{h_t} = \tilde{\phi}$ for all histories h_t. For an example of such a construction, assume that $Y = \{y', y''\}$ and let $\tilde{\phi}(y') = p'$, $\tilde{\phi}(y'') = p''$, and $\tilde{\phi}(y', y'') = 1 - p' - p''$. For $t \geq 0$, let $(Y_\tau)_{\tau=0}^t$ be a sequence of sets with each y_τ equal to either $\{y'\}$, $\{y''\}$ or $\{y', y''\}$. With any such sequence, define the set A by

$$\omega \in A \Longleftrightarrow [\omega_\tau \in Y_\tau, \tau = 0, \dots, t].$$

Let $\mathcal{E}$ be the σ-algebra generated by such sets. To specify the credence function, it suffices to specify its values on the set of generating sets. let A be such a set, and let $m'(A)$ and $m''(A)$

agents' credence functions are equivalent. Suppose the state ω is such that the limit $n' = \lim_{t\to\infty} \frac{1}{t}\mathbf{1}_{y'}$ exists, as do similar limits n'' and n''' for y'' and y'''. Then each agent's beliefs will converge to a credence function under which outputs are drawn independently according to the capacity from the set $\{\tilde{\phi}_n\}$ that maximizes

$$\tilde{\phi}(\{y', y''\})^{n'+n''}\tilde{\phi}(\{y', y'''\})^{n'+n'''}\tilde{\phi}(\{y'', y'''\})^{n''+n'''}.$$

■

The obstacle to merging is thus not the appearance of nonBayesian conjectures per se. Rather, merging arguments break down when credence functions incorporate different "levels" or "styles" or reasoning, such as mixtures of Bayesian and nonBayesian conjectures, with nothing in the updating process ensuring that the agents agree in their mixtures of these various styles of reasoning.

From a different direction, agents who are Savage Bayesian, and who cease trading on the strength of merging beliefs, can still exhibit behavioral anomalies:

Example 4 Suppose $Y = \{y', y'', y'''\}$ and X is a singleton. Consider a finite collection $\{\tilde{\phi}_n\}$ of capacities whose Möbius transforms put weight only on the sets $\{y'\}$ and $\{y'', y'''\}$. Each agent has a credence function corresponding to a full support probability distribution over the set $\{\tilde{\phi}_n\}$, though these distributions differ across agents, with independent draws from the selected capacity then governing the subsequent realizations. The agents are Savage Bayesians, with equivalent credence functions. Once again, assuming the relevant limits exist, the agents' beliefs will merge.

Notice that the agents in this case can fall prey to the Ellsberg paradox. An agent whose behavior is governed by the limiting beliefs will prefer to bet on $\{y'\}$ rather than $\{y''\}$, but will prefer to bet on $\{y'', y'''\}$ rather than $\{y', y'''\}$. As noted in Section 2.3.2, the key to this appearance of the Ellsberg paradox is that some of these events are not measurable for this agent. In effect, the agent has been unsuccessful in capturing relevant aspects of the state and excluding irrelevant ones when constructing her Savage-Bayesian state space.[29] ■

be the number of times the sets $\{y'\}$ and $\{y''\}$ appear in the sequence $(Y_\tau)_{\tau=0}^t$ associated with the set A. Then we let

$$\tilde{\phi}(A) = p'^{m(A)}p''^{m'(A)}(1-p'-p'')^{n-m'(A)-m''(A)}.$$

Such beliefs satisfy Epstein and Schneider's [21] rectangularity condition.

[29] We could similarly construct a two-urn version of the Ellsberg paradox. Intuitively, we think of one urn that is known to contain half red and half black balls, and one urn containing red and black balls in an unknown proportion. We can than think of four outputs, denoted by RR, RB, BR, and BB. In the absence of any further information, the credence function induces a capacity in each period given by

$$\begin{aligned}\tilde{\phi}(\{RR, RB\}) &= 0.5\\ \tilde{\phi}(\{BR, BB\}) &= 0.5.\end{aligned}$$

An asset that pays 1 only in output RR is worth nothing a priori, as in an asset that pays 1 only in output RB, whereas an asset paying off in both outputs is valuable.

How do we summarize these results? Merging is essentially a Bayesian phenomenon. Familiar arguments will typically force the Bayesian components of agents' reasoning to merge, but we may not have analogous forces working on the nonBayesian component, and there are no forces pushing together the balance between Bayesian and nonBayesian reasoning.

5.3 Revising Credence Functions

Section 5.2 indicates that once we move beyond Bayesian beliefs, we can expect little from merging style arguments. Perhaps there are other learning processes that will bring agents' beliefs together? Indeed, once we move outside the small-worlds setting of Bayesian reasoning, there is no reason to expect agents to learn only via a counterpart of Bayes' rule. And once we recognize that a small-worlds setting with Bayesian reasoning is something constructed by the agents, we might expect small-worlds agents who are sufficiently surprised by the realized state to reconsider their construction. This section accordingly considers a process in which agents begin with given credence functions, but then may switch to new credence functions as their histories identify alternative and better-performing credence functions.

How do we square this analysis with Section 2.5, where we insisted (on the strength of Proposition 7 in Appendix A.1) that we need never consider agents who switch between credence functions? The key observation behind Proposition 7 is that if the agents are inclined to shift toward better-performing credence functions as their experience accumulates, then they can build such shifts into their original credence functions. We could thus conduct all of the following analysis in terms of fixed credence functions that are conditioned on experience, but never adjusted. However, if the agents share similar views of what such "shifting toward better credence functions" means, then they should build some common structure into their initial credence functions. We can thus view the exercise in this section as seeking the common elements that we expect (perhaps eventually) to see in agents' credence functions, with these common elements most effectively isolated by examining a credence-function updating process.

We simplify the notation by temporarily taking the set X to be a singleton, relaxing this assumption in Section 5.3.5.

5.3.1 Likelihoods for Credence Functions

Assessments of credence functions must be driven by some tradeoff between the strength of the predictions they make and the success of those predictions.[30] Our first step is to develop a measure of predictive success. If the credence

[30] One might think or simply ranking credence functions by now well they have predicted previous outcomes. However, for any history h_t, there exist an infinite number of credence functions whose predictions for the first t periods exactly match h_t. Building on this observation, Gilboa and Samuelson [34] identify conditions under learning in such an environment that is guided solely by past success will perform no better than chance.

function ϕ is Bayesian, then history h_t gives the familiar likelihood

$$\prod_{\tau=0}^{t-1}\left(\sum_{y\in Y}\mathbf{1}_y^\tau\tilde{\phi}_{h_\tau}(\{y_\tau\})\right),$$

where $\tilde{\phi}_{h_\tau}$ is the capacity induced by credence function ϕ in period τ given history h_τ, and $\mathbf{1}_y^\tau$ is an indicator function identifying the value of y in period τ. It is convenient to work with the log likelihood and then to render comparisons more convenient by taking averages. Let

$$\pi(\phi,h_t)=\frac{1}{t}\sum_{\tau=0}^{t-1}\log_2\left(\sum_{y\in Y}\mathbf{1}_y^\tau\tilde{\phi}_{h_\tau}(\{y_\tau\})\right).$$

We can interpret higher values as identifying credence functions that do a better job of explaining the data, with the maximum of 0 corresponding to a credence function that attaches probability 1 to the realized history.

We might interpret the likelihood score as reflecting a market in which each credence function uses a fixed budget to buy a portfolio of histories in proportions matching the probabilities the credence function attaches to the histories. Each credence function's payoff is proportional to the quantity it owns of the asset corresponding to the realized history. Alternatively, this likelihood corresponds to the commonly studied logarithmic or self-information loss function in prediction problems (Cesa-Bianchi and Lugosi [13, Chapters 9-10]).

Extending this notion to capacities raises the question of how to evaluate weight attached to a nonsingleton subset of Y. Suppose the Möbius transform ϕ^m (suppressing the history subscript) attaches positive probability to the set $\{y',y''\}$, and the output y' is realized. How much "credit" should the credence function get for attaching probability to $\{y',y''\}$? Presumably, we do not want to proceed as if all such probability was attached to the outcome y', since the set $\{y',y''\}$ is a less precise prediction than is $\{y'\}$. If $Y=\{y',y''\}$, we could argue that attaching weight to the set $\{y',y''\}$ conveys *no* information, and should be ignored. However, Y may be larger, and it may be quite informative to attach weight to the set $\{y',y''\}$, even if less informative than dividing this weight among the singletons. The following example illustrates the issues that arise in assessing predictions of different precision.

Example 5 Let $Y=\{y',y''\}$ and let X be a singleton. Let $\{\tilde{\phi}_1,\ldots,\tilde{\phi}_n\}$ be a set of capacities over Y. Each agent's credence function treats the draws of y_t as independent draws from one of the capacities $\{\tilde{\phi}_1,\ldots,\tilde{\phi}_n\}$.

The agents' credence functions differ only in terms of the prior distributions over which of the capacities $\{\tilde{\phi}_1,\ldots,\tilde{\phi}_n\}$ is the relevant one. Now, as the state unfolds, the familiar process of deleting falsified conjectures will automatically shift weight among the various credence functions underlying the capacities $\{\tilde{\phi}_1,\ldots,\tilde{\phi}_n\}$, without any need for explicitly calculating likelihoods.

Suppose we observe a history in which the outputs y' and y'' have appeared n' and n'' times, respectively. Then the weight placed on capacity $\tilde{\phi}_k$ is proportional to

$$\rho(\tilde{\phi}_k)[\tilde{\phi}_k(\{y'\}) + \tilde{\phi}_k(\{y', y''\})]^{n'}[\tilde{\phi}_k(\{y''\}) + \tilde{\phi}_k(\{y', y''\})]^{n''},$$

where $\rho(\tilde{\phi}_k)$ is the prior probability placed on capacity $\tilde{\phi}_k$. The following calculation is then immediate. If we have any two capacities $\tilde{\phi}_1$ and $\tilde{\phi}_2$ that attach the same relative weights to the sets $\{y'\}$ and $\{y''\}$, then the capacity attaching the larger weight to $\{y', y''\}$ will eventually swamp the former. Uninformative conjectures fare well in an updating process based on eliminating refuted conjectures. More generally, the eliminating-refuted-conjectures updating procedure has a built-in tendency to select capacities that make weaker predictions, in the sense that they place more weight on larger subsets of Y. Indeed, the trivial capacity that attaches all weight to set Y will always overwhelm all others. ■

With these considerations in mind, we let $|Y'|$ denote the cardinality of $Y' \subseteq Y$ and generalize our definition of $\pi_\theta(\phi, h_t)$ to:

$$\pi_\theta(\phi, h_t) = \frac{1}{t}\sum_{\tau=0}^{t-1}\sum_{y\in Y}\left(\mathbf{1}_y^\tau \log_2\left(\sum_{Y'\subset Y : y\in Y'}\theta(y, Y')\phi_{h_\tau}^m(Y')\right)\right), \tag{6}$$

where $\theta(y, Y')$ (with $\theta : Y \times \mathcal{P}(Y) \to (0, 1]$) is defined for elements $y \in Y'$ to equal 1 if $Y' = \{y\}$, and otherwise $\theta(y, Y') < \frac{1}{|Y'|}$. Intuitively, the function θ ensures that when (for example) $\phi^m(\{y', y'', y'''\}) > 0$, the agent can say "I'll wait until I see whether one of these three states realizes, say y', and then I'll take credit for predicting it by adding some of the probability $\phi^m(\{y', y'', y'''\})$ to the probability I've already attached to the singleton $\{y'\}$, but I'll only be able to add $\theta(y', \{y', y'', y'''\})\phi^m(\{y', y'', y'''\})$ credit for it." In terms of our asset market, we now imagine the credence function being able to buy assets corresponding to all subsets of Y rather than just the singleton subsets. However, the assumption that $\theta(y, Y') < \frac{1}{|Y'|}$ when Y' is not a singleton ensures that the credence function faces a risk-return tradeoff—a larger subset is safer in the sense that it is more likely that the realized state will come from this subset, but pays a smaller return. An implication is that if outputs are drawn independently across periods according to a stationary distribution that puts probability $q(y)$ on output y (with $\sum_{y\in Y} q(y) = 1$), then the highest likelihood corresponds to a Bayesian belief that matches the data generating process. In this sense, our likelihood evaluations build in a premium on more precise beliefs.[31] However, a credence function that attaches weights only to singletons, but does so in a misspecified way, can perform worse than a function that attaches weight to some nonsingletons.

[31] In a similar vein, Holland's [44] genetic algorithm places a premium on more precise rules.

5.3.2 What Can We Hope to Accomplish?

How should we think about learning credence functions? One possibility is to formulate a prior distribution over credence functions, to be combined with likelihoods in a calculation reminiscent of Bayes' rule to obtain a posterior distribution over credence functions. However, we regard this approach as rather jarringly at odds with reasoning based on credence functions in the first place. An agent willing to be a Bayesian in her choice of credence functions should presumably simply be Bayesian in the formulation of her beliefs about the outcome chosen from Y in each period, obviating the need to learn about credence functions. We are accordingly unwilling to simply build a Bayesian model on a different space.

Perhaps the most common alternative when thinking about nonBayesian learning is to maximize the worst-case performance. In order to talk about the worst-case performance of a credence function, we need some method of aggregating sequences of likelihood scores. We work with a limiting criterion, asking that credence functions eventually become informative.

Given a credence function ϕ, let

$$\underline{\pi}_t(\phi) = \inf_{\omega \in \Omega} \limsup_{t \to \infty} \pi_\theta(\phi, h_t(\omega)).$$

We say that a credence function ϕ is worst-case optimal if it maximizes $\underline{\pi}_t$ over the set of credence functions.

Let ϕ^* be a credence function that, after every history h_t, generates a capacity that is Bayesian, and attaches probability $1/|Y|$ to every element of Y. We refer to this as the *uninformative* credence function, since it never ventures any opinion as to the likely output.

Lemma 3 *The uninformative credence function is worst-case optimal.*

Proof. It is immediate that the uninformative credence function ϕ^* satisfies

$$\pi_\theta(\phi^*, h_t(\omega)) \geq -\frac{1}{|Y|}, \quad \forall t, \omega.$$

Now consider any other credence function $\hat{\phi}$. We generate a state recursively. Given a sequence $((x_0, y_0), \ldots, (x_{t-1}, y_{t-1}))$, let (x_t, y_t) solve

$$\min_{(x_t, y_t) \in X \times Y} \pi_\theta(\hat{\phi}, ((x_0, y_0), \ldots, (x_t, y_t))).$$

Then it is immediate that given credence function $\hat{\phi}$, this state generates a likelihood that never exceeds $-1/|Y|$, and falls strictly short of $-1/|Y|$ should the credence function $\hat{\phi}$ ever generate a capacity different than that of ϕ^*. ■

The proof shows that the uninformative credence function satisfies a rather strong notion of worst-case optimality, since for any alternative credence function $\hat{\phi}$, we can find a single state with the property that $\hat{\phi}$ fares worse than the uninformative credence function after every history at which their induced beliefs differ.

The implication we draw from this result is that if one's objective is to learn in all circumstances, then one can never learn at all. One must choose some circumstances under which to learn, and some circumstances under which to surrender all hopes of learning.[32] It is no surprise that we cannot hope for a universal learning result. Much of the theory of statistical learning is built on the assumption that the data are generated by an exchangeable process (e.g., Hastie, Tibshirani and Friedman [42]). Alternatively, the analysis of a prediction problem typically begins by identifying a collection of experts, with the learning process then designed to ensure the agent does well relative to the performance of the given experts (cf. Cesa-Bianchi and Lugosi [13]).

We might reasonably hope for effective learning in a set of particularly salient environments. In constructing the latter, we focus on a basic tradeoff. People believe they understand some aspects of their environment well enough to offer deterministic explanations. Other aspects are sufficiently inscrutable as to prompt refuge in random explanations. Hence, we seek a learning process that works well for both deterministic and random processes. Obviously, this task would be straightforward if an oracle could first identify the data generating process as being either deterministic or random. Much of the complication arises out of the fact that the agent must assess the likelihood of these possibilities from the unfolding of the state itself.

5.3.3 Preliminaries

Our first step is to identify deterministic and random processes. Given our temporary assumption that X is a singleton (and hence can be ignored), a state ω is a sequence $(y_t)_{t=0}^{\infty}$. For each such state and each $y \in Y$, let $n(y,t)$ be the number of times the observation y has occurred in periods $0,\dots,t-1$. Then:

Definition 1

[1.1] A state ω is computable if there exists a Turing machine that, for any $t \in \mathbb{N}$, if its input tape includes (only) the entry t (in binary encoding), then the machine halts and produces y_t as its output.[33]

[1.2] A state is eventually constant if there exists a time t and an output $\overline{y}$ such that for all $\tau \geq t$, $y_\tau = \overline{y}$.

[1.3] A state has a random limit if, for all $y \in Y$, there is a number $q(y) \in [0,1]$ such that

$$\frac{n(y,t)}{t} \xrightarrow[t\to\infty]{} q(y).$$

[32] This result is reminiscent of a series of "no free lunch" theorems due to Wolpert and Macready [77]. The underlying idea is straightforward. An informative credence function must sometimes commit itself to a prediction. But we can then construct a state in which this prediction fails, yielding a result worse than no prediction at all.

[33] This definition of computability, from Minsky [59, p. 159], reflects Turing's original vision, but is now somewhat archaic, though useful for our purposes.

A state is random if (i) it has a random limit and (ii) every recursively defined subsequence has the same random limit.

We introduce the concept of an eventually-constant state to highlight one of the complications in the analysis. An eventually-constant state is computable and is also (perhaps somewhat counterintuitively) random. More generally, some random states are arbitrarily close to computable states. Classifying states as random or computable can then be difficult.

The following, with proof in Appendix A.6, collects some basic facts:

Lemma 4

[4.1] There are countably many computable states.

[4.2] There are countably many eventually-constant states.

[4.3] A random state is computable if and only if it is eventually constant.

[4.4] There are uncountably many random states, and hence uncountably many noncomputable random states.

[4.5] There are uncountably many states that neither have a random limit nor are computable.

[4.6] If the values $(y_t)_{t=0}^{\infty}$ are drawn from an exchangeable measure, then with probability one the measure will select a random state.

This formulation of a random state is taken from von Mises [75]). Every random state has a random limit, but the converse fails. For example, the sequence (01010101 . . .) has a random limit but is not random. von Mises adopted the name "random" in anticipation that that such states would satisfy any conceivable property of randomness. As Martin-Löf [56] explains, Ville [74] showed that this was not the case, leading Martin-Löf [57] to formulate a definition of randomness based on Kolmogorov complexity that has the requisite properties. We first make use of the limits introduced in Definition 1.3, prompting us to begin with von Mises' formulation. As Remark 3 explains, we would like to work with a set of states on which the convergence is uniform, which will be the case for a set of states of probability arbitrarily close to one if states are drawn from an exchangeable measure. We then invoke Doob's [19] result that Lemma 4.6 can be strengthened, so that with probability one, the state drawn by such a measure satisfies the remaining required characteristics of randomness.

5.3.4 A Learning Result

We consider two types of explanations agents may bring to bear in understanding the world. One is that they may think in terms of deterministic theories that predict precisely what may happen. The second is that they may be less ambitious and seek to explain the world only in random terms. Different agents may have differing proclivities to appeal to one type of explanation over another—some may be quick to see patterns in the data that suggest deterministic explanations, others may see only confusion and randomness.

We let Φ be a countable set of credence functions that are entertained by the agents as possible explanations for the data generating process. The most obvious motivation for restricting attention to a countable set of credence functions is that the agents entertain only computable credence functions, and we accordingly assume Φ contains at least the computable credence functions. We will assume that agents have a ranking of credence functions that reflects various preferences over explanations, such as their proclivity to detect deterministic patterns in the data. In each period, the agents combine information from this ranking and information from the observed explanatory power of the various credence functions to choose one such function. Let $\Gamma : H \to \Phi$ be a *learning rule* that, for each history $h_t \in H = \cup_{t=0}^{\infty} H_t$ selects a credence function to be applied in period t. We imagine each agent being characterized by such a rule.

Two rules Γ and Γ' will each generate collections of credence functions, one for each history. We say that the rules Γ and Γ' (or, equivalently, the induced sequence of beliefs), ζ-converge ($\zeta > 0$) at state ω if they satisfy the following condition: for every $Y' \subset Y$ and $z > 0$, the limsup of the proportion of periods in which the induced capacities $\tilde{\phi}$ and $\tilde{\phi}'$ attach beliefs to Y' that differ by more than z, denoted by Z, satisfies $z < 1 - 2^{-\zeta/Z}$. More precisely, for $z > 0$, define (letting the $0/0$ term corresponding to $t = 0$ be defined arbitrarily)

$$Z(z) = \max_{Y' \subset Y} \left[\text{limsup}_{t \to \infty} \frac{1}{t} \# \left\{ \tau \;\middle|\; \left| \tilde{\phi}_{h_\tau}(Y') - \tilde{\phi}'_{h_\tau}(Y') \right| > z \right\} \right].$$

Hence, $Z(z)$ is the highest (limiting) relative frequency of the periods for which $\tilde{\phi}_{h_t}$ and $\tilde{\phi}'_{h_t}$ differ by more than z (on any given subset of outcomes Y'). Then the beliefs induced by Γ and Γ' ζ-converge at state ω if they satisfy the following condition for all $z > 0$:

$$z < 1 - 2^{-\zeta/Z(z)}.$$

Beliefs ζ-converge for all ζ exceeding

$$\zeta^* = \inf_{z>0} \left[Z(z) \log_2 \frac{1}{1-z} \right].$$

The interpretation here is that ζ- convergence allows the beliefs attached to a set Y' to differ relatively more, if they do so in a smaller proportion of periods. In one limit, these beliefs can differ by as much as 1 ($z = 1$) if they do so in 0 proportion of periods ($Z = 0$), and the allowed difference in beliefs decreases as the proportion of periods increases. Decreasing ζ moves the function downward, imposing a more stringent convergence condition. As long as there is some curvature in utility functions, trade will wither away in the face of the ζ-convergence of beliefs.

Remark 3 explains the asymmetry in the statement of the following proposition, with measure-theoretic considerations entering only in the second part.

Proposition 5 *For any $\varepsilon > 0$ and $\zeta > 0$, there exists a class of learning rules such that, if each agent is characterized by a learning rule in this class, then:*

[5.1] If the realized state ω is computable, then beliefs will ζ-converge with the Dirac measure δ_ω, as well as with each other.

[5.2] For any exchangeable measure λ on Ω, there is a set of states of λ-probability at least $1 - \varepsilon$ on which beliefs will ζ-converge with λ, as well as with each other.

Notice that we do not assert that the agents have prior beliefs corresponding to λ, or have any beliefs. We use the device of an exchangeable measure to identify a subset of the set of random states.[34] The result indicates that each agent's belief must grow close to that of the data generating process, and hence the agents' beliefs similarly grow close to one another.

Remark 2 (Intuition for the Proof) The point of departure for the argument is the idea that an agent chooses a credence function with a relatively high likelihood. However, this alone will not be useful, as for every finite history, there will be an infinite number of credence functions that predict that history exactly, and so have the maximal likelihood score. We respond by asking the agent to first form an assessment of whether she should think in terms of deterministic or random theories. She does so by calculating the conditional Kolmogorov complexity of the history she has observed. If this complexity is small, then the agent thinks in terms of deterministic theories, made operational by being quite demanding in terms of the likelihood score she requires in order to adopt a credence function. If the complexity is large, she entertains the possibility that the state may be random, and is less demanding in terms of likelihood, sufficiently so that an *iid* process qualifies as an acceptable explanation.

Once the agent has set a likelihood target and identified the credence functions that have met this target sufficiently often, there will typically be many remaining explanations.[35] How does she choose among them? One possibility would be to choose randomly, but Gilboa and Samuelson [34] explain why this can lead to ineffective prediction. Instead, we follow Gilboa and Samuelson [34] in assuming that agent has a ranking of theories, and chooses according to this ranking from among those that exhibit an acceptable likelihood. ■

Remark 3 (Speed versus Precision) Proposition 5.2, addressing the case in which λ is exchangeable, shows that for any ε and $\zeta > 0$, we can find a learning process such that with probability at least $1 - \varepsilon$, beliefs are eventually within ζ. We can find learning processes that make ε and ζ arbitrarily small, but at the cost that learning will take longer. We could devise learning rules

[34] Grant, Meneghel and Tourky [37] make use of the condition of conditional exchangeability (that, for every t, the sequences $(\omega_0, \ldots, \omega_t, \omega_{t+1})$ and $(\omega_0, \ldots, \omega_t, \omega_{t+2})$ have identical distributions). Conditional exchangeability retains the salient features of exchangeability. For example, Kallenberg [49, Proposition 2.1] shows that a process that is both stationary and conditionally exchangeable is also exchangeable. Fortini, Petrone and Sporysheva [25, Section 2] show that any sequence of conditionally exchangeable random variables converges to a sequence of exchangeable random variables.

[35] In deciding whether a credence function satisfies the resulting likelihood threshold, the agent requires that the credence function do so not only in the current period, but has also done so for a number of periods. Gilboa and Samuelson [34] provide an example of why such a feature is needed.

that push ε and ζ to zero, at the cost of both considerable complication and of sacrificing all sense of how long learning might take.

An interesting open question concerns the characterization of learning schemes that optimally balance this trade-off, while also taking into account the necessity of distinguishing computable from random states. A first difficulty is that the objective function for such an analysis is not obvious. A common approach would be to calculate expected payoffs with respect to some measure over states, but our large-worlds viewpoint forbids this, while we have seen that worst-case criteria are also of no use.

Relatedly, one might wonder why we express Proposition 5.2 in terms of an exchangeable measure, instead of simply addressing random states. Random states exhibit converging frequencies of observations, allowing us to establish properties much like those of Lemmas 5 and 7 below. However, Lemma 5 establishes a uniform convergence result (on a set of large measure) that we would lose if we required simply that the state in question would be random. We could instead define a subset of the set of random states whose limits converge uniformly. The measurability formulation in Proposition 5.2 is one convenient way of doing so. ■

Proof of Proposition 5. An agent is characterized by a relation $\succsim$ on Φ. This relation may capture all sorts of considerations, including notions of plausibility and simplicity. If $\phi' \succ \phi''$, we will view the agent as "preferring" credence function ϕ' to ϕ'' (with $\succ$ being the asymmetric part of $\succsim$, with the latter interpreted as "at least as good"), in the sense that the agent will appeal to credence function ϕ' rather than ϕ'' unless the history provides some reason to reverse this preference. We assume that $\succsim$ is complete, transitive, and, importantly, has finite upper contour sets.[36]

Let $\lambda^q \in \Delta Y$ be a probability measure over Y (recalling our maintained assumption that X is a singleton). For each such measure, let ϕ^q denote the credence function (i.e., the measure on $\mathcal{A}$) that corresponds to a sequence of independent draws governed by the measure λ^q. We let $q(y)$ be the probability attached to $y \in Y$ by λ^q and let q denote the vector of such probabilities.

The credence function ϕ^q and the measure λ^q are formally different objects, but contain the same information, and so we shall use them interchangeably when it causes no confusion. In particular, we will feel free to replace ϕ^q with λ^q as the first argument of the likelihood π_θ. We will also abuse notation by using the symbol π again in defining

$$\pi_\theta(\lambda^q, q) = \sum_{y \in Y} q(y) \log_2 q(y).$$

De Finetti's [16] theorem ensures that, for an exchangeable measure λ, we can view the process determining ω as one of first choosing $q \in \Delta Y$ and then

[36] One implication is that the views of the agents captured by their orders $\succsim$ cannot be too different. In particular, for any given credence function, the difference in the rankings of this credence function in the orders $\succsim$ corresponding to two agents must be finite.

choosing ω from ϕ^q. In the exchangeable case, we will accordingly sometimes write the state as (q, ω).

A state ω induces a sequence of partial histories $(h_t)_{t=0}^{\infty}$. For each such partial history, we let $q(h_t)$ denote the empirical distribution of the realizations in this partial history and let $K(h_t|t)$ denote its conditional (upon knowing the length of the sequence) Kolmogorov complexity.

We collect some preliminary results. Appendix A.7 notes that straightforward law-of-large-numbers arguments prove:

Lemma 5 *Let λ be exchangeable and let $\varepsilon > 0$ be given. Then:*

[5.1] There exists a function $\delta(t) : \mathbb{N} \to \mathbb{R}_+$, with $\lim_{t\to\infty} \delta(t) = 0$, such that,

$$\lambda\{(q, \omega) : \pi_\theta(\lambda^{q(h_t(\omega))}, h_t(\omega)) - \pi_\theta(\lambda^q, h_t(\omega)) \leq \delta(t) \quad \forall t \in \mathbb{N}\} \geq 1 - \varepsilon/2.$$

[5.2] There exists T such that

$$\lambda\{(q, \omega) : |\pi_\theta(\lambda^{q(h_t(\omega))}, h_t(\omega)) - \pi_\theta(\lambda^q, q)| < \varepsilon/4 \quad \forall t \geq T\} \geq 1 - \varepsilon/4.$$

To interpret the first result, suppose that the state is generated by a sequence of independent draws from the measure λ^q, and that the agent has observed the history h_t. The likelihood of this history under the measure λ^q is given by $\pi_\theta(\lambda^q, h_t)$. This history would have an even higher likelihood under the measure $\lambda^{q(h_t)}$, i.e., under the measure that draws outputs in precisely the proportions observed in the history h_t. However, for long histories, it is very likely that the proportions $q(h_t)$ are very close to q, and hence the likelihood $\pi_\theta(\lambda^{q(h_t)}, h_t)$ cannot be too much larger than $\pi_\theta(\lambda^q, h_t)$. Lemma 5.1 makes this precise—for any $\varepsilon > 0$ and any t, with high probability (i.e., probability $1-\varepsilon/2$) $\pi_\theta(\lambda^{q(h_t)}, h_t)$ cannot exceed $\pi_\theta(\lambda^q, h_t)$ by more than $\delta(t)$.

Appendix A.8 establishes a key characteristic computable states:

Lemma 6 *Let ω be computable. Then*

$$\lim_{n\to\infty} \frac{K(h_t|t)}{t} = 0.$$

We can establish a link between the likelihood (or relative entropy) and the Kolmogorov complexity of random states. The following is a straightforward adaption of the corresponding result in Cover and Thomas [15, Theorem 14.5.3]) for the case of Bernoulli random variables:

Lemma 7 *Let ω be drawn from the distribution λ^q for some q. Then*

$$\frac{K(h_t|t)}{t} \to -\pi_\theta(\lambda^q, q),$$

where the convergence is in probability.

We now specify an agent's learning rule. Fix $\alpha \in (0, 1)$ for this agent.

1. In each period t, the agent calculates $\pi_\theta(\lambda^{q(h_t)}, q(h_t))$ and $K(h_t|t)$. Then the agent constructs a function $f(\tau, h_t)$ for $\tau \leq t$ by setting

$$f(\tau, h_t) = \begin{cases} 0 & \text{if } \frac{K(h_t|t)}{t} \leq -\alpha\pi_\theta(\lambda^{q(h_t)}, q(h_t)) \\ \pi_\theta(\lambda^{q(h_\tau)}, q(h_\tau)) - \delta(\tau) - \zeta/2\Theta I & \text{otherwise,} \end{cases}$$

   where I is the number of agents, $\alpha \in (0,1)$, and $\Theta \geq 1$ will be set later.

   The interpretation is that $f(\tau, h_t)$ gives the "likelihood demands", made by the agent in (retroactively) assessing a candidate credence function's performance in every period $\tau \leq t$, after observing history h_t. Intuitively, if the history has low conditional Kolmogorov complexity, then the agent entertains the idea that there is a plausible deterministic theory that explains the data, and is quite demanding in terms of the likelihood the agent requires in order to consider a theory acceptable. In particular, if the complexity of the history is sufficiently low, the agent entertains only credence functions whose predictions match the history exactly. If the history has high complexity, then the agent concludes the data are likely to reflect random elements, and sets a lower standard for the likelihood required to consider a credence function acceptable. If the history is sufficiently complicated, the agent sets a likelihood standard that admits an exchangeable process, effectively concluding that the data are unpredictable.

   We have made f a step function, but could easily replace this with a smoother function.

2. The agent can then be thought of as forming the set $\tilde{\Phi}(h_t)$ of credence functions ϕ to be considered. First, for each $\tau \leq t$, the agent forms the set of credence functions $\Phi^\dagger_\tau(h_t)$ consisting of all those ϕ that satisfy,

$$\pi_\theta(\phi, q(h_\tau)) \geq f(\tau, h_t).$$

   Then the agent takes the intersection of these credence functions over τ, giving

$$\tilde{\Phi}(h_t) = \cap_{\tau=1}^{t} \Phi^\dagger_\tau(h_t).$$

   The agent asks for credence functions that match the data well not only in the current period, but in every past period. This naturally raises a concern about whether the intersection might be empty. However, for any history h_t, there is always a credence function that attaches all of its weight to a single output that matches the h_t precisely. This credence function gives a log-likelihood of zero, which equals or exceeds $f(\tau, h_t)$ for every τ, which in turn ensures that this credence function appears in every $\tilde{\Phi}^\dagger_\tau(h_t)$.

3. The agent then chooses from the resulting set $\tilde{\Phi}(h_t)$ according to the relation $\succsim$.

This gives us our candidate for a learning process. We now argue as follows.

[Proposition 5.1] First, let the state ω be computable. Then as t increases, the relative conditional Kolmogorov complexity $K(h_t|t)/t$ of the history converges to 0, as does $(1/\alpha)K(h_t|t)/t$. We now consider two sets of periods.

- The first set includes all those periods in which $-\alpha\pi_\theta(\lambda^{q(h_t)}, q(h_t)) \geq K(h_t|t)/t$. Intuitively, this is the expected case in which the agent recognizes that the state has low conditional Kolmogorov complexity and hence is unlikely to be random, and so the agent turns to deterministic explanations.

  In all periods along this sequence, the relevant value of $f(\tau, h_t)$, for all $\tau \leq t$, is given by 0. Hence, in each period t, $\tilde{\Phi}(h_t)$ consists of all those credence functions that have generated a likelihood value of 0 in each period, i.e., all those credence functions that have predicted a single element of Y in each previous period, with that prediction matching the realized element in that period. One such credence function ϕ^ω, identifies the realized state and is contained in $\Phi^\dagger_\tau(h_t)$ for every $\tau \leq t$, and hence is contained in $\tilde{\Phi}(h_t)$. No credence function ranked below ϕ^ω under $\succsim$ is relevant, and we may as well eliminate such credence functions from $\tilde{\Phi}(h_t)$. There are finitely many credence functions that are preferred to ϕ^ω under $\succsim$. Hence, the relevant subset of the set $\tilde{\Phi}(h_t)$ is decreasing in t, and the agent must eventually settle on a single credence function and use that credence function in every subsequent period, with the function predicting behavior identical to ω.

- The second set includes all those periods in which $-\alpha\pi_\theta(\lambda^{q(h_t)}, q(h_t)) < K(h_t|t)/t$. Intuitively, this is an unusual case in which the conditional Kolmogorov complexity of the sequence is shrinking, but so is the entropy of the sequence, allowing the agent to continually entertain the hypothesis that the data are generated by an *iid* process that is nearly degenerate.

  In all periods along this sequence, the relevant likelihood target in each period $\tau \leq t$ is $\pi_\theta(\lambda^{q(h_t)}, q(h_t)) - \delta(\tau) - \zeta/2\Theta I$. Once again, one credence function, ϕ^ω, identifies the realized state and is contained in $\Phi^\dagger_\tau(h_t)$ for every $\tau \leq t$, and hence is contained in $\tilde{\Phi}(h_t)$. Again, no credence function ranked below ϕ^ω under $\succsim$ is relevant, and we may as well eliminate such credence functions from $\tilde{\Phi}(h_t)$. There are finitely many credence functions that are preferred to ϕ^ω under $\succsim$. The agent must again eventually settle on a single credence function ϕ and use that credence function in every subsequent period. In this case, however, all we know is that the credence function ϕ eventually employed satisfies (since $K(h_t|t)/t$ and hence $-\alpha\pi_\theta(\lambda^{q(h_t)}, q(h_t))$ converge to 0, and so does $\delta(\tau)$)

$$\pi_\theta(\phi, q(h_t)) > -\zeta/\Theta I. \tag{7}$$

- We then have the result that the each agent (at least eventually, as $-\alpha\pi_\theta(\lambda^{q(h_t)}, q(h_t)) - \delta(\tau)$ converges to 0) settles on a single credence function, possibly satisfying $\pi_\theta(\phi, q(h_t)) = 0$ and at least satisfying $\pi_\theta(\phi, q(h_t)) > -\zeta/\Theta I$. We need to argue that beliefs ζ-converge. Suppose there are two agents (so $I = 2$, with the extension to more agents being immediate). To maximize the divergence in beliefs while satisfying this constraint, in each period τ in which beliefs differ, we should set one credence function to attach probability 1 to the realized y_τ, while the other attaches some probability $1 - z_\tau$ to y_τ. Let us first assume we always set agent 1's probability equal to 1 in such a period. In order for agent 2 to satisfy (7), we must have

$$\frac{1}{t}\sum_{\tau=0}^{t-1} \log_2(1 - z_\tau) > -\zeta/2\Theta,$$

or, equivalently,

$$\prod_{\tau=0}^{t-1}(1 - z_\tau) > 2^{-t\zeta/2\Theta}.$$

This implies the desired result. In particular, if beliefs agree in all but Z proportion of periods, in which case they differ by z, this is

$$(1 - z)^{tZ} > 2^{-t\zeta/2\Theta},$$

or

$$z < 1 - 2^{-\zeta/2\Theta Z}.$$

This gives us the proportion of times beliefs can differ by more than z because agent 1 attaches probability 1 to an event and agent 2 attaches a smaller probability. The agents can of course reverse roles, and so the total number of times beliefs can differ must be doubled, giving

$$z < 1 - e^{-\zeta\Theta Z},$$

which, because $\Theta \geq 1$, is as desired.

[Proposition 5.2] Now suppose that the state is generated by an exchangeable measure. Then (by de Finetti's [16] result) there exists some parameter $q \in \Delta Y$ such that the realized state is generated by a sequence of *iid* draws from the measure λ^q.

Suppose first that the distribution from which q is drawn places positive probability on a degenerate distribution. Then we can condition on the event that such a degenerate distribution is drawn, and apply to the resulting degenerate distribution the argument from the proof of Proposition 5.1.

We may accordingly assume that q is nontrivial. Then Lemmas 5.2 and 7 ensure that there exists a value T such that, for all $t \geq T$, with probability at least $1 - \varepsilon/2$ we have $\frac{K(h_t|t)}{t} > -\alpha\pi_\theta(\lambda^{q(h_t)}, q(h_t))$. We hereafter restrict attention to such t and to the probability $1 - \varepsilon/2$ event that $\frac{K(h_t|t)}{t} > -\alpha\pi_\theta(\lambda^{q(h_t)}, q(h_t))$. We then note that for any t and every $\tau \leq t$, the relevant value of $f(\tau, h_t)$ is

$$\pi_\theta(\lambda^{q(h_\tau)}, q(h_\tau)) - \delta(\tau) - \zeta/2\Theta I.$$

Lemma 5 ensures that, with probability at least $1 - \varepsilon/2$, we have, for all $\tau \leq t$,

$$\pi_\theta(\lambda^q, q(h_\tau)) \geq \pi_\theta(\lambda^{q(h_\tau)}, q(h_\tau)) - \delta(\tau).$$

One would like to then conclude that the credence function ϕ^q is contained in $\tilde{\Phi}(h_t)$ for all t. However, the credence function ϕ^q need not be computable and hence may not be contained in Φ. Fortunately, exploiting the presence of $\delta(\tau)$, there exists a computable $\hat{q}$ (with $\hat{q}$ close to q) with

$$\pi_\theta(\lambda^q, q(h_\tau)) \geq \pi_\theta(\lambda^{q(h_\tau)}, q(h_\tau)) - \delta(\tau) - \zeta/2\Theta I,$$

and hence we have that $\phi^{\hat{q}}$ is contained in $\tilde{\Phi}(h_t)$ for all t.

Moreover, the set of credence functions preferred to $\phi^{\hat{q}}$ under $\succsim$ is finite.

The likelihood target $\pi_\theta(\lambda^{q(h_t)}, q(h_t)) - \delta(\tau) - \zeta/2\Theta I$ varies as does t, and hence we cannot conclude that the sequence $\tilde{\Phi}(h_t)$ is a descending sequence. However, because $q(h_t)$ converges with probability one to q and $\delta(\tau)$ approaches 0, it must be the case that for every credence function ϕ preferred under $\succsim$ to $\phi^{\hat{q}}$, it is either the case that ϕ is eventually in $\tilde{\Phi}(h_t)$ for all sufficiently large t, or is eventually excluded from $\tilde{\Phi}(h_t)$ for all large t. This ensures that agent i will eventually hit upon a credence function ϕ_i and thereafter use this credence function in every subsequent period.

Now we argue that beliefs ζ-converge. We again present the argument for two agents, with the generalization being straightforward. The previous arguments ensure that agents i and j each individually converge to particular credence functions and use them thereafter. If agents i and j hit upon the same credence function, whether ϕ^q or not, then obviously beliefs ζ-converge. However, their learning rules may differ—they may have different orders $\succsim$, they may have different parameters α, or different functions $\delta(t)$ (and hence different functions f).

Suppose they hit upon different credence functions, and assume en route to a contradiction that beliefs do not ζ-converge. Hence, there exists at least one value $y' \in Y$, and an infinite sequence of periods $(t_1, t_2, \ldots,)$ (with the property that the relative frequency of such periods does not converge to zero, a qualification hereafter taken for granted) at which the probabilities the agents attach to output y differ by at least ζ. Then there must also exist at least one agent, say i, and an infinite sequence, also denoted by $(t_1, t_2, \ldots,)$, at which the probabilities the agent and q attach to y' differ by at least $\zeta/2$.

The credence function ϕ_i has a likelihood $\pi_\theta(\phi_i, h_t)$ that converges to $\pi_\theta(\lambda^q, q) - \zeta/2\Theta$ (since $I = 2$). For sufficiently large Θ, this is consistent with the existence of an infinite sequence of periods $(t_1, t_2, \ldots,)$ at which the probabilities the agent and q attach to y' differ by at least $\zeta/2$ only if there also exists for i an output $y \in Y$ and an infinite sequence of periods (also denoted by $(t_1, t_2, \ldots,)$ such that the probabilities attached to output y by ϕ_i *exceeds* that attached to y by ϕ^q by at least some $\hat{\theta} > 0$, and in which the frequency with which the output y is realized is greater than the probability attached to y by the measure q. But then this sequence of periods could be used to construct the type of gambling system that Doob [19] shows is impossible for a random state, giving us a contradiction and ensuring that beliefs must ζ-converge. ■

Remark 4 (Persistent Trade) The state may be neither random nor computable. In this case, we can expect agents to switch back and forth between random theories and deterministic ones. There is no reason to expect different agents to make the various switches at the same time, and hence the agents will often have different beliefs, and will often have cause to trade.

More precisely, while the relative entropy of a sequence generated by *iid* draws converges in probability to the entropy of the underlying distribution (Cover and Thomas [15, Theorem 14.5.3]), this is not true in general. Consider a sequence whose conditional Kolmogorov complexity does not converge to the underlying entropy. Then, there must exist a value ε such that for arbitrarily large values of n, we have

$$\frac{K(h_t|t)}{t} \leq -\pi_\theta(\lambda^{q(h_t)}, q(h_t))$$

(with Theorem 14.2.5 of Cover and Thomas [15] precluding the possibility that the conditional Kolmogorov complexity is consistently *above* $-\pi_\theta(\lambda^{q(h_t)}, q(h_t)))$. We then need only consider two agents whose values of α, introduced in step 1 of the description of the learning algorithm, are such that $\frac{K(h_t|t)}{t}$ frequently lies between the two agents' values of $\alpha\pi_\theta(\lambda^{q(h_t)}, q(h_t))$.[37] In each such period, one agent will induce a random prediction and one a deterministic prediction of the state, ensuring that there is an advantageous trade. ■

There thus exists a rich collection of states exhibiting persistent trade. How does this explanation differ from the small-worlds case of equivalent Bayesian priors, where a judiciously drawn state can again confound learning? In the small-worlds setting, the agents must be surprised by such an outcome, since they attach zero probability to the collection of all such states. In the current setting, there is no comparable sense in which such states are surprising—agents may well have experience with beliefs that do not converge, all the while knowing that there is no learning process that can prevent all such occurrences. Indeed, given the trade-off between the coverage of a learning rule and the speed or learning, agents will *prefer* rules that do not always lead to the convergence of beliefs.

Remark 5 (Styles of Explanation) We can interpret agents whose learning rules feature different values of α as having different dispositions for invoking deterministic or random explanations, or different "styles of explanation". An agent whose value of α is relatively large is one who is relatively prone to adduce deterministic explanations. An agent with a smaller value of α is skeptical of

[37]For example, we can construct such states by defining a generating procedure that alternatives between periods in which independent, equiprobable draws are take from $\{y', y''\}$ along with periods in which y' and y" alternate deterministically.

deterministic explanations, and more readily concludes that the data are beyond comprehension, prompting refuge in random explanations.

Proposition 5 shows that if one of these positions is unambiguously correct—the data are either generated by a deterministic, computable process, or represent independent draws from identical distributions—then both agents will eventually come to such a view and trade will (asymptotically) vanish. Depending on the nature of the data generating process, one agent may come readily to the corresponding view of the world while the other is dragged, kicking and screaming, to a similar view. In the course of these adjustments, we may often see disagreements that lead to trade, but trade must eventually dissipate. If neither position is unambiguously correct, then different styles of reasoning can fuel persistent trade. ■

5.3.5 Generalization

This section generalizes Proposition 5 in two directions—we let X be nontrivial, and we allow the agents to entertain a wider class of explanations. Let $\chi : \cup_{t=0}^{\infty} (X \times Y)^t \to \mathbb{K}$ be a computable function that assigns each history to an element of a countable set $\mathbb{K}$. We refer to an element k of $\mathbb{K}$ as a category, and χ as a categorization.

Categories are a device allowing the agents to consider more sophisticated explanations of the data. We can think of Proposition 5 as the special case in which there is only one category, so that the agents can entertain (only) the possibilities that the state is deterministic or represents a sequence of *iid* draws. By letting $\mathbb{K}$ contain one category for each element of Y, the agent can capture the possibility that the output is governed by a first-order Markov process that is independent of the signals X. Taking $\mathbb{K}$ to contain one category for each element of $X \times Y$ allows the possibility that the state is a first-order Markov process whose transitions are conditioned on the most recent realization of the signal x as well as output y. Other categories allow the possibility of higher-order Markov processes, outputs that depend on signals, time-dependent processes, and various other possibilities.

Notice that each history is assigned to only one category. This initially appears limiting, since one would presumably want to assign various period-2 histories (for example) to a category representing an *iid* process, as well as to a category representing a first-order Markov process, and so on. However, the former is a special case of the latter, and hence it suffices to include only the latter possibility. In general, the inclusion of a data-generating process ensures that any of its simplifications are also included.

For a category $k \in \mathbb{K}$, we say that the state is k-computable if the mapping from histories in $\chi^{-1}(k)$ to outputs induced by the state is computable. We say that a sequence of histories $h_{t_1}, h_{t_2}, \ldots$ is consistent if each history in the sequence is a prefix of its successor. We say a measure λ is k-exchangeable if, for any collection sequence of consistent $h_{t_1}, h_{t_2}, \ldots$ of histories drawn from $\chi^{-1}(k)$, the distributions of outputs conditional on these histories is exchangeable.

Proposition 6 *Fix a categorization χ. For any $\varepsilon > 0$ and $\zeta > 0$, there exists a class of learning rules such that, if each agent is characterized by a learning rule in this class, then:*

[6.1] If the state is k-computable for some $k \in \mathbb{K}$, then beliefs along any consistent sequence $h_{t_1}, h_{t_2}, \ldots$ of histories drawn from $\chi^{-1}(k)$ will ζ-converge.

[6.2] For any $k \in \mathbb{K}$, k-exchangeable measure λ and any sequence $h_{t_1}, h_{t_2}, \ldots$ of histories drawn from $\chi^{-1}(k)$, there is a set of states that have probability at least $1-\varepsilon$ under λ, conditioned on the realization of $h_{t_1}, h_{t_2}, \ldots$, on which beliefs will ζ-converge.

The proof is a straightforward category-by-category application of the arguments behind Proposition 5.

This extends the learning result of Proposition 5 to specifications of countable complexity. Notice, however, that for this extension to be effective, the agents must have a common understanding of the set of categories $\mathbb{K}$. For example, suppose $Y = \{y', y''\}$ and the data are generated by independent draws across periods, with the probability of y' being 1/3 in even periods and 2/3 in odd periods. An agent whose set of categories $\mathbb{K}$ allows even and odd periods to be distinguished will converge to beliefs matching these probabilities. An agent who fails to make distinctions will converge to beliefs attached probability 1/2 to y' in each period. Once again, we see that information can give rise to trade among agents who do not share a common view of the world.

5.3.6 Indiscriminate Learning

The learning process described in Propositions 5-6 describes an agent intent not only on making good predictions, but understanding the world. The agent begins with an initial ordering over credence functions, responds to likelihood scores, makes hypotheses as to whether the world is deterministic or random—all the sorts of things one would expect in a generalization of Bayesian reasoning. At the opposite end of the spectrum from Bayesian reasoning are methods such as reinforcement learning (e.g., Bush and Mosteller [12], appearing contemporaneously with Savage's [68] formalization of Bayesian decision making) that make no pretence of understanding the world, and yet can often predict quite well. We consider here two classes of learning rules that bear a similar relationship to our generalization of Bayesian learning.

To keep this comparison simple, let X be a singleton and consider an agent who in each period t observes the history of outputs $y_0, \ldots, y_{t-1}$ and must choose a mixture from a finite subset of ΔY, interpreted as a prediction of the period-t output.

First, let T denote the set of periods $\{0, 1, 2, 4, 8, 16, 32 \ldots\}$. The agent enumerates the computable states. In each period $t \in T$, the agent chooses the lowest-numbered computable state in her enumeration that is consistent with the history she has observed, and uses this state to make predictions until it is falsified. Once her lowest-ranking computable state is falsified, her prediction in each subsequent period is given by the empirical frequencies observed along the history, until the next element of T arrives. Then, she chooses the next computable state, and the process continues.

If the state is indeed computable, then the agent will eventually come to a description of the state in her enumeration, and will thereafter predict perfectly. If the state is generated by a draw from an exchangeable measure, then with probability one the proportion of periods the agent uses the empirical frequencies as her prediction will converge to one. Moreover, her predictions during these periods will converge to the frequencies corresponding to the underlying data-generating measure. One detects few traces of understanding the world in this process, but it nonetheless predicts well.

Second, Foster and Vohra [27] established the remarkable result that there exists a strategy mapping from histories into mixtures over ΔY ensuring that for any $\varepsilon > 0$, for every sequence of outputs, the agent is ε-calibrated.[38]

The obvious response to these initial calibration results is that matching the empirical frequency does not always ensure that the agent is predicting well. The standard example is the outcome $yy'yy'yy'yy'yy'yy'\ldots$. An agent who predicts that the probability of the observation y is $1/2$ is perfectly calibrated given this outcome. However, it seems clear that the agent is missing some relevant information.

Sandroni, Smorodinsky and Vohra [67], building on Lehrer [52], introduce the idea of a checking rule. A (forecast-based) checking rule is a function that maps, for every t, from the history $(y_0, \ldots, y_{t-1})$ and the period-t choice p_t of the agent, into $\{0, 1\}$, indicating whether the rule is active in period t or not. For example, a rule might apply if the past four realizations were y' and the agent predicted a probability of y' greater than .8, or if the empirical frequency of y' is sufficiently close to 0.6 and the agent names a probability sufficiently close to .4. The central result of Sandroni, Smorodinsky and Vohra [67] is the following. Choose any countable collection of checking rules. Then there is a learning rule ensuring that, for every checking rule, the agent is calibrated on the collection of periods that the checking rule is active. Sandroni, Smorodinsky and Vohra then note that there are only countably many computable functions, and so there are only countably many criteria one could apply for whether the agent is predicting "correctly". By making each such criterion a checking rule, we can be assured that no computable test could lead to the conclusion that the agent is not forecasting correctly.

The application of checking rules ensures that the agent will not respond to outcomes such as $yy'yy'yy'yy'yy'yy'\ldots$ by continuing to predict $1/2$, with the "only even periods" and "only odd periods" checking rules instead forcing the agent to recognize this pattern. More importantly for our purposes, Sandroni, Smorodinsky and Vohra's [67] result can be exploited to ensure that if the state is computable, then (taking refuge in a rough statement) the agent will eventually predict nearly perfectly, and if the state is drawn from an exchangeable measure

[38] Roughly, the agent is ε-calibrated if, for every prediction p the agent makes infinitely often, the empirical frequency of the realizations in those periods in which the agent predicts p are within ε of p. Foster and Vohra [27] provide an algorithm that describes such a strategy, and Foster [26] provides a simpler argument.

λ, then with arbitrarily high probability (under λ) the agent will eventually make predictions arbitrarily close to the measure over Y chosen by λ.[39] Hence, the agent predicts correctly under the circumstances that Proposition 6 ensures correct predictions.

Both methods described in this section give salutary asymptotic results, just as do Propositions 5-6. Our agents will in addition at least believe that they are doing the best they can in terms of prediction in the meantime, whereas the preceding learning rules do not allow such an interpretation. Indeed, the calibration results require that agents sometimes mix over forecasts. It is less clear how to interpret mixed predictions in terms of the agent's quest to understand the world she faces. It is then perhaps unsurprising that the most prominent application of calibration results has been to argue that agents can predict well, even with *no* understanding of the process they face.

6 Discussion

Our points of departure for this work are the familiar agreeing-to-disagree and no-trade theorems. The literature contains many variations on the points that disagreement or trade can arise when people hold different beliefs, or process information differently. Our model joins this list. What have we added?

The first component of our two-fold answer is conceptual. We agree that it is quite demanding to assume that people have identical prior beliefs and information-updating rules. However, the arguments leading one to doubt such fortuitous agreement should also lead one to doubt that a small worlds model is appropriate, prompting our large-worlds model. We expect this to be useful in applications beyond no-trade theorems.

The second component of our answer is more substantive. We view a no-trade counter argument as most compelling if one can argue that we expect agents to *persistently* engage in trade, despite opportunities to learn to agree. Persistent trade cannot arise in the familiar world of Bayesian agents with different (but equivalent) priors, forcing us to appeal to a large-worlds alternative. We then find that agents with different models of their interactions can reasonably be expected to come to similar views of the world, and cease trading, *if* the data generating process is sufficiently simple, such as being either deterministic or emanating from *iid* draws. In effect, we can expect agents to detect (at least approximately) deterministic processes and come to agree on them—people no

[39] For each computable state ω and each element $u \in Y$, let there be a checking rule that is "on" in period t if the history is consistent with ω and the period-t value of the computable state is the y. Hence, for each computable state we have $|Y|$ checking rules. Let there also be the "always on" checking rule. If the state is computable, then the agent must be calibrated on the sequence of periods in which the state takes value y, for each $y \in Y$. But this can be the case only if in an arbitrarily high proportion of such periods the agent makes a prediction arbitrarily close to putting probability one on y. If the state is drawn from an exchangeable measure, then the "always on" checking rule ensures that in an arbitrarily high proportion of periods, the agent must make a prediction arbitrarily close to the measure over Y chosen by λ.

longer argue about the motions of the planets. We can also expect people to recognize and agree on processes that are too complicated to explain deterministically, but stationary enough to explain as random—these are the conditions under which insurance markets flourish. However, if the data generating process is more confounding, then we cannot be sure of eventual agreement.

In the latter case, we can expect agents described by our learning process to sometimes exhibit abrupt shifts in beliefs, as their preferred explanation shifts from one type of explanation to another. Khaw, Stevens and Woodford [50] report results from a learning experiment that exhibit analogous behavior (with their Figure 2 providing a good illustration), though the subjects in their experiment are attempting to track the state of a hidden Markov process determining the parameters of conditionally stationary random draws rather than shifting between deterministic and random explanations.[40]

These results lead us to a more complete understanding of when we might expect to see trade, or when we are more likely to see more trade. When facing *iid* processes, such as simulated roulette wheels in a laboratory setting, we should see trade only if people are risk seeking, or have different risk attitudes that they have not already been able to accommodate with state-contingent trades. Facing a more complicated process, such as a financial market, we may well see trade, even after agents have written trades to implicitly insure their differences in risk attitudes. In particular, no one seriously argues that we should model financial markets as deterministic. At the same time, Reinhard and Rogoff [63] make a convincing case that throughout the last millennium, people have consistently rejected the view that financial markets are governed by stationary processes, moving us outside the purview of Propositions 5-6. Moreover, we might expect to especially see such trade after unexpected events have forced at least some agents to reevaluate their models of the market. Giacomini, Skreta and Turén [32] argue that these are the conditions under which professional forecasters are most likely to reconsider their models.

A Appendices

We provide foundations for each of the three components of our model. Proposition 7 plays an important role in our analysis, as we explain in Section 2.5.

A.1 Updating about the Data Generating Process

Our updating rule is the counterpart of Bayesian updating—we eliminate conjectures that have been refuted, retain those that have not been refuted (while eliminating any refuted states from those conjectures), and renormalize.

Our updating rule imposes no restrictions on posterior beliefs beyond minimal consistency requirements. In particular, we establish a version of Kolmogorov's theorem (Billingsley [4, Theorem 36.1, p. 517] for our model. Let

[40] The subjects in Khaw, Stevens and Woodford's [50] experiment had the advantage of being told the data generating process, putting them immediately into a small world.

(h_t, Y') denote that subset of $(X \times Y)^{t+1}$ (recalling that the first period is numbered 0) consisting of elements of the form (h_t, y) for some $y \in Y'$.

Proposition 7 *Let $(\tilde{\psi}_t)_{t=0}^{\infty}$ be a sequence of totally monotone capacities on $(X \times Y)^{t+1}$ satisfying the consistency conditions that for any t and for any set $S_t \subseteq (X \times Y)^{t+1}$,*

$$\tilde{\psi}_{t+1}(S_t) = \tilde{\psi}_t(S_t).$$

Then there exists a credence function ϕ on $(\mathcal{A}, \mathcal{E})$ such that for any h_t and for any set $Y' \subseteq Y$,

$$\tilde{\phi}_{h_t}(Y') \sim \tilde{\psi}_t((h_t, Y')).$$

Proof. Epstein and Wang [22, Theorem D.2, p. 1369] establish a generalization of Kolmogorov's extension theorem. As they note [22, p. 1370], a special case of their theorem gives an extension theorem for convex (including, as a special case, totally monotonic) capacities. Hence, there is a totally monotone capacity $\tilde{\Psi}$ on Ω with (for $[S_t] = \{\omega \in \Omega | (\omega(0), \ldots, \omega(t)) \in S_t\}$)

$$\tilde{\Psi}([S_t]) = \tilde{\psi}_t([S_t])$$

for any $S_t \subseteq (X \times Y)^{t+1}$, for any t. This implies

$$\tilde{\Psi}([h_t, Y']) = \tilde{\psi}_t((h_t, Y'))$$

for any history h_t and set $Y' \subseteq Y$. The capacity $\tilde{\Psi}$ thus generates the requisite beliefs, and we then need only show that there is a credence function that does to. Such a credence function is the Möbius transform of the totally monotone capacity $\tilde{\Psi}$. ∎

Remark 6 (Any Posterior from Any Prior) A special case of Proposition 7 is the familiar "any posterior from any prior" property for Bayesian updating. In particular, let Ω be a (for convenience) finite state space, let ρ be a prior belief on Ω, and let ρ' be a putative posterior on Ω, with the support of ρ' contained in the support of ρ. Then we can introduce the signal space $\{t', t''\}$ and construct a signal structure such that signal t' yields the posterior ρ'. Hence, Bayesian updating is also permissive with respect to posteriors.

The preceding observation allows us to state a version of the no-trade theorem to which we refer in Section 4.1:

Proposition 8 *Consider an exchange economy with state space Θ and (at least weakly) risk averse agents. Let the prior ρ on Θ be Bayesian and let the allocation z be efficient. Then z is also efficient for any Bayesian belief ρ' on Θ whose support is contained in ρ.*

Proof. Suppose that z is efficient with respect to distribution ρ but not distribution ρ'. Then we can introduce the signal space $\{t', t''\}$ and construct a distribution g over $\Theta \times \{t', t''\}$ whose marginal distribution on y is given by ρ, with signal t' giving rise to posterior ρ' and with signal t'' giving rise to some posterior ρ''.[41] Let z' be an allocation that, given posterior ρ', is strictly Pareto superior to allocation z. Let z'' be an allocation, possibly equal to allocation z, that given posterior ρ'', is weakly Pareto superior to allocation z. Then given the distribution g, the trade allocation given by z' if signal t' occurs and z'' if signal t'' occurs also strictly Pareto improves on z. Now taking the expected value of this trade (with respect to the signal distribution) leads to a trade that conditions only on Θ and Pareto improves on z given the distribution ρ, a contradiction. ■

■

Dempster [18] defines a rule for updating plausibility and uses this to derive a rule for updating his notion of belief. As we note in Gilboa, Samuelson and Schmeidler [35, p. 1407], the resulting rule for updating Dempster's notion of belief is equivalent to our updating rule. Similarly, Dempster [18] formulates a "rule of combination" for combining the information contained in two credence functions. It is straightforward to show that if one of the credence functions is equivalent to conditioning on an event, then the combination of the two credence functions gives a result equivalent to his (and hence our) rule for updating beliefs.[42]

[41] For example, let $k > 0$ be a number such that, for all $\omega \in \Omega$ we have

$$k\rho'(\omega) < \rho(\omega).$$

Since the support of ρ' is contained in the support of ρ, such a k exists. Now let the joint distribution of states and signals be given by, for all $\omega \in \Omega$,

$$\begin{aligned} g(\omega, t') &= k\rho'(\omega) \\ g(\omega, t'') &= \rho(\omega) - k\rho'(\omega). \end{aligned}$$

This is clearly a probability distribution whose marginal distribution on the set Ω is the prior ρ. Upon observing signal t', the posterior is

$$\frac{g(\omega, t')}{\sum_{\tilde{\omega}\in\Omega} g(\tilde{\omega}, t')} = \frac{k\rho'(\omega)}{\sum_{\tilde{\omega}\in\Omega} k\rho'(\tilde{\omega})} = \rho'(\omega).$$

[42] Gilboa and Schmeidler [36] define a class of updating rules for nonadditive priors, with a special case being the Dempster-Shafer rule for updating belief. All of these rules correspond to Bayes' rule when the prior is additive. Gilboa and Schmeidler [36] also define a class of rules for updating *sets* of additive priors, which consist of first excluding some priors, and then updating the remainder individually according to Bayes' rule. One method for choosing which priors to exclude is to drop all those that do not attain maximum likelihood, given the history. Footnote 46 explains an equivalence in decision making between using a unique nonadditive prior and using multiple additive priors. Gilboa and Schmeidler [36] show that using the maximum-likelihood-exclusion-then-updating rule on multiple priors is equivalent to using the Dempster-Shafer belief updating rule for (unique) nonadditive priors, extending the equivalence of footnote 46 to updating.

A.2 Forming Beliefs about y_t

Let E (generic for event) be a subset of Ω and let ϕ be a credence function. Then we define

$$\overline{\phi}(E) = \phi(\{A \in \mathcal{E} | A \cap E \neq \emptyset\})$$

and

$$\underline{\phi}(E) = \phi(\{A \in \mathcal{E} | A \subseteq E\}).$$

Dempster [18] (see also Shafer [72]) defines the first of these to be the *plausibility* of the event E and the second to be the *belief* of E. Dempster suggests that the plausibility is the largest probability one could conceive of attaching to E, and the belief is the smallest probability one could conceive of attaching to E.

The definition of the capacity $\tilde{\phi}_{h_t}$ given in (1) is the counterpart of Dempster's notion of belief—$\mathcal{A}(h_t, Y')$ is the set of surviving (given h_t) conjectures whose prediction for the current period is a *subset* of Y'. Of Dempster's two extremes, why work with belief rather than plausibility, or why not work with some capacity in the middle?

A capacity $\tilde{\phi}_{h_t}$ on Y is totally monotone if, for any set $Y' \subseteq Y$, the Möbius transform $\phi^m_{h_t}$ on Y corresponding to the capacity $\tilde{\phi}_{h_t}$ attaches positive weight to Y', i.e.:

$$\phi^m_{h_t}(Y') := \sum_{Y'' \subseteq Y'} (-1)^{|Y'|-|Y''|} \tilde{\phi}_{h_t}(Y'') \geq 0. \tag{8}$$

A positive Möbius transform can be viewed as a probability distribution over the power set of Y, with the discrete σ-algebra.

We then recover in our setting the following version of a standard result (cf. Shafer [72, pp. 37-39]):[43]

Proposition 9 *For any history h_t, the capacity $\tilde{\phi}_{h_t}$ defined by the belief function given in* (1) *is totally monotone.*

The proof is a mildly tedious but straightforward calculation that, given (1) and (8), leads to

$$\phi^m_{h_t}(Y') \quad = \quad \frac{\phi([h_t, Y'])}{\phi(\{A \in \mathcal{A} | \emptyset \neq A \cap [h_t] \subseteq [h_t, Y]\})}. \tag{9}$$

Total monotonicity then follows from (9) and the fact that ϕ is a probability measure over $\mathcal{A}$, and hence takes positive values. Working with belief functions thus ensures total monotonicity. Total monotonicity in turn captures our view of conjectures as representing arguments that might be brought to bear in favor of an event. In particular, a calculation analogous to that yielding (9) gives

$$\tilde{\phi}_{h_t}(Y') \quad = \quad \sum_{Y'' \subseteq Y'} \phi^m_{h_t}(Y'').$$

[43] In a similar vein, Gilboa, Samuelson and Schmeidler [35, Proposition 4] show that a qualitative capacity (representing an ordinal ranking) is p-monotone (the analog of total monotonicity for qualitative capacities) if it is derived from a belief function.

We can view $\phi^m_{h_t}(Y'')$ (for $Y'' \subseteq Y'$) as the marginal contribution made to the weight attached to the set Y' by those conjectures supporting the set Y''. Total monotonicity ensures that the marginal contribution of any conjecture supporting the set Y' is positive. The same is not true of plausibility.[44]

Instead of belief, we might work with capacities that lie between those given by belief and those given by plausibility, some of which also preserve total monotonicity. To see what distinguishes belief, consider a credence function ϕ and history h_t, with the induced capacity $\tilde{\phi}$ (suppressing the history subscript). Let $\tilde{\Psi}$ denote the set of capacities $\tilde{\psi}$ that attach to each subset $\tilde{Y} \subseteq Y$ a value $\tilde{\psi}(\tilde{Y})$ that lies between the belief and plausibility of $\tilde{Y}$. We can induce a partial order $\leq$ on $\tilde{\Psi}$ by saying that $\tilde{\psi} \leq \tilde{\tilde{\psi}}$ if, for every $Y' \subseteq Y$, we have $\tilde{\psi}(Y') \leq \tilde{\tilde{\psi}}(Y')$.

Given a Möbius transform ϕ^m, say that the Möbius transform ψ^m is obtained from ϕ^m via a reduction if there is a set $Y' \subseteq Y$, a partition $(Y_1, \ldots, Y_n)$ of Y', and a vector $(\alpha_1, \ldots, \alpha_n) \in \mathbb{R}^n_+$ such that

$$\begin{array}{rcl}
\psi^m(Y_1) & = & \phi^m(Y_1) + \alpha_1 \\
\vdots & \vdots & \vdots \\
\psi^m(Y_n) & = & \phi^m(Y_n) + \alpha_n \\
\psi^m(Y') & = & \phi^m(Y') - \alpha_1 - \ldots - \alpha_n \geq 0,
\end{array}$$

with $\psi^m(Y'') = \phi^m(Y'')$ for any $Y'' \subseteq Y$ not contained Y'. Hence, the reduction chooses a subset Y', reduces the probability attached to Y' by ϕ^m by the amount $\alpha_1 + \ldots + \alpha_n$, and distributes this probability among the subsets of Y'. Intuitively, weight in the Möbius transform moves from a relatively large set Y' to smaller subsets of Y'. We write $\tilde{\phi} \succsim \tilde{\psi}$ if ψ^m can be obtained from ϕ^m via a finite sequence of reductions.

The following is now a straightforward calculation:

Proposition 10 *Fix a credence function ϕ and history h_t, inducing (via Dempster's notion of belief) the capacity $\tilde{\phi}$, and its Möbius transform ϕ^m. Then:*
[10.1] $\tilde{\psi} \in \tilde{\Psi} \Longrightarrow \tilde{\psi} \geq \tilde{\phi}$.
[10.2] $\tilde{\psi} \in \tilde{\Psi} \Longleftrightarrow \tilde{\phi} \succsim \tilde{\psi}$.

To interpret this result, recall that $\tilde{\Psi}$ is the set of all capacities that are bounded by Dempster's notions of belief and plausibility. Proposition [10.1] confirms that $\tilde{\phi}$, being belief, is a minimal element of this set. Proposition [10.2] indicates that every other element $\tilde{\psi}$ of the set can be obtained from $\tilde{\phi}$ by pushing probability

[44] For the simplest such example, let $Y = \{y', y''\}$ and let $\tilde{\phi}$ be given by

$$\tilde{\phi}(\{y'\}) = \tilde{\phi}(\{y''\}) = \tilde{\phi}(\{y', y''\}) = \frac{1}{3}.$$

Then plausibility gives the capacity

$$\tilde{\phi}(\{y'\}) = \tilde{\phi}(\{y''\}) = \frac{2}{3}, \qquad \tilde{\phi}(\{y'.y''\}) = 1,$$

which fails total monotonicity.

in the Möbius transform ϕ downward from larger sets to smaller sets. Hence, of all the totally monotone capacities that lie between Dempster's notions of belief and plausibility, we are working with the most conservative, in the sense that the agent is least assertive in how the weight attached to a conjecture associated with a set Y' should be apportioned among the subsets of Y'. We find this choice intuitive: whenever the agent is willing to interpret weights attached to a set Y' as also lending support to strict subsets of Y', then we expect the agent to instead have build this interpretation into her credence function. If we are then to respect this construction, we must derive beliefs from credence functions via Dempster's notion of belief.

A.3 Decisions

Schmeidler [70] provides a foundation for the decision criterion given by (2). Let $\succsim$ be a preference order on $\mathcal{L}$ (the set of functions from Y into probability distributions over the set of consequences Z). Schmeidler [70] shows:[45]

Proposition 11 *The preference order $\succsim$ satisfies weak order, comonotonic independence, continuity, monotonicity and nondegeneracy if and only if there exist an affine (and unique up to affine transformations) function $\tilde{u}$ on $\mathcal{L}$ and a unique capacity $\tilde{\phi}$ on Y such that*

$$\ell \succsim \ell' \iff \int_Y \tilde{u}(\ell(y))d\tilde{\phi}(y) \geq \int_Y \tilde{u}(\ell'(y))d\tilde{\phi}(y).$$

This criterion matches (2). We thus have a relaxation of the axioms of Anscombe and Aumann [1] that (taking nondegeneracy for granted) replaces independence with comonotonic independence, and is equivalent to Choquet expected utility maximization.

Why might one prefer comonotonic independence to independence? For any function $\ell \in \mathcal{L}$, let $\mathcal{Y}(\ell)$ be the coarsest σ-algebra with respect to which ℓ is measurable. Then if ℓ and ℓ' are comonotonic, we have $\mathcal{Y}(\ell) = \mathcal{Y}(\ell')$. Hence, when we approach a problem in comonotonic independence by first ranking two comonotonic acts, and then by considering mixtures of these two acts with a third comonotonic act, this mixing requires the consideration of no new events. This is not the case with independence, where the mixtures can require considering the measures of sets that were previously unevaluated. One may well be comfortable with stronger restrictions on comonotonic acts than acts in general.[46]

[45] The conditions on $\succsim$ are familiar, and we refer to the original for definitions, with the possible exception of comonotonic independence. Two acts are comonotonic if $\ell(y) \succsim \ell(y') \Rightarrow \ell'(y) \succsim \ell'(y')$. Comonotonic independence holds if the usual independence axioms holds for all triples of acts that are pairwise comonotonic. Hence, for pairwise comonotonic ℓ, ℓ' and ℓ'' and $\alpha \in (0,1)$, we have $\ell \succ \ell' \Longrightarrow \alpha\ell + (1-\alpha)\ell'' \succ \alpha\ell' + (1-\alpha)\ell''$. Shervish, Seidenfeld and Kadane [69, Sections 5 and 6] (see Wakker [76] for a similar argument) examine the extent to which axioms characterizing preferences are preserved as states are grouped together to reduce the size of a model.

[46] Gilboa and Schmeidler [36] show that the preferences represented by maximizing the

Remark 7 (State-independent Utility) The utility function u (cf. (3)) in this result is state independent, unlike the utility functions that appear in Milgrom and Stokey's [58] presentation of the no-trade theorem and Morris' [60] examination of different prior beliefs, and much of the applied literature. It is without loss of generality to work without state independent utilities. By appropriately defining an allocation, one can always represent a utility function as being state independent, perhaps in the process working with a model in which some allocations arise only in some states. Given sets of states Θ and consequences Z' and a state-dependent utility function, we can take the set of consequences to be $Z = Z' \times \Theta$, and then restrict attention to allocations that always map state θ into an element of $Z' \times \{\theta\}$. We can then define a state-independent utility function on Z. However, this device can sometimes become quite cumbersome, and it is often more convenient to work with state-dependent utility. We readily move between state dependent and state independent utility functions as is convenient.[47] ■

A.4 Proof of Lemma 1

[1.1] Statement [1.1] is immediate from the strict concavity of the utility functions—if $z_{ik} > z_{i\ell}$, then strict concavity ensures that $u_i'(z_{ik}) < u_i'(z_{i\ell})$, so balancedness requires $u_i'(z_{jk}) < u_i'(z_{j\ell})$, at which point strict concavity ensures $z_{jk} > z_{j\ell}$, delivering strict comonotonicity.

[1.2] We next show that balanced allocations are efficient. Fix the capacity $\tilde{\phi}$ and its Möbius transform ϕ^m. Let $\mathcal{Z}$ be the set of feasible allocations, i.e., the set of functions given by

$$\left\{ z : I \times Y \to \mathbb{R}_+ \text{ such that } \forall y \in Y, \sum_{i \in I} z_{iy} = \sum_{i \in I} e_{iy} \right\}.$$

Let $\mathcal{P}(Y)$ denote the set of nonempty subsets of Y. Then define the set $\mathcal{W}$ to be

$$\left\{ w : I \times \mathcal{P}(Y) \to \mathbb{R}_+ \text{ such that } \forall Y' \in \mathcal{P}(Y), \sum_{i \in I} w_{iY'} \leq \min_{y \in Y'} \left\{ \sum_{i \in I} e_{iy} \right\} \right\}.$$

Choquet integral of a utility function u (cf. (3)) against a capacity $\tilde{\phi}$ are equivalent to those represented by the maxmin of the expected utility u against the set of multiple priors given by the core of $\tilde{\phi}$. (For a capacity $\tilde{\phi}$, say that the measure p is in the core of $\tilde{\phi}$ if $p(A) \geq \tilde{\phi}(A)$ for all A. Every totally monotonic capacity is convex ($\phi(Y \cup Y') + \phi(Y \cap Y') \geq \phi(Y) + \phi(Y')$), and convex capacities have nonempty cores.) Conversely for any maxmin representation with a convex set of priors, we can derive a capacity from the set of priors such that the preferences are represented by a Choquet integral against this capacity.

[47] Binmore [7, p. 351] invokes Aesop's principle (cf. footnote 9) to argue that utility should be modeled as state independent. However, we interpret Aesop's principle as requiring that preferences over a set of objects should be independent of the probabilities the agent attaches to those objects, and that preferences over objects should remain unaffected as the consumption set of feasible objects expands or contracts, but as making no comment as to the identity (i.e., only consequences, or consequences and states) of the objects over which preferences are defined.

We will refer to an element of $\mathcal{W}$ as an a $\mathcal{W}$-allocation, a moniker justified by thinking of an alternative economy in which the elements of $\mathcal{P}(Y)$ are outputs. Now define the function $\mathbf{w} : \mathcal{Z} \to \mathcal{W}$ by letting

$$\mathbf{w}(z)_{iY'} = \min_{y \in Y'} z_{iy}.$$

We need to verify that $\mathbf{w}$ indeed maps into $\mathcal{W}$, which follows from noting that for any allocation z, any set $Y' \in \mathcal{P}(Y)$ and any $y' \in Y'$, we have

$$\sum_{i \in I} \boldsymbol{w}(z)_{iY'} \leq \sum_{i \in I} \min_{y \in Y'} \{z_{iy}\} \leq \sum_{i \in I} z_{iy'} = \sum_{i \in I} e_{iy'}$$

and hence (since this holds for any $y' \in Y'$), we have $\sum_{i \in I} \boldsymbol{w}(z)_{iY'} \leq \min_{y \in Y'} \{\sum_{i \in I} e_{iy}\}$, as needed. We now note that, from the definition of the Choquet integral, we have, for any $z \in \mathcal{Z}$ and any capacity $\tilde{\phi}$ with associated Möbius transform ϕ^m,

$$\int_Y u_i(z_{iy}) d\tilde{\phi} = \sum_{Y' \in \mathcal{P}(Y)} u_i(\mathbf{w}(z)_{iY'}) \phi^m(Y').$$

Finally, we can define a $\mathcal{W}$-allocation as being balanced, analogously to our definition for allocations in Z. We then note that if the allocation z is balanced, then it is strictly comonotonic (Proposition 1.1) and hence the set $\arg\min_{y \in Y} z_{iy}$ is identical across agents, and so the $\mathcal{W}$-allocation $\mathbf{w}(z)$ is balanced.

Now fix a balanced allocation z and capacity $\tilde{\phi}$, with associated Möbius transform ϕ^m. The arrival of information gives rise to a new capacity $\tilde{\psi}$, with associated Möbius transform ψ^m. We need to show that the allocation z is still efficient. This is equivalent to showing that given the probability distribution ψ^m on the set $\mathcal{P}(Y)$, there is no $w \in \mathcal{W}$ that Pareto dominates $\mathbf{w}(z)$.

Putting these considerations together, it suffices to show that the balanced $\mathcal{W}$-allocation z on the output set $\mathcal{P}(Y)$ is efficient for any probability distribution. In particular, it is then efficient for the probability distributions ϕ^m and ψ^m. But this is immediate from the facts that (i) an allocation problem is efficient if and only if it can be obtained as the solution to the problem of maximizing a weighted sum of the agents' utilities, (ii) the assumption that the utility functions u_i are differentiable and concave ensures that the first-order conditions for this maximization problem are sufficient, and (iii) balancedness ensures that these first-order conditions are satisfied.[48]

[1.3] This result follows from the observations that finished the previous argument—(i) an allocation problem is efficient if and only if it can be obtained as the solution to the problem of maximizing a weighted sum of the agents' utilities, (ii) the assumption that the utility functions u_i are differentiable and concave ensures that the first-order conditions for this maximization problem are sufficient, and (iii) for any set of weights, there is a balanced allocation solving the first-order conditions. ■

[48]Similar ideas lie behind Proposition 8 in Appendix A.1.

A.5 Proof of Lemma 2

Let z' be an efficient but unbalanced allocation, and let z^* be the balanced allocation yielding payoffs identical to those of z'. Let $\lambda \in (0,1)$. For every i, we have

$$
\begin{aligned}
& \sum_{Y' \in \mathcal{P}(Y)} u_i(\min_{y \in Y'} \{\lambda z'_{iy} + (1-\lambda) z^*_{iy}\}) \phi^m(Y') \\
\geq & \sum_{Y' \in \mathcal{P}(Y)} (u_i(\lambda \min_{y \in Y'} \{z'_{iy}\} + (1-\lambda) \min_{y \in Y'} \{z^*_{iy}\}) \phi^m(Y') \\
\geq & \ \lambda \sum_{Y' \in \mathcal{P}(Y)} u_i(\min_{y \in Y'} \{z'_{iy}\}) \phi^m(Y') + (1-\lambda) \sum_{Y' \in \mathcal{P}(Y)} u_i(\min_{y \in Y'} \{z^*_{iy}\}) \phi^m(Y') \\
= & \sum_{Y' \in \mathcal{P}(Y)} u_i(\min_{y \in Y'} \{z'_{iy}\}) \phi^m(Y') \\
= & \sum_{Y' \in \mathcal{P}(Y)} u_i(\min_{y \in Y'} \{z^*_{iy}\}) \phi^m(Y'),
\end{aligned}
$$

where the first term is the payoff from the feasible allocation $\lambda z' + (1-\lambda) z^*$, the first inequality follows from moving the minimizations inside the convex combination, the next inequality from the concavity if u, and the equalities follow from the observation that the final two terms are the (equal) payoffs from allocations z' and z^*. Because z' and z^* are distinct, it must be the case that as λ varies from zero to one, resources are shifted between agents without affecting payoffs, i.e., there is slack in the allocation z'. ■

A.6 Proof of Lemma 4

[4.1] This follows because there are countably many Turing machines.

[4.2] These statements follow immediately from the definitions.

[4.3]

It is immediate that an eventually-constant state is random. Suppose the state is not eventually constant. Then there exists an output $y' \in Y$ and an infinite sequence of periods $\{t'_1, t'_2, \ldots\}$ in which the output is y', and there also exists an output y'' and an infinite sequence of periods $\{t''_1, t''_2, \ldots\}$ in which the output is y''. If the state is computable, then the sequences $\{t'_1, t'_2, \ldots\}$ and $\{t''_1, t''_2, \ldots\}$ can be recursively defined. Because these sequences have different random limits, the state is not random.

[4.4] Church [14] shows that the cardinality of the set of random states is that of the continuum, and the fact that there are only countably many computable states then ensures that there is continuum of noncomputable, random states. Second, Doob [19] shows that if $Y = \{y', y''\}$ and each period's value is drawn independently, with equiprobable realizations, then with probability 1 the state is random (and has limiting proportion $1/2$ of y' realizations). Church [14] notes that if each period's value is drawn independently with probability $p \in [0,1]$ of output y', then with probability 1 the state drawn will be random and will have limiting proportion p of y' realizations. The extension to large sets Y is immediate, and from these observations, it follows that if the measure λ corresponds to an exchangeable process, then with probability 1 the state will be random.

[4.5] A state may neither have a random limit nor be computable, and there are a continuum of such states. To construct a collection of such states, let $Y = \{y', y''\}$ and hence $\Omega = \{y', y''\}^\infty$, and let $(n_k)_{k=0}^\infty$ be the sequence of numbers defined by letting $n_0 = 1$ and

$$n_k = 3 \sum_{\ell=1}^{k-1} n_\ell .$$

Then, let $F : \Omega \to \Omega$ be the function defined by letting, for each state ω, $F(\omega)$ be that state in which the first n_0 realizations equal ω_0, the next n_1 realizations equal ω_1, next n_2 realizations equal ω_2, and so on. Then F is a 1-1. Hence, there is obviously a continuum of states in $F(\Omega)$, and thus necessarily a continuum of noncomputable such states. Moreover, each such state is random only if the original state from which it is constructed has either only finitely many y' values or finitely many y'' values in its specification, i.e., is eventually constant. There are in turn a countable number of such states, and so we have a continuum of nonrandom, noncomputable states.

[4.6] Doob [19] shows that if $Y = \{y', y''\}$ and each period's value is drawn independently, with equiprobable realizations, then with probability 1 the state is random (and has limiting proportion 1/2 of y' realizations). Church [14] notes that if each period's value is drawn independently with probability $p \in [0, 1]$ of output y', then with probability 1 the state drawn will be random and will have limiting proportion p of y' realizations. The extension to large sets Y is immediate, and from these observations, it follows that if the measure λ corresponds to an exchangeable process, then with probability 1 the state will be random. ■

A.7 Proof of Lemma 5

[5.1] We have

$$\pi_\theta(\lambda^{q(h_t(\omega))}, h_t(\omega)) - \pi_\theta(\lambda^q, h_t(\omega)) = \sum_{y \in Y} q_y(h_t(\omega)) \left(\log_2 q_y(h_t(\omega)) - \log_2 q_y \right). \tag{10}$$

It follows from the strong law of large numbers that, conditional on q, the proportions $q_y(h_t(\omega))$ converge to q_y almost surely. Hence, for each q, (10) converges to zero λ^q-almost surely, and so (10) converges to zero λ-almost surely. Then, by Egorov's theorem, there is a set of measure at least $1 - \varepsilon/2$ on which (10) converges to zero uniformly. Hence, we can find a function $\delta(t)$ with $\lim_{\delta \to \infty} 0$ such that, with probability at least $1 - \varepsilon/2$, we have

$$\pi_\theta(\lambda^{q(h_t(\omega))}, h_t(\omega)) - \pi_\theta(\lambda^q, h_t(\omega)) \leq \delta(t) \quad \forall t \in \mathbb{N},$$

which is the desired result.

[5.2] This follows from a similar argument. ∎

A.8 Proof of Lemma 6

Because ω is computable, there exists a Turing machine that responds to input t by producing the output ω_t. A history h_t can then be described by a program that includes the specification of the Turing machine, an instruction for the machine to produce the sequence of outputs it would produce in response to the inputs $(1, \ldots, \tilde{t})$ for some $\tilde{t}$, and a specification that $\tilde{t} = t$. The complexity of the first two components is finite and independent of t, while that specification of t can be done with complexity at most $2\log_2(t)$. We thus have that $K(h_t|t)/(t) \leq (k + 2\log_2(t))/(t)$ for some k, which converges to 0 as t grows. ∎

References

[1] Francis J. Anscombe and Robert J. Aumann. A definition of subjective probability. *Annals of Mathematical Statistics*, 34(1):199–205, 1963.

[2] Kenneth J. Arrow. *Essays on the Theory of Risk Bearing.* Markham, Chicago, 1971.

[3] Robert J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976.

[4] Patrick Billingsley. *Probability and Measure.* John Wiley and Sons, New York, anniversary edition, 2012.

[5] Antoine Billot, Alain Chateauneuf, Itzhak Gilboa, and Jean-Marc Tallon. Sharing beliefs: Between agreeing and disagreeing. *Econometrica*, 68(3):685–694, 2000.

[6] Kenneth G. Binmore. *Rational Decisions.* Princeton University Press, Princeton, 2009.

[7] Kenneth G. Binmore. A miminal extension of Bayesian decision theory. *Theory and Decision*, 109(3):341–362, 2016.

[8] David Blackwell and Lester Dubins. Merging of opinions with increasing information. *The Annals of Mathematical Statistics*, 33(3):882–886, 1962.

[9] Lawrence E. Blume, David A. Easley, and Joseph Y. Halpern. Redoing the foundations of decision theory. In *Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR2006)*, pages 14–24, 2006.

[10] Adam Brandenburger, Eddie Dekel, and John Geanakoplos. Correlated equilibrium with general information structures. *Games and Economic Behavior*, 4(2):182–201, 1992.

[11] Isabelle Brocas, Juan D. Carrillo, Stephanie W. Wang, and Colin F. Camerer. Imperfect choice or imperfect attention: Understanding strategic thinking in private information games. *Review of Economic Studies*, 81(3):944–970, 2014.

[12] Robert R. Bush and Frederick Mosteller. *Stochastic Models for Learning.* John Wiley and Sons, New York, 1955.

[13] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning and Games.* Cambridge University Press, Cambridge, 2006.

[14] Alonzo Church. On the concept of a random sequence. *Bulletin of the American Mathematical Society*, 46(2):130–135, 1940.

[15] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory.* John Wiley and Sons, New York, second edition, 2006.

[16] Bruno de Finetti. Funzione caratteristica di un fenomeno aleatorio. *Atti della Academia Nazionale dei Lincei Rendiconti, Class di Scienze Fisiche, Matematiche e Naturali*, 4:86–133, 1930.

[17] Bruno de Finetti. La prevision: Ses lois logiques, ses sources subjectives. *Annales de l'Institute Henri Poincare*, 7(1):1–68, 1937.

[18] Arthur P. Dempster. Upper and lower probabilities induced by a multivaued mapping. *Annals of Mathematical Statistics*, 38(2):325–339, 1967.

[19] Joseph L. Doob. Note on probability. *Annals of Mathematics*, 37(2):363–367, 1936.

[20] James Dow, Vincente Madrigal, and Sergio Ribeiro da Costa Werlang. Preferences, common knowledge and speculative trade. Working paper number 149, London Business School, 1990.

[21] Larry G. Epstein and Martin Schneider. Recursive miltiple-priors. *Journal of Economic Theory*, 113(1):1–31, 2003.

[22] Larry G. Epstein and Tan Wang. "Beliefs about beliefs" without probabilities. *Econometrica*, 64(6):1343–1373, 1996.

[23] Erik Eyster and Matthew Rabin. Cursed equilibrium. *Econometrica*, 73(5):1623–1672, 2005.

[24] Erik Eyster, Matthew Rabin, and Dimitri Vayanos. Finanical markets where traders neglect the informational content of prices. *Journal of Finance*, 74(1):371–399, 2018.

[25] Sandra Fortini, Sonia Petrone, and Polina Sporysheva. On a notion of partially conditionally identically distributed sequences. *Stochastic Processes and Their Applications*, 128(3):819–846, 2018.

[26] Dean Foster. A proof of calibration via Blackwell's approachability theorem. *Games and Economic Behavior*, 29(1-2):73–78, 1999.

[27] Dean Foster and Rakesh V. Vohra. Asymptotic calibration. *Biometrika*, 85(2):379–390, 1998.

[28] Kenneth R. French. Presidential address: The cost of active investing. *Journal of Economic Finance*, 63(4):1537–1573, 2008.

[29] John Geanakoplos. Game theory without partitions, and applications to speculation and concensus. Cowles Foundation discussion paper no. 914, Yale University, 1989.

[30] John Geanakoplos and Heraklis Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 27(2):192–200, 1982.

[31] Paolo Ghirardato. Coping with ignorance: Unforseen contingencies and non-additive uncertainty. *Economic Theory*, 17(2):247–276, 2001.

[32] Raffaella Giacomini, Vasiliki Skreta, and Javier Turén. Models, inattention and Bayesian updates. *American Economic Journal: Macroeconomics*, 2017.

[33] Itzhak Gilboa, Stefania Minardi, Larry Samuelson, and David Schmeidler. States and contingencies: How to understand Savage without anyone being hanged. Working paper, Tel Aviv University, HEC Paris, Yale University, and Tel Aviv University, 2019.

[34] Itzhak Gilboa and Larry Samuelson. Subjectivity in inductive inference. *Theoretical Economics*, 7(2):183–216, 2012.

[35] Itzhak Gilboa, Larry Samuelson, and David Schmeidler. Dynamics of inductive inference in a unified model. *Journal of Economic Theory*, 148(4):1399–1432, 2013.

[36] Itzhak Gilboa and David Schmeidler. Updating ambiguous beliefs. *Journal of Economic Theory*, 59(1):33–49, 1993.

[37] Simon Grant, Idione Meneghel, and Rabee Tourky. Learning under unawareness. Australian National University, 2017.

[38] Yoram Halevy. Trade between rational agents as a result of asymmetric information. Caress working paper 98-03, University of Pennsylvania, 1998.

[39] Yoram Halevy. The possibility of speculative trade between dynamically consistent agents. *Games and Economic Behavior*, 46(1):189–198, 2004.

[40] John C. Harsanyi. Games with incomplete information played by Bayesian players, parts I-III. *Management Science*, 14:159–182, 320–334, 486–502, 1967–68.

[41] John C. Harsanyi. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge University Press, Cambridge, Massachusetts, 1977.

[42] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning, Second Edition*. Springer, New York, 2017.

[43] Geoffrey Heal and Antony Millner. Reflections: Uncertainty and decision in climate change economics. *Review of Environmental Economics and Policy*, 8(1):120–137, 2014.

[44] John H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, 1975.

[45] David Hume. *An Enquiry Concerning Human Understanding*. Clarendon Press, Oxford, 1748.

[46] Philippe Jehiel. Analogy-based expectation equilibrium. *Journal of Economic Theory*, 123(2):81–104, 2005.

[47] Philippe Jehiel. Investment strategy and selection bias: An equilibrium perspective on overoptimism. *American Economic Review*, 108(6):1582–1597, 2018.

[48] Philippe Jehiel and Frédéric Koessler. Revisiting games of incomplete information with analogy-based expectations. *Games and Economic Behavior*, 62(2):533–557, 2008.

[49] Olav Kallenberg. Spreading and predictive sampling in exchangeable sequences and processes. *The Annals of Probability*, 16(2):508–534, 1988.

[50] Mel Win Khaw, Luminita Stevens, and Michael Woodford. Discrete adjustment to a changing environment: Experimental evidence. Columbia University and University of Maryland, 2017.

[51] Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

[52] Ehud Lehrer. Any inspection is manipulable. *Econometrica*, 69(5):1333–1347, 2001.

[53] Duncan Luce and Howard Raiffa. *Games and Decisions*. John Wiley and Sons, New York, 1957.

[54] Chenghu Ma. A no-trade theorem under Knightian uncertainty with general preferences. *Theory and Decision*, 51(2):173–181, 2000.

[55] Michael Magill and Martine Quinzii. *Theory of Incomplete Markets*. MIT Press, Cambridge, 2002.

[56] Per Martin-Löf.

[57] Per Martin-Löf. The definition of random sequences. *Information and Control*, 9(6):602–619, 1966.

[58] Paul R. Milgrom and Nancy Stokey. Information, trade, and common knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.

[59] Marvin Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, Englewood Cliffs, NJ, 1967.

[60] Stephen Morris. Trade with heterogeneous prior beliefs and asymmetric information. *Econometrica*, 62(6):1327–1347, 1994.

[61] Sujoy Mukerji. Understanding the nonadditive probability decision model. *Economic Theory*, 9(1):23–46, 1997.

[62] Sujoy Mukerji and Hyun Song Shin. Equilibrium departures from common knowledge in games with non-additive expected utility. *Advances in Theoretical Economics*, 2(1):Article 2, 2002.

[63] Carmen M. Reinhart and Kenneth S. Rogoff. *This Time is Different: Eight Centures of Financial Folly.* Princeton University Press, Princeton, 2009.

[64] Barry Ritholtz. *Bailout Nation.* John Wiley and Sons, Hoboken, New Jersey, 2009.

[65] Stephen A. Ross. Intertemporal asset pricing. In S. Bhattacharya and G. Constantinides, editors, *Theory of Valuation: Frontiers in Modern Financial Theory, Volume 1*, pages 85–96. Rowman and Littlefield, Totowa, N.J., 1989.

[66] Dov Samet. Ignoring ignorance and agreeing to disagree. *Journal of Economic Theory*, 52(1):190–207, 1990.

[67] Alvaro Sandroni, Rann Smorodinsky, and Rakesh V. Vohra. Calibration with many checking rules. *Mathematics of Operations Research*, 28(1):141–153, 2003.

[68] Leonard J. Savage. *The Foundations of Statistics.* Dover Publications, New York, 1972. Orignally 1954.

[69] Mark J. Schervish, Teddy Seidenfeld, and Joseph B. Kadane. State-dependent utilities. *Journal of the American Statistical Association*, 85(11):840–847, 1990.

[70] David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–587, 1989.

[71] James K. Sebenius and John Geanakoplos. Don't bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association*, 78(382):424–426, 1983.

[72] Glenn Shafer. *A Mathematical Theory of Evidence.* Princeton University Press, Princeton, 1976.

[73] Sylvain Sorin. Merging, reputation, and repeated games with incomplete information. *Games and Economic Behavior*, 29(1/2):274–308, 1999.

[74] Jean Ville. *Etude Critique de la Notion de Collectif.* Gauthier-Villars, Monographies des Probabilits, no. 3, Paris, 1939.

[75] Richard von Mises. *Probability, Statistics, and Truth.* Allen and Unwin, 1928. Second revised English edition, 1961.

[76] Peter P. Wakker. Nonexpected utility because of incompleteness of models: Rereading Savage's small worlds. Mimeo, CenTER, University of Tilburg, 1996.

[77] David H. Wolpert and William G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, 1997.

[78] Alexander Zimper. Half empty, half full, and why we can *agree to disagree* forever. *Journal of Economic Behavior and Organization*, 71(2):283–299, 2009.