Uncertainty and Disagreement in Equilibrium Models*

Nabil I. Al-Najjar† Eran Shmaya‡

October 2012;
This version: November 2012

Abstract

Leading equilibrium concepts require agents' beliefs to coincide with the model's true probabilities and thus be free of systematic errors. This implicitly assumes a criterion that tests beliefs against the observed outcomes generated by the model. We formalize this requirement in stationary environments. We show that there is a tension between the requirements that beliefs can be tested against systematic errors, on the one hand, and allowing agents to disagree or be uncertain about the long-run fundamentals. We discuss the implications of our analysis in the contexts of asset pricing and dynamic games.

*We thank Luciano De Castro for his comments, and Luciano Pomatto and Jaber Zarzadeh for their research assistance.

† Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208.

‡ Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208.

Contents

1 Introduction

Leading equilibrium concepts such as Bayesian-Nash equilibrium and rational expectations equilibrium incorporate the idea that agents' beliefs about future outcomes coincide with the model's true probabilities.[1] A natural motivation for this requirement is that competitive pressures in economic and strategic interactions create strong incentives to avoid systematic forecast errors.

Requiring beliefs to be free of systematic errors implicitly assumes that there is a criterion or test to verify that they are consistent with observed outcomes. This paper formalizes this requirement and examines its implications for the sort of equilibria one might expect to arise. Roughly, an equilibrium is *testable* if every non-equilibrium belief can be rejected with positive probability by a test that compares that belief with observed outcomes. Testability is motivated by, and is a stylized representation of statistical tests that outside observers, such as an econometrician, might use to take theoretical models to data. Our focus, however, is on the economic and strategic implications of testability rather than the empirics of testing per se.

We first take the perspective of an outside observer who models an infinite-horizon economy as a stochastic process P. This process describes the assumed evolution of exogenous variables, as well as the optimizing behavior of agents and the resulting equilibrium outcomes. Leading equilibrium concepts assume that agents' beliefs agree with the true process P. We start by taking the perspective of an outside observer who directly assumes that the system is in equilibrium, without explicitly modeling the agents' optimization. Our central assumption is that P is stationary. Stationarity is a natural assumption, widely used in dynamic stochastic models for its tractability and conceptual appeal. It is also necessary for the application of many standard empirical methods. We illustrate the role and limitations of stationarity in a number of important contexts, such as consumption-based asset pricing,

[1] This is also the case in a self-confirming equilibrium (Fudenberg and Levine (1993a)), although players there may have incorrect beliefs about the consequences of their actions off the equilibrium path. See Section 5 for discussion. Non-equilibrium concepts like rationalizability do no impose a connection between beliefs and the model's objective probabilities.

Markov perfect equilibria, and Bayesian-Nash equilibrium in Markovian environments.

Our first main result is that requiring an equilibrium to be testable is equivalent to any one of the following three properties. First, an equilibrium P is testable if and only if it precludes disagreement: any alternative belief Q that is consistent with observed outcomes with probability one must be equal to P. Second, an equilibrium is testable if and only if agents' beliefs preclude structural uncertainty about long-run fundamentals, in the sense that no realization of past observations can change their opinions about the model's long-run properties. Intuitively, the absence of structural uncertainty means that agents have nothing further to learn from data about structural parameters. Learning may still play a role, but only for predicting short-run outcomes. Third, we connect testability with the properties of moment conditions used in the econometrics of non-linear dynamic stochastic models. We show that an equilibrium is testable if and only if the empirical moment conditions identify the model's true parameters asymptotically.[2]

We interpret these findings as highlighting a connection (and potential tension) between compelling desiderata of dynamic stochastic models. Testing agents' equilibrium beliefs against actual observations simply reflects the view that the concept of equilibrium should have observable implications, and not just be taken as an article of faith. A non-testable equilibrium is one for which there exists alternative beliefs that cannot be rejected with positive probability, regardless of the amount of data. Testability is also naturally related to empirical estimation of equilibrium models, a link formalized by its equivalence to the consistency of empirical moment conditions. Given the wide use of moment conditions in empirical work, we view consistency as another desirable property a model should have.

Our main theorem says that testability precludes disagreement and structural uncertainty about long-run fundamentals, two properties that are increasingly important in modeling economic and financial phenomena. An im-

[2] We focus on the generalized method of moments (GMM) introduced by Hansen (1982) which estimates a model by replacing theoretical moment conditions with their empirical counterparts. GMM includes standard techniques, such as linear regression and two-stage least squares, as special cases. See Section 3.3 for a detailed formal discussion.

plication of agreement, in the form of common or concordant priors, is the no trade theorems pioneered by Milgrom and Stokey (1982). No trade theorems are widely regarded to be inconsistent with observed investors' behavior and volume of trade.[3] A natural response to this conflict between empirical evidence and theoretical predictions is to weaken the common prior assumption so agents can disagree on how to interpret information.[4] Relatedly, structural uncertainty about long-run fundamentals and its gradual resolution through learning are increasingly viewed as relevant to understanding a number of empirical puzzles. In their survey of the subject, Pástor and Veronesi (2009) suggest that "[m]any facts that appear baffling at first sight seem less puzzling once we recognize that parameters are uncertain and subject to learning."[5] Our theorem indicates a potential tension between disagreement and structural-uncertainty, on the one hand, and the usual conceptual justification for equilibrium and its empirical estimation, on the other.

So far, the stochastic process P in the theorem represents an outside observer's model of a dynamic economic system. To interpret this process as an equilibrium requires that agents' decisions are optimal given their beliefs. A natural question is whether agents in a stationary environment can be uncertain about its long-run structural parameters. Proposition 4.1 shows, under additional assumptions about beliefs and payoffs, that structural uncertainty tends to lead to non-stationary behavior. Stated differently, a stationary equilibrium model must be testable and is thus inconsistent with agents' disagreement or structural uncertainty. Intuitively, if an agent's optimal behavior depends on his beliefs about structural parameters, then this behavior will change as this structural uncertainty is resolved through learning. This

[3] Kandel and Pearson (1995) is an early empirical study suggesting that agents disagree on the interpretation of information. Hong and Stein (2007) survey the literature on the trade volume puzzle.

[4] The case for the possible role of disagreement in explaining this and other puzzles is the subject of a large and growing literature. See the Hong and Stein (2007) survey. Ross (1989), for instance, notes that: "It seems clear that the only way to explain the volume of trade is with a model that is at one and the same time appealingly rational and yet permits divergent and changing opinions in a fashion that is other than ad hoc."

[5] Lewellen and Shanken (2002), Brav and Heaton (2002), and Weitzman (2007) are early examples of works that introduce learning about fundamentals to explain equity premia, risk-free rates, excess volatility, and predictability of returns. See Pástor and Veronesi's survey for a comprehensive review of the literature.

result is in the spirit of Nachbar (2005)'s theorem on the tension between learning and optimization in repeated games. We discuss Nachbar's work in Section 5.

Statistical testing is a natural bridge between empirical methods and theoretical solution concepts.Although leading equilibrium concepts require agents' beliefs to agree with the model's true probabilities, models differ in how this requirement is implemented. In a Bayesian-Nash equilibrium (with a common prior), agents are not required to know the true parameters of the model, only the probabilities used by Nature to select these parameters. A more demanding approach is the rational expectations equilibrium concept where beliefs are assumed to coincide with the objective empirical frequencies implied by the model. Beliefs in a rational expectations equilibrium are testable, and thus preclude disagreement or structural uncertainty. In a Bayesian-Nash equilibrium, on the other hand, beliefs may accommodate structural uncertainty and disagreement. The ideal is to "have it both ways," *i.e.,* maintain the flexibility of the Bayesian-Nash model without sacrificing testability and empirical identification of rational expectations. Our results point out to difficulties in reaching such compromise.

The paper proceeds as follows. In Section 2 we introduce the central concept of testability. We discuss three settings in asset pricing and game theory to motivate the main concepts and make connections to the literature. Section 3 introduces the three main properties an equilibrium model might have: disagreement, structural uncertainty, and the consistency of empirical moment conditions. Section 4 states the main results and discusses some extensions and qualifications. Finally, Section 5 provides additional connections to the literature.

2 Testing Beliefs

2.1 Preliminaries

We consider an infinite horizon model, with time periods denoted $n = 0, 1, \ldots$. In period n, an outcome s_n in a finite set S is realized.[6] Let $H = S \times S \times \cdots$ denote the set of infinite histories. The set of histories of length n is denoted H_n, and the (finite) algebra of events generated by these histories is $\mathcal{H}_n$. Stochastic processes are probability distributions on the set of infinite histories H, with the σ-algebra $\mathcal{H}$ generated by the finite-horizon events $\cup_{n=1}^{\infty} \mathcal{H}_n$.

A probability distribution P on H will stand for either the true (exogenous or equilibrium) process, or for an agent's belief about that process. Which of these interpretations is being referred to will be clear from the context. Our focus is on *stationary distributions*. Recall that a probability distribution P is stationary if, for every k, its marginal distribution on the algebra $\mathcal{H}_l^{l+k}$ generated by coordinates $l, \ldots, l+k$ does not depend on l. See, for example, Stokey and Lucas (1989). Let $\mathcal{P}$ denote the set of stationary distributions. We discuss non-stationarity in Section 4.2.

2.2 Tests

Equilibrium concepts, such as Nash equilibrium and rational expectation equilibrium, require agents' beliefs to coincide with the true equilibrium process. Consider an outside observer (a modeler, or an econometrician) who assumes that agents hold equilibrium beliefs and observes a realization of the process. For this equilibrium assumption to be refutable, it is natural to require that there is an objective criterion, or mechanism to determine whether agents' beliefs are indeed consistent with observed outcomes. We formulate this property in terms of statistical tests. In the next section we characterize some of the implications of testability in terms of what equilibria must look like.

Fix a stationary stochastic process P which we interpret as the "equilib-

[6] The finiteness of S is convenient to avoid inessential technical complications.

rium" process. This process describes the evolution of the model's exogenous variables (*e.g.,* income and technology shocks) as well as endogenous variables (*e.g.,* strategies, consumption levels, or asset pricing). A statistical *test* is any function

$$T : \mathcal{P} \times H \to \{0, 1\}$$

that takes as input a distribution P and a history h and returns a yes/no answer. Interpret an outcome $T(P, h) = 1$ to mean "history h is consistent with the process P;" otherwise, h is inconsistent with P. The set:

$$T_P \equiv \{h : T(P, h) = 1\}$$

is the set of all observations consistent with P. We may interpret T_P as the empirical predictions of P relative to the test T: if P is correct, then the observed sequence of outcomes must be in T_P.

In this paper we assume that the outcome of a test can depend on the entire infinite realization of the process. Real-world statistical tests can depend on only a finite (and usually small) number of observations. Nevertheless, asymptotic tests are a useful idealization to clarify the statistical and theoretical properties of finite tests. Asymptotic testing makes testability an easier hurdle for an equilibrium theory to overcome: an equilibrium theory that is not testable under any asymptotic test is certainly not testable using tests that depend on a fixed finite number of observations.[7] The asymptotic testing formalism is also simpler and, in some cases, results in stronger theorems.

The *Type I error* of a test T at a process P is the number $1 - P(T_P)$ representing the probability that the test rejects P when it is true. Statistical models usually require tests with low Type I error. In our asymptotic testing setting, it is natural to consider tests that are *free of Type I error*, *i.e.,* ones for which $P(T_P) = 1$ for every P. This makes for shaper statements and interpretation of the results. For tests based on finite observations, small but positive Type I errors are more natural. We discuss finite tests in Section 4.4.

[7] A finite test is a test that is measurable with respect $\mathcal{H}_n$ for some finite n. The set of finite tests is obviously a subset of the set of all tests. See Section 4.4 for further discussion.

Definition 1. *A stationary distribution P is testable if for every stationary $Q \neq P$ there is a test $T^\star$ such that*

1. *$T^\star$ is free of Type I error;*

2. *$Q\big(T^\star_P\big) < 1$.*

We can interpret the definition in two ways. First, think of P as the theory held by an outside observer, such as a theorist or an econometrician, of the data generating process implied by an economic model. The theory P is testable if any other different, and therefore incorrect, theory Q can be proven wrong with positive probability under P. The alternative theory could represent an agent's subjective belief about his environment or an outside observer's alternative model. A failure of this requirement means that there is a distinct theory $Q \neq P$ that cannot be distinguished from P under *any* Type I error free test. In statistical language, testability requires that for any such Q, there is a test $T^\star$ for Q that has *some* power against P. That is, the Type II error of this test is not 100%.

We may alternatively interpret P as the beliefs of an agent within the model. Equilibrium requires agents to have unquestioning faith in their equilibrium beliefs. Such agents have no use for tests that check whether their beliefs are right or wrong (since they are convinced they know the truth). On the other hand, agents less certain about their knowledge may view the testability of their beliefs as a desirable criterion.

Our notion of testability is meant to shed light on the implications of equilibrium rather than as a guide for designing practical testing procedures. Actual empirical tests must tolerate some Type I error while we require our asymptotic tests to be free of such errors. On the other hand, testability is an easier burden for a theory to satisfy in our setting since we do not restrict tests to use finite data or to have high power. In summary, testability retains some essential features of real world empirical tests, but abstracts from others in order to focus on theoretical questions about equilibrium beliefs.

2.3 The Role of Stationarity

Our interest in the stationarity of the process P stems from two sources. First, stationarity is an important assumption in economic models and in their empirical estimation.

Second, stationarity enables us to identify a non-trivial set of testable beliefs. To see this, let P be the distribution of a sequence of i.i.d. fair coin tosses. We will see later that the stationary i.i.d. distribution P is testable according to Definition 1. Consider now the non-stationary alternative belief Q under which the first coin is biased $(1/3, 2/3)$ and the other coins are fair (and thus coincides with P). Any Type I error-free test at P must pass Q with probability 1. In fact, if we extend Definition 1 by removing the stationarity assumption, then the only distributions that will be seen as testable are Dirac's atomic measures on the realizations, i.e. those that give rise to no uncertainty at all ![8]

2.4 Examples

2.4.1 Asset Pricing

As a first illustration, consider the canonical consumption-based asset pricing model. The primitive is a stochastic process $\{c_n\}$, interpreted as the consumption of a representative agent with a time separable utility $\sum_{t=1}^{\infty} \delta^n u(c_n)$. Here, u is a (differentiable) period utility and $\delta \in [0, 1)$ is the discount factor.[9] Assume a finite outcome space rich enough to model the consumption process and any collection of asset returns or other variables of interest.[10]

[8] Indeed, Dirac's atomic measures on realizations are the extreme point of $\Delta(S^{\mathbb{N}})$. Therefore, if P is not Dirac's atomic measure then it can be written as $P = \lambda P' + (1-\lambda)P''$ for some beliefs P, P'' such that $P' \neq P''$ and some $0 < \lambda < 1$. Therefore, every Type I error-free test must have the property that $P'(T_P) = P''(T_P) = 1$. If $Q = \mu P' + (1-\mu)P''$ for some $0 < \mu < 1$ and $\mu \neq \lambda$ then $Q \neq P$ but $Q(T_P) = \mu P'(T_P) + (1-\mu)P''(T_P) = 1$ for every such test.

[9] The consumption process can be either exogenous, as in Lucas (1978)-style endowment economies, or the outcome of an unmodeled optimization and is taken as given. See Cochrane (2005) for a textbook exposition and discussion of these issues.

[10] We assumed, for technical convenience, that the outcome space is finite. In asset pricing theory, it is more natural to consider continuous outcome spaces. We can either extend the model to continuous spaces, or assume S to be some appropriately fine grid.

The marginal rate of substitution between consumption in periods $n+1$ and n is the random variable

$$m_n \equiv \delta \, \frac{u'(c_{n+1})}{u'(c_n)}, \tag{1}$$

also known as the stochastic discount factor.

We assume that m_n is stationary with distribution P.[11] The process $\{m_n\}$ determines the equilibrium rates of returns of all assets. Specifically, we consider a collection of assets, each represented by a stochastic process $\{R_n\}$ giving the gross rate of return at time t of a dollar invested in that asset at time $t-1$. Equilibrium requires that the return process of any asset satisfy the standard Euler equation $E_P\big[m_{n+1}\, R_{n+1} \,|\, h^n\big] = 1$, for all positive probability histories h^n.

The model postulates P as the true process governing how the model's variables evolve. Our formal notion of testability attempts to give an operational meaning to the statement "P is the true process." Suppose that the modeler, or the representative agent, is confronted by an arbitrageur with an alternative theory $Q \neq P$ of the stock market. Testability captures the intuition that the modeler's theory P must have some observable implications that can be used to "prove Q wrong." Since even a wrong theory can give, by accident, more accurate predictions than the true theory, the statement "proving Q wrong" must be qualified.[12] The notion of testability accomplishes this by requiring the modeler (or the agent) to produce, for any competing theory $Q \neq P$, an objective criterion that does not fail P, but that has some power against Q.

The process P in this example describes the evolution of a system in equilibrium. It draws no distinction between exogenous variables, such as the discount factor process $\{m_n\}$, and endogenous market-determined asset

[11] Note that this does not require that $\{c_n\}$ is stationary. The stationarity of $\{m_n\}$ can be derived from primitive assumptions about the consumption process and the utility function. For example, in a lognormal consumption growth model, while consumption is non-stationary, the discount factor under power utility is stationary. Here we assume stationary directly.

[12] A theory that predicts a fair coin to land *Heads* in a given toss is more accurate than the true theory (that the coin is fair) about half of the time!

returns that depend on agents' beliefs.[13]

An agent who believes in an alternative theory $Q \neq P$ that cannot be rejected by any test under P cannot be dismissed on objective grounds as wrong or irrational. Of course, injecting agents with belief Q in the model will likely lead to observably distinct behavior that is inconsistent with the original process P. The impact of introducing non-equilibrium beliefs depends on the specifics of the model, and is not something we examine here.

2.4.2 Markov Perfect Equilibria

Our next example is a dynamic game with a Markovian state variable. There is a finite set of states, X, and players, I. The set of action profiles is $A = A^1 \times \cdots \times A^I$, where A^i is player i's set of actions. In accordance with our notation, we write $S = X \times A$ to denote the outcome of the game in any given round.

The horizon is infinite, with time periods denoted $n = 0, 1, \ldots$. The distribution of period $n+1$'s state given period n's state x_n and action profile a_n is determined by a commonly known time-invariant transition function $\pi : X \times A \to \Delta(X)$. In each period $n \geq 1$, players choose action profile $a_n \in A$ then a state x_n is realized from the distribution $\pi(x_{n-1}, a_n)$. We assume for simplicity that $\pi(x, a)$ has full support for every outcome (x, a). The initial outcome (x_0, a_0) is chosen according to some distribution p on S. A strategy for player i specifies his mixed action at every stage as a function of past states and action profiles. A strategy is Markovian if players' actions depend only on the current state. Thus, a Markovian strategy for player i is of the form $\sigma^i : X \to \Delta(A)$.

A Markovian strategy profile σ and a transition function π induces a unique stationary Markovian distribution P over $\Delta(S^{\mathbb{N}})$ that describes the steady state of the play. The steady state distribution does not depend on the initial choice of (x_0, a_0).

Agent i maximizes discounted expected utility, with period utility $u^i(x, a)$

[13] What is considered endogenous and exogenous is, obviously, model-dependent. We follow the common practice in asset pricing theory by assuming that consumption is exogenous, although a richer model would treat it as endogenous.

and discount factor δ. A profile σ is a Markov perfect equilibrium (MPE) if for each player i, σ^i is optimal against the profile of strategies of his opponents, σ^{-i}.[14] The MPE model is widely used in applied work to model, among other things, industry dynamics following the seminal work of Ericson and Pakes (1995). An attractive feature of this model is that it lends itself naturally to empirical implementation, an issue we shall return to below.[15]

We will show later that an MPE distribution P on S^∞ is testable: given an alternative stationary theory $Q \neq P$ of how the game evolves, we can construct a Type I error free statistical test such that Q is rejected with positive probability under P. We cannot in general conclude that a false theory Q will necessarily be rejected with *high* probability. For example, if P assigns high probability $1-\epsilon$ to Q and probability ϵ to some $Q' \neq Q$, then Q is accepted most of the time.

A variant of the above model is one where we assume that an analyst observes the players' actions and some function $f : X \to X'$ of the state.[16] For example, in empirical models it is common to assume that players condition their actions on disturbances that are unobservable by the analyst. Every MPE distribution induces a belief of the analyst over $(A \times X')^{\mathbb{N}}$. The analyst' belief is typically not a Markov process but is still stationary. We will show later that it is also testable.

3 Disagreement and Structural Uncertainty

We introduce three properties a dynamic economic model may have: the possibility of disagreement, absence of structural uncertainty, and the feasibility of empirical estimation. The main theorem of this paper, Section 4, will show that all three properties are equivalent to testability.

[14] As usual, all deviations are allowed, not just to Markovian ones.

[15] In empirical models, strategies usually depend also on unobservable disturbances.

[16] In particular, if f is constant then the analyst only see the players actions but not the state of nature.

3.1 Disagreement

A standard assumption in economic models is that agents share a common prior (Aumann (1976 and 1978)) and, as a result, a common interpretation of information. This assumption has been questioned on a number of grounds. First, that agents may disagree about how to interpret information seems both intuitive and consistent with the basic axioms of rationality. Second, the absence of disagreement leads to paradoxical theoretical conclusions, such as the no-trade theorems (Milgrom and Stokey (1982)), that are difficult to reconcile with reality. Third, as noted in the introduction, there is large empirical evidence on volume of trade that seems difficult to reconcile with common interpretation of information.[17]

A natural way to introduce disagreement is to assume that agents have different prior beliefs about the underlying uncertainty. Heterogenous prior models are used to generate realistic trade volumes and to account for other asset pricing anomalies. See Hong and Stein (2007) and Pástor and Veronesi (2009) for surveys.

An agent's uncertainty about the true process P is formally represented by a distribution $\mathcal{Q} \in \Delta(\Delta(H))$ on the set of stochastic processes. To any belief $\mathcal{Q}$ corresponds a process $Q \in \Delta(H)$ with identical distribution on sample paths.[18] Since there is no difference between the belief $\mathcal{Q}$ and the process Q, we use the latter to describe beliefs.

Definition 2. *Two beliefs* Q_1, Q_2 *are* compatible *if for every event* B, $Q_1(B) = 1$ *if and only if* $Q_2(B) = 1$.

Compatibility of beliefs is the property known as mutual absolute continuity, a condition that appears in the seminal paper by Blackwell and Dubins (1962) and introduced to the study learning in repeated games by Kalai

[17] Whether the motives for trading under common priors (unanticipated liquidity and rebalancing needs) are sufficient to explain observed volumes is an important empirical question which is beyond the scope of this paper. In their survey, Hong and Stein (2007) suggest that these motives are far too small to explain the 51 trillion dollar volume of trade in equity in 2005, say. See their paper for additional examples.

[18] Define $Q(A) \equiv \int_{\Delta(H)} P(A)\mathcal{Q}(\mathrm{d}P)$. Note that since the set of stationary distributions is closed and convex, if a belief $\mathcal{Q}$ is concentrated on stationary distributions, then its reduction Q is also stationary.

and Lehrer (1993). Blackwell and Dubins showed that compatible beliefs "merge," in the sense of generating the same predictions in the limit as data increases.[19]

Compatibility is a common requirement in heterogenous belief models. It has the interpretation that, while beliefs may initially disagree, their disagreement must eventually vanish. Incompatible beliefs must continue to disagree even in the limit, with infinite data.

Definition 3. *A stationary belief P* precludes disagreement *if for every stationary belief Q compatible with P, we have $Q = P$.*

3.2 Structural Uncertainty

An intuitive distinction is often made between situations where agents have "structural uncertainty about long-run fundamentals" and ones where there is no such uncertainty.

How should uncertainty about fundamentals be defined? Intuitively, an agent who understands the structure of his environment believes he learned all that can be learned. Uncertainty about fundamentals suggests an agent who believes (or behaves as if) he does not know all the long-run properties of the process and that additional learning about them is possible. As a simple example, suppose that there are just two candidate processes, P_1 and P_2, both of which is i.i.d. For an agent who knows that the true process is P_1, say, no amount of new observations can change his beliefs about the probability of future outcomes. In this case, there is no structural uncertainty, and new information has no predictive value. On the other hand, it is intuitive to think of an agent who believes that the true process is P_1 with probability α and P_2 with probability $1 - \alpha$ as someone who is uncertain about the structure, and for whom additional observations are potentially valuable.

This simple intuition breaks down in more general settings. Suppose, that P_1 and P_2 are, instead, two non-i.i.d. Markov processes. If $\alpha \in (0, 1)$, the agent is uncertain about the true process, and information is as valuable as before. The problem is that information is also valuable even if the agent

[19] See Blackwell and Dubins (1962) or Kalai and Lehrer (1993) for a formal statement of this well-known result.

knew that the true process is P_1, say. The reason is that with a Markovian process, the outcome of one period is informative about outcomes of future periods even when the process is known.

Intuitively, the idea of "structural uncertainty about long-run fundamentals" suggests lack of knowledge about the long-run properties of the the process, not outcomes in the 'near' future. In the Markovian example above, for an agent who knows P_1 additional observations are informative about the near future, but have no impact on his beliefs about the probability of outcomes in the distant future.

To make this distinction formal, fix a function $f : S^k \to \mathcal{R}^m$, with finite k, and define $f_n \equiv f(s_{n-k+1}, \ldots, s_n), n = k, k+1, \ldots$. We shall think of $f_k, f_{k+1}, \ldots$ as a payoff stream where, in each period n, a payment f_n is made that depends on the realization of the past k outcomes according to the (stationary) formula f.[20] For example, f_k may represent the history-dependent dividend paid by an asset in period k and $f_k, f_{k+1}, \ldots$ is the stream of such payments. Another example is that f takes values in a finite set of actions $\{a_1, \ldots, a_m\}$ in some vector space and represents the actions of an opponent who uses some Markovian strategy.

For a sequence $(x_k, x_{k+1}, \ldots)$ of elements $\mathcal{R}^m$, define the limiting average:

$$V(x_k, x_{k+1}, \ldots) = \lim_{n\to\infty} \frac{1}{n} \sum_{i=k}^{n} x_i,$$

whenever the limit exists. We shall abuse notation and refer to $f_k, f_{k+1}, \ldots$ as f. Let $V(f)$ be the random variable whose value $V(f)(h)$ at a history h is the limiting average of $f_k, f_{k+1}, \ldots$ generated by f at h.[21]

In the asset pricing example (Section 2.4.1), $V(f)(h)$ represents the long-run average return of an asset along a history h. In the dynamic game example (Section 2.4.2), $V(f)$ takes values in the m unit simplex Δ^m, with $V(f)(h)$ representing the limiting empirical frequency of actions played.

Definition 4. P displays no structural uncertainty *if for every bounded func-*

[20] To allow history dependence of the payoff stream, payments start in period k.

[21] This random variable is well-defined with probability 1 for any stationary P and function f (Proposition A.2).

tion f and finite history h^{t-1} with $P(h^{t-1}) > 0$

$$E_{P(\cdot\,|\,h^{t-1})} V(f) = E_P V(f).$$

We could replace the limit-of-averages criterion in the definition with discounted payoffs, where the discount increases to 1. What is important is that information relevant for short-run predictions is irrelevant for the long run, thus capturing the role of information that changes beliefs about the long-run properties of the process. The limit-of-average criterion simplifies the statement of the property and the results.

3.3 Empirical Identification

In empirical studies, the implications of an economic model are usually summarized by *moment conditions.* Formally,

Definition 5. *A* moment condition *is a bounded continuous function*[22]

$$f : Z \times S^k \to \mathcal{R}^q \tag{2}$$

where k and q are positive integers, and $Z \subset \mathcal{R}^m$ is a compact set.

We say that f identifies *P if there is a unique $\bar{z} = \bar{z}(P, f)$ such that*

$$E_P f\,(\bar{z}, s_1, \ldots, s_k) = 0.$$

Intuitively, a moment condition represents a finite-horizon trait, or feature of the underlying model P. As an example, assume that $S = \{0, 1\}$, $Z = [0, 1]$, and $k = q = 1$. Let $f : Z \times S \to \mathcal{R}$ be the moment condition $f(z, s) \equiv z - s$. For any stationary distribution P, we have $E_P\, f(z, s) = 0$ precisely at $\bar{z}_P \equiv P(s_k = 1)$. The number $\bar{z}_P$ is a reduced form parameter that reflects one aspect of the distribution but reveals nothing, for instance, about any intertemporal correlation P might have. This moment condition identifies the marginal of P on any coordinate. Suppose now that P_1 is i.i.d. while P_2 is a process whose realization is $0, 1, 0, \ldots$ or $1, 0, 1, \ldots$, with probability 0.5 each. Then, $E_{P_1} f(z, s) = E_{P_2} f(z, s) = 0$ if and only if $z = 1/2$. These

[22] Continuity has bite only with respect to Z, since S is assumed to be discrete.

two stationary processes are indistinguishable from the perspective of the moment condition f, even though one is perfectly correlated over time, while the other is independent.

Moment conditions underlie many statistical techniques used in econometric practice, such as the generalized method of moments, introduced in Hansen (1982), which generalizes many standard techniques. The idea is to estimate the true $\bar{z}_P$ by the element $\hat{z} \in Z$ that minimizes the empirical estimate of $E_P f$. Formally, given f and length of data $n \geq k$, define the empirical average

$$F_n(z, h) = \frac{1}{n-k} \sum_{t=k+1}^{n} f\left(z, s_{t-k}, \ldots, s_t\right),$$

viewed as $q \times 1$ column vector, with transpose denoted $F_n^\top$. The generalized moment estimator is the random variable

$$\hat{z}_n(h) \in \underset{z \in Z}{\operatorname{argmin}} \; F_n(z, h)^\top F_n(z, h).^{23} \tag{3}$$

Definition 6. *A stationary distribution P can be* empirically identified *if $\hat{z}_n \xrightarrow{P} \bar{z}$ for every moment condition that identifies P.*[24]

The property that P can be empirically identified from data means that all of the implications of P can be recovered, via moment conditions, from observing the evolution of the process. Alternatively, if P' cannot be empirically identified, then there must be an implication of P', in the sense that $E_{P'} f\left(\bar{z}, s_1, \ldots, s_k\right) = 0$, such that $\bar{z}$ cannot be recovered, even asymptotically as data grows without bound.

[23] The generalized method of moments estimator looks instead at the quadratic form $F(z,h)^\top Q\, F(z,h)$, where Q is an appropriately chosen $q \times q$ matrix. Hansen (1982) shows how to select Q optimally. We focus here on the baseline case where Q is the identity for simplicity.

[24] That is, $P\{h : |\hat{z}_n(h) - \bar{z}| > \alpha\} \to 0$ for every α as $n \to \infty$. Hansen shows that consistency obtains for any ergodic P. For this argument, it is enough to consider k=1. But it is not in general possible to prove the converse without using higher k's.

4 Characterization of Testable Beliefs

4.1 Main Result

Our first main result relates the four concepts introduced earlier:

Theorem 1. *For any stationary process P, the following four statements are equivalent:*

1. *P is testable;*
2. *P precludes disagreement;*
3. *P precludes structural uncertainty;*
4. *P can be empirically identified.*

Theorem 1 divides stationary models into testable and untestable ones. Disagreement and structural uncertainty, two properties with conceptual appeal and strong empirical support, conflict with testability and moments-based empirical methods. As we note in Section 5.4 below, the theorem does not say that disagreement and structural uncertainty have no observable implications on behavior. It only says that when beliefs disagree, it is not possible to determine which belief is correct or to objectively test the properties of the equilibrium process. Long-run structural uncertainty will usually have implications on agents' decisions that can be measured empirically.

Many familiar models in the literature correspond to the testable case. A first example is i.i.d. processes over outcomes. A second example is an irreducible Markov chain. Recall that a Markov transition function is given by an $S \times S$ matrix $(\pi(t|s))_{s,t \in S}$ such that $\sum_t \pi(t|s) = 1$ for every s. We say that π is *irreducible* if for every states $s, t \in S$ there exists some n such that $\pi^n(t|s) > 0$ where π^n is the n-th power of π. Every Markov transition π and every $p \in \Delta(S)$ induces a distribution P over $S^{\mathbb{N}}$ such that

$$P(s_0, s_1, \ldots, s_n) = p(s_0) \cdot \pi(s_1|s_0) \cdot \cdots \cdot \pi(s_n|s_{n-1}).$$

If π is irreducible then there exists a unique p that makes P stationary. Call any such distribution P a *(stationary) irreducible Markov process*. This definition extends naturally to processes with some memory k.

As an example of a non-Markovian stationary process, consider a *hidden Markov process*. Such processes are given by some Markov process on a set S of *unobservable outcomes*, a set T, and a function $f : S \to T$ where $f(s)$ is the observation made given the (unobservable) outcome s. The distribution induced over $T^{\mathbb{N}}$ is stationary but not necessarily Markovian of finite memory.

The following facts will be useful in dealing with examples below:

Fact: the conditions of Theorem 1 are satisfied for every irreducible Markov process (with finite memory k), and every hidden Markov process where the underlying Markov process is irreducible.

We now revisit the earlier examples in light of the theorem and the above fact. The Markov perfect equilibrium model induces an irreducible Markov process of outcomes and is therefore testable by the Fact above. The Markov perfect equilibria with partially observed states is a hidden Markov model and the underlying Markov process over states is irreducible, and so testable.

4.2 Learning and Non-Stationarity

Theorem 1 assumes an equilibrium processes that is stationary. As discussed earlier, stationarity is an important property to consider given the role it plays in empirical and theoretical work. We also noted that there may be significant conceptual difficulties in defining testing when stationarity fails.

Here we discuss another subtle aspect of stationarity, namely its connection with learning. A process P describing a dynamic stochastic model in equilibrium must incorporate agents' learning about the equilibrium process itself. Start with an equilibrium process that is stationary but non-testable. The theorem implies that agents must be uncertain about some long-run structural aspects of the process. As agents observe the process unfold, their beliefs about these long-run fundamentals change as a result of learning. Changes in beliefs typically lead to changes in the actions agents take. But since a process in equilibrium must also describe how agents' behavior evolve over time, it seems likely that non-testable equilibrium processes will be non-stationary.

This informal argument suggests that non-testability tends to imply non-stationarity. We will show by example that this is not always the case. Connecting testability and stationarity requires assumptions about how agents respond to changes in their beliefs to rule out that behavior remains stationary even though beliefs change.

To clarify these and other issues, we consider another class of examples:

4.3 Examples with Learning

4.3.1 Passive Learning

Consider the MPE setting of Section 2.4.2 except that agents are uncertain about the transition π. Pakes and Ericson (1998) study a model along these lines with the goal of empirically testing for the implications of agents' passive learning on industry dynamics.

The formalism here closely follows that of the MPE setting. There is a finite set of outcomes $S = X \times A$, where X and A represent the states and action profiles, respectively. We continue to use the same notation to denote strategies and utilities. There are two departures from the MPE model. First, agents may be uncertain about the transition. Formally, agents have beliefs $\mu_i, i \in I$, with a finite and common support $\Pi = \{\pi_1, \ldots, \pi_L\}, L \geq 1$, representing their uncertainty about the transition. Second, learning is passive in the sense that agent's actions do not influence the evolution of the exogenous process. This rules out agents' active experimentation and keeps the theoretical and empirical analysis tractable. Formally, we require that each $\pi_l \in \Pi$ takes the form $\pi_l(x_n)$, and so depends only on x_n.

The evolution of the system is described by a stochastic process P on S^∞. Use P_X to denote its marginal on the exogenous variables X^∞ and let $P_X(\pi_l)$ be the stationary distribution on X^∞ induced by the transitions $\pi_l \in \Pi$. Agent i believes that the exogenous variables evolve according to the stationary distribution $P_X(\mu_i) = \mu_i(\pi_1)P_X(\pi_1) + \cdots + \mu_i(\pi_l)P_X(\pi_l)$ which no agent can influence.

We can apply Theorem 1 either to the entire process P or just to its marginal P_X. We distinguish three nested cases:

Case 1: $L = 1$, *so agents know the true transition* π.

Case 2: $L > 1$ *and agents share a common prior* ($\mu_i = \mu$, *for all* i).

Case 3: $L > 1$ *and agents have different priors* ($\mu_i \neq \mu$, *for all* i).

Case 1 is a special case of the MPE introduced earlier.[25] The marginal process P_X is Markovian, testable and displays no structural uncertainty. The same holds for the full equilibrium process P under an MPE.

Turning to Case 2, we model decisions using the concept of Bayesian-Nash equilibrium as solution concept. Fix one such equilibrium and let P denote the distribution it generates. In any such equilibrium, the marginal process P_X on exogenous variables is the stationary process $P_X(\mu)$. Applying Theorem 1 to P_X, it is easy to verify that it is non-testable.[26]

The fact that the marginal process P_X is stationary and non-testable makes it "likely" that the equilibrium process P as a whole is not stationary. Intuitively, when agents learn about the long-run fundamentals (the true π), their beliefs and decisions will change in a non-stationary fashion. The next theorem formalizes this intuition:

Proposition 4.1. *Suppose that agents have a common prior* μ *with support* $\{\pi_1, \ldots, \pi_L\}, L > 1$. *Then there exists action sets and utilities,* A^i, u^i, *for* $i = 1, \ldots, I$, *such as the resulting game admits no Markov equilibrium.*

When players know the Markov transition, their beliefs about the next state depends only on their current state. Their equilibrium behavior then depends only on the current state, and the play is that of a perfect Markov equilibrium. On the other hand, when there is uncertainty, the current state is no longer sufficient to determine agents' belief about the next outcome, and behavior may depend on the entire history of past observations.

Making this argument formal requires that changes in beliefs have an impact on actions. This is why the theorem is stated in terms of *some* specification of the payoff functions.[27]

---

[25] Special because we require passive learning

[26] We can verify this, using the theorem, by noting that agents are uncertain about the long-run fundamentals (*i.e.,* which π is the true process).

[27] One can define a more general equilibrium concept, stationary equilibrium, which is a strategy profile that renders the joint stochastic process of states and player actions stationary. Similar theorem like Theorem 4.1 applies for this equilibrium.

Uncertainty about the long-run fundamentals implies that early observations have a persistent effect on what is observed in the long run. This implication is exploited by Pakes and Ericson (1998) to construct a statistical test of whether observed industry dynamics is consistent with firms learning about their environment.

We finally note a connection between testability and disagreement. In Case 2, the process P_X is non-testable yet agents agree on it. As argued earlier, justifying agreement that μ is the true process is difficult when there can be no statistical test, based on observing the exogenous variables, that can disprove alternative beliefs $\mu' \neq \mu$. If one views testability as a (necessary) condition for agreement, then Case 2 seems artificial, and Case 3 where agents disagree is the more compelling alternative.

4.3.2 Bayesian Nash Equilibrium

Next we consider learning in a repeated game with incomplete information using the setup of Kalai and Lehrer (1993). For simplicity, assume there are just two players, each with a finite number of actions. A type of a player is his payoff matrix. Assume that there are m possible types for each player, and denote by $T^i = \{D^i_1, \ldots, D^i_m\}$ the set of types of player i. A type profile is realized from common prior distribution μ on $T^1 \times T^2$. After types are realized, each player observes his own type. The standard solution concept for this class of games is that of a Bayesian Nash equilibrium. Kalai and Lehrer (1993, Theorem 2.1) show, roughly, that conditional on (D^1, D^2) being the true type profile, the play of the game will eventually be close to that of a Nash equilibrium of the true game.

We consider an outside analyst who observes the actions of the players but not the realized types. Say that a type D of player i is *committed to action a* if playing a each period is a dominant strategy for player i in the repeated game when his payoff matrix is D.

We illustrate the relationship between testability, stationarity, and learning in three examples chosen for their tractability. In Example 1 types are completely correlated, so players have no uncertainty.

Example 1. *Assume that μ is uniform on $\{(D^1_j, D^2_j)\}_{j=1,\ldots,m}$. Then players* *know the type profile but an observer has uncertainty. Consider the Bayes Nash equilibrium according to which, when the type profile is* $\left(D_j^1, D_j^2\right)$ *players repeatedly play some Nash equilibrium of that game.*

The equilibrium profile in this example is stationary (in fact, exchangeable) but will typically not be testable. This is because there is uncertainty about the long-run distribution of play—unless the support of μ is concentrated on type profiles that choose the same actions. In our next example, players are uncertain about the game so they learn something about their opponent' type as they observe the evolution of play. However, the payoffs are such that learning does not affect the players behavior.

Example 2. *Assume that* μ *has full support over* $T^1 \times T^2$ *and that all types are committed (to different actions). In a Bayes-Nash equilibrium, each player always plays his commitment strategy.*

The equilibrium profile in Example 2 is, trivially, stationary, but, when some types of a player are committed to different actions, it is not testable.

In our final example, players have uncertainty and the example is such that learning must affect behavior. In this case, no Bayesian Nash equilibrium induces a stationary play.

Example 3. *Assume that:*

1. μ *has full support over* $T^1 \times T^2$.
2. *There exists types* $D, D' \in T^1$ *of player 1 which are committed to different actions* a, a'.
3. *There exists type* D'' *of player 2 such that the set* $B(a)$ *of best responses of player 2 under* D'' *to* a *and the* $B(a')$ *of best responses of player 2 under* D'' *to* a' *are disjoint.*

We have the following result:

Proposition 4.2. *A Bayes Nash equilibrium of the game described in Example 3 cannot be stationary.*

The proof is in the Appendix. Roughly speaking, this follows from the fact that in every equilibrium in the game with payoff matrices (D, D''), the players action profile is in $\{a\} \times B(a)$, and that in every equilibrium in the game with payoff matrices (D', D''), the players action profile is in $\{a'\} \times B(a')$. Kalai and Lehrer's theorem implies that players learn to play as in equilibrium of the game with the realized types, so that there is a positive probability that from some period onwards the action profile played is always in $\{a\} \times B(a)$ and a positive probability that from some period onwards the action profile played is always in $\{a'\} \times B(a')$. Because of stationarity, this should be the case already for the actions at period 0, and because these actions are independent given the types, it follows that at day 0 there is a positive probability that the action will be in $\{a\} \times B(a')$ which is a contradiction.

4.4 Finite Horizon Testing

The properties in Theorem 1 are formulated in terms of infinite outcome sequences. The definitions have natural analogues for finite horizons. Take, for instance, the concept of testability. Call a test T *finite* if there exists an integer n such that T depends only on the outcome of the first n periods.[28] Tests used in statistical practice are obviously finite. It is easy to show the following:

Claim. *If P is testable then for every stationary Q and every $\alpha > 0$ there exists a finite horizon test T^* such that:*

1. *The Type I error of T^* is smaller than α.*
2. $Q(T^*(P)) < 1 - \alpha$

Formulating our result in a finite horizon setting requires close attention to the amount of data available to the test, an issue that did not arise in the asymptotic testing context. The length of the horizon n needed depends on the complexity of the intertemporal structure of the processes being tested. To illustrate, consider the following simple example:

[28] Formally, T is measurable with respect to $\mathcal{H}_n$.

Example 4. *Let* $S = \{0, 1\}$ *and let* P *be the steady state distribution of a Markov process with transition*

- $\pi(0|0) = \pi(1|1) = 1 - \epsilon$ *and*
- $\pi(1|0) = \pi(0|1) = \epsilon$

for a small $\epsilon < \frac{1}{n}$*. Since* P *is an irreducible Markov chain, it follows from Fact 1 above that* P *is testable. Let* Q *be the process that puts unit mass on the sequence* $(0, 0, \dots)$*. Then for every test* T *of horizon* n *and Type I error smaller than* $\alpha = 0.10$*, we have* $Q(T^*(P)) = 1$*.*

The problem in this example is that the process P is highly persistent. If we are given a very short horizon, this persistence can easily confound P with the constant process Q.

Testability with such short horizon has little bite. This should not be surprising: no statistical technique based on such limited data can reject the process Q with probability more than 0.5. Relatedly, the GMM estimator of the one dimensional marginal of P will fail to converge to its correct over such small horizon.

In summary, finite data is a limiting factor to statistical methods in general, and not specifically our model. To apply the notion of testability in finite horizon, either the class processes has to be restricted, or the horizon extended (or both). In our asymptotic testing framework, we impose no restrictions beyond stationarity, but require the horizon to be infinite. Alternatively, we could limit n to be finite, but then we must restrict the class of processes to ensure fast enough rate of convergence to make statistical testing feasible.[29] We find the asymptotic approach to give a clearer picture of the implications of testability.

[29] A natural way to proceed is to require the process to be mixing with a uniform mixing rate. This, indeed is the common assumption made in practice. Vector autoregressive processes, for example, belong to this class. See also our discussion of Pakes and Ericson (1998) below.

5 Discussion and Related Literature

In addition to the literature discussed earlier, this paper is also related to other important strands of literature.

5.1 Testing Strategic Experts

A recent literature examines the problem of testing an expert whose goal is to pass the test regardless of his knowledge of the underlying process. The expert is strategic in the sense that he selects a randomized forecast in order to manipulate the test. A number of papers, including Sandroni (2003), Olszewski and Sandroni (2008), Shmaya (2008), provide general conditions under which a strategic expert can manipulate any test using randomized predictions. Positive results establishing the existence of non-manipulable tests appear in Dekel and Feinberg (2006), Olszewski and Sandroni (2009), Al-Najjar, Sandroni, Smorodinsky, and Weinstein (2010). See Olszewski (2012) for a survey.

In our paper, agents do not strategically randomize their forecasts to manipulate the test. Rather, they hold subjective beliefs about their environment, and choose optimally given these beliefs. Our focus are the implications of requiring that erroneous (non-equilibrium) beliefs can be rejected by data under the true model.

5.2 Bayesian Learning in Games

Sections 2.4.2 and 4.3.2 illustrate the connection with the literature on learning in games. The connection we make between testability, long-run uncertainty, and stationarity is related to Nachbar (2005)'s results for learning in games. In his motivating example (pp. 459-460) two players play an infinitely repeated game, where the stage game has just two actions A, B for each player. Player $i = 1, 2$, believes player $j \neq i$ follows an i.i.d. strategy where A is played with probability q_i each period. Player i's prior belief about the strategy of the opponent is a probability distribution μ_i over $[0, 1]$. In our terminology, each player faces a stationary environment described by his belief about the strategy of his opponent. In this special setting, there is long-run uncertainty if and only if μ does not put unit mass on one value of q_i. Proposition 4.1 shows that when players update their beliefs, their behavior will not in general be i.i.d., and is therefore not consistent with the players's priors about their opponents' strategies.

Nachbar obtains a general result for infinitely repeated game strategies where the assumption of stationarity is not natural. In this paper, we take the perspective of a single decision maker and our arguments rely heavily on stationarity. Our concepts of testability and empirical identification also have no counterpart in Nachbar's work.

5.3 Self-confirming Equilibrium

The process P in this paper describes the evolution of all variables in the model, including agents' decisions. Beliefs are in equilibrium if their predictions about future outcomes coincide with P. In our analysis, the process P puts no restrictions on agents' predictions at counter-factual events, *i.e.,* events that have zero probability under P. This is important in games, where a player's decisions depend on his beliefs about what his opponents will do at all events, including those not observed in equilibrium. Although the process P describes all what an outside observer can hope to see, it is insufficient to explain what players do. This distinction is captured by the notion of self-confirming equilibrium introduced by Fudenberg and Levine (1993a) where players are assumed to know the true distribution on observed outcomes, but may have incorrect off-path conjectures.

Fudenberg and Levine (1993b) motivate the concept of self-confirming equilibrium as the result of a steady-state learning process. Related to our discussion of disagreement, Dekel, Fudenberg, and Levine (2004) argue that this steady-state interpretation of self-confirming equilibrium is difficult to reconcile with heterogenous beliefs about the state of nature.

5.4 Testing Beliefs vs. Testing Behavior

Beliefs about a stationary exogenous process may be non-testable yet have important observable consequences on behavior. Consider the context of asset pricing: although the beliefs of agents with structural uncertainty are untestable, their behavior will be influenced by their learning, and will typically have observable implications. Lewellen and Shanken (2002) provide an asset pricing model where investors' learning returns might appear predictable and excessively volatile even though prices react efficiently to information. Lewellen and Shanken also note that learning may confound empirical testing of asset pricing models. Another example, in a Markovian model of industry dynamics, Pakes and Ericson (1998) develop econometric tests to detect firms' learning based on whether the first few outcomes have a persistent impact on their posterior beliefs. Evidence for learning may be found by looking for non-stationary behavior (the fact that early observations have persistent correlation with the long-run evolution of the model).

A Proofs

A.1 Mathematical Preliminaries

A function $f : S^{\mathbb{N}} \to \mathbb{R}$ is finitely-based if there is an integer k such that for any history h, f depends only on the first k coordinates of h (that is, f is $\mathcal{H}^k$-measurable). Abusing notation, we express finitely-based functions in the form $f : S^k \to \mathbb{R}$. We use the following lemma:

Lemma A.1. *For any pair of distributions $P, Q \in \Delta(H), P \neq Q$ there exists a finitely-based function f such that $E_P f \neq E_Q f$.*

Given a stationary P and any (Borel)function $g : S^{\mathbb{N}} \to \mathcal{R}$, the *ergodic theorem* states that the limit

$$\tilde{g} = \lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} g\left(s_k, s_{k+1}, \dots\right) \tag{4}$$

exists P-a.s. and that $E_P\tilde{g} = E_P g$. If the limit $\tilde{g}$ is constant (i.e., $\tilde{g} = E_P g$ almost surely) for every g then P is called *ergodic*. It is sufficient for ergodicity that the limit be constant for every finitely-based function g.

The following proposition states the implication of the ergodic theorem for finitely based functions. In the terminology of Section 3, the (random) limiting average payoff of a finitely-based f is well-defined:

Proposition A.2. *Let P be stationary and $f : S^k \to \mathbb{R}$ be finitely-based. Then*

$$V(f)(h) = \lim_{n\to\infty} \frac{1}{n} \sum_{i=k}^{k+n-1} f(s_{i-k}, \dots, s_{i-1})$$

exists for P-almost every history $h = (s_0, s_1, \dots)$, and $E_P V(f) = E_P f$. Moreover, if P is ergodic then $V(f(h))$ is P-almost surely the constant $E_P f$.

The set of all stationary distributions P over $S^{\mathbb{N}}$ is convex and weak*-compact. We denote by $\mathcal{E} \subset \mathcal{P}$ the set of all ergodic distributions over $S^{\mathbb{N}}$. The *ergodic decomposition theorem* states that there exists a Borel function $\varepsilon : S^{\mathbb{N}} \to \mathcal{E}$ such that $\mu\left(\varepsilon^{-1}(\mu)\right) = 1$ for every ergodic μ and such that

$$P(E) = \int \mu(E)\bar{P}(\mathrm{d}\mu) \tag{5}$$

for every event $E \subseteq S^{\mathbb{N}}$ where $\bar{P} \in \Delta(\mathcal{E})$ is the push-forward of P under ε (i.e., $\bar{P}(E) = P(\varepsilon^{-1}(E))$ for every Borel subset E of $\mathcal{E}$). It follows from the ergodic decomposition theorem that the extreme points of this set are the ergodic distributions. That is, P is ergodic if and only if there exists no stationary distributions R', R'' such that $R' \neq R''$ and $P = \lambda R' + (1-\lambda)R''$ for some $0 < \lambda < 1$.

A.2 Proof of Theorem 1

A.2.1 P is ergodic $\implies$ P displays no structural uncertainty.

Assume that P is ergodic and let $f : S^k \to \mathbb{R}$ be bounded. By Proposition A.2 $V(f)$ is almost surely the constant $E_P V(f)$. It follows that $E_P(V(f)|E) = E_P V(f)$ for every event E with $P(E) > 0$. In particular, $E_{P(\cdot|h^{t-1})}V(f) = E_P V(f)$ for every finite history h^{t-1} with $P(h^{t-1}) > 0$.

A.2.2 P displays no structural uncertainty $\implies$ P is testable.

Assume that P displays no structural uncertainty. Let $f : S^k \to \mathcal{R}$ be bounded. From the martingale convergence theorem it follows that

$$E_{P(\cdot|h^{t-1})}V(f) \xrightarrow[t\to\infty]{} V(f)(h)$$

for P-almost every h. If P displays no structural uncertainty then

$$E_{P(\cdot|h^{t-1})}V(f) = E_P V(f) = E_P f,$$

where the second equality follows from Proposition A.2. Therefore $V(f)(h) = E_P f$ for P-almost every h.

Let T_f be a test such that $T_f(P) = \{h | V(f)(h) = E_P f\}$ and $T_f(Q) = \Omega$ for every stationary $Q \neq P$. By the argument above it follows that T is Type I error free. Moreover if $Q(T_f(P)) = 1$ for some stationary distribution Q then $V(f)(h) = E_P f$ for Q-almost every h, and therefore $E_Q f = E_Q V(f) = E_P f$, where the first equality follows from Proposition A.2.

If $Q \neq P$ is any stationary distribution, then by Proposition A.1 there exists a bounded function $f : S^k \to \mathcal{R}$ such that $E_P f \neq E_Q f$. Then it follows that T_f is Type I error free and that $Q(T_f(P)) < 1$, as desired.

A.2.3 P is testable $\implies$ P precludes disagreement.

Let P be testable and let Q be any stationary belief such that $Q \neq P$. We claim that Q is not compatible with P. Indeed, let T^* be a Type I error free test such that $Q(T^*(P)) < 1$. Since T^* is Type I error free, we get that $P(T^*(P)) = 1$. Therefore P and Q are not compatible, as desired.

A.2.4 P precludes disagreement $\implies$ P is ergodic.

Assume by contradiction that P is not ergodic so that $P = \lambda R' + (1-\lambda)R''$ for some stationary beliefs $R' \neq R''$ and $0 < \lambda < 1$. Since $P(B) = \lambda R'(B) + (1-\lambda)R''(B)$ for every event B it follows that $P(B) = 1$ if and only if $R'(B) = R''(B) = 1$. Let $0 < \mu < 1$ such that $\mu \neq \lambda$ and let $Q = \mu R + (1-\mu)R''$. Then $Q(B) = 1$ if and only if $R'(B) = R''(B) = 1$ and so $Q(B) = 1$ if and only if $P(B) = 1$. Thus P and Q are compatible, and since $R \neq R''$ and $\lambda \neq \mu$ it follows that $P \neq Q$, contradicting the assumption that P precludes disagreement.

A.2.5 P is ergodic $\implies$ P can be empirically identified

This part of the theorem is a slight extension of Hansen (1982)'s result on the consistency of the GMM estimator. First we restate his result in our notation:

Proposition A.3. *Assume that some ergodic distribution P is identified by a moment condition (2) with $k = 1$. Let $\hat{z}_n$ be the GMM estimator given by (3). Then $\hat{z}_n \xrightarrow{P} \bar{z}$.*

We need to adapt this result to arbitrarily finite k. Let P be ergodic and let f be a moment condition given by (2):

$$f : Z \times S^k \to \mathcal{R}^q.$$

Consider the new set of outcomes $\bar{S} = S^k$. Define the probability distribution $\bar{P}$ over $\bar{S}^{\mathbb{N}}$ as the push-forward of P under the map $(s_0, s_1, s_2, \dots) \in S^{\mathbb{N}} \mapsto (\bar{s}_0, \bar{s}_1, \dots)$ where

$$\bar{s}_n = (s_n, s_{n+1}, \dots, s_{n+k-1}).$$

Then $\bar{P}$ is ergodic and the moment conditions f on P translates to a moment condition on $\bar{P}$ with $k = 1$. Thus, by Proposition A.3 it follows that the GMM estimator is consistent. Since this is the case for every moment condition f it follows that P can be empirically identified.

A.2.6 P can be empirically identified $\implies$ P is ergodic

Let $P \in \mathcal{P}$ be stationary and empirically identified. Fix an integer k and let $Z = \Delta(S^k)$. Define the moment condition $f : Z \times S^k \to \mathbb{R}^{S^k}$ be given by

$$f(z, h)[h'] = z[h'] - \delta_{h,h'}$$

for every $z \in \Delta(S^k)$ and $h' \in S^k$, where $\delta_{h,h'} = \begin{cases} 1, & \text{if } h = h' \\ 0, & \text{otherwise.} \end{cases}$. If μ is a stationary distribution then f identifies μ since the unique $\bar{z}$ that satisfies $E_\mu f(\bar{z}, h) = 0$ is given by $\bar{z} = \Pi^k(\mu)$ where $\Pi^k : \mathcal{P} \to \Delta(S^k)$ is such that $\Pi^k(\mu)$ is the distribution on k-tuples induced by μ for every stationary distribution μ.

It follows from A.2.5 that for every ergodic $\mu \in \mathcal{E}$

$$\hat{z}_n \xrightarrow{P} \Pi^k(\mu) \tag{6}$$

under μ, where z_n is the GMM estimator.

We claim that P,

$$\hat{z}_n \xrightarrow{P} \Pi^k \circ \varepsilon. \tag{7}$$

where $\varepsilon : S^{\mathbb{N}} \to \mathcal{E}$ is the ergodic decomposition function introduced earlier satisfying 5. Indeed, for every $\alpha > 0$

$$P\left(\left\{\omega \in S^{\mathbb{N}} \middle| \, |\, \hat{z}_n(\omega) - \Pi^k \circ \varepsilon(\omega)\, | > \alpha\right\}\right) = \\ \int \bar{P}(\mathrm{d}\mu)\mu\left(\left\{\omega \in S^{\mathbb{N}} \middle| \, |\, \hat{z}_n(\omega) - \Pi^k(\mu)\, | > \alpha\right\}\right) \xrightarrow[n\to\infty]{} 0$$

where the equality follows from the definition of ε and the limit from the bounded convergence theorem and (6).

On the other hand,

$$\hat{z}_n \xrightarrow{P} \Pi^k(P). \tag{8}$$

since P can be empirically identified. It follows from (7) and (8) that $\Pi^k \circ \varepsilon = \Pi^k(P)$, P-a.s. Since this is true for every k it follows that ε is constant P-a.s., i.e., that P is ergodic.

A.3 Proof of Proposition 4.1

Since $\pi_1, \dots, \pi_L$ are not identical, there exists some $s^* \in S$ such that $\pi_1(s^*) \neq \pi_2(s^*)$. Therefore there exists some $u \in \mathbb{R}^S$ and $\alpha \in \mathbb{R}$ such that

$$\pi_1(s^*) \cdot u < \alpha < \pi_2(s^*) \cdot u$$

where $\cdot$ is the inner product in $\mathbb{R}^S$.

Assume that at every period n each player has two actions: a safe action that gives payoff α regardless of the next state, and a risky action that gives payoff $u[s_{n+1}]$. Thus, a player's payoff depends only on her own action and not on the opponent's actions. In this game, every equilibrium strategy σ is such that after every history h the player chooses the safe action if the conditional expectation of $u[s_{n+1}]$ is smaller than α and the risky action if the conditional expectation of $u[s_{n+1}]$ is greater than α. Since after sufficiently long periods the player learns the ergodic transition q it follows that there are arbitrary long histories $(s_0, s_1, \dots, s_n)$ with $s_n = s^*$ after which the player must play the safe action and arbitrary long histories after which the player must play the risky action. Thus, the equilibrium cannot be Markovian.

A.4 Proof of Proposition 4.2

The proof uses the following lemma, which follows immediately from the ergodic theorem.

Lemma A.4. *Let S be a finite set of outcomes and let P be a stationary distribution over $S^{\mathbb{N}}$. Then P-almost every realization $\omega = (s_0, s_1, \dots) \in S^{\mathbb{N}}$ has the property that s_0 appears infinitely often in ω.*

Proof: By the ergodic decomposition we can assume w.l.o.g. that P is ergodic. Let $p \in \Delta(S)$ be the marginal one-period distribution of P, i.e. $p[s] = P(\{\omega = (s_0, s_1, \dots) | s_0 = s\})$ for every $s \in S$. Then it follows from the ergodic theorem that, for P-almost every realization ω, every $s \in S$ appears in ω with frequency $p[s]$. In particular, if $p[s] > 0$ then s appears infinitely often in ω. This implies that s_0 appears infinitely often in ω for almost every realization $\omega = (s_0, s_1, \dots) \in S^{\mathbb{N}}$, as desired. ∎

To prove the proposition, let $\epsilon > 0$ be small enough such that:

1. The only subgame-perfect ϵ-equilibrium in the repeated game with incomplete information (D, D'') is such that player 1 always play a and player 2 always play an action from $B(a)$.

2. The only subgame-perfect ϵ-equilibrium in the repeated game with incomplete information (D', D'') is such that player 1 always plays a' and player 2 always play an action from $B(a')$.

Let $\eta > 0$ sufficiently small. It follows from Kalai Lehrer's Theorem that in every equilibrium of the game with incomplete information, if the types are (D, D'') there is a time from which, with probability at least $1 - \eta$, player 2 always plays an action from $B(a)$. It follows from Lemma A.4 that according to the equilibrium profile, conditioned on the types being, (D, D''), player 2 plays an action in $B(a)$ with probability at least $1 - \eta$ (since almost every realization of the play path on which he doesn't play an action in $B(a)$ cannot have the property that he plays an action from $B(a)$ from some point onward.) This implies that according to his equilibrium strategy, player 2 must play an action from $B(a)$ with probability $1 - \eta$ at day 0 if his type is D''. Similarly, player 2 must play an action from $B(a')$ with probability $1 - \eta$ at day 0 if his type is D''. Since $B(a) \cap B(a'') = \emptyset$ we get a contradiction.

References

Al-Najjar, N. I., A. Sandroni, R. Smorodinsky, and J. Weinstein (2010): "Testing Theories with Learnable and Predictive Representations," *Journal of Economic Theory*, 145, 2203–17.

Aumann, R. (1976): "Agreeing to disagree," *The annals of statistics*, 4(6), 1236–1239.

Blackwell, D., and L. Dubins (1962): "Merging of opinions with increasing information," *Ann. Math. Statist.*, 33, 882–886.

Brav, A., and J. Heaton (2002): "Competing theories of financial anomalies," *Review of Financial Studies*, 15(2), 575.

Cochrane, J. (2005): *Asset pricing*. Princeton university press.

Dekel, E., and Y. Feinberg (2006): "Non-Bayesian Testing of an Expert," *Review of Economic Studies*, 73, 893–906.

Dekel, E., D. Fudenberg, and D. Levine (2004): "Learning to play Bayesian games," *Games and Economic Behavior*, 46(2), 282–303.

Ericson, R., and A. Pakes (1995): "Markov-perfect industry dynamics: A framework for empirical work," *The Review of Economic Studies*, 62, 53–82.

Fudenberg, D., and D. K. Levine (1993a): "Self-confirming equilibrium," *Econometrica*, 61(3), 523–545.

——— (1993b): "Steady state learning and Nash equilibrium," *Econometrica*, 61(3), 547–573.

Hansen, L. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, pp. 1029–1054.

Hong, H., and J. Stein (2007): "Disagreement and the stock market," *The Journal of Economic Perspectives*, 21(2), 109–128.

Kalai, E., and E. Lehrer (1993): "Rational Learning Leads to Nash Equilibrium," *Econometrica*, 61, 1019–1045.

Kandel, E., and N. Pearson (1995): "Differential interpretation of public signals and trade in speculative markets," *Journal of Political Economy*, 103(4), 831–872.

Lewellen, J., and J. Shanken (2002): "Learning, Asset-Pricing Tests, and Market Efficiency," *The Journal of finance*, 57(3), 1113–1145.

Lucas, R. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, pp. 1429–1445.

Milgrom, P., and N. Stokey (1982): "Information, trade and common knowledge," *Journal of Economic Theory*, 26(1), 17–27.

Nachbar, J. H. (2005): "Beliefs in repeated games," *Econometrica*, 73(2), 459–480.

Olszewski, W. (2012): "Calibration and Expert Testing," in *Handbook of Game Theory*, vol. IV.

Olszewski, W., and A. Sandroni (2008): "Manipulability of Future-Independent Tests," *Econometrica*, 76, 1437–66.

——— (2009): "A Nonmanipulable Test," *Annals of Statistics*, 37, 1013–39.

Pakes, A., and R. Ericson (1998): "Empirical Implications of Alternative Models of firm Dynamics," *Journal of Economic Theory*, 79, 1–45.

Pástor, L., and P. Veronesi (2009): "Learning in Financial Markets," *Annu. Rev. Financ. Econ.*, 1(1), 361–381.

Ross, S. (1989): "Intertemporal asset pricing," *Theory of Valuation: Frontiers in Modern Financial Theory*, 1, 85–96.

Sandroni, A. (2003): "The Reproducible Properties of Correct Forecasts," *Internat. J. Game Theory*, 32(1), 151–159.

Shmaya, E. (2008): "Many Inspections are Manipulable," *Theoretical Economics*, 3, 367–82.

Stokey, N., and R. Lucas (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press.

Weitzman, M. (2007): "Subjective Expectations and Asset-Return Puzzles," *The American Economic Review*, 97(4), 1102–1130.