Uncertainty and Disagreement in Equilibrium Models

Nabil I. Al-Najjar & Eran Shmaya
Northwestern University *Tel Aviv University*

RUD, Warwick, June 2014

Forthcoming: *Journal of Political Economy*

Motivation

- Equilibrium concepts, such as Nash equilibrium and rational expectation equilibrium, explicitly assume that agents know the "true" and "objective" stochastic process P governing their environment
- Motivation:
  - There is a "true" process P
  - P is not metaphysical: any $Q \neq P$ can be "proven wrong"
  - In particular, $\implies$ no disagreement
- Sargent 2008 describes (critically)

  > *A rational expectations equilibrium asserts that the same model is shared by (1) all of the agents within the model, (2) the econometrician estimating the model, and (3) nature, also known as the data generating mechanism.*

This ought to be disturbing to a subjectivist

Outline

Examine the idea that *equilibrium probabilities* are “objective,”

- Operationalize the idea of “objectivity”
- “objective” = “statistically testable”
- Relate to ongoing debates in Macro-Finance-Econometrics
- Formulate the theoretical question
  - untied to a specific economic context
  - .. in a way a (decision) theorist can relate to

Philosophy

Two polar traditions:

1. Rational Expectations: beliefs are anchored in empirical frequencies
   - Econometrically tractable
   - ..but rules out “uncertainty” and disagreement
2. Subjectivist Tradition: a belief is “a state of mind”
   - Consistent with heterogeneity and disagreement
   - ..but difficult to implement in practice (anything goes?)

The Hope:
Merge the two traditions into a *subjectivist equilibrium theory*

This paper shows why this is very hard;
better understanding of the sources of difficulty

Dynamical Economic Systems

- A dynamic model is a stochastic process P
  - Finite set of period outcomes S
  - Sequences of outcomes

$$H = S \times S \times \cdots$$

  - P is a distribution on H
- Information: finite initial histories h^{n-1}
- Standard topology and measures
  - Borel sigma-algebra on products topologies
- Finite S is technically convenient but not critical

Stationarity

Stationarity

For every integer m, the marginal distribution of P on

$$S_{k+1} \times \cdots S_{k+m}$$

does not vary with $k = 1, 2, \ldots$

- Important in theoretical and econometric models
- Later we relax it
- $\mathcal{P}$ is the set of stationary distributions

Example: Asset Pricing Models

- Stochastic consumption process $\{c_n\}$ of the representative agent
  - (Assume rich enough finite outcome space)
- $\{c_n\}$ is exogenously given
  - Asset prices and returns are determined by the model
- Time separable utility $\sum_{t=1}^{\infty} \delta^n u(c_n)$

Example: Asset Pricing Models

- The *stochastic discount factor* (SDF) is the process:

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

- *Our primitive is a process P on $\{m_n\}$.*
- The equilibrium return of any asset is determined via the Euler equation:

$$E_P\left[m_{n+1}\, R_{n+1} \,|\, h^n\right] = 1$$

where $\{R_n\}$ is an asset's (stochastic) gross rate of return

Example: Markov Perfect equilibrium

- Exogenous variables evolve according to a Markovian transition π
- Notation too heavy; see the paper
- Ingredients:
  - Firms use Markovian strategies
  - Known π+ Markovian strategies $\implies$ Markovian P
- P has "every right" to be called the true objective theory, shared by agents and outside observers
  - Real "sense" that probabilities are "objective"
  - ..hence no room for disagreement
  - No uncertainty about the long-run
  - Standard econometrics "works"

Tests: Idea

- Say an outside observer (a modeler, econometrician) assumes that agents hold correct equilibrium beliefs
- Verifying this requires an objective criterion that compares beliefs with an observed realization of the process
- Idea to formalize:

  "If P is incorrect and the correct alternative Q is presented, P can be proven wrong with some positive probability"

Tests: Formal Definition

A statistical *test* is a 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

- The set of all histories consistent with P:

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

- These are the empirical predictions of P relative to test T: if P is correct, then the observed sequence of outcomes must be in T_P

Type I Errors

- The *Type I error* of a test T at P is the number:

$$1 - P(T_P)$$

- In our asymptotic setting, it is natural to consider tests that are *free of Type I error*:

$$P(T_P) = 1, \quad \forall P$$

- With finite observations, small but positive Type I errors must be allowed
- Asymptotic setting makes for shaper statements and interpretation of the results

Testability

Definition

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. $Q(T^\star_P) < 1$

- $Q(T^\star_P)$ = probability that P is confirmed when false
  - Just the Type II error of $T^\star$
- Definition says: for any Q there is *some* test for which Type II error is not 100% (*i.e.*, has *some* power)

Testability: Intuition

- Imagine we propose P as an equilibrium model of asset prices
- Suppose P is not testable
- Then someone can suggest another model $Q \neq P$ such that no matter what Type I error-free test we use, there is no way to prove P false
- P has no testable implications against some alternative Q
- *We think of testability as a desirable property of a model*

Disagreement

- There is a strong intuition that people hold different beliefs about their environments
- Differing beliefs perhaps reflect persistent imperfections in learning
- Empirical paradoxes and theoretical responses:
  - Trade volume puzzles
  - Speculative trade
  - One response: models of heterogeneous priors

Precluding Disagreement

Definition

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

Mutual absolute continuity

Definition

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

P accommodates disagreement if people can hold different compatible beliefs

Uncertainty about Long-Run Fundamentals

What do we mean by long-run fundamentals?

- Suppose P is a Markov process
- There is uncertainty about next period outcome ..and the period next, and next..
- But there is no uncertainty about the long-run distribution
  - Formally, Markov processes are mixing
  - There is uncertainty about near horizon events but not about the long-run
- *Value of information:* additional observations will not help me better predict the long-run
  - Implication: I already know it

No-Long Run Uncertainty

For a sequence $(x_k, x_{k+1}, \ldots)$ 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

No-Long Run Uncertainty

- For a function $f : S^k \to \mathbb{R}^m$, with finite k, define:

$$f_n \equiv f\left(s_{n-k+1}, \ldots, s_n\right), n = k, k+1, \ldots$$

- $f_k, f_{k+1}, \ldots$ is a payoff stream that depends on past realizations via the (stationary) formula f
- $f \equiv \{f_k, f_{k+1}, \ldots\}$
- $V(f)$ is the random variable defined by:

$$V(f)(h) = \lim_{n\to\infty} \frac{1}{n} \sum_{i=k}^{n} f_k$$

i.e., the limiting average of the stream $f_k, f_{k+1}, \ldots$ generated along h

No-Long Run Uncertainty

Definition

P displays no long-run structural uncertainty if for every f and history h^{t-1} with $P(h^{t-1}) > 0$

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

- We may entertain doubts about the short-run outcomes, but not about the long-run behavior of the process

Moment Conditions: Asset Pricing Revisited

- Moment conditions are “finite-horizon features of P”
- Example: the Euler equation of asset pricing

$$E_P\left[m_{n+1}\, R_{n+1} \mid h^n\right] - 1 = 0$$

- Assume CRRA utility, with unknown parameter γ
- Fundamental parameters to be estimated:

$$Z = \{\delta, \gamma, \text{ moments of } P \ldots\}$$

- Empirical moment condition

$$\frac{1}{n}\sum_{n=1}^{N}\left[\frac{c_n^{\gamma}}{c_{n+1}^{\gamma}}\, R_{n+1} \mid \text{ information}\right] - 1 = 0$$

- Moment estimation underlie many statistical techniques used in econometric practice
  - Generalized method of moments (GMM) introduced in Hansen (1982)
  - GMM itself covers many techniques such as linear and non-linear regressions, VAR, ... etc

Moment Conditions and Statistical Identification

Definition

A moment condition *is a bounded continuous function*

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

- *k and q are positive integers*
- *$Z \subset \mathbb{R}^m$ is compact*

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

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

Empirical Identification

In GMM, starting with a *theoretical* moment condition, one estimates the true $\bar{z}_P$ by the element

$$\hat{z} \in Z$$

that minimizes the *empirical* moment condition

(Details are standard: see the paper or a time series textbook)

Definition

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

- This says 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'

$$E_{P'} f(\bar{z}, s_1, \ldots, s_k) = 0$$

  such that $\bar{z}$ cannot be recovered, even asymptotically as data grows without bound

Main Result

Theorem

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.*

The Testable Case

Useful result:

Proposition

The conditions of the Theorem are satisfied for

- Every irreducible memory k Markov process
- Every hidden Markov process where the underlying Markov process is irreducible.

This case covers:

- All MPE models in IO, PE, etc..
- Parametric consumption-based asset pricing models
  - Lucas-style asset pricing model
  - Exogenous and endogenous variables are Markovian

By the Theorem, they preclude disagreement and uncertainty about long run fundamentals

Simple intuition

Skip slide

- Two stationary Markov processes, indexed by transitions

$$\{\pi_1 \neq \pi_2\}$$

  - Simplest possible example: i.i.d. processes

- Each P^{π_i} is testable (by the proposition)
- Given P^{π_i}, easy to see that:
  - There can be no disagreement (any stationary P' compatible with P^{π_i} must coincide with it)
  - Given P^{π_i} there is no uncertainty about the long-run
  - GMM is consistent

Uncertain beliefs

Introduce doubts about which of $\{\pi_1, \pi_2\}$ is the true model

$$\mu \in \hat{\Delta} \equiv \text{ interior } \Delta\{\pi_1, \pi_2\}$$

Disagreement and uncertainty about the long-run are possible

However: *Let μ, μ' be any pair of beliefs in $\hat{\Delta}$ and T any Type I error-free test. Then*

$$P^{\mu}\{T(\mu', h) = 1\} = 1$$

An agent who believes the data is generated by μ also believes that no zero-Type-I-error-test can disprove the wrong belief μ'

How should Rational Expectations be defined in this case?

Proof

Skip slide

Zero Type I error means:

$$1 = P^{\mu'}\{T(\mu', h) = 1\}$$

$$= P^{\pi_1}\{T(\mu', h) = 1\}\mu'(\pi_1) + P^{\pi_2}\{T(\mu', h) = 1\}\mu'(\pi_2).$$

$$\Longrightarrow$$

$$P^{\pi_1}\{T(\mu', h) = 1\} = P^{\pi_2}\{T(\mu', h) = 1\} = 1$$

$$\Longrightarrow$$

$$P^{\mu}\{T(\mu', h) = 1\} = P^{\pi_1}\{T(\mu', h) = 1\}\mu(\pi_1) + P^{\pi_2}\{T(\mu', h) = 1\}\mu(\pi_2)$$

$$= 1$$

The other equality is proved similarly.

Concluding Observations

- There is an emerging consensus that injecting subjectivity in equilibrium analysis is essential to explaining theoretical and real-world puzzles
- Much of economic theory is based on the idea that there is a ‘true’ objective process (deep down we are objectivists)
- My view is that this commitment is like enter into a “bargain with the devil,” eventually exacting a substantial price
- My view also is that this has little to do with the axioms of decision theory
- Question: is there a way to develop subjectivist statistics?

Additional slides
not shown in the talk.

Non-Parametric Models

- Empirical models that make no explicit assumptions about the structure of P
- Use GMM to estimate moments of P
- These model assume that GMM is consistent
- By the Theorem, these models assume that the underlying process is testable and thus preclude disagreement and structural uncertainty

The Non-Testable Case

- An implication of that theorem is that in stationary non-testable environment there is always something to learn about the long-run fundamentals.
- As agents learn from observing the process
- Changing beliefs can lead to non-stationary decisions, undermining the stationarity of the process

Goal: introduce limited form of non-stationarity so we can still leverage the Theorem

Learning, Stationarity, and Optimization

- Write $S = X \times A$
  - Action profiles A at time n
  - Exogenous variables X not affected by these actions
- Evolution of the system
  - stochastic process P on $H = (X \times A)^\infty$
  - P_X the marginal of P on X^∞
  - P_A the marginal on A^∞
- Require P_X to be stationary, but not necessarily the entire process P
- Makes it possible to apply the Theorem restricted to P_X
- Beliefs about the exogenous variables are stationary, but actions (and beliefs about actions) are not necessarily so

Passive Learning

- MPE setting except that agents are uncertain about the transition π of exogenous variables
- Pakes and Ericson (1998) study a model along these lines
- Their goal is to empirically test for the implications of agents' passive learning on industry dynamics.
- Agents have beliefs $\mu_i, i \in I$, with a finite and common support $\Pi = \{\pi_1, \ldots, \pi_L\}, L \geq 1$
- Agents learn, but learning is passive
  - agents' actions do not influence exogenous variables
  - Rules out agents' active experimentation
  - Keeps the theoretical and empirical analysis tractable
- Formally, each $\pi_l \in \Pi$ takes the form $\pi_l(x_n)$

Passive Learning

- 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

$$P_X(\mu_i) = \mu_i(\pi_1)P_X(\pi_1) + \cdots + \mu_i(\pi_l)P_X(\pi_l)$$

- Agents believe none can influence this distribution

How about actions?

Learning and Optimization

Consider three cases:

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

All of P falls under the testable case

Case 2: *$L > 1$ and agents share a common prior ($\mu_i = \mu$, for all i) with support $\{\pi_1, \ldots, \pi_L\}$.*

Case 3: *$L > 1$ and agents have different priors ($\mu_i \neq \mu_j$, for some i, j) with common support $\{\pi_1, \ldots, \pi_L\}$.*

We show that under assumptions about payoffs that P_A may be inconsistent with *any* stationary behavior

The Econometrician's Perspective

- Pakes and Ericson (1998) exploit this to conduct tests for stationarity
- An econometrician making the usual assumptions of rational expectations (e.g. assume GMM or testability) will get very strange results looking at the data
- This issue is endemic (was pointed out in asset pricing models, Lewellen-Shanken (2002), Brav-Heaton (2002), Weitzman (2005) among many others)

Testing Beliefs vs. Testing Behavior

- When beliefs disagree or display uncertainty about long-run fundamentals, the theorem says that it is not possible to devise an objective test to refute alternative beliefs Q
- It not mean that disagreement and structural uncertainty can have observable implications on behavior and market outcomes
- The Theorem does imply that in this case one would need tools that look very different form GMM

The Econometrician's Perspective *Redux*

Pakes and Ericson (1998) develop econometric techniques in a Markovian setting to detect firms' learning

- Main issue is that beliefs are not directly observable
- Need to look for *indirect* manifestations of subjective uncertainty
- Idea: look for whether early observations have a persistent impact on outcomes

Other example can be found in the context of asset pricing

But all this is *very hard*. GMM is so nice.. when it works!

Bayesian Games

- This part is informal; see the paper for details
- Nature moves to select a stationary process for fundamentals
  - Industry shocks in an IO or trade model
- Players strategies are part of an equilibrium
  - In particular, they best respond to their beliefs about the evolution of fundamentals and strategies of others
- Unless the game is “weird” when players learn, their behavior is no longer stationary

Idea of the Proof

- The four properties above characterize ergodic distributions

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

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

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*.

Idea of the Proof

- Ergodic Decomposition Theorem: Every stationary process is a mixture of ergodic processes.
- For an ergodic process P, the decomposition is trivia
  - P has distinct support from any other ergodic process
    - Testability
  - Empirical averages coincide with the theoretical probabilities
    - Moment conditions
  - If P is a non-trivial mixture, then there is uncertainty about the true ergodic component
    - There is uncertainty about long-run fundamental properties of the process