40th conference on

The mathematics of operations research

Landelijk Netwerk Mathematische Besliskunde

Lunteren, January 14, 2015

Learning, Predictability and Stability in Large Games

Ehud Kalai

Northwestern University

Predictability and stability is critical for the well functioning of social systems.

Producers and consumers need reliable prices in order to plan their activities.

Reliable driving patterns are important for proper delivery of goods, proper planning of roads by traffic engineers, transportation of passengers, etc..

Reliable demand and supply is important in health delivery systems, in the provision of new technology, etc. etc..

Stability means that in every period of play, once the period's outcome is observed, the players have no incentives to deviate from the plans that lead to the outcome, (with hindsight knowledge of the outcome, they would have no regret over their choices). This is stability within a period, which does not imply that the play is constant over different periods.

When such **hindsight stability** fails, in situations where players are able to revise their choices, the period's outcome may become chaotic. For example, consider the game of driving during the rush hour.

This presentation

Part A: Motivation and Review, Kalai, Econometrica (2004)

Review hindsight stability in **one shot** games with **independent types,**

New material, Kalai and Shmaya (2013 and 2014)

Part B: Markov perfect equilibrium in an **imagined-continuum model** of a repeated population game: A simple ϵ-fully rational equilibrium, in a game which is too complex to play otherwise

Part C: **Learning** and **hindsight Stability**.

Part D: **Price stability** in a repeated **Cournot game**.

This presentation

Part A: Motivation and Review, Kalai, Econometrica (2004)

Review hindsight stability in **one shot** games with **independent types,**

Example: hindsight instability in small vs large coordination game

Example

Simultaneous-move Computer choice game: independent types

Players: i = 1,2,…,n; each chooses an action: $a^i = \mathcal{PC}$ or $a^i = \mathcal{M}$

Player's types: iid Pr(t^i = $\mathcal{PC}$) = Pr(t^i = $\mathcal{M}$) = .50

Individual's payoff: $u^i = \text{prop}_{j \neq i}\,(a^j = a^i)^{1/3} + 0.2\delta_{a^i = t^i}$,

i.e., (the proportion of opponents he matches)$^{1/3}$
+ 0.2 if he chooses his computer type (0 otherwise).

Choose your favorite computer ($a^i = t^i$) is a *Nash* equilibrium.

It is "asymptotically **hindsight stable**" (as the number of players increases).

Even more, it is "asymptotically **structurally robust**": it remains an equilibrium in all extensive-game alterations that (1) start with the same initial information, (2) preserve the players' strategic possibilities and (3) do not alter the players' payoffs.

For example, it survives under **sequential play** (no herding),
and, more generally, under **general dynamic play**: **revision of choices, information leakage, cheap talk, delegation possibilities**, and more..

Modelling Partially-specified games

Students choosing computers on the web

Instructions: "Go to web site xyz before Friday and click in your choice $\mathcal{PC}$ or $\mathcal{M}$." Types, choices, and payoffs as before.

Need to know: who are the players? the order of play? monitoring? communications? commitments? delegations? revisions?... **Impossible**

But under structural robustness: any equilibrium of the one-shot simultaneous-move game, (e.g., choose your favorite computer) remains equilibrium no matter how you answer the above.

Price formation in Shapley Shubik market games

Hindsight stability $\rightarrow$ price stability

Kalai Econometrica (2004): In n-player one-shot simultaneous-move Bayesian games with independent types:

many players,
semi-anonymity,
continuity &
~~Independent types~~

So hindsight stability is important

But it fails when player types are correlated, common in economic interaction.

All Nash equilibria are asymptotically **hindsight stable**

~~**Price stability** in market games~~

~~All Nash equilibria are asymptotically **structurally robust**~~

Hindsight stability fails with correlated types

will be used repeatedly in the following slides

Computer choice game with correlated types.

Players: i = 1,2,…,n, each chooses $\mathcal{PC}$ or $\mathcal{M}$.

Unknown **state of nature**: the computer with better overall features is:

s = $\mathcal{PC}$ or s = $\mathcal{M}$ with prob .50 , .50

Player types: iid conditional on s: $\Pr(t^i = s) = 0.7$

Payoffs: as before.

Answer to follow:
If the number of players is large, with the exception of a finite number of chaotic learning periods, all periods are asymptotically hindsight stable.

E

play they all know the state of
types are (conditionally) independent.
ggests the study taken next

What happens with hindsight stability in large **repeated games** with **correlated types?**

The Model

A Bayesian repeated game with

1. a large but unknown number of players n, who
2. have fixed types, correlated through an unknown state of nature s, and
3. have imperfect monitoring.

Computing **standard Bayesian best response is too demanding.**

The Players need priors over:

- The number of players.
- The state of the nature.
- The types of the players.

Prior to every period the players must update their beliefs over:

- The number of players.
- The state of nature.
- The types of the players.
- The history of actions selected in all previous periods.

The model and equilibrium presented next eliminate these.

The Model

In our imagined continuum model, each individual player thinks of himself as negligible in a game with a continuum of players;
but as game theorists we perform a correct probabilistic analysis of what really happens with such n players.

Computing standard Bayesian best response is too demanding.

The Players need priors over:

- The number of players.
- The state of the nature.
- The types of the players.

Prior to every period the players must update their beliefs over:

- The number of players.
- The state of nature.
- The types of the players.
- The history of actions selected in all previous periods.

The model and equilibrium presented next eliminate these.

The Model

In our imagined continuum model, each individual player thinks of himself as negligible in a game with a continuum of players;
but as game theorists we perform a correct probabilistic analysis of what really happens with such n players.

This imagined continuum model is different from the standard continuum-of-players game (a la Schmeidler (1973)) in which both, the players **and the game theorist** pretend that there is a continuum of players.

In addition to simplifying the Bayesian analysis of the players, in the imagined continuum model the game theorist retains a **coherent probability structure** (lost in the continuum of players game) needed for the analysis of events (e.g., merging, hindsight stability, …).

The Model

In our imagined continuum model, each individual player thinks of himself as negligible in a game with a continuum of players;
but as game theorists we perform a correct probabilistic analysis of what really happens with such n players.

We restrict ourselves to anonymous and symmetric games of proportions.
A repeated illustrative example will be:

Repeated computer choice game with correlated types.

Prior probabilities:

One (unknown) **state of nature** s,
$\theta_0(s)$ is the known prior prob that the state is s.

$\theta_0(S = \mathcal{PC}) = \theta_0(S = \mathcal{M}) = 1/2$,
each computer is equally likely to have the better overall features

Players' privately known **types** are distributed by a conditional iid,
τ_s(t) is the prob of a player being of type t when the state is s.

$\tau_{\mathcal{PC}}(t = \mathcal{PC})$ = .7 and $\tau_{\mathcal{PC}}(t = \mathcal{M})$ = .3:
when the state is $\mathcal{PC}$, independently of the others every player has prob of 0.7 to be a $\mathcal{PC}$ type and 0.3 to be a $\mathcal{M}$ type.

Symmetrically, $\tau_{\mathcal{M}}$(t=$\mathcal{M}$) = .7 and $\tau_{\mathcal{M}}(t = \mathcal{PC})$ = .3

The Stage Game, played in periods k = 0,1,2,…:

Actions available to every player are denoted by $a \epsilon A$

a = $\mathcal{PC}$ or a = $\mathcal{M}$, chooses $\mathcal{PC}$ or chooses $\mathcal{M}$.

$e_k(t, a)$ is the **empirical proportion** of players who are of type t and choose the action a.

$e_k(\mathcal{PC}, \mathcal{PC})$ the proportion of players who like $\mathcal{PC}$ and chose $\mathcal{PC}$,
$e_k(\mathcal{PC}, \mathcal{M})$ the proportion of players who like $\mathcal{PC}$ but chose $\mathcal{M}$, etc.

A random **outcome** x, is chosen with prob $\chi_{s,e}(x)$ and made public at the end of the period.

A sample with replacement of J computer users is taken $\boldsymbol{x} = x(\mathcal{PC})$ is the sample proportion of $\mathcal{PC}$ users. $x(\mathcal{M}) \equiv 1 - x(\mathcal{PC})$.

$u(t, a, x)$ is the **period payoff** of a player of type t who took the action a in a period with the outcome x.

u = (the proportion of users in the sample that his choice matches)$^{1/3}$
+ 0.2 if his choice is the same as his type, $a = t$.

The Repeated game

The repetition/payoff structure

Can be:

Infinitely repeated with discounting.

Finitely repeated with the average of the periods payoffs.

Any function that is continuous and strictly monotonic in the periods payoffs. Will elaborate.

The repeated computer choice game is infinitely repeated with individual discount parameters.

Strategies and equilibrium terminology

We restrict the players to simplified symmetric equilibrium that does not depend on n, but we do a full Bayesian asymptotic analysis of such n players, as $n \to \infty$.

A **common strategy** F is a symmetric profile in which all the players play F.

F is **Markov**, if it depends only on the player's type and the "public-belief" over the unknown state, will elaborate.

In the **imagined play path**, periods' random empirical distributions of types & actions are replaced by their deterministic conditional expectations. The only uncertainty is about the unknown state s.

A **Markov strategy** is a function $\mathrm{F}: \Delta(S) \times \mathrm{T} \to \Delta(A)$

$F_{\theta,t}(a)$ is the probability of choosing the action a by a player of type t in periods in which the *Markov state* is θ, where θ is the prob distribution that describes the public belief about the unknown state s, to be described.

An $\boldsymbol{\alpha}$ threshold strategy : With prob 1:

Choose your type of computer t in periods with

$$\alpha < \theta(t) \quad ,$$

choose the other computer t^c in periods with

$$\theta(\mathrm{t}) \leq \alpha$$

The **public beliefs** (in the imagined game) about s under a common strategy F

The **initial** public belief is θ_0 .

Continuing inductively, in a period that starts with the imagined public belief θ:

the expected values from the continuum game.

1. to every state s associate the (imagined) deterministic **empirical distribution**

$$\mathrm{d}_\theta(t,a) = \tau_s(t) \cdot F_{\theta,t}(a), \quad \text{and}$$

2. after observing the period outcome x, compute the **posterior** public belief by Bayes rule

$$\hat{\theta}_{\theta,x}(s) \equiv \frac{\theta(s) \cdot \chi_{s,d_\theta}(x)}{\sum_{s'} \theta(s')\chi_{s',d_\theta}(x)}$$

Suppose a period's **prior public belief** is $\theta(s = \mathcal{PC}) = .6$, $\theta(s=\mathcal{M}) = .4$ and that under $F_{\theta,t}$ each player chooses his computer type with probability 1.

1. The **imagined** empirical distribution $d_\theta(t,a)$ is

For s=$\mathcal{PC}$

$$d_\theta(\mathcal{PC},\mathcal{PC}) = .7\cdot 1 = .7$$
$$d_\theta(\mathcal{PC},\mathcal{M}) = .7\cdot 0 = .0$$
$$d_\theta(\mathcal{M},\mathcal{PC}) = .3\cdot 0 = 0$$
$$d_\theta(\mathcal{M},\mathcal{M}) = .3\cdot 1 = .3$$

For s=$\mathcal{M}$

$$d_\theta(\mathcal{PC},\mathcal{PC}) = .3\cdot 1 = .3$$
$$d_\theta(\mathcal{PC},\mathcal{M}) = .3\cdot 0 = .0$$
$$d_\theta(\mathcal{M},\mathcal{PC}) = .7\cdot 0 = 0$$
$$d_\theta(\mathcal{M},\mathcal{M}) = .7\cdot 1 = .7$$

Recall, the period outcome x($\mathcal{PC}$) is the proportion of $\mathcal{PC}$ users in a sample with replacement of J computer users from the population.

2. Suppose that in J=20 observations there were 11 $\mathcal{PC}$ users, then the **posterior public belief** is :

$$\hat{\theta}_{\theta,11/20}(s = \mathcal{PC}) = \frac{0.6\cdot \mathrm{B}_{20,0.7}(11)}{0.6\cdot \mathrm{B}_{20,0.7}(11) + 0.4\cdot \mathrm{B}_{20,0.3}(11)} = 0.89$$

$$\hat{\theta}_{\theta,11/20}(s = \mathcal{M}) = \frac{0.4\cdot \mathrm{B}_{20,0.3}(11)}{0.6\cdot \mathrm{B}_{20,0.7}(11) + 0.4\cdot \mathrm{B}_{20,0.3}(11)} = 0.11$$

$B_{20,0.7}(11)$ is binomial prob of 11 successes in 20 trials with success prob 0.7

Private beliefs in a period with public belief θ of a player of type t is

$$\theta^{(t)}(s) \equiv \frac{\theta(s) \cdot \tau_s(\mathrm{t})}{\sum_{s`} \theta(s') \cdot \tau_{s'}(\mathrm{t})}$$

Period expected payoff (in the imagined game) of a type t who chooses the action a when the common strategy is F and the public belief is θ

$$u_F(\theta, t, a) \equiv \sum_x \boxed{\left(\sum_s \theta^{(t)}(s) \cdot \boxed{\chi_{s,d_\theta}(x)} \right)} u(t, a, x)$$

The probability that the player assigns to the outcome x

The probability that the player assigns to the outcome x, for a given s,

Computed with the public belief θ, not with $\theta^{(t)}$, because the strategies of the players are conditioned on the public belief, θ.

Recall: $\mathrm{d}_\theta(t, a) = \tau_s(t) \cdot F_{\theta,t}(a)$,

Definition: A common strategy F is a **Markov equilibrium** (in the imagined game) if for every public belief $\theta\epsilon\Delta(S)$ and every type $t\epsilon T$, if $F_{\theta,t}(a) > 0$, then a maximizes $u_F\ (\theta, t, a)$.

Kalai & Shmaya (2013) defines equilibrium (in the imagined game) more generally, without the Markov property and myopicity, and show that:

- Myopicity is a result (not an assumption): every equilibrium (in the imagined game) is myopic.
- When the number of players is large:

1. Period probabilities in the imagined game are approximately the same as the real probabilities.
2. Equilibria in the imagined game are standard $\boldsymbol{\varepsilon}$ - Nash equilibria in the real n-person game.

Definition: A common strategy F is a **Markov equilibrium** (in the imagined game) if for every public belief $\theta \epsilon \Delta(S)$ and every type $t \epsilon T$, if $F_{\theta,t}(a) > 0$, then a maximizes $u_F\ (\theta, t, a)$.

Due to myopicity:

- Markov equilibrium is applicable to other repetition / payoff structures.
- It is applicable to games with new players entering and players exiting the game at different periods, provided that all players know the current period belief about the unknown state. For example, with overlapping generations, we get learning and stability across generations of players.
- Existence is a simple matter since it has to be shown only for the stage games.

Definition: A common strategy F is a **Markov equilibrium** (in the imagined game) if for every public belief $\theta \epsilon \Delta(S)$ and every type $t \epsilon T$, if $F_{\theta,t}(a) > 0$, then a maximizes $u_F\ (\theta, t, a)$.

<u>Information about the number of players</u>

The players need no such information.

The game theorist must only know that the number is sufficiently large, but does not have to know its precise value.

For the common strategy G in which every player chooses her favorite computer, define $\boldsymbol{\alpha} = \theta(\mathcal{PC})$ at which $u_G\ (\theta, \mathcal{PC}, \mathcal{PC}) = u_G\ (\theta, \mathcal{PC}, \mathcal{M})$,

$\boldsymbol{\alpha}$ is the **tipping value**: when everybody chooses her favorite computer, how low must the prior on $\boldsymbol{S} =$ PC be to make PC types choose M.

$\boldsymbol{\alpha} <$ ½; it is the same for $\mathcal{PC}$ and for $\mathcal{M}$.

The common α threshold strategy is a Markov equilibrium.

- It is trivial to check: direct from the definition.
- It is easy to play: start by choosing your computer type, continue by updating the public belief and following the threshold rule.
- Coordination: from some time on they will all use the same computer.
- It is a real ε Nash equilibrium if the number of players is large.
- If in addition the sample size is large, they will all be using the better computer from the second period on.

Hindsight Stability

Definition of asymptotic Hindsight Stability

Definition (Kalai 2004): consider a common Markov strategy F. *Period* $\boldsymbol{k}$ *is asymptotically* **hindsight stable** up to (ε,ρ) if with sufficiently many players,

$$\Pr\left(\begin{array}{l} \textit{after observing the outcome of the k-th period,} \\ \textit{by a unilateral revision of his action} \text{ some player} \\ \text{can improve his period payoff by more than } \boldsymbol{\varepsilon} \end{array} \right) \le \boldsymbol{\rho}$$

The real probability, computed by the game theorist for the real n-person process.

Definition: *u* has Lipschitz constant L if for all a, x and x' :

$$|u(t,a,x) - u(t,a,x')| < L \cdot d(x,x')$$

Hindsight Stability Theorem

Theorem: For every $\epsilon, \rho > 0$ there is an integer K such that in every Markov equilibrium F and every $d > 0$, all periods except at most K are asymptotically hindsight stable up to

$$[\ d + 2Q_G(d/L) + \epsilon\ ,\ 2Q_G(d/L) + \rho\ ]\ .$$

The lack of concentration of the worst possible uncertainty in the publicly reported outcome: The measure of the set of outcomes x that cannot fit into a ball of diameter < d/L in the worst case (over all s and e).

Corrolary: If the public signal has standard deviation σ. Then for every ε >0 there is a finite integer K such that in every imagined Markov equilibrium F all but at most K periods are asymptotically hindsight stable up to:

$$[2\varepsilon + 8\left(\frac{L\sigma}{\varepsilon}\right)^2, \varepsilon + 8\left(\frac{L\sigma}{\varepsilon}\right)^2]$$

With the α threshold equilibrium in the rpt computer choice game,

For any $\epsilon, \rho > 0$, in all but a finite number of periods the equilibrium is asymptotically hindsight stable up to

$$\left[2\epsilon + \frac{1}{J\epsilon^2} \ , \ \frac{1}{J\epsilon^2} + \rho \right].$$

Recall, J = sample size when sampling for the proportion of PC users.

- Reducing the uncertainty in the publicly observed outcomes (bigger sample size J in our example) improves hindsight stability.
- But with substantial uncertainty in the observed outcomes, hindsight **in**stability is unavoidable, regardless of the number of players.

Rough intuition about the proof

Consider first the |T| *imagined processes*, in which the t-types hold deterministic beliefs about the probabilities of the period empirical distributions of types and actions, $\mathrm{d}_\theta(t, a)$.

- Merging, under the automatic grain of truth, implies that with high probability, except for a finite number of learning periods, the forecasted probabilities over the outcome of the periods are appx accurate. (Fudenberg-Levin, Sorin, Kalai-Lehrer), i.e., the same as would be forecasted with knowledge of the unknown state.
- High concentration (small variance in our example) of the outcome distribution, combined with the fact that the empirical distributions in the imagined processes are deterministic conditional on the states, implies that with high probability at the non-learning periods they predict the realized period outcomes (not just its probability).

So in the imagined processes, in all non-learning periods players will have no regret over their chosen actions, thus hindsight stability holds.

So in the imagined processes, in all non-learning periods players will have no regret over their chosen actions, thus hindsight stability holds.

But what about in the real process, in which the players observe **the randomly realized** real outcomes?

Building on Kalai (2005), Kalai and Shmaya (2013) show that when the number of players is large and outcome probabilities are continuous, real probabilities of period events are approximated well by the probabilities in the imagined process. Thus appx hindsight stability holds with (real) high probability in the non-learning periods.

Remarks:

1. **Hindsight stability is a result of "no further learning" from some time on.** Similar to multi arm bandit problems, the players do not necessarily learn the real state of nature, or even learn to play "as if" they know it.

2. **On the rate of getting to stability**: We know from Sorin (1999) that the number of chaotic periods is monotone in the size of the grain of truth, which is bounded below in our population game. Thus the number of chaotic periods is bounded above.

Price Stability in Repeated Cournot Game

Production possibilities and prices:

$i = 1,2, \ldots, n,$ producers
$k = 0,1,2, \ldots,$ periods
${a^i}_k = 0 \ or\ 1$ possible production levels of i in period k.

$P_k = \frac{1}{2} - A_k + \epsilon$, period-k price,
the publicly announced outcome x_k in the model.

$A_k =$ average (per producer) production,

$\epsilon \sim N(0, \sigma)$ is a random uncertainty, iid across periods (e.g., information shock, noisy traders, …).

Implicitly:There are n buyers. At price x each demands:
$½ - x + \epsilon$

Prior probabilities:

s = *easy* or *difficult* , equally-likely **unknown state of fundamentals**.

t_i privately known, fixed **producer types**, generated by conditional iid's:

$$\Pr(t_i = eff \mid s = easy) = 3/4 \qquad \Pr(t_i = eff \mid s = diff) = 1/4$$

$$\Pr(t_i = ineff \mid s = easy) = 1/4 \qquad \Pr(t_i = ineff \mid s = diff) = 3/4$$

Per-unit production cost: 0 for *efficient* types and 1/8 for *inefficient* types.

The play: In every period, based on her starting info, each producer chooses a production level, the realized price P_k becomes publicly known.

Period payoffs:

For players who produce *0* it is *0*.

For players who produced 1, depending on type $= \begin{cases} P_k, & \text{if } eff \\ P_k - \frac{1}{8}, & \text{if } ineff \end{cases}$

A Unique Markov Equilibrium in period k

Let $\theta_K = \theta_K$ (easy) denote the prior public belief prob that the product is easy to produce. Recall $\theta_0 = 1/2$.

Regim 1 $\theta_k \geq \frac{7+\sqrt{33}}{16} \approx 0.8$: The inefficient types are idle,
the efficient types produce w.p. $p = (4\theta_k + 2)/(8\theta_k + 1)$

Regim 2 $\frac{35-\sqrt{649}}{64} < \theta_k < \frac{7+\sqrt{33}}{16}$: there are unique $0 < q < p < 1$
the inefficient players produce wp q and the efficient ones wp p.

Regim 3 $0.15 \approx \frac{35-\sqrt{649}}{64} \geq \theta_k$: all the efficient ones produce,
the ineff ones produce wp $q = (3 - 6\theta_k)/(18 - 16\,\theta_k)$.

The unique Markov equilibrium of the repeated game

Start in Regime 2, keep updating and playing the corresponding regimes.

Assymptotic Price stability in all but a finite number of chaotic learning periods, with sufficiently many players,

$$\Pr\left(\begin{matrix}\text{at the equilibrium price } \textit{some player can} \\ \textit{improve his payoff by more than } \mathbf{10\sigma^{2/3}}\end{matrix}\right) \leq \mathbf{9\sigma^{2/3}}$$

A Unique Compressed Markov Equilibrium

Let θ_K denote the Bayesian updated prob, prior to the st…
(based on the history of prices), that the st…

Regim 1 $\theta_k \geq \frac{7+\sqrt{33}}{\ldots}$ …

… $(\ldots\theta_k + 1)$

…ere are unique $0 < q < p < 1$
…ient players produce wp q and the efficient ones wp p.

Regim 3 $0.15 \approx \frac{35-\sqrt{649}}{64} \geq \theta_k$: all the efficient ones produce,
the ineff ones produce wp $q = (3 - 6\theta_k)/(18 - 16\,\theta_k)$.

Significant price stability in large market games is obtained in all but a finite number of chaotic learning periods, provided that random uncertainty is only a small component of the price.

The unique Markov equilibrium of the repeated game

Start in Regime 2, keep updating and playing the corresponding regimes.

Assymptotic Price stability in all but a finite number of chaotic learning periods, with sufficiently many players,

$$\Pr\left(\begin{array}{l}\text{at the equilibrium price } \textit{some player can} \\ \textit{improve his payoff by more than } \mathbf{10\sigma^{2/3}}\end{array}\right) \leq \mathbf{9\sigma^{2/3}}$$

Summary

Conceptual contribution:

Hindsight stability

- An important property for equilibrium in strategic interaction.
  Examples: Partially-specified games played on the web.
  Price stability in market games.

- In Markov equilibrium of large Bayesian repeated games, with the exception of a finite number of chaotic learning periods, all periods are asymptotically hindsight stable. True even if the player types are correlated through unknown fundamentals.

Methodological contribution:

Imagined continuum model:

- A hybrid of the continuum and asymptotic models of large games.
- A well defined probability space for Bayesian analysis.
- Avoids problematic computational issues of standard Bayesian equilibrium.
- Equilibrium strategies are independent of the number of players.
- Allows the existence of Markov equilibria.
- Easy to compute ε Nash equilibrium (ε best response) with many players.

To reduce the number of unstable periods:

The number of unstable chaotic periods may be reduced by improving the information of the players in at least two ways:

First, we know from Sorin 1999, that the number of learning periods, in which stability is violated, depends on the accuracy of the players' initial beliefs about the unknown fundamental. i.,e, the size of the grain of truth. Thus,

- starting with better public information about unknown fundamentals is likely to reduce the number of unstable periods.

The other source of instability is the lack of predictability of the periods' publically announced outcome. For example small sample size in the computer choice game, noise traders and unpredictable shocks to demand in the repeated production game . Thus

- less random uncertainty in the publically reported period outcomes is likely to reduce the number of unstable periods.

To reduce the number of unstable periods:

The number of unstable chaotic periods may be reduced by improving the information of the players in at least two ways:

First, we know from Sorin 1999, that the number of learning peri in which stability is violated, depends on the accuracy of th initial beliefs about the unknown fundamental. i.,e th grain of truth. Thus,

- starting with better public infor ntals is likely to reduce the nu

The other so of predictability of the periods' public example small sample size in the ertainty traders and unpredictable shocks to d production game . Thus

- ndom uncertainty in the publically reported period outcomes s likely to reduce the number of unstable periods.

Indeed, as we hear often form stock market analysts,

More educated traders and increased transparency in reported market information reduce instability

A question for future research:

What happens if the fundamental state changes, slowly, and you only observe random outcomes that depend on the changing state?

Need a mathematical theorem on Bayesian learning (merging), in a hidden but slowly changing Markov chain.

Thank you!