Confronting Theory with Experimental Data and vice versa
Lecture II: Social Learning

Academia Sinica

Mar 3, 2009

Background

- Social learning describes any situation in which individuals learn by observing the behavior of others.

- Several economic theories explain the existence of uniform social behavior:

  - benefits from conformity

  - sanctions imposed on deviants

  - network / payoff externalities

  - social learning.

The canonical model of social learning

- A set of players N, a finite set of actions $\mathcal{A}$, a set of states of nature Ω, and a common payoff function

$$U(a, \omega)$$

where $a \in \mathcal{A}$ is the action chosen and $\omega \in \Omega$ is the state of nature.

- Player i receives a private signal $\sigma_i(\omega)$, a function of the state of nature ω, and uses this private information to identify a payoff-maximizing action.

The canonical assumptions

– Bayes-rational behavior

– Incomplete and asymmetric information

– Pure information externality

– Once-in-a-lifetime decisions

– Exogenous sequencing

– Perfect information

Direct methodological extensions

- Caplin & Leahy (*AER* 1994), Chamley & Gale (*ECM* 1994)

- Avery & Zemsky (*AER* 1999), Chari & Kehoe (*JET* 2004)

- Çelen & Kariv (*GEB* 2004), Smith & Sørensen (2008)

- Bala & Goyal (*RES* 1998), Gale & Kariv (*GEB* 2004), Acemoglu et al. (2008)

The model of BHW (JPE 1992)

- There are two decision-relevant events, say A and B, equally likely to occur *ex ante* and two corresponding signals a and b.

- Signals are informative in the sense that there is a probability higher than 1/2 that a signal matches the label of the realized event.

- The decision to be made is a prediction of which of the events takes place, basing the forecast on a private signal and the history of past decisions.

- Whenever two consecutive decisions coincide, say both predict A, the subsequent player should also choose A even if his signal is different b.

- Despite the asymmetry of private information, eventually every player imitates her predecessor.

- Since actions aggregate information poorly, despite the available information, such herds / cascades often adopt a suboptimal action.

- Anderson and Holt (*AER* 1997) investigate the social learning model of BHW experimentally.

- They report that "rational" herds / cascades formed in most rounds and that about half of the cascades were incorrect.

- Extensions: Hung and Plott (*AER* 2001), Kübler and Weizsäcker (*RES* 2004), Goeree, Palfrey, Rogers and McKelvey (*RES* 2007).

The model of Smith and Sørensen (ECM 2000)

- Two phenomena that have elicited particular interest are *informational cascades* and *herd behavior*.

  - Cascade: players 'ignore' their private information when choosing an action.

  - Herd: players choose the same action, not necessarily ignoring their private information.

- Smith and Sørensen (2000) show that with a continuous signal space herd behavior arises, yet there need be no informational cascade.

The model of Çelen and Kariv (GEB 2004)

<u>Signals</u>

- Each player $n \in \{1, ..., N\}$ receives a signal θ_n that is private information.

- For simplicity, $\{\theta_n\}$ are independent and uniformly distributed on $[-1, 1]$.

<u>Actions</u>

- Sequentially, each player n has to make a binary irreversible decision $x_n \in \{0, 1\}$.

Payoffs

- $x = 1$ is profitable if and only if $\sum_{n \leq N} \theta_n \geq 0$, and $x = 0$ is profitable otherwise.

Information

- Perfect information

$$\mathcal{I}_n = \{\theta_n, (x_1, ..., x_{n-1})\}$$

- Imperfect information

$$\mathcal{I}_n = \{\theta_n, x_{n-1}\}$$

The decision problem

– The optimal decision rule is given by

$$x_n = 1 \text{ if and only if } \mathbb{E}\left[\sum_{i=1}^{N} \theta_i \mid \mathcal{I}_n\right] \geq 0.$$

Since $\mathcal{I}_n$ does not provide any information about the content of successors' signals, we obtain

$$x_n = 1 \text{ if and only if } \theta_n \geq -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid \mathcal{I}_n\right].$$

The cutoff process

– For any n, the optimal strategy is the *cutoff strategy*

$$x_n = \begin{cases} 1 & if \quad \theta_n \geq \hat{\theta}_n \\ 0 & if \quad \theta_n < \hat{\theta}_n \end{cases}$$

where

$$\hat{\theta}_n = -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid \mathcal{I}_n\right]$$

is the optimal history-contingent cutoff.

– $\hat{\theta}_n$ is sufficient to characterize the individual behavior, and $\{\hat{\theta}_n\}$ characterizes the social behavior of the economy.

A three-agent example

A three-agent example

A three-agent example under perfect information

A three-agent example under imperfect information

The case of perfect information

The cutoff dynamics follows the cutoff process

$$\hat{\theta}_n = \begin{cases} \frac{-1+\hat{\theta}_{n-1}}{2} & \text{if} \quad x_{n-1} = 1 \\ \frac{1+\hat{\theta}_{n-1}}{2} & \text{if} \quad x_{n-1} = 0 \end{cases}$$

where $\hat{\theta}_1 = 0$.

A sequence of cutoffs under perfect information

A sequence of cutoffs under perfect information

Informational cascades

- $-1 < \hat{\theta}_n < 1 \ \forall n$ so any player takes his private signal into account in a non-trivial way.

Herd behavior

- $\{\hat{\theta}_n\}$ has the martingale property by the Martingale Convergence Theorem a limit-cascade implies a herd.

The case of imperfect information

The cutoff dynamics follows the cutoff process

$$
\hat{\theta}_n = \begin{cases} -\dfrac{1+\hat{\theta}_{n-1}^2}{2} & \text{if} \quad x_{n-1} = 1 \\ \dfrac{1+\hat{\theta}_{n-1}^2}{2} & \text{if} \quad x_{n-1} = 0 \end{cases}
$$

where $\hat{\theta}_1 = 0$.

A sequence of cutoffs under imperfect and perfect information

A sequence of cutoffs under imperfect and perfect information

Informational cascades

- $-1 < \hat{\theta}_n < 1 \; \forall n$ so any player takes his private signal into account in a non-trivial way.

Herd behavior

- $\{\hat{\theta}_n\}$ is not convergent and the divergence of cutoffs implies divergence of actions.

- Behavior exhibits periods of uniform behavior, punctuated by increasingly rare switches.

Takeaways

- The dynamics of social learning depend crucially on the extensive form of the game.

- Longer and longer periods of uniform behavior, punctuated by (increasingly rare) switches.

- A succession of fads: starting suddenly, expiring easily, each replaced by another fad.

- Why do markets move from 'boom' to 'crash' without settling down?

The experiments of Çelen and Kariv (AER 2004, ET 2005)

- In market settings, we observe behavior but not beliefs or private information.

- In the laboratory, we can elicit subjects' beliefs and control their private information.

- Test the model's predictions and study the effects of variables about which our existing theory has little to say.

Perfect information
(Mean cutoffs when all predecessors acted alike)

Percentage of concurring, neutral and contrary decision points

The distribution of contrary subjects

Imperfect information
(Mean cutoffs in concurring decisions)

Imperfect information
(Mean cutoffs)

The econometric analysis

- At each decision turn n, with probability p_n a player is rationally, and with probability $1 - p_n$ he is noisy.

- The cutoff of a noisy player is a random draw from a distribution function G_n with support $[-1, 1]$ and mean $\tilde{\theta}_n$.

- Others cannot observe whether a player behavior is noisy, but the sequences $\{p_n\}$ and $\{G_n\}$ are common knowledge.

The estimated cutoff process

– The cutoff dynamics of rational players follow the process

$$\hat{\theta}_n = \hat{\theta}_{n-1} - \begin{cases} \dfrac{10+(1-p_{n-1})\tilde{\theta}_{n-1}+p_{n-1}\hat{\theta}_{n-1}}{2} & \text{if } x_{n-1}=A, \\ \dfrac{-10+(1-p_{n-1})\tilde{\theta}_{n-1}+p_{n-1}\hat{\theta}_{n-1}}{2} & \text{if } x_{n-1}=B, \end{cases}$$

where $\hat{\theta}_1 = 0$.

– The estimated parameters for the first decision-turn are employed in estimating the parameters for the second turn, and so on.

Sequences of cutoffs under perfect information
(Theory and estimated)

Sequences of cutoffs under perfect information
(Theory and estimated)

Localized conformity

- One of the most striking regularities of (human) society is localized conformity.

- In many social and economic situations, individuals are influenced by the decisions of others.

- The commonest examples occur in everyday life, as in choosing a fashionable restaurant or a popular movie.

- Similar influences also affect technology adoption and asset market decisions.

Why should individuals behave in this way?

N - finite set of agents

Ω - set of states of nature

$\sigma_i(\omega)$ - private signal

$\mathcal{A}$ - finite set of actions

$U(a, \omega)$ - common payoff function

- Incomplete and asymmetric information

  – Agents are uncertain about the event and the information about it is shared asymmetrically among them.

- Pure information externality

  – Each agent's payoff depends only on the state of nature and his own action.

Social learning

- Agents use their information to identify a payoff-maximizing action so the choice of action reflects that information.

- By observing an agent's action, it is possible to learn something about his information and make a better decision.

- In social settings, where agents can observe one another's actions, it is rational for them to learn from one another.

- *Social learning* occurs when individuals learn by observing the behavior of others.

What have we learned from Social Learning?

The striking uniformity of social behavior is an implication of social learning:

- Despite the asymmetry of information, agents rationally 'ignore' their own information and 'follow the herd'.

- Despite the available information, so-called *herd behavior* and *informational cascades* often result in an inefficient choice.

- Mass behavior is fragile, in the sense that small shocks may cause behavior to shift suddenly and dramatically.

The standard social-learning model

- The standard model has special features that are restrictive and deserve further examination:

  - Perfect information

  - Once-in-a-lifetime decisions

- The dynamics of social learning depend crucially on the extensive form of the game.

Sequential social-learning model:
Well heck, if all you smart cookies agree, who am I to dissent?

Imperfect information:
Which way is the wind blowing?!

The social network model

- Agents are bound together by a *social network*, a complex of relationships that brings them into contact with other agents.

- Markets are characterized by agents connected by complex, multilateral information networks.

- The network is represented by a family of sets $\{N_i\}$ where N_i denotes the set of agents $j \neq i$ who can be observed by agent i.

- Agents choose actions simultaneously and revise their decisions as new information is received.

Related literature

- The theoretical paper most closely related is Bala and Goyal (1998). The models differ in two ways:

    - Boundedly rational agents

    - Observing payoffs as well as actions.

- A model of *social experimentation* (with a multi-armed bandit) rather than social learning.

Outline

- Theory

  - The theoretical implications of the parametric model tested in the laboratory.

- Experiment

  - Interpret the data generated by the experiments using the theoretical framework.

Overview of results

- Theory

  - Learning occurs through diversity but is eventually replaced by uniformity.

  - Uniformity of actions is the robust feature of connected social networks.

  - The network architecture affects the efficiency and dynamics of social learning.

- Experiment

  - A variety of different networks and informational treatments provides a reasonable test of the theory.

  - The theoretical approach is indeed relevant for the interpretation of the data.

  - A structural econometric model based on the theory does quite well in explaining the experimental data.

Equilibrium

A weak perfect Bayesian equilibrium consists of a sequence of random variables $\{X_{it}\}$ and σ-fields $\{\mathcal{F}_{it}\}$ such that for each $i = 1, ..., n$ and $t = 1, 2, ...,$

(i) $X_{it} : \Omega \to \mathcal{A}$ is $\mathcal{F}_{it}$-measurable,

(ii) $\mathcal{F}_{it} = \mathcal{F}\left(\sigma_i, \{X_{js} : j \in N_i\}_{s=1}^{t-1}\right)$, and

(iii) $E[U(x(\omega), \omega)] \leq E[U(X_{it}(\omega), \omega)]$, for any $\mathcal{F}_{it}$-measurable function $x : \Omega \to \mathcal{A}$.

Asymptotic properties

- The welfare-improvement principle

  - Agents have perfect recall, so expected utility is non-decreasing over time. This implies that equilibrium payoffs form a submartingale.

- The imitation principle

  - In a connected network, asymptotically, all agents must get the same average (unconditional) payoffs.

Convergence Let $\{X_{it}, \mathcal{F}_{it} : i = 1, ..., n, t = 1, 2, ...\}$ be an equilibrium. For each i, define $V_{it}^* : \Omega \to \mathbf{R}$ by

$$V_{it}^* = E[U(X_{it}, \cdot)|\mathcal{F}_{it}].$$

Then $\{V_{it}^*\}$ is a submartingale with respect to $\{\mathcal{F}_{it}\}$ and there exists a random variable $V_{i\infty}^*$ such that V_{it}^* converges to $V_{i\infty}^*$ almost surely.

Connectedness Let $\{X_{it}, \mathcal{F}_{it}\}$ be the equilibrium and let V_{it}^* be the equilibrium payoffs. If $j \in N_i$ and j is connected to i then $V_{i\infty}^* = E[V_{j\infty}^* | \mathcal{F}_{i\infty}]$.

Imitation Let i and j be two agents such that $j \in N_i$ and j is connected to i. Let E^{ab} denote the measurable set on which i chooses a infinitely often and j chooses b infinitely often. Then $V_{i\infty}^a(\omega) = V_{i\infty}^b(\omega)$ for almost every ω in E^{ab}.

- Apart from cases of disconnectedness and indifference, diversity of actions is eventually replaced by uniformity.

- This is the network-learning analogue of the herd behavior found in the standard social learning model.

- The convergence properties of the model are general but many important questions about learning in networks remain open.

- Identify the impact of network architecture on the efficiency and dynamics of social learning.

A three-person example

- The network consists of three agents indexed by $i = A, B, C$. The neighborhoods $\{N_A, N_B, N_C\}$ completely define the network.

- Uncertainty is represented by two equally likely events $\omega = -1, 1$ and two corresponding signals $\sigma = -1, 1$.

- Signals are informative in the sense that there is a probability $\frac{2}{3}$ that a signal matches the event.

- With probability q an agent is informed and receives a private signal at the beginning of the game.

- At the beginning of each date t, agents simultaneously guess $a_{it} = -1, 1$ the true state.

- Agent i receives a positive payoff if his action $a_{it} = \omega$ and zero otherwise.

- Each agent i observes the actions a_{jt} chosen by the agents $j \in N_i$ and updates his beliefs accordingly.

- At date t, agent i's information set I_{it} consists of his private signal, if he observed one, and the history of neighbors' actions.

Complete   *Star*   *Circle*

A   A   A
C   B   C   B   C   B

A   A
C   B   C   B

Learning dynamics

- Learning is 'simply' a matter of Bayesian updating but agents must take account of the network architecture in order to update correctly.

- If all agents choose the same action at date 1, no further information is revealed at subsequent dates (an absorbing state).

- We can trace out possible evolutions of play when there is diversity of actions at date 1.

- The exact nature of the dynamics depends on the signals and the network architecture.

Complete network

$$N_A = \{B, C\}, N_B = \{A, C\}, N_C = \{A, B\}$$

	A	B	C
t/σ	1	0	0
1	1	-1	-1
2	-1	-1	-1
3	-1	-1	-1
4	-1	-1	-1
...	...	...	...

Star network

$$N_A = \{B, C\}, N_B = \{A\}, N_C = \{A\}$$

	A	B	C
t/σ	1	0	0
1	1	-1	-1
2	-1	1	1
3	1	1	-1
4	1	1	1
...	...	...	...

Circle network

$$N_A = \{B\}, N_B = \{C\}, N_C = \{A\}$$

	A	B	C
t/σ	1	0	0
1	1	-1	-1
2	1	-1	1
3	1	1	1
4	1	1	1
...	...	...	...

Takeaways

- Convergence to a uniform action tends to be quite rapid, typically occurring within two to three periods.

- Significant differences can be identified in the equilibrium behavior of agents in different networks.

- Even in the three-person case the process of social learning in networks can be complicated.

- Because of the lack of common knowledge, inferences agents must draw in order to make rational decisions are quite subtle.

Experimental design

- Each experimental session consisted of 15 independent rounds and each round consisted of six decision-turns.

- The network structure and the information treatment $(q = \frac{1}{3}, \frac{2}{3}, 1)$ were held constant throughout a given session.

- The ball-and-urn social learning experiments paradigm of Anderson and Holt (1997).

- A serious test of the ability of a structural econometric model based on the theory to interpret the data.

Selected data
(star network under high-information)

t/σ	A	B	C
	1	0	0
1	1	−1	−1
2	−1	1	1
3	−1	−1	−1
4	−1	−1	−1
5	1	1	−1
6	1	1	1

Herd behavior

Herd behavior is characterized by two related phenomena:

  – *Stability*: the proportion of subjects who continue to choose the same action.

  – *Uniformity*: a score function that takes the value 1 if all subjects act alike and takes the value 0 otherwise.

Uniformity will persist and lead to herd behavior if stability takes the value 1 at all subsequent turns.

Stability under high-information
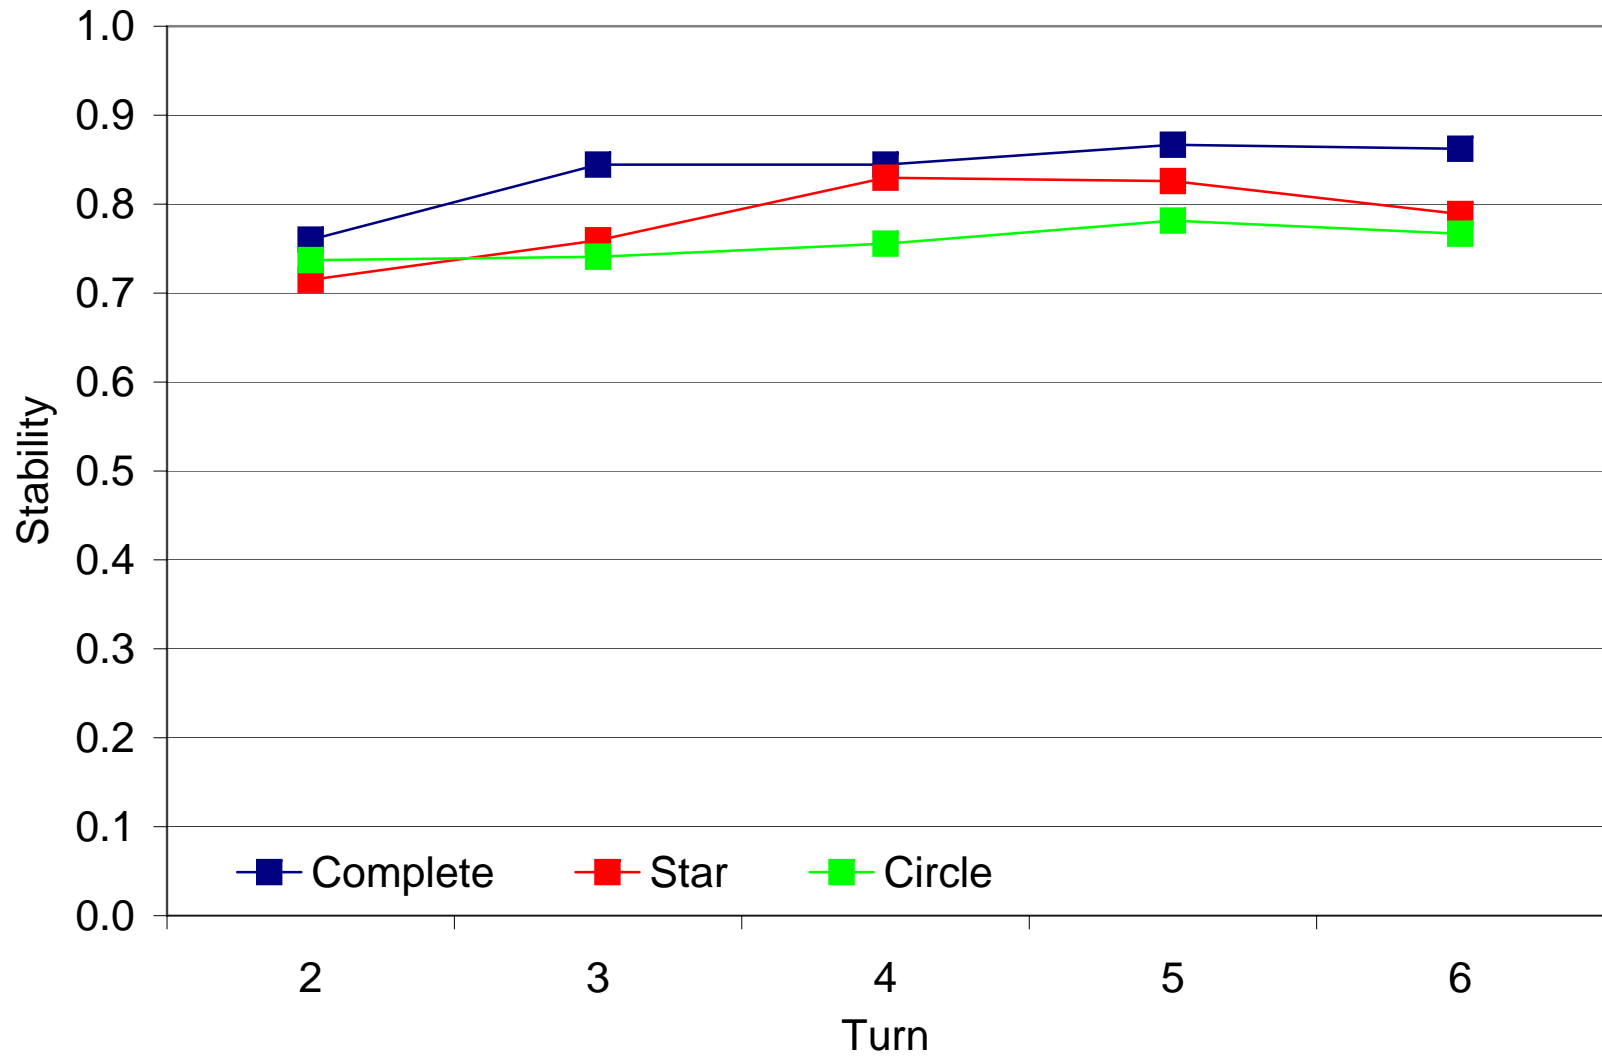
Uniformity under high-information

Herd behavior under high-information

Informational efficiency

The efficiency of decisions is measured in two ways:

- *actual efficiency* $= \dfrac{\pi_a - \pi_r}{\pi_e - \pi_r}$

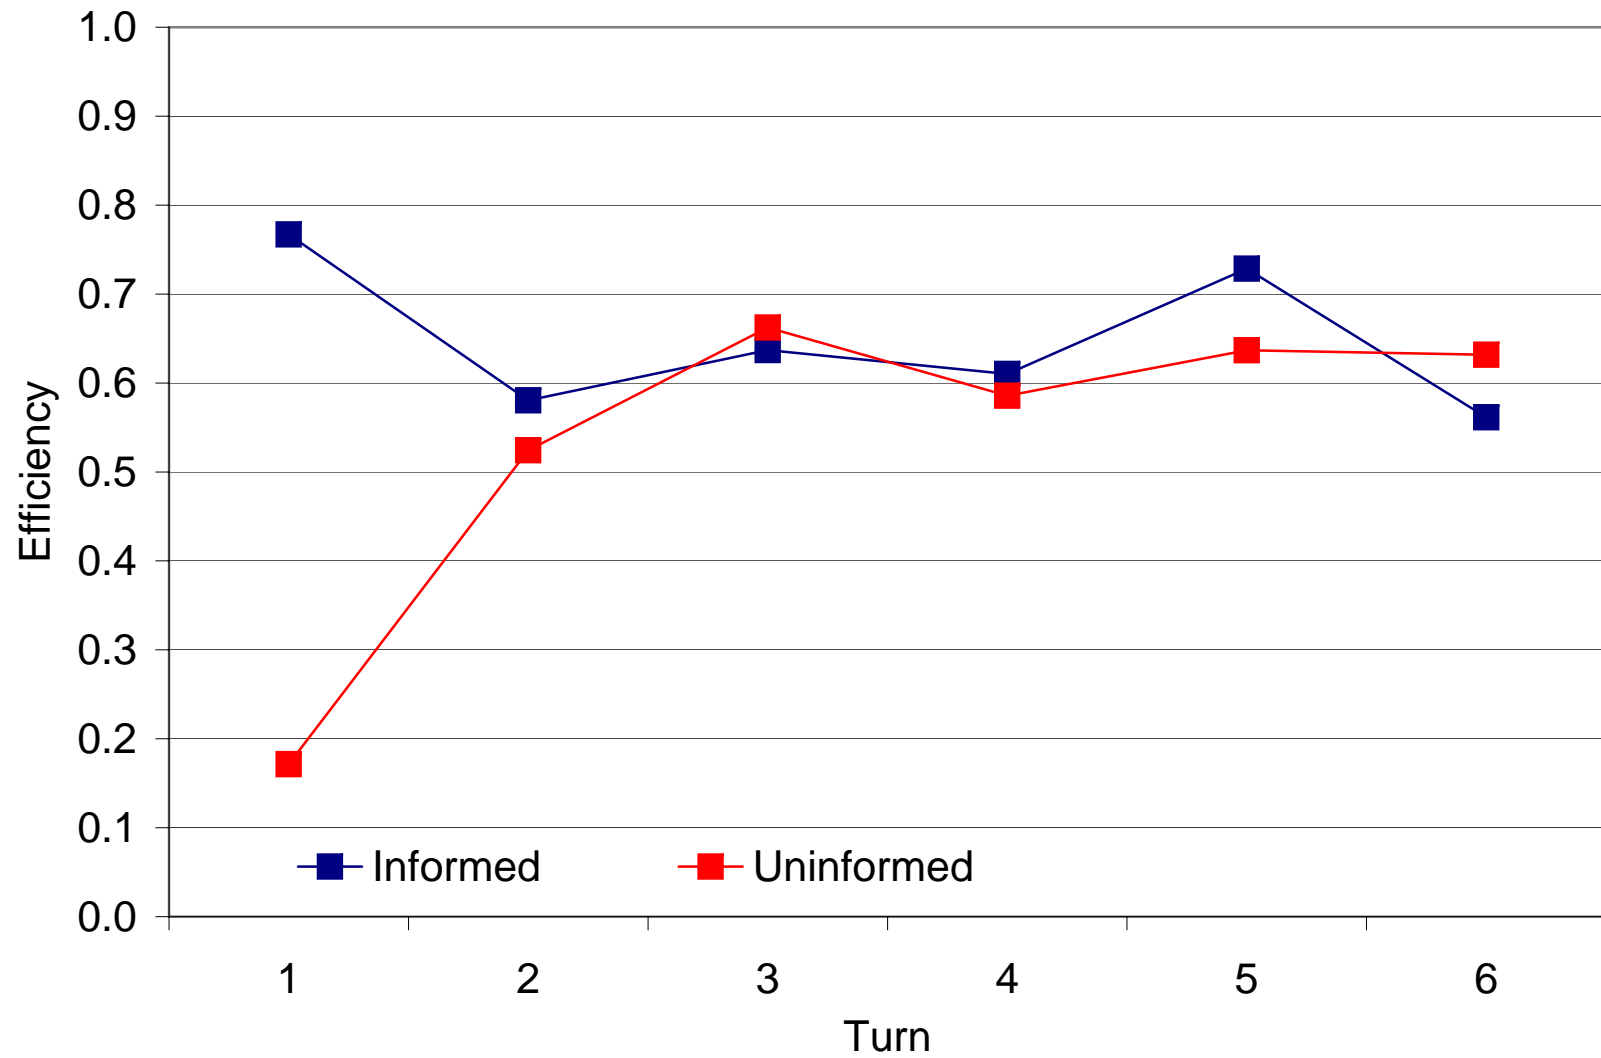- *private-information efficiency* $= \dfrac{\pi_p - \pi_r}{\pi_e - \pi_r}$

The net actual (private) efficiency $\pi_a - \pi_r \ (\pi_p - \pi_r)$ as a fraction of the net pooled efficiency $\pi_e - \pi_r$.

Efficiency under high-information

Efficiency in the circle network under high-information

Summary of experimental data

- A strong tendency toward herd behavior and a marked efficiency of information aggregation.

- There are significant differences between the behavior of different networks and information treatments.

- Differences might be explained by the symmetry or asymmetry of the network or the information treatment.

- There is some variation across networks and treatments but the error rates are uniformly fairly low.

Quantal response equilibrium (QRE)

- Mistakes are made and this should be taken into account in any theory of rational behavior.

- The payoff from a given action is assumed to be a weighted average of the theoretical payoff and a logistic disturbance.

- The "weight" placed on the theoretical payoff is determined by a regression coefficient.

- The recursive structure of the model enables to estimate the coefficients of the QRE model for each decision-turn sequentially.

The logit equilibrium can be summarized by a choice probability function following a binomial logit distribution:

$$\Pr\left(a_{it} = 1 | I_{it}\right) = \frac{1}{1 + \exp\left(-\beta_{it} x_{it}\right)}$$

where β_{it} is a coefficient and x_{it} is the difference between the expected payoffs from actions 1 and -1.

The regression coefficient β will be positive if the theory has any predictive power.

- Use the estimated coefficient from turn t to calculate the theoretical payoffs from the actions at turn $t + 1$.

- The behavioral interpretation is that subjects have rational expectations and use the true mean error rate.

- The parameter estimates are highly significant and positive, showing that the theory does help predict the subjects' behavior.

- A series of specification tests shows that the restrictions of the QRE model are confirmed by the data.

Beta = 2.591 , Bw = 0.339 , Complete (t=4 , high)

Beta = 2.566 , Bw = 0.326 , Circle (t=5 , high)

Beta = 2.565 , Bw = 0.354 , Star-type B&C (t=5 , high)

Beta = 3.118 , Bw = 0.407 , Star-type A (t=3 , high)

The beta time-series under full-information

The beta time-series under high-information

The beta time-series under low-information

The beta time-series in the circle network

Concluding remarks

- Use the theory to interpret data generated by experiments of social learning in three-person networks.

- The family of three-person networks includes several architectures, each of which gives rise to its own distinctive learning patterns.

- The theory, modified to include the possibility of errors, adequately accounts for large-scale features of the data.

- A strong support for the use of models as the basis for structural estimation and the use of QRE to interpret experimental data.