LEARNING BY SIMILARITY IN COORDINATION PROBLEMS

Jakub Steiner
Colin Stewart

CERGE-EI

Charles University
Center for Economic Research and Graduate Education
Academy of Sciences of the Czech Republic
Economics Institute

WORKING PAPER SERIES (ISSN 1211-3298)
Electronic Version

324

Working Paper Series 324
(ISSN 1211-3298)

Learning by Similarity in Coordination Problems

Jakub Steiner
Colin Stewart

CERGE-EI
Prague, April 2007

ISBN 978-80-7343-123-5 (Univerzita Karlova. Centrum pro ekonomický výzkum a doktorské studium)
ISBN 978-80-7344-112-8 (Národohospodářský ústav AV ČR, v.v.i.)

Learning by Similarity in Coordination Problems*

Jakub Steiner[†]
University of Edinburgh

Colin Stewart[‡]
Yale University

March 30, 2007

Abstract

We study a learning process in which subjects extrapolate from their experience of similar past strategic situations to the current decision problem. When applied to coordination games, this learning process leads to contagion of behavior from problems with extreme payoffs and unique equilibria to very dissimilar problems. In the long-run, contagion results in unique behavior even though there are multiple equilibria when the games are analyzed in isolation. Characterization of the long-run state is based on a formal parallel to rational equilibria of games with subjective priors. The results of contagion due to learning share the qualitative features of those from contagion due to incomplete information, but quantitatively they differ.

Abstrakt

Studujeme process učení v němž hráči extrapolují svou zkušenost z podobných minulých strategických situací k nynějšímu problému. V případě koordinačních her vede tento proces učení k šíření ustáleného chování z her s extrémními pravidly a nepřímo k selekci rovnovážných stavů i ve velmi rozdílných hrách. Z dlouhodobého hlediska je rovnovážný stav jedinečný, ačkoliv, kdyby byla každá hra zkoumaná odděleně, rovnovážných stavů by bylo několik. Při výpočtu rovnovážného stavu využíváme podobnost s koordinačními hrami s asymetrickou informací.

Keywords: Similarity, learning, contagion, case-based reasoning, global games, coordination, subjective priors.

---

*We are grateful to Philippe Jehiel, George Mailath, Stephen Morris, Ben Polak, Larry Samuelson, Avner Shaked, organizers of the VI Trento Summer School in Adaptive Economic Dynamics, and seminar participants at the University of Edinburgh, PSE Paris, Stanford University, Yale University, and the Econometric Society meetings in Minneapolis and Vienna. Jakub Steiner benefited from the grant "Stability of the Global Financial System: Regulation and Policy Response" during his research stay at LSE. While working on this paper, Jakub Steiner was also supported by research center grant No. LC542 of the Ministry of Education of the Czech Republic implemented at CERGE-EI—the joint workplace of the Center for Economic Research and Graduate Education, Charles University, Prague, and the Economics Institute of the Academy of Sciences of the Czech Republic.
[†]email: jsteiner@cerge-ei.cz
[‡]email: colin.stewart@yale.edu

1 Introduction

In standard models of learning, players repeatedly interact in the *same* game, and use their experience from the history of play to myopically optimize in each period. In many cases of interest, decision-makers are faced with *many different* strategic situations, and the number of possibilities is so vast that a particular situation is virtually never experienced twice. The history of play may nonetheless be informative when choosing an action, as previous situations, though different, may be similar to the current one. A tacit assumption of standard learning models is that players extrapolate from their experience of previous interactions similar to the current one.

The central message of this paper is that such extrapolation has important effects: *similarity-based* learning can lead to contagion of behavior across very different strategic situations. Two situations that are not directly similar may be connected by a chain of intermediate situations, along which each is similar to the neighboring ones. One effect of this contagion is to select a unique long-run action in situations that would allow for multiple steady states if analyzed in isolation. For this to occur, the extrapolations at each step of the similarity-based learning process need not be large; in fact, the contagion effect remains even in the limit as extrapolation is based only on increasingly similar situations.

We focus here on the application of similarity-based learning to coordination games. Consider, as an example, the class of 2×2 games $\Gamma(\theta)$ in Table 1 parameterized by a fundamental, θ. Action I, interpreted as investing, is strategically risky, as its payoff depends on the action of the opponent. The safe action, NI, gives a constant payoff of 0. For extreme values of θ, the game $\Gamma(\theta)$ has a unique equilibrium as investing is dominant for $\theta > 1$, and the safe action is dominant for $\theta < 0$. When θ lies in the interval $(0, 1)$, the game has two strict pure strategy equilibria.

The contagion effect can be sketched without fully specifying the learning process, which we postpone to Section 3. Two myopic players interact in many rounds in a game $\Gamma(\theta_t)$, with θ_t selected at random in each round. Roughly, we assume that players estimate payoffs for the game $\Gamma(\theta)$ on the basis of past experience with fundamentals similar to θ, and that two games $\Gamma(\theta)$ and $\Gamma(\theta')$ are viewed by players as similar if the difference $|\theta - \theta'|$ is small.

	I	NI
I	θ, θ	$\theta - 1, 0$
NI	$0, \theta - 1$	$0, 0$

Table 1: Payoffs in the Example of Section 2.

Since investing is dominant for all sufficiently high fundamentals, there is some $\overline{\theta}$ above which players eventually learn to invest. Now consider a fundamental just below $\overline{\theta}$, say $\overline{\theta} - \varepsilon$. At $\overline{\theta} - \varepsilon$, investing may not be dominant, but players view some games with values of θ above $\overline{\theta}$ as similar. Since the opponent has learned to invest in these games, strategic complementarities in payoffs increase the gain from investing. When ε is small, this increase outweighs the potential loss from investing in games below $\overline{\theta}$, where the opponent may not invest. Thus players learn to invest in games with fundamentals below, but close to $\overline{\theta}$, giving a new threshold $\overline{\theta}'$ above which both players invest.

Repeating the argument with $\overline{\theta}$ replaced by $\overline{\theta}'$, investment continues to spread to games with smaller fundamentals, even though these are not directly similar to games in the dominance region. The process continues until a threshold fundamental θ is reached at which the gain from investment by the opponent above θ is exactly balanced by the loss from non-investment by the opponent below θ. Not investing spreads contagiously beginning from low values of the fundamental by a symmetric process. Together, these processes meet at the same threshold, giving rise to a unique long-run outcome, provided that similarity drops off quickly in distance.[1]

Contagion effects have previously been studied in local interaction and incomplete information games. In local interaction models, actions may spread contagiously across members of a population because each member has an incentive to coordinate with her neighbors in a social network (e.g. Morris (2000)). In incomplete information games with strategic complementarities (global games), actions may spread contagiously across types because private information gives rise to uncertainty about the actions of other players (Carlsson and van Damme 1993). Unlike these models, contagion through learning depends neither on any network structure nor on high orders of reasoning about the beliefs of other players. The

[1]In other words, players place much more weight on values of the fundamental very close to the present one when forming their payoff estimates.

contagion is driven solely by a natural solution to the problem of learning one's own payoffs when the strategic situation is continually changing. This problem is familiar from econometrics, where one often wishes to estimate a function of a continuous variable using only a finite data set. The similarity-based payoff estimates used by players in our model have a direct parallel in the use of kernel estimators by econometricians. Moreover, the use of such estimates for choosing actions is consistent with the case-based decision theory of Gilboa and Schmeidler (2001), who propose similarity-weighted payoff averaging as a general theory of decisions under uncertainty.

While the learning model we have described is one of complete information, the same reasoning applies when, as in the global game model, players imperfectly observe the value of the fundamental. In order to directly compare the process of contagion through learning to that from incomplete information, players in the general model of Section 3 observe private signals of the fundamental that may be noisy. The fundamental and signals are independently drawn in each round. From the history of play, players have experience with realized payoffs for signals similar to, but different from, their current signal. They estimate the current payoffs based on the payoffs of similar types in the past.

The main tool for understanding the result of contagion through learning is a formal parallel to rational play in a modified version of the game. This modified game differs from the original game only in the priors: players eventually behave *as if* they incorrectly believe their own signal to be noisier than it actually is, while holding correct beliefs about the precision of the other players' signals. More precisely, players learn not to play strategies that would be serially dominated in the modified version of the game (see Theorem 3.1).

This result enables us to solve the modified game by extending the techniques of Carlsson and van Damme (1993), further developed by Morris and Shin (2003). With complete information, the original game has a continuum of equilibria, but contagion leads to a unique learning outcome when similarity is concentrated on nearby fundamentals. With small noise in observations of the fundamental, the underlying game has a unique equilibrium as a result of contagion from incomplete information. In this case, there is also a unique learning outcome when similarity is concentrated, but this outcome depends on the relative size of the noise compared to the concentration of the similarity. In particular, the process of contagion through learning does not generally coincide with that of contagion from incomplete information. However, the qualitative features of these processes agree, as both converge to play of symmetric threshold strategies, and give rise to comparative statics of the same sign.

After an illustrative example in the following section, Section 3 describes the general learning model, and characterizes its long-run behavior in terms of the modified game. Section 4 fully identifies the long-run state in the limit of small noise and narrow similarity distributions, and examines comparative statics. Section 5 reviews the related literature.

2 Example

This section presents an example to illustrate in more detail the process of contagion through learning before describing the general model in Section 3.

The underlying family of coordination problems consists of the 2-player games in Table 1. We denote by $u(a^i, a^{-i}, \theta)$ the payoff to choosing action a^i when the opponent chooses action a^{-i}, and the fundamental is θ.

The game is played repeatedly in periods $t \in \mathbb{N}$, with the fundamental θ_t drawn independently across periods according to a uniform distribution on an interval $[-b, 1+b]$, where $b > 0$. Each realization θ_t is perfectly observed by both players, who play a myopic best response to their beliefs in each period. Beliefs are based on players' previous experience, but since θ is drawn from a continuous distribution, players (almost surely) have no past experience with the current game $\Gamma(\theta_t)$, and must extrapolate from their experience playing different games. In each period, players directly estimate payoffs as a weighted average of historical returns in which the weights are determined by the similarity between the current and past fundamentals. Strategic considerations play no role in these estimates: players treat the past actions of their opponents as given. Thus following any history $\left\{\theta_s, a_s^1, a_s^2\right\}_{s<t}$, the estimated payoff to player i from choosing action a^i given the fundamental θ_t is

$$\frac{\sum_{s<t} g(\theta_s - \theta_t) u(a^i, a_s^{-i}, \theta_s)}{\sum_{s<t} g(\theta_s - \theta_t)}, \tag{1}$$

where $g(\cdot)$ is the *similarity function* determining the relative weight assigned to past cases.

For this example, suppose that $g(\cdot)$ is the density corresponding to a uniform distribution on the interval $[\tau\frac{c-1}{2}, \tau\frac{c+1}{2}]$, where $c \in [-1, 1]$ and $\tau \in (0, b]$. Beliefs may be chosen arbitrarily if the history contains no fundamental similar to θ_t, that is, if $\sum_{s<t} g(\theta_s - \theta_t) = 0$.

The learning process is stochastic, but suppose that the empirical distribution of realized cases may be approximated by the probability distribution over θ (this idea is formalized in Section 3 below). By focusing on the most extreme strategies remaining for the opponent at each stage of the learning process—those involving investment at the most or the fewest fundamentals—we may bound the payoff estimates independently of the precise evolution of the opponent's strategy. Accordingly, consider a fixed strategy $l : \Theta \to \{0, 1\}$ of the opponent, where 1 is associated with investing and 0 with the safe action. Upon observing the fundamental θ, the player forms estimates of the true payoffs to investing $u(1, l(\theta), \theta) = \theta + l(\theta) - 1$, and to the safe action $u(0, l(\theta), \theta) = 0$. Similarity-based learning leads to payoff estimates

$$\int_{\Theta} \left(\theta' + l(\theta') - 1\right) g(\theta' - \theta) d\theta' \tag{2}$$

from choosing to invest, and 0 from the safe action. The expression (2) is formally equivalent to the conditional expectation $E[\Theta' + l(\Theta') - 1|\theta]$ when θ is an imprecise signal of θ', with noise distributed according to density $g(\cdot)$. Thus, in the longrun, the similarity-based learner behaves *as if* her observation of θ were not the true fundamental, but only a noisy signal.

Let $\tilde{\theta}(\theta) = \int_{\Theta} \theta' g(\theta' - \theta) d\theta'$ denote the posterior expected value $E[\Theta'|\theta]$ of the fundamental after observing the signal θ under this "virtual signal" interpretation of the payoff estimates. Players (eventually) learn to invest at those values of θ for which $\tilde{\theta}(\theta)$ lies above 1, because the estimated payoff is positive even if the opponent has never invested.

Next, consider some θ for which $\tilde{\theta}(\theta) = 1 - \alpha$, with $\alpha > 0$ small relative to τ. Suppose that a sufficiently long time has passed since the completion of the first learning stage as to make this earlier history negligible. Close to half of the similarity weight at this θ will be assigned to past cases θ' with $\tilde{\theta}(\theta') > 1$. In these cases, the opponent always invests, causing the estimated benefit to investing in these cases to outweigh the estimated loss for smaller values of the fundamental where the opponent may not invest. Players therefore learn to invest at fundamentals θ with $\tilde{\theta}(\theta) > 1 - \overline{\alpha}$, for some $\overline{\alpha} > 0$.

Following this second learning stage, we may apply the same argument for fundamentals θ with $\tilde{\theta}(\theta) = 1 - \overline{\alpha} - \alpha$ for small $\alpha > 0$. Iterating the argument in this way, players eventually learn to invest for ever lower fundamentals. The process continues until a threshold $\overline{\theta}$ is reached at which the gain in the estimated return to investing due to the opponent investing above $\overline{\theta}$ is exactly offset by the loss in this return if the opponent chooses the safe action below $\overline{\theta}$.

The same reasoning applies to the safe action beginning from low fundamentals, giving rise to a threshold $\underline{\theta}$ below which both players choose the safe action. The threshold $\underline{\theta}$ satisfies the same offsetting-payoff condition as $\overline{\theta}$. Since the virtual estimate $\tilde{\theta}(\theta)$ is increasing in θ, this condition is satisfied for a unique fundamental θ^*. Thus the two contagion processes meet at the same threshold $\overline{\theta} = \underline{\theta} = \theta^*$. At this threshold, players are indifferent between their two actions given their long-run payoff estimates. Thus we have

$$\tilde{\theta}\left(\theta^*\right) - 1 + \int_{\Theta} l_{\theta^*}(\theta')g(\theta' - \theta^*)d\theta' = 0,$$

where $l_{\theta^*}(\theta')$ is the threshold strategy with threshold θ^*.

The threshold θ^* depends on the shape of the similarity function. The threshold type's estimate of the likelihood that her opponent invests is equal to the similarity weight the type assigns to higher fundamentals: in this case, $\frac{c+1}{2}$. The long-run threshold therefore solves $\tilde{\theta}\left(\theta^*\right) + \frac{c-1}{2} = 0$.

In Section 3, we introduce the general model of similarity-based learning. In addition to more general payoff functions, and general similarity functions $g^i(\cdot)$, we allow for incomplete information in the observation of the fundamental. Long-run behavior is influenced by both the true error in players' signals and the virtual error arising from the use of extrapolation in similarity-based learning.

3 The Learning Model

The model is comprised of an *underlying game*, which shares much of the structure of global games, together with a dynamic process by which players form beliefs about their payoffs as a function of the observed history. We begin by describing the underlying game.

Two players play a common value game Γ_σ. In this game, a state θ is drawn from a connected space $\Theta \subseteq \mathbb{R}$ according to the continuous distribution $\Phi(\cdot)$ with density $\phi(\cdot)$. Each player i then receives a possibly noisy signal $x^i \in X^i \subseteq \mathbb{R}$ of the state θ given by $x^i = \theta + \sigma\epsilon^i$, where ϵ^i is drawn from a continuous distribution with density $f^i(\cdot)$. These draws are independent across players. The parameter σ governing the precision of the signals is assumed to be nonnegative; in particular, we consider not only the incomplete information case of $\sigma > 0$, but also the complete information case of $\sigma = 0$ in which the state θ becomes common knowledge before the players choose their actions. Letting $p\left(\theta, x^1, x^2\right)$ denote the probability density associated with the combination $\left(\theta, x^1, x^2\right)$ when $\sigma > 0$, we have[2]

$$p\left(\theta, x^1, x^2\right) = \phi(\theta) f^1\left(\frac{x^1 - \theta}{\sigma}\right) f^2\left(\frac{x^2 - \theta}{\sigma}\right).$$

Each player has two actions, 0 and 1 (these correspond, respectively, to the actions NI and I in the Introduction above). Payoffs depend only on the state θ and the action profile. To economize on notation, we normalize the payoff from action 0 to be equal to 0 in every state θ, and write $u(\theta, l)$ for the expected payoff from choosing action 1 in state θ when the opponent chooses action 1 with probability l.[3] More generally, $u(\theta, l)$ represents the *difference* in payoffs from choosing action 1 instead of action 0 given θ and l.

We place the following restrictions on the payoffs throughout:

A1. $u(\theta, l)$ is increasing in θ.

A2. $u(\theta, l)$ is non-decreasing in l.

A3. Uniform limit dominance: there exists some $\overline{\theta}$ and $\varepsilon > 0$ such that $u(\theta, l) > \varepsilon$ whenever $\theta \geq \overline{\theta}$ and $u(\theta, l) < -\varepsilon$ whenever $\theta \leq -\overline{\theta}$ (for all $l \in [0, 1]$).

A4. Bounded payoffs: there exists some $V \in \mathbb{R}$ such that $|u(\theta, l)| < V$ uniformly for all $(\theta, l) \in \Theta \times [0, 1]$.

Assumptions A2 and A3 are quite standard in global games (see, e.g., Morris and Shin (2003)).

[2]When $\sigma = 0$, we have $x^1 = x^2 = \theta$, and the density is simply $\phi(\theta)$.

[3]The reason for defining payoffs so as to allow a non-linear dependence on the probability distribution over the opponent's action is to facilitate the move to a model with a continuum of players, in which l represents the share of the population choosing action 1.

For simplicity, Assumption A1 strengthens the usual nondecreasing payoffs assumption. This additional strength is needed only to guarantee that the function $\tilde{m}(x, x)$ defined below possesses a unique root; the corresponding uniqueness condition is assumed directly in the earlier literature.

The learning process is based on the idea that players estimate their possible payoffs based on play in past situations, with more similar situations being assigned greater weight in this estimate. More precisely, fixing $\sigma \geq 0$, we suppose that the game Γ_σ is played in each period $t = 1, 2, \ldots$, with the fundamental and signals drawn independently across periods. Let θ_t, x_t^i, and a_t^i denote, respectively, the payoff-relevant state, player i's signal, and player i's action in period t.

Regardless of her own action, each player i learns at the end of each period t the payoff $u\left(\theta_t, a_t^j\right)$ that she received, or would have received, from choosing action 1. The assumption that players learn counterfactual payoffs from actions that they have not chosen simplifies the analysis by ensuring that initial beliefs do not prevent players from learning. We offer two interpretations of this assumption. In certain applications, the counterfactual payoff may be directly observable from public reports in the media. Thus, for example, in a currency attack, even those who have not participated learn about the outcome of the attack. Alternatively, one may suppose that in each period, players have a small but fixed probability of choosing their action at random, independently of the history of play, either by mistake or for the purpose of experimentation. As this error probability becomes small, the long-run outcomes approach those of our model.

Each player i is endowed with a similarity function $g^i : \mathbb{R} \longrightarrow \mathbb{R}$ that depends only on the difference between two types, so that the weight placed by type x on experience as type x' is given by $g^i(x' - x)$. The similarity function is assumed to be non-negative everywhere and integrable, and we will normalize it to be a probability density function. Following a history $h_t = \left(\theta_s, x_s^1, x_s^2, a_s^1, a_s^2\right)_{s=1,\ldots,t}$, type x^i of player i forms the *estimated return* to action 1 by

$$r\left(x^i; h_t\right) := \frac{\sum_{s=1}^t u\left(\theta_s, a_s^{-i}\right) g^i\left(x_s^i - x^i\right)}{\sum_{s=1}^t g^i\left(x_s^i - x^i\right)} \tag{3}$$

whenever the denominator on the right-hand side is non-zero. Player i chooses action 1 in period t if and only if this estimated return is positive. This formulation captures the notion that players form estimates of payoffs based on their experience with similar types, and, as usual in the literature on learning in games, behave myopically based on these estimates. We place no restrictions on behavior when $\sum_{s=1}^{t} g^i\left(x_s^i - x^i\right) = 0$,[4] and all of our results hold for any specification of behavior or beliefs at these histories.

This learning process is a form of case-based decision theory, as formulated by Gilboa and Schmeidler (2001). Alternatively, the model has a cognitive interpretation based on Billot et al. (2005), who describe and axiomatize a belief formation process according to which a statistician estimates the probability of an outcome on the basis of its frequency among previous cases, where these cases are weighted by their similarity to the present one. Our players can be viewed as statisticians satisfying the axioms of Billot et al., who, after forming beliefs, maximize their expected payoffs.[5] The key axioms in both Gilboa and Schmeidler (2001) and Billot et al. (2005) preclude learning of the similarity function, which is consistent with our modelF in which similarity is exogenous.

The informational requirements of the learning process are modest. In particular, players need not have any initial knowledge of their own payoff function, nor must they observe their opponent's actions, payoffs, or types. It is even possible for players to follow this process without knowing that they are involved in strategic interaction, as they are simply forming estimates of the optimal action based on their own payoff history in similar situations.

To simplify the analysis, the payoff estimates $r\left(x^i; h^t\right)$ place equal weight on all past observations regardless of how much time has elapsed. More generally, one could suppose that observations are discounted over time according to a non-increasing sequence $\delta(\tau) \in (0,1]$ by modifying equation (3) to include an additional factor of $\delta(t-s)$ in each sum. In the undiscounted model, the convergence results presented below rely on the property that changes in payoff estimates in a single period become negligible once players have accumulated enough experience. Since this property continues to hold as long as the series $\sum_{\tau=0}^{\infty} \delta(\tau)$ diverges, we conjecture that all of our results hold in this more general setting. If, on the other hand, this sum converges, then the situation becomes more complicated, as the learning

[4]That is, when the history of play contains no cases similar to the present one.

[5]This connection is subject to the caveat that Billot et al. assume a finite outcome space, whereas the outcome space is infinite here.

process will not converge in general. It is therefore not possible for the long-run behavior to agree with that of the undiscounted process in every period. However, as long as memory is "sufficiently long," we expect this agreement to occur in a large fraction of periods. For example, if memory is discounted exponentially, so that $\delta(\tau) = \rho^\tau$ for some $\rho \in (0, 1)$, then we expect play to be consistent with our results most of the time when ρ is close to 1.

The following technical assumptions are required for the analysis:

A5. Each similarity function $g^i(\cdot)$ is bounded by some M^i.

A6. Each similarity function $g^i(\cdot)$ is uniformly continuous.[6]

Since $g^i(x)$ is a probability density function, Assumption A6 implies that similarity tends to zero in distance; that is, $\lim_{x\to\infty} g^i(x) = \lim_{x\to-\infty} g^i(x) = 0$. Otherwise, there must exist some $\epsilon > 0$ such that for each M there is some $x > M$ for which $g^i(x) > \epsilon$. By the uniform continuity of $g^i(\cdot)$, this implies that there exist infinitely many disjoint intervals of fixed length on which $g^i(\cdot)$ is everywhere greater than $\frac{\epsilon}{2}$, contradicting that $g^i(\cdot)$ is integrable.

Let $p^i_x(\cdot)$ denote the marginal density corresponding to the distribution of the signal of player i. With incomplete information, we have $p^i_x(x) = \int_\Theta \phi(\theta) f^i\left(\frac{x-\theta}{\sigma}\right) d\theta$, and with complete information, $p^i_x(x) = \phi(x)$.

A7. The marginal densities $p^i_x(\cdot)$ are continuous.

Note that in the incomplete information case, Assumption A7 allows for discontinuities in the densities $\phi(\cdot)$ and $f^i(\cdot)$; for example, this assumption holds if the discontinuities of both of these densities are topologically isolated.

In order to ensure that learning occurs everywhere in finite time, we assume:

A8. Either the state and type spaces are compact, or each similarity function $g^i(\cdot)$ has full support on the real line.

3.1 Long-run Characterization

The learning process described above converges, in a sense that will be made precise below, to the set of strategies surviving IEDS in a game with subjective priors that we will refer to

[6] All of the results hold if instead we suppose that for each i, there exists some x^i below which g^i is non-decreasing, and above which it is non-increasing.

as the modified game. Whereas the underlying game describes the actual situation in which the players interact, the modified game describes a virtual situation in which rational players would exhibit the same behavior as the learning players of our model (in the longrun).

In order to motivate the formulation of beliefs in the modified game, consider the incomplete information case ($\sigma > 0$). Recall that under the specified learning dynamics, behavior is determined by the sign of the numerator of the estimated return in (3). Against a fixed strategy $a^j(x')$ of the opponent, the expected value of this numerator is proportional to

$$\frac{\int_{X^i}\int_{X^j}\int_{\Theta} u\left(\theta, a^j(x')\right) g^i(x-x^i)p(\theta,x,x')d\theta dx' dx}{\int_{X^i}\int_{X^j}\int_{\Theta} p(\theta,x,x')g^i(x-x^i)d\theta dx' dx}$$
$$= \int_{X^j}\int_{\Theta} u\left(\theta, a^j(x')\right) q^i(\theta, x'|x^i)d\theta dx', \quad (4)$$

where

$$q^i\left(\theta, x^j|x^i\right) = \frac{\int_{X^i} p\left(\theta,x,x^j\right) g^i\left(x-x^i\right)dx}{\int_{\Theta}\int_{X^j}\int_{X^i} p\left(\theta,x,x'\right) g^i\left(x-x^i\right) dx dx' d\theta}. \quad (5)$$

Note that the right-hand side of (4) is precisely the expected payoff to player i from playing action 1 given the posterior beliefs $q^i(\theta, x'|x^i)$, suggesting that the long-run behavior under the learning dynamics should correspond to rational behavior given these subjective beliefs.

The *modified game* is identical to the underlying game, except that the beliefs of type x^i of player i are given in the incomplete information case by $q^i\left(\theta, x^j|x^i\right)$. In the complete information case ($\sigma = 0$), the beliefs of each type x^i assign probability one to the event $x^j = \theta$, and correspond to the density

$$q^i\left(\theta|x^i\right) = \frac{\phi(\theta)g^i\left(\theta - x^i\right)}{\int_{\Theta}\phi(\theta)g^i\left(\theta-x^i\right)d\theta}. \quad (6)$$

An equivalent definition of the modified game specifies the subjective priors from which the posterior beliefs q^i may be derived. These priors correspond to an incorrect model of signal formation on the part of each player. To be precise, consider a model in which, after θ is drawn according to the correct distribution, player i's signal is formed in a two-stage process. The first stage of this process is the same as for the signal in the underlying game; that is, a noisy signal $\tilde{x} = \theta + \sigma\tilde{\epsilon}$ of θ is generated by drawing $\tilde{\epsilon}$ according to the density

$f^i(\cdot)$. What player i observes, however, is a noisy signal x of $\tilde{x}$ drawn according to the density $g^i(\tilde{x} - x)$. The beliefs q^i correspond to this two-stage process when player i holds the correct beliefs about her opponent's signal, namely that it arises from only the first stage of this process. The effect of learning by similarity in the long-run may therefore be viewed as if players add noise to their own signals, but not to that of their opponents. This interpretation explains in part why many of the results discussed below are close, but not identical to those of the standard global games literature. In particular, when there is complete information in the underlying global game, this form of subjective noise may lead to a unique equilibrium in the same way that adding small noise does in global games with rational players.

Given any game with subjective priors, we may define (interim) dominated strategies in the same way as for Bayesian games with common priors.[7] In fact, we will require a slightly stronger form of dominance in which the payoff difference exceeds some fixed $\pi \geq 0$. To define this formally, let $u^i\left(\theta, a^i, a^{-i}\right)$ denote the payoff to player i from the action profile $\left(a^i, a^{-i}\right)$ when the fundamental is θ. The action $a^i \in \{0,1\}$ is π*-dominated for type* x^i against a set S^{-i} of strategies for the opponent if there exists some other action $\tilde{a}^i$ such that

$$E_{q^i(\theta,x^{-i}|x^i)} u^i\left(\theta, \tilde{a}^i, s^{-i}\left(x^{-i}\right)\right) - E_{q^i(\theta,x^{-i}|x^i)} u^i\left(\theta, a^i, s^{-i}\left(x^{-i}\right)\right) > \pi$$

for all $s^{-i} \in S^{-i}$.[8] In words, the expected payoff of type x^i based on her posterior beliefs could be increased by more than π by playing a different action, regardless of the strategy of the opponent. We call a strategy π-dominated for player i if it specifies a π-dominated action for some type. As usual, we will say simply that s^i is dominated if it is π-dominated with $\pi = 0$.

The need to consider π-domination instead of ordinary strict domination arises because of the difference between estimated returns following finite histories and their long-run expectations. In the proof of Theorem 3.1 below, we show that for any $\pi > 0$, estimated payoffs under the learning process almost surely eventually lie within π of the corresponding expected

[7] Since no other notion of domination will be employed here, we henceforth drop the term "interim" and refer simply to "dominated strategies."

[8] The notion of π-domination should not be confused with the unrelated concept of p-dominance that has appeared in the literature on higher-order beliefs.

payoffs in the modified game. It follows that actions that are π-dominated in the modified game will (almost surely eventually) not be played under the learning process. The following lemma shows that considering π-domination for arbitrary $\pi > 0$ suffices to prove the result for $\pi = 0$, that is, for strict domination.

Lemma 3.1. *Suppose that action $a \in \{0,1\}$ is serially dominated for type $\overline{x}$ of player i. Then there exists some $\pi > 0$ such that action a is serially π-dominated for type $\overline{x}$.*

The idea of the proof, which is relegated to the appendix, is that as π is made to decrease toward zero, smaller sets of strategies survive iterated elimination of π-dominated strategies (IEπDS). The lemma states that the set obtained in the limit is equal to that from ordinary IEDS. Suppose this is not the case, and consider any type x for which they differ. The argument proceeds by induction on the round in which elimination occurs under ordinary IEDS for this type, call it N. We show that by choosing π sufficiently small, the set of types on which IEπDS differs from IEDS in the first $N-1$ rounds can be made to have an arbitrarily small measure, and hence this difference has an arbitrarily small impact on the possible expected payoffs received by type x.

Given any subset $\alpha \subseteq \{0,1\}$, let $X^i(\alpha) \subseteq X^i$ denote the set of types of player i for which α is precisely the set of serially undominated actions in the modified game. The main result of this section, given in the following theorem, shows that, in the longrun, players will not play strategies that are serially dominated in the modified game.

Theorem 3.1. *The probability that play under the specified dynamics is consistent with IEDS in the modified game approaches one as time tends to infinity. Moreover, on any compact set of types not intersecting $X(\{0,1\})$, convergence almost surely occurs in finite time.*

Proof. See appendix. □

Using the strong law of large numbers, it is relatively straightforward to show that for a given type against a fixed strategy, the long-run payoff estimate is equal to the expected payoff in the modified game. The main difficulty in the proof of the preceding theorem arises because, in order for the analogue of IEDS to occur under the learning dynamics, infinite sets of types must "eliminate" actions in finite time. Accordingly, the proof demonstrates that it is possible to reduce the problem to one involving a finite state space while introducing only an arbitrarily small error in the payoff estimates.

4 Limit Results and Comparative Statics

4.1 Narrow Similarity and Small Noise

By Theorem 3.1, applying IEDS in the modified game allows us to identify strategies that may survive in the long-run of the learning process. We therefore shift our attention in this section to the solution of the modified game.

¿From this point on, we focus on the case in which the densities $\phi(\cdot)$ and $f^i(\cdot)$ have full support on the real line, and assume that the game and the learning process are symmetric with respect to players; that is, $f^1(\cdot) = f^2(\cdot)$ and $g^1(\cdot) = g^2(\cdot)$. Since it follows that the subjective beliefs in the modified game take the same form, we drop the player index from $q^i(\cdot)$. Like the underlying game, the modified game does not have a unique equilibrium in general. However, the techniques developed for global games with rational players (see, e.g., Morris and Shin (2003)) can be extended to show that uniqueness arises as long as the noise and the similarity weights are both sufficiently concentrated on a narrow interval. In order to make this precise, we introduce a similarity parameter $\tau \in \mathbb{R}_{++}$, and replace the similarity function $g(x' - x)$ with $\frac{1}{\tau} g\left(\frac{x'-x}{\tau}\right)$. Decreasing τ increases the similarity weight given to types x' close to x.

The proof of the following proposition closely follows that of the corresponding result in Morris and Shin (2003). They show that in the limit as σ tends to zero, the essentially unique serially undominated strategy for each player is defined by the threshold θ^* solving $\int_0^1 u(\theta, l)dl = 0$, with action 0 taken by types below θ^*, and action 1 by types above θ^*. The following result identifies the long-run solution under the learning dynamics in the limit with both small noise ($\sigma \to 0$) and narrow similarity ($\tau \to 0$), while holding the ratio $\frac{\sigma}{\tau}$ fixed.

Proposition 4.1. *For any $\delta > 0$, there exists $\overline{\gamma} > 0$ such that for any $\gamma \in (0, \overline{\gamma})$, if the strategy $s(x)$ survives IEDS in the modified game $\Gamma^m(\tilde{\sigma}\gamma, \tilde{\tau}\gamma)$, then $s(x) = 0$ for $x < \theta^* - \delta$*

and $s(x) = 1$ *for* $x > \theta^* + \delta$*, where, for* $\sigma > 0$*,* θ^* *solves*

$$\int_0^1 u(\theta, l) dH(l) = 0 \tag{7}$$

and $H(\cdot)$ *is defined by*

$$H(l) = \int_{\Xi} g(\xi) \left(1 - F\left(\frac{\tilde{\tau}}{\tilde{\sigma}}\xi + F^{-1}(1-l)\right)\right) d\xi, \tag{8}$$

and for $\sigma = 0$*,* θ^* *solves*

$$G(0)u(\theta, 0) + (1 - G(0))\, u(\theta, 1) = 0, \tag{9}$$

where $G(\cdot)$ *is the distribution function corresponding to the density* $g(\cdot)$*.*

The expression $1 - F\left(\frac{\tilde{\tau}}{\tilde{\sigma}}\xi + F^{-1}(1-l)\right)$ in (8) is the equilibrium belief over l in the underlying game for a player observing a signal at distance $\sigma\xi$ from the threshold. A player in the modified game observing the threshold signal x is uncertain over the true value x', and thus her belief $H(l)$ is an average of the rational beliefs induced by signals close to the threshold.

Proof. The proof in the complete information case is similar to, but simpler than that for the incomplete information case, and is omitted.

Let $\sigma = \tilde{\sigma}\gamma$, $\tau = \tilde{\tau}\gamma$, and $q_\theta(\theta|x)$ denote the marginal density associated with the subjective beliefs $q(\theta, x'|x)$ given σ and τ. Define

$$\tilde{m}_{\sigma,\tau}(x, k) \equiv \int_{\Theta} q_\theta(\theta|x) u\left(\theta, 1 - F\left(\frac{k - \theta}{\sigma}\right)\right) d\theta, \tag{10}$$

which is the expected payoff to action 1 for type x in the modified game $\Gamma^m(\sigma, \tau)$ when the opponent plays a threshold strategy with threshold k. Step 1 consists of showing that action 0 is serially dominated for $x > \overline{\theta}^*$ and action 1 is serially dominated for $x < \underline{\theta}^*$, where $\overline{\theta}^*$ and $\underline{\theta}^*$ are, respectively, the maximal and minimal roots of $\tilde{m}_{\sigma,\tau}(x, x) = 0$. The proof of step 1 is essentially the same as the relevant portion of the proof of Proposition 2.1 in Morris and

Shin (2003), and we therefore do not repeat it here.[9]

Step 2 consists of expressing $\tilde{m}_{\sigma,\tau}(x,k)$ in terms of $m_\sigma(x,k)$ of the underlying standard global game, defined by

$$m_\sigma(x,k) \equiv \int_\Theta p(\theta|x)u\left(\theta, 1-F\left(\frac{k-\theta}{\sigma}\right)\right)d\theta,$$

where $p(\theta|x)$ denotes the objective conditional distribution in the underlying game (given σ). The only difference between this and $\tilde{m}_{\sigma,\tau}(x,k)$ is that $m_\sigma(x,k)$ is computed with the use of objective conditional probabilities $p(\theta|x)$, whereas $\tilde{m}_{\sigma,\tau}(x,k)$ uses the subjective beliefs $q_\theta(\theta|x)$.

We have

$$q_\theta(\theta|x) = \int_X p(\theta|x')q_{x'}(x'|x)dx',$$

where

$$q_{x'}(x'|x) = \frac{p_x(x')\frac{1}{\tau}g(\frac{x'-x}{\tau})}{\int_X p_x(\tilde{x}')\frac{1}{\tau}g(\frac{\tilde{x}'-x}{\tau})d\tilde{x}'}.$$

Substituting this expression and interchanging the order of integrals in (10), we obtain

$$\tilde{m}_{\sigma,\tau}(x,k) = \int_X q_{x'}(x'|x)m_\sigma(x',k)dx',$$

which completes step 2.

Step 3 consists of computing the limits $\lim_{\tau\to 0} q_{x'}(x'|x)$ and $\lim_{\sigma\to 0} m_\sigma(x,k)$. Note that, since $\int_{x-\varepsilon}^{x+\varepsilon}\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx' = \int_{\frac{-\varepsilon}{\tau}}^{\frac{\varepsilon}{\tau}} g(z)\,dz$, given any $\delta > 0$ and $\varepsilon > 0$, there exists some $\tau > 0$ such that $\int_{x-\varepsilon}^{x+\varepsilon}\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx' > 1-\delta$. In particular, for any function $\psi(\cdot)$ that is continuous at x, we have

$$\lim_{\tau\to 0}\int_X \psi(x')\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx' = \psi(x).$$

[9] The basic idea is that if action 1 has been eliminated for all types below $k < \underline{\theta}^*$, then $\tilde{m}_{\sigma,\tau}(x,k) < 0$ for x sufficiently close to k, indicating that there are further types for which action 1 can be eliminated. There is a slight complication in that $\tilde{m}_{\sigma,\tau}(x,k)$ may not be increasing in x. As a result, the inductive procedure of their proof may eliminate fewer strategies than IEDS, which poses no problem for the claim of Step 1 here. In addition, the fact that $\tilde{m}_{\sigma,\tau}(x,k)$ is decreasing in k for each x suffices to guarantee the monoticity of the sequences used in Morris and Shin's proof.

It follows that

$$\lim_{\tau\to 0}\tau q_{x'}\left(x+\tau\xi|x\right)=\lim_{\tau\to 0}\frac{p_x\left(x+\tau\xi\right)g(\xi)}{\int_X p_x\left(\tilde{x}'\right)\frac{1}{\tau}g(\frac{\tilde{x}'-x}{\tau})d\tilde{x}'}=g\left(\xi\right),$$

and for each x, convergence is uniform on compact subsets of Ξ since p_x and g are continuous.

Morris and Shin (2003, Appendix A) show that

$$m_\sigma(x,k)=\int_{l=0}^{1}u(k-\sigma F^{-1}(l),l)d\Psi_\sigma(l;x,k),$$

where $\Psi_\sigma(\cdot;x,k)$ is a distribution function over the interval $[0,1]$. Moreover, as $\sigma\to 0$,

$$\Psi_\sigma(l;x+\sigma\xi,x)\to 1-F(\xi+F^{-1}(1-l)) \tag{11}$$

uniformly.

Step 4 consists of taking the limit as $\gamma\to 0$ and combining the limits from Step 3. Accordingly, we have

$$\begin{aligned}
\tilde{m}(x,x) &= \lim_{\gamma\to 0}\int_X q_{x'}(x'|x)m_{\tilde{\sigma}\gamma}(x',x)dx' \\
&= \lim_{\gamma\to 0}\int_X q_{x'}(x'|x)\int_{l=0}^{1}u(x-\tilde{\sigma}\gamma F^{-1}(l),l)d\Psi_{\tilde{\sigma}\gamma}(l;x',x)dx' \\
&= \lim_{\gamma\to 0}\int_\Xi \tilde{\tau}\gamma q_{x'}(x+\tilde{\tau}\gamma\xi|x)\int_{l=0}^{1}u(x-\tilde{\sigma}\gamma F^{-1}(l),l)d\Psi_{\tilde{\sigma}\gamma}(l;x+\tilde{\tau}\gamma\xi,x)d\xi \\
&= \int_\Xi g\left(\xi\right)\int_{l=0}^{1}u(x,l)d\left(\lim_{\gamma\to 0}\Psi_{\tilde{\sigma}\gamma}(l;x+\tilde{\tau}\gamma\xi,x)\right)d\xi \\
&= \int_{l=0}^{1}u(x,l)d\left(\int_\Xi g\left(\xi\right)\left(\lim_{\gamma\to 0}\Psi_{\tilde{\sigma}\gamma}(l;x+\tilde{\tau}\gamma\xi,x)\right)d\xi\right),
\end{aligned}$$

with the interchanging of the limit and the integral justified by uniform convergence on compact subsets of $\Xi\times(0,1)$.

Substituting the limit from (11) gives

$$\tilde{m}(x,x)=\int_{l=0}^{1}u(x,l)d\left(\int_\Xi g\left(\xi\right)\left(1-F(\frac{\tilde{\tau}}{\tilde{\sigma}}\xi+F^{-1}(1-l))\right)d\xi\right).$$

Thus defining

$$H(l) = \int_{\Xi} g(\xi) \left(1 - F(\frac{\tilde{\tau}}{\tilde{\sigma}}\xi + F^{-1}(1-l))\right) d\xi,$$

we have

$$\tilde{m}(x,x) = \int_{l=0}^{1} u(x,l) dH(l).$$

Since $u(x,l)$ is increasing in x, it follows that the equation $\tilde{m}(x,x) = 0$ has at most one root.

It follows from uniform limit dominance that there exist signals $\underline{x}$ and $\overline{x}$ such that action 1 is dominated for $x < \underline{x}$ and action 0 is dominated for $x > \overline{x}$. Thus we may restrict attention to signals in some compact set $\overline{X}$, on which $\tilde{m}_{\tilde{\sigma}\gamma,\tilde{\tau}\gamma}(x,x)$ converges to $\tilde{m}(x,x)$ uniformly. Given any neighborhood N of the unique root x^* of $\tilde{m}(x,x)$, there exists some $\varepsilon > 0$ such that the absolute value of $\tilde{m}(x,x)$ is uniformly bounded away from zero outside of N. Choosing $\overline{\gamma} > 0$ small enough so that whenever $\gamma < \overline{\gamma}$, $\tilde{m}_{\tilde{\sigma}\gamma,\tilde{\tau}\gamma}(x,x)$ is within ε of $\tilde{m}(x,x)$ everywhere on $\overline{X}$, guarantees that $\tilde{m}_{\tilde{\sigma}\gamma,\tilde{\tau}\gamma}(x,x)$ has no root in $\overline{X} \setminus N$. □

Proposition 4.1 can be generalized from to 2 to N players in a straightforward way.[10] Let the payoff to investment be $\tilde{u}(\theta,k)$ under a pure strategy profile in which k players invest. Let

$$u(\theta,l) \equiv \sum_{k=1}^{N} \binom{N-1}{k-1} l^{k-1}(1-l)^{N-k}\tilde{u}(\theta,k) \tag{12}$$

be the expected payoff to investment if each player independently randomizes and invests with probability l. Suppose that $\tilde{u}(\theta,k)$ is non-decreasing in k, which implies that $u(\theta,l)$ is non-decreasing in l. With this definition of payoffs, Proposition 4.1 applies to the N player case unchanged, and the limit long-run threshold is characterized by equation (7) with the distribution $H(l)$ satisfying (8). Moreover, let us rewrite the pure strategy profile payoff as $\tilde{u}(\theta,k) = v\left(\theta, \frac{k}{N}\right)$, where $v : \mathbb{R} \times [0,1] \to \mathbb{R}$. Then $\lim_{N\to\infty} u(\theta,l) = v(\theta,l)$, as all the weight in the summation in (12) becomes concentrated at $k = l$. In particular, as N grows large, the long-run threshold converges to the solution of

$$\int_0^1 v(\theta,l) dH(l) = 0,$$

[10] The generalized proofs require only an expansion of notation, and are available upon request.

just as in the two player case.

For each ratio $\rho = \frac{\tilde{\tau}}{\tilde{\sigma}}$, let θ^*_ρ denote the threshold defined by equation (7) when ρ is finite, and by (9) when $\rho = \infty$ (that is, when $\tilde{\sigma} = 0$). In addition, let θ^*_{gg} denote the threshold corresponding to the unique serially undominated strategy of the underlying game in the small noise limit. In general, the long-run threshold θ^*_ρ differs from the global game prediction θ^*_{gg}, and is sensitive to both the noise and similarity distributions. The following corollary identifies sufficient conditions under which the quantitative predictions agree.

Corollary 4.1. *Suppose that the noise and similarity distributions are symmetric about* 0*, and the payoff function $u(\theta, l)$ is linear in l. Then for $\tilde{\sigma} > 0$, the long-run threshold $\theta^*_{\frac{\tilde{\tau}}{\tilde{\sigma}}}$ is identical to the equilibrium threshold θ^*_{gg} of the underlying global game.*

Proof. By the symmetry of the noise distribution, we have $F^{-1}(l) = -F^{-1}(1-l)$, which implies that

$$H(1-l) = \int_\Xi g(\xi)\left(1 - F\left(\frac{\tilde{\tau}}{\tilde{\sigma}}\xi - F^{-1}(1-l)\right)\right) d\xi.$$

Substituting $\tilde{\xi} = -\xi$ and using the symmetries $g(\xi) = g(-\xi)$ and $1 - F(-y) = F(y)$ gives

$$H(1-l) = \int_\Xi g(\tilde{\xi})\left(F\left(\frac{\tilde{\tau}}{\tilde{\sigma}}\tilde{\xi} + F^{-1}(1-l)\right)\right) d\tilde{\xi},$$

which is equal to $1 - H(l)$. Hence the distribution $H(\cdot)$ is symmetric about $1/2$, which, together with the linearity of $u(\theta, l)$ in l, implies that

$$\int_0^1 u(\theta, l) dH(l) = \int_0^1 u(\theta, l) dl$$

for every θ. □

The conditions of Corollary 4.1 are strong. Although symmetry of both the error distribution and the similarity function guarantees symmetric beliefs over l, the thresholds in the two models may not agree if payoffs depend non-linearly on l.

Proposition 4.1 allows for the computation of the long-run outcome as τ and σ approach zero while holding $\frac{\sigma}{\tau}$ fixed. The following proposition states that as $\tilde{\tau}$ becomes small relative to $\tilde{\sigma}$, the equilibrium of the underlying game emerges as the long-run outcome of the learning model; on the other hand, as $\tilde{\sigma} > 0$ becomes small relative to $\tilde{\tau}$, the long-run outcome approaches that obtained under complete information. The latter result contrasts sharply with the rational model, in which the predictions vary discontinuously at $\sigma = 0$.

Proposition 4.2. *Suppose that the densities $f(\cdot)$ and $\phi(\cdot)$ are bounded, and that each has only finitely many discontinuities on any compact set. Then*

1. *For $\tilde{\sigma} > 0$, the threshold $\theta^*_{\frac{\tilde{\tau}}{\tilde{\sigma}}}$ of the learning model tends to the equilibrium threshold θ^*_{gg} as $\frac{\tilde{\tau}}{\tilde{\sigma}}$ tends to zero; that is, $\lim_{\rho \to 0} \theta^*_\rho = \theta^*_{\mathrm{gg}}$*

2. *Suppose in addition that, for each θ, $u(\theta, l)$ is continuous in l at 0 and 1. Then the incomplete information threshold $\theta^*_{\frac{\tilde{\tau}}{\tilde{\sigma}}}$ tends to the complete information threshold θ^*_∞ as $\frac{\tilde{\tau}}{\tilde{\sigma}}$ tends to infinity; that is, $\lim_{\rho \to \infty} \theta^*_\rho = \theta^*_\infty$.*

Proof. See Appendix. □

4.2 Comparative Statics

The predictions of the learning model are more ambiguous than those of the global games model because the long-run threshold depends on the similarity function which is unknown to an outside observer. Yet the comparative statics with respect to many parameters of practical interest are unambiguous and have the same sign as in the global game model.

Consider the learning process characterized by $\tilde{\tau}$, $\tilde{\sigma}$ in the limit $\gamma \to 0$, let the payoff function $u(\theta, l; z)$ depend on an exogenous parameter z, and assume throughout this subsection that it is continuously differentiable with respect to θ and z.

Proposition 4.3. *If the sign of $\frac{\partial u}{\partial z}(\theta, l; z)$ is the same for all θ and l, then*

$$\mathrm{sign}\left(\frac{\partial \theta^*_{\frac{\tilde{\tau}}{\tilde{\sigma}}}}{\partial z}\right) = \mathrm{sign}\left(\frac{\partial \theta^*_{\mathrm{gg}}}{\partial z}\right) = -\mathrm{sign}\left(\frac{\partial u}{\partial z}\right),$$

independent of $\tilde{\tau}$, $\tilde{\sigma}$ and $g(\cdot)$.

Proof. By Proposition 4.1, the long-run threshold $\theta^*_{\frac{\tilde{\tau}}{\tilde{\sigma}}}$ is the solution to

$$\int_0^1 u(\theta, l; z) dH(l) = 0.$$

By the implicit function theorem,

$$\frac{\partial \theta^*_{\frac{\tilde{\tau}}{\tilde{\sigma}}}}{\partial z} = -\frac{\int_0^1 u_z(\theta, l; z) dH(l)}{\int_0^1 u_\theta(\theta, l; z) dH(l)}.$$

The denominator is positive by Assumption A1, that returns are increasing in θ. The sign of the numerator is equal to sign $\left(\frac{\partial u}{\partial z}\right)$. The sign does not depend on the distribution $H(\cdot)$ and thus is independent of $\tilde{\tau}$, $\tilde{\sigma}$ and $g(\cdot)$, and therefore equal to the sign obtained when $H(\cdot)$ corresponds to the uniform distribution on $[0, 1]$. □

Proposition 4.3 may be applied to many comparative statics analyses found in applications. For instance, though the size of the effect depends on details of the model, the long-run threshold always increases with the outside option value[11], and decreases with the measure of players.[12] Heinemann, Nagel, and Ockenfels (2004) experimentally study both of these comparative statics effects, and confirm the qualitative predictions of the global game model, and thus also those of the present model.

Next we study comparative statics with respect to $F(\cdot)$ and $g(\cdot)$. Unlike in the global game theory, the long-run threshold in the learning model generally depends on the error distribution. However, translations of this distribution have no effect on the threshold. To see this, consider a change in the error distribution from $F(\cdot)$ to $\tilde{F}(\epsilon) = F(\epsilon - \mu)$ for $\mu \in \mathbb{R}$. The distribution $H(\cdot)$ given by (8) in Proposition 4.1 becomes

$$\begin{aligned}
\tilde{H}(l) &= \int_\Xi g(\xi) \left(1 - \tilde{F}\left(\frac{\tilde{\tau}}{\tilde{\sigma}}\xi + \tilde{F}^{-1}(1-l)\right)\right) d\xi \\
&= \int_\Xi g(\xi) \left(1 - F\left(\frac{\tilde{\tau}}{\tilde{\sigma}}\xi + F^{-1}(1-l) + \mu - \mu\right)\right) d\xi \\
&= H(l).
\end{aligned}$$

The long-run thresholds under the two error distributions are therefore identical.

Now fix $F(\cdot)$ and consider two similarity functions $g(\cdot)$ and $\tilde{g}(\cdot)$ such that the distribution corresponding to g first-order stochastically dominates that corresponding to $\tilde{g}$. In this case,

[11] A game with payoff $u(\theta, l)$ in which the value of the outside option was raised from 0 to z can be re-normalized to a game in which the outside option is 0, but $\tilde{u}(\theta, l) = u(\theta, l) - z$.

[12] Consider the limit of continuum of players. Increasing the measure of players from 1 to z is equivalent to keeping the measure constant but changing the payoff function to $\tilde{u}(\theta, l) = u(\theta, lz)$.

we say that g is *more optimistic* than $\tilde{g}$, as a player characterized by g unambiguously assigns more weight to similar higher signals than a player characterized by $\tilde{g}$.

Proposition 4.4. *If g is more optimistic than $\tilde{g}$, then the long-run threshold of players learning according to g is weakly lower than that of players learning according to $\tilde{g}$. The inequality is strict if $u(\theta, l)$ is strictly increasing in l.*

Proof. Let $H(\cdot)$ and $\tilde{H}(\cdot)$ be the distributions given by (8) in Proposition 4.1 corresponding to g and $\tilde{g}$ respectively. Since $1 - F\left(\frac{\tilde{\tau}}{\tilde{\sigma}}\xi + F^{-1}(1-l)\right)$ is decreasing in ξ, the first-order stochastic domination of $\tilde{g}$ by g implies that $H(l) < \tilde{H}(l)$ for every $l \in (0, 1)$; in other words, H first-order stochastically dominates $\tilde{H}$. But then since $u(\theta, l)$ is non-decreasing in l, we have

$$\int_0^1 u(\theta, l) dH(l) \geq \int_0^1 u(\theta, l) d\tilde{H}(l)$$

for each θ, and therefore $\theta^* \leq \tilde{\theta}^*$. □

4.3 The Environmental Multiplier

In this subsection, we analyze the impact of the prior distribution using an example of the learning process with *non-vanishing* σ and τ. We keep the structure of the example compatible with the setting of Chapter 3 of Morris and Shin (2003), in which the authors examine the strategic influence of public information. This allows us to compare the influence of the prior via the strategic link studied by Morris and Shin to that via learning.

The underlying game Γ_σ of this example is characterized by the payoff function in Table 1, the distribution of fundamentals $\theta \sim N(y, \omega^2)$, and the distribution of error terms $\sigma\epsilon^i \sim N(0, \sigma^2)$. Players are characterized by their similarity function $\frac{1}{\tau}g(\frac{x'-x}{\tau})$, which we take to be the density function of $N(0, \tau^2)$.

Applying Theorem 3.1, the long-run behavior is consistent with IEDS in the modified game $\Gamma^m_{\sigma,\tau}$. We use the procedure utilized in the proof of Proposition 4.1 according to which the solution of $\Gamma^m_{\sigma,\tau}$ reduces to solving $\tilde{m}_{\sigma,\tau}(x, x) = 0$. The normality of the distributions and of the similarity function allows us to express $\tilde{m}_{\sigma,\tau}(x, k)$ analytically. For this purpose, we explicitly express the subjective probability distribution of $X^{-i}|x^i$. In the first step we compute the subjective probability distribution of $\Theta|x$. In estimating θ, each player processes two normally distributed signals, the public signal y and the private signal x. Each player subjectively evaluates x as $\theta + \sigma\epsilon + \tau\xi$ where $\sigma\epsilon \sim N(0, \sigma^2)$ and $\tau\xi \sim N(0, \tau^2)$. Thus, ignoring the public signal y, the subjective $\Theta|x$ would be distributed as $N(x, \sigma^2 + \tau^2)$. Finally, incorporating the public signal y,

$$\Theta|x \sim N\left(\frac{y(\sigma^2+\tau^2)+x\omega^2}{\omega^2+\sigma^2+\tau^2}, \frac{\omega^2(\sigma^2+\tau^2)}{\omega^2+\sigma^2+\tau^2}\right). \tag{13}$$

Players in the modified game have correct beliefs about the conditional distribution of $X^{-i}|\theta$. The subjective belief $X^{-i}|x$ thus consists of a sum of the normal random variable in (13) and $N(0, \sigma^2)$, which gives

$$X^{-i}|x \sim N\left(\underbrace{\frac{y(\sigma^2+\tau^2)+x\omega^2}{\omega^2+\sigma^2+\tau^2}}_{A(x)}, \underbrace{\frac{\omega^2(\sigma^2+\tau^2)+\sigma^2(\omega^2+\sigma^2+\tau^2)}{\omega^2+\sigma^2+\tau^2}}_{B^2}\right).$$

Recall that the function $\tilde{m}_{\sigma,\tau}(x, k)$ is the subjective expected return in the modified game, given that the opponent's threshold is k. Hence, for the payoffs in Table 1, $\tilde{m}_{\sigma,\tau}(x, k) = A(x) - F\left(\frac{k-A(x)}{B}\right)$, and the threshold x^* is the root of

$$\tilde{m}_{\sigma,\tau}(x, x) = A(x) - F\left(\frac{x-A(x)}{B}\right) = 0. \tag{14}$$

Keeping σ and τ fixed, the left hand side of (14) is strictly increasing in x for sufficiently large ω, as can be verified by explicit computation of $\frac{d}{dx}\left(A(x) - F\left(\frac{x-A(x)}{B}\right)\right)$. We assume below that ω is sufficiently large, which rules out multiple equilibria.

We are now able to analyze the comparative statics of the threshold x^* with respect to the public signal y. Consider first the comparative statics under the limit $\tau \to 0$. Applying the implicit function theorem to equation (14) gives

$$\frac{\partial x^*}{\partial y} = -\frac{\sigma^2 + f((x-y)D)D(\omega^2+\sigma^2)}{\omega^2 - f((x-y)D)D(\omega^2+\sigma^2)}, \tag{15}$$

where $D = \frac{1}{\sqrt{(\omega^2+\sigma^2)(\frac{2\omega^2}{\sigma^2})+1}}$. This result is equivalent to the result in Morris and Shin (2003, Section 3.1). The coincidence is a consequence of Proposition 4.2, which states that the learning process converges to the equilibrium profile in the global game Γ_σ for $\tau \ll \sigma$.

If the rational players were to ignore strategic considerations and only process the information in the public signal, the effect would be of size $\frac{\partial x^*}{\partial y} = \frac{\sigma^2}{\omega^2}$. But the actual effect in (15) is larger due to the strategic behavior in the global game model. Although (15) also holds in the learning model (in the limit $\tau \to 0$), the interpretation must differ, as the players do not directly process public information, nor are they capable of any strategic reasoning. The need for a different interpretation is even more pronounced if we consider the limit as $\sigma \to 0$ holding $\tau > 0$ fixed. In this case

$$\frac{\partial x^*}{\partial y} = -\frac{\tau^2 + f((x-y)E)E(\omega^2+\sigma^2)}{\omega^2 - f((x-y)E)E(\omega^2+\sigma^2)}, \tag{16}$$

where $E = \frac{\tau}{\omega}\frac{1}{\sqrt{\omega^2+\tau^2}}$. The public signal plays no informational role in the limit as players observe the fundamental perfectly. However, the outcome of learning varies with y because y defines the environment in which the learning takes place. Increasing y corresponds to an improvement in the environment, and thus, *ceteris paribus*, improves players' experiences. Higher experienced returns translate into higher estimated returns; consequently, x^* must decrease in order to keep the threshold player indifferent between the two actions. We summarize the difference in interpretations by renaming the "public information" multiplier to be the "environmental" multiplier for the purposes of our model.

The difference in the interpretations of the multipliers in the global game and learning models stems from the fact that the reasoning of players is entirely deductive in the global game model, whereas it is entirely inductive in the learning model. Both of these assumptions are extreme. Consider a publicly announced change in the prior from $\Phi(\theta)$ to $\tilde{\Phi}(\theta)$ at time t. According to the global game theory, players *immediately* and substantially adjust their behavior. While the learning model also predicts a large impact on the behavior, it predicts that there will be *no* immediate reaction; the adjustment occurs only in the long-run. Some combination of the two models could lead to less extreme predictions involving an instantaneous reaction combined with partial inertia.

5 Related Literature

Processes of learning from similar games have been examined in several papers, which typically define similarity by an equivalence relation on a given set of games. LiCalzi (1995) provides sufficient conditions for convergence of fictitious play with similarity-based learning in 2×2 games. Germano (2004) considers rules that specify a strategy for each game in a given set $\mathcal{G}$. Rules are subject to stochastic evolutionary selection, and those that do not survive IEDS almost surely disappear in the longrun. Stahl and Van Huyck (2002) demonstrate learning from similar games experimentally. Subjects repeatedly interacted in stag-hunt games randomly drawn from a particular set, with the set being varied under two different treatments. The observed long-run behavior in a particular game contained in both sets varied across treatments, indicating that subjects were influenced by their experience playing different games.

Jehiel and Koessler (2006) study steady states of learning processes in incomplete information games. Let Ω be the set of states of the world. Learning by each player is governed by a partition $\mathcal{A}^i$ of Ω: when learning an opponent's action in state $\omega \in \Omega$, player i aggregates the history of opponent's strategy in all states in the set $\mathcal{A}^i(\omega)$.[13] Jehiel and Koessler apply their equilibrium concept to a global game, assuming the coarsest similarity partition, according to which each player completely disregards the circumstances under which her opponent choses an action. The main predictions of our model arise at the opposite extreme, in which only cases from a small neighborhood of the current case are given significant weight. Another important difference, however, lies in Jehiel and Koessler's formulation of similarity as a partition, which prevents actions from spreading contagiously across types.

Argenziano and Gilboa (2005) consider coordination problems drawn from a finite set. Players perfectly observe the current problem and form beliefs about their opponents' strategies by aggregating their experience in similar past games. When games with dominant actions are sufficiently rare, the long-run outcome of learning depends on historical accidents.

Once most of the work on the present paper was completed, we discovered a paper by

[13] In a related paper, Jehiel and Samet (2004) suppose that players use partitions of their actions spaces in order to estimate payoffs directly. While this approach is similar to that of our model, their paper is focused on very different issues.

Carlsson (2004) that proposes a learning model closely related to the one studied here. In his model, players use similarity to estimate their opponent's strategy in two-player, complete information global games. Carlsson offers an informal argument to suggest that the learning process can be approximated by the best-response dynamics of a modified game. Theorem 3.1 formalizes the corresponding result for binary action global games when similarity is instead used to estimate payoffs directly. Carlsson's focus is on providing evolutionary foundations for the global game equilibrium, which agrees with the long-run outcome of learning under the conditions of his model. Our analysis suggests that under more general conditions (in particular, allowing for incomplete information in the learning process), the main qualitative predictions of the learning and global game models coincide, although the outcomes may differ quantitatively.

Milgrom and Roberts (1990) study supermodular games, of which global games are a special case, and show that only serially undominated strategies are played in the long-run under a large class of adaptive dynamics. These dynamics, however, require that players adjust to the full strategies of their opponents. In games with large type spaces, where play of the game (at most) reveals the actions assigned by strategies to the particular types that are drawn, such dynamics are difficult to justify. The use of similarity in learning can be seen as generating "close to" adaptive dynamics, as reflected in the modified serially undominated result of Theorem 3.1.[14]

An alternative approach to learning in binary-action supermodular games is offered by Beggs (2005), who proposes a class of adaptive learning rules where players are restricted to using monotone (threshold) strategies. The threshold evolves based on payoffs from similar types, with similarity weights becoming increasingly concentrated on nearer types over time. Under stronger restrictions on similarity than those imposed here, the threshold strategies converge to an equilibrium of the game.

[14] In addition, both Samuelson and Zhang (1992) and Nachbar (1990) specify classes of learning processes under which players learn not to play serially dominated strategies; however, both papers assume finite sets of pure strategies.

6 Conclusion

The main difficulty in formulating a learning model for games with large type spaces is that players must learn optimal behavior in many contingencies despite having relatively limited experience. To enable learning in such games, we have supposed that players extrapolate from their experience in past cases in which their type was similar to the current one. This approach allows for learning even if interactions arise only rarely relative to the size of the type space. In environments with strategic complementarities, this similarity-based learning process leads to contagion of actions across types.

Contagion through learning shares the main qualitative features of contagion from incomplete information. Players learn to play symmetric threshold strategies, and the comparative statics predictions share the same sign. Quantitatively, however, these two processes generally lead to different outcomes. This difference is captured by the subjective priors of the modified game, as compared to the objective priors of the usual incomplete information model. With objective priors and small noise, players always believe with probability $\frac{1}{2}$ that their opponent has received a signal greater than their own. With subjective priors, this probability generally depends on the priors.

The extrapolations used by players in similarity-based learning typically lead to biases in payoff estimates away from their true values. As similarity becomes more heavily concentrated on nearby states, these biases disappear, but their impact on behavior does not. Narrowly concentrated similarity is analogous to the bandwidth of a kernel estimator vanishing. Long-run estimates are consistent under general conditions as long as the estimated function is not changing. In a strategic setting, however, payoff estimates depend on the strategies of the other players, which in turn depend on their own payoff estimates. Since short-run biases in these estimates are unavoidable, their effects may persist over time even if the long-run estimates are unbiased. Thus contagion through learning persists even as the biases in similarity-based estimates vanish.

A well-known formal equivalence exists between Bayesian games and local interaction models (see Morris (1997)). Under this equivalence, types correspond to members of a population, and posterior beliefs about the types of other players correspond to probabilities of being matched with the corresponding members of the population. This equivalence readily extends to similarity-based learning. Learning payoffs from certain other types in a Bayesian game is equivalent to learning payoffs from certain other players in a local interaction model. In this setting, the modified game result indicates that the outcomes of learning may be viewed as equilibria of a modified local interaction game. The subjective priors of the modified game in the Bayesian setting correspond to subjectivity concerning the structure of the network in the local interaction setting. In these subjective networks, players generally believe that interactions are asymmetric: it may be that player i's payoff depends on the action of player j, but not vice versa.

We have explored the long-run outcomes of similarity-based learning only in a particular class of games. However, this learning process may be applied more generally to the broad class of games with large type spaces, in which standard learning models fail. We conjecture that, under general conditions, players will learn not to play serially dominated strategies for sufficiently concentrated similarity. Such a result would extend the theorem of Samuelson and Zhang (1992) for finite games.

A Appendix

Proof of Lemma 3.1. First note that the uniform continuity of $g^i(\cdot)$ implies that the beliefs $q^i\left(\cdot|x^i\right)$ are uniformly continuous in x^i for every $\left(\theta, x^j\right)$, and hence, for a fixed strategy of the opponent, expected payoffs also vary continuously in the player's own type.

Uniform limit dominance implies that, for some $\pi > 0$, there exist types $\underline{x}$ and $\overline{x}$ such that action 1 is π-dominated for $x < \underline{x}$ and action 0 is π-dominated for $x > \overline{x}$. Thus it suffices to prove the result for types on any compact interval $[b, c]$.

Step 1: First we show that given any $\varepsilon > 0$, there exists some $\delta > 0$ such that changing the opponent's strategy on a set of Lebesgue measure at most δ changes the expected payoff of any type by at most ε.

For any two strategies s, s', let $\mu\left(s, s'\right) \in \mathbb{R}_+ \cup \{\infty\}$ denote the Lebesgue measure of the set of types on which s and s' differ. Let $U\left(s, x\right)$ denote the expected payoff received by type x when playing action 1 against an opponent who plays strategy s. Given any $\varepsilon > 0$ and any type x, define

$$\delta(x;\varepsilon) = \inf_{\substack{s,s' \\ |U(s,x)-U(s',x)|\geq\varepsilon}} \mu(s,s').$$

In words, a measure of at least $\delta(x;\varepsilon)$ of the opponent's types must change their actions in order to induce a payoff change of at least ε for type x when choosing action 1. Dropping the ε from the notation, clearly $\delta(x) > 0$ everywhere, so if we show that $\delta(\cdot)$ is continuous, then it must attain a strictly positive minimum on the compact interval $[b,c]$. Accordingly, suppose that there is a discontinuity of size at least $\eta > 0$ in $\delta(\cdot)$ at some type x_0 (that is, there does not exist any $\gamma > 0$ such that $|\delta(x) - \delta(x_0)| < \eta$ whenever $|x - x_0| < \gamma$). Suppose that for every $\gamma > 0$ there exists some $x \in (x_0 - \gamma, x_0 + \gamma)$ such that $\delta(x) > \delta(x_0) + \eta$ (the argument is similar if instead $\delta(x_0) > \delta(x) + \eta$). Let s and s' be strategies for the opponent such that $\mu(s,s') < \delta(x_0) + \frac{\eta}{2}$ and $U(s,x_0) - U(s',x_0) \geq \varepsilon$. Note that since $\mu(s,s')$ is finite, either $s \neq 1$ or $s' \neq 0$. Thus there either exists some strategy s'' such that $\mu(s,s'') \leq \frac{\eta}{2}$ and $U(s'',x_0) > U(s,x_0)$ or there exists some strategy such that $\mu(s',s'') \leq \frac{\eta}{2}$ and $U(s',x_0) > U(s'',x_0)$. Suppose the former (the argument for the other case is similar). Then we have $U(s'',x_0) - U(s',x_0) > \varepsilon$, and by the continuity of $U(s',\cdot)$ and $U(s'',\cdot)$, there exists some neighborhood N of x_0 such that $U(s'',x) - U(s',x) > \varepsilon$ whenever $x \in N$, contradicting the definition of η since $\mu(s'',s') \leq \mu(s,s') + \mu(s,s'') < \delta(x_0) + \eta$.

Step 2: Given any type x for which action 0 (say) is eliminated in the Nth round, there exists some $\pi(x) > 0$ such that the expected payoff for playing action 1 is at least $\pi(x)$ whenever the opponent plays an action consistent with $N-1$ rounds of elimination. From Step 1, it suffices to show that given any $\delta > 0$, there exists some $\pi > 0$ small enough such that $N-1$ rounds of elimination of π-dominated strategies differs from $N-1$ rounds of elimination of dominated strategies on a set of types of measure at most δ.

Consider any positive sequence $\pi_1, \pi_2, ...$ such that $\lim_{n\to\infty} \pi_n = 0$. Fix a set S of strategies for the opponent that contains a unique "worst-case" strategy, that is, contains a strategy s_0 with the property that $s_0(x) = 0$ whenever $s(x) = 0$ for some $s \in S$ (note that the set of strategies surviving N rounds of iterated deletion of π-dominated strategies satisfies this property for any $\pi \geq 0$ and any N). Let $X(n)$ denote the set of types that receive an expected payoff greater than π_n when playing action 1 against any strategy in S (equivalently, against s_0), and let $\overline{X}$ denote the set of types for which action 1 is dominant against the set S. Then $X(n)$ is a monotone sequence of sets that increases to $\overline{X}$ in the limit, for otherwise $\overline{X} \setminus \lim_{n\to\infty} X(n)$ is non-empty, and any type contained in it cannot receive a positive payoff when playing action 1 against s_0, contradicting the definition of $\overline{X}$.

We now proceed by induction on N. The result is trivial for $N = 1$. For $N > 1$, assume the result to be true for $N - 1$, that is, assume that given any $\delta > 0$, there exists some $\pi > 0$ for which $N - 2$ rounds of elimination of π-dominated strategies differs from $N - 2$ rounds of elimination of dominated strategies on a set of types of measure at most δ. For each n and $\pi \geq 0$, let $S_n(\pi)$ denote the set of strategies remaining for the opponent after n rounds of iterated deletion of π-dominated strategies. Note that for each n, $S_n(\pi)$ is non-decreasing in π in the sense that $s \in S_n(\pi)$ implies $s \in S_n(\pi')$ whenever $\pi' > \pi$.

Given $\delta > 0$, choose $\pi' > 0$ small enough so that the set of types for which a given action is dominated but not π' -dominated against $S_{N-2}(0)$ has measure at most δ. By Step 1, there exists some $\delta' > 0$ such that changing the actions of at most a measure of δ' of the opponent's types changes a player's expected payoff by at most $\frac{\pi'}{2}$. But then by the inductive hypothesis, we can choose $\pi'' > 0$ such that each element of $S_{N-2}(\pi'')$ differs from one of $S_{N-2}(0)$ on a set of types of measure at most δ'. Consider $\pi = \min\left\{\frac{\pi'}{2}, \pi''\right\}$. We need to show that $s \in S_{N-1}(\pi)$ implies that s differs from a member of $S_{N-1}(0)$ on a set of types of measure at most δ. Consider any type x. Since $S_{N-2}(\pi) \subseteq S_{N-2}(\pi'')$, the payoff received by x from playing action 1 against any member of $S_{N-2}(\pi)$ is within $\frac{\pi'}{2}$ of that from some member of $S_{N-2}(0)$. But then if either action is π' -dominated for type x against $S_{N-2}(0)$, it must be $\frac{\pi'}{2}$-dominated, and therefore π-dominated against $S_{N-2}(\pi)$, as needed. □

Proof of Theorem 3.1. We give the proof only for the incomplete information case, as that for the complete information case is essentially the same, except that instead of partitioning $\Theta \times X^1 \times X^2$, it suffices to partition Θ alone. We begin with the case of a compact state space, then show how the argument can be extended to non-compact spaces if the similarity function has full support.

By the lemma, it suffices to prove the result for IEπDS in the modified game for any $\pi > 0$. We proceed by induction on the round of deletion, n, fixing $\pi > 0$. For $n = 0$ there is nothing to prove. Suppose for induction that there almost surely comes a time after which each player i only plays strategies in the set $S^i(n-1)$ of those consistent with $n-1$ rounds of IEπDS.

Suppose that action 1 is π-dominated for type $\tilde{x}^i$ against $S^j(n-1)$ in the modified game (for $j \neq i$). Then

$$\int_\Theta \int_{X^j} u\left(\theta, l\left(x'\right)\right) q^i\left(\theta, x' | \tilde{x}^i\right) dx' d\theta < -\pi$$

for all strategies $l(\cdot) \in S^j(n-1)$. That is,

$$\frac{\int_\Theta \int_{X^j} \int_{X^i} u\left(\theta, l\left(x'\right)\right) p\left(\theta, x, x'\right) g^i\left(x - \tilde{x}^i\right) dx dx' d\theta}{\int_\Theta \int_{X^j} \int_{X^i} p\left(\theta, x, x'\right) g^i\left(x - \tilde{x}^i\right) dx dx' d\theta} < -\pi \tag{17}$$

for every $l(\cdot) \in S^j(n-1)$. Defining

$$\pi' := \inf_{x^i} \pi \int_\Theta \int_{X^j} \int_{X^i} p\left(\theta, x, x'\right) g^i\left(x - x^i\right) dx dx' d\theta,$$

the compactness of the type space implies that $\pi' > 0$. Since $u(\theta, l)$ is nondecreasing in l for every θ, inequality (17) implies that

$$\int_\Theta \int_{X^j} \int_{X^i} u\left(\theta, \bar{l}(x')\right) p\left(\theta, x, x'\right) g^i\left(x - \tilde{x}^i\right) dx dx' d\theta < -\pi', \tag{18}$$

where $\bar{l}(\cdot)$ is the strategy in $S^j(n-1)$ that chooses action 1 for every type for which this action has not been eliminated.

Let $\Theta = [b, c]$ be the payoff-relevant state space, which, along with the type spaces, is assumed to be compact. Given $\delta > 0$, partition each X^i and Θ into a finite number of subintervals of length at most δ. We will denote these partitions by $P_\delta\left(X^i\right)$ and $P_\delta(\Theta)$ respectively. To simplify the notation below, we assume that the partition $P_\delta\left(X^j\right)$ may be chosen so that $\bar{l}(\cdot)$ is $P_\delta\left(X^j\right)$-measurable, and for $\mu^j \in P_\delta\left(X^j\right)$, we will write $\bar{l}\left(\mu^j\right)$ for $\bar{l}\left(x^j\right)$ for $x^j \in \mu^j$. Otherwise, if it is not possible to choose the partition in this way, note that, since expected payoffs are continuous in types, $\bar{l}^{-1}(0)$ is an open set in X^j. We may therefore choose the partition in such a way that only an arbitrarily small measure of types of player j lie in elements of $P_\delta\left(X^j\right)$ on which $\bar{l}(\cdot)$ is not constant. This small measure of types almost surely (henceforth a.s.) has an arbitrarily small impact on player i's payoff estimates in the long-run, and so will only affect the following argument by introducing an additional arbitrarily small error term.

For any combination $\left(\rho, \mu^1, \mu^2\right) \in P_\delta(\Theta) \times P_\delta\left(X^1\right) \times P_\delta\left(X^2\right)$, and any $\eta > 0$, the strong law of large numbers guarantees that there will a.s. come a time after which the fraction of earlier periods t having $\left(\theta_t, x_t^1, x_t^2\right) \in \left(\rho, \mu^1, \mu^2\right)$ is within η of the probability associated with the event $\left(\rho, \mu^1, \mu^2\right)$. Since the number of such events is finite, there will almost surely come a time after which this is true for all $\left(\rho, \mu^1, \mu^2\right) \in P_\delta(\Theta) \times P_\delta\left(X^1\right) \times P_\delta\left(X^2\right)$.

We want to show that since the estimated payoff to type $\tilde{x}^i$ from action 1 in the modified game is less than $-\pi$, the estimated payoff under the learning dynamics will a.s. eventually lie below zero; hence this type learns to play action 0. By the induction hypothesis, any finite history in which the opponent played strategies outside $S^j(n-1)$ will a.s. eventually have arbitrarily small weight in player i's payoff estimates. Thus it suffices to show that, a.s., eventually,

$$\sum_{s=1}^{t-1} u\left(\theta_s, \bar{l}\left(x_s^j\right)\right) g^i\left(x_s^i - \tilde{x}^i\right) < 0$$

since the denominator on the right-hand side of (3) is (eventually) positive. Accordingly, we have

$$\begin{aligned}
&\sum_{s=1}^{t-1} u\left(\theta_s, \bar{l}\left(x_s^j\right)\right) g^i\left(x_s^i - \tilde{x}^i\right) \\
&\leq (t-1)\left(\begin{array}{l}
\sum_{\substack{\rho,\mu^1,\mu^2 \\ u\left(\sup(\rho),\bar{l}\left(\mu^j\right)\right)<0, \\ \Pr\left(\rho,\mu^1,\mu^2\right)>0}} \left(\Pr\left(\rho, \mu^1, \mu^2\right) - \eta\right) u\left(\sup(\rho), \bar{l}\left(\mu^j\right)\right) \inf_{x\in\mu^i} g^i\left(x - \tilde{x}^i\right) \\
+\sum_{\substack{\rho,\mu^1,\mu^2 \\ u\left(\sup(\rho),\bar{l}\left(\mu^j\right)\right)>0}} \left(\Pr\left(\rho, \mu^1, \mu^2\right) + \eta\right) u\left(\sup(\rho), \bar{l}\left(\mu^j\right)\right) \sup_{x\in\mu^i} g^i\left(x - \tilde{x}^i\right)
\end{array}\right),
\end{aligned}$$

where η is chosen to be sufficiently small so that each term $\Pr\left(\rho, \mu^1, \mu^2\right) - \eta$ in the first sum on the right-hand side is positive, which is possible since the partition is finite. We want to show that for sufficiently small η and δ, the expression inside the parentheses is negative. First, letting $\rho(\theta)$ denote the element of $P_\delta(\Theta)$ containing θ, and similarly for $\mu^i(x)$ and

$\mu^j(x)$, define the step function

$$\xi^i\left(\theta,x,x^j;\tilde{x}^i\right)=\begin{cases}u\left(\sup\left(\rho\left(\theta\right)\right),\bar{l}\left(x^j\right)\right)\inf_{x'\in\mu^i(x)}g^i\left(x'-\tilde{x}^i\right) & \text{if } u\left(\sup\left(\rho\left(\theta\right)\right),\bar{l}\left(x^j\right)\right)<0\\ u\left(\sup\left(\rho\left(\theta\right)\right),\bar{l}\left(x^j\right)\right)\sup_{x'\in\mu^i(x)}g^i\left(x'-\tilde{x}^i\right) & \text{if } u\left(\sup\left(\rho\left(\theta\right)\right),\bar{l}\left(x^j\right)\right)\geq 0.\end{cases}$$

Since the set $P_\delta(\Theta)\times P_\delta\left(X^1\right)\times P_\delta\left(X^2\right)$ is finite, and both u and g^i are bounded, there exists some finite K such that the integral $\int_\Theta\int_{X^j}\int_{X^i}\xi^i\left(\theta,x,x';\tilde{x}^i\right)p\left(\theta,x,x'\right)dxdx'd\theta$ is within ηK of the bracketed expression. Thus by choosing η sufficiently small (given the partitions), it suffices to show that $\int_\Theta\int_{X^j}\int_{X^i}\xi^i\left(\theta,x,x';\tilde{x}^i\right)p\left(\theta,x,x'\right)dxdx'd\theta<0$.

Since $u(\cdot,\cdot)$ is increasing in its first argument and defined on a compact set, given any $\lambda>0$, there may exist only finitely many discontinuities of $u(\cdot,1)$ of size at least λ (that is, for which there exists no $\delta>0$ such that $u(\theta,1)-u(\theta',1)\in(-\lambda,\lambda)$ whenever $\theta-\theta'\in(-\delta,\delta)$), and similarly for $u(\cdot,0)$. Thus given any $\varepsilon>0$, the partition $P_\delta(\Theta)$ may be chosen so that $u(\theta,1)-u(\theta',1)\in(-\varepsilon,\varepsilon)$ and $u(\theta,0)-u(\theta',0)\in(-\varepsilon,\varepsilon)$ whenever θ and θ' lie in the same element of $P_\delta(\Theta)$. Similarly, since each $g^i(\cdot)$ is uniformly continuous, given any $\varepsilon>0$, we may choose the partition $P_\delta\left(X^i\right)$ with δ small enough so that $g^i\left(x-\tilde{x}^i\right)-g^i\left(x'-\tilde{x}^i\right)\in(-\varepsilon,\varepsilon)$ whenever $x,x'\in\mu^i$ for some $\mu^i\in P_\delta\left(X^i\right)$. That is, letting

$$\varepsilon=\sup_{\substack{\theta,\theta'\\ \rho(\theta)=\rho(\theta')}}\max\left\{u\left(\theta,1\right)-u\left(\theta',1\right),u\left(\theta,0\right)-u\left(\theta',0\right)\right\}$$

$$\text{and}\quad\varepsilon'=\sup_{\substack{\tilde{x},x,x'\\ \mu^i(x)=\mu^i(x')}}g^i\left(x-\tilde{x}\right)-g^i\left(x'-\tilde{x}\right),$$

we may choose the partitions in such a way as to make ε and ε' arbitrarily small positive numbers.

Recalling that $g^i(\cdot)$ is bounded by M^i, and $u(\cdot,\cdot)$ is bounded by V, we have

$$\left|\xi^i\left(\theta,x,x';\tilde{x}^i\right)-u\left(\theta,\bar{l}\left(x'\right)\right)g^i\left(x-\tilde{x}^i\right)\right|<\varepsilon M^i+\varepsilon'V+\varepsilon\varepsilon',$$

and therefore

$$\int_\Theta \int_{X^j} \int_{X^i} \xi^i\left(\theta, x, x'; \tilde{x}^i\right) p\left(\theta, x, x'\right) dx dx' d\theta$$
$$\leq \int_\Theta \int_{X^j} \int_{X^i} u\left(\theta, \bar{l}\left(x'\right)\right) p\left(\theta, x, x'\right) g^i\left(x - \tilde{x}^i\right) dx dx' d\theta + \varepsilon M^i + \varepsilon' V + \varepsilon\varepsilon',$$

which, by (18), is negative for ε and ε' sufficiently small.

We have shown that if action 1 is π-dominated against $S^j(n-1)$ for type $\tilde{x}^i$ in the modified game, then there will almost surely come a time after which $\tilde{x}^i$ will not play this action under the learning dynamics. Furthermore, the same payoff approximations apply to any type x^i of player i, with the only difference being a shift in the arguments of the similarity function. This completes the proof of the inductive step. The symmetric argument proves the corresponding result for action 0. This completes the proof of the first statement when the state space is compact.

For the second statement, note that for any compact set S of types not intersecting $X(\{0,1\})$, there exists some $\pi > 0$ and $n \in \mathbb{N}$ such that after n rounds of elimination of π-dominated strategies (in the modified game), only the serially undominated action remains for each type in S. We have shown that under these conditions, there will almost surely come a time after which types in S play only their serially undominated actions, as needed.

If the state space is not compact, instead of repeating the proof for the compact case, we show only how the argument can be modified by the introduction of an arbitrarily small error term in the payoff estimates.

Given $\delta > 0$, we must show that there will a.s. be a period after which the probability measure of the set of player i's types playing actions consistent with IEDS in the modified game is at least $1 - \delta$. Consider some interval $[b, c]$ such that $\Pr\left(x^i \in [b, c]\right) > 1 - \delta$, and choose any $x \in [b, c]$. We want to show that for any $\varepsilon > 0$ there exist $\underline{x}, \overline{x}, \underline{\theta}, \overline{\theta}$ such that there will almost surely be some period T for which

$$\left| \frac{1}{\sum_{s=1}^{t-1} g^i\left(x_s^i - x\right)} \sum_{\substack{s=1,\dots,t-1 \\ \left(\theta_s, x_s^i, x_s^j\right) \notin \left[\underline{\theta}, \overline{\theta}\right] \times \left[\underline{x}, \overline{x}\right] \times \left[\underline{x}, \overline{x}\right]}} u\left(\theta_s, a_s^j\right) g^i\left(x_s^i - x\right) \right| < \varepsilon \qquad (19)$$

whenever $t > T$; that is, the contribution to the estimated payoff of those draws $\left(\theta_s, x_s^i, x_s^j\right)$ outside the compact space $\left[\underline{\theta}, \overline{\theta}\right] \times [\underline{x}, \overline{x}] \times [\underline{x}, \overline{x}]$ can be made arbitrarily small by an appropriate choice of $\underline{x}, \overline{x}, \underline{\theta}, \overline{\theta}$. The proof for the compact case may then proceed for types in the interval $[b, c]$ by partitioning the set $\left[\underline{\theta}, \overline{\theta}\right] \times [\underline{x}, \overline{x}] \times [\underline{x}, \overline{x}]$ and allowing for an additional arbitrarily small error term in the resulting bounds.[15]

The sum in the numerator may be naturally divided into parts according to whether θ_s lies below, in, or above $\left[\underline{\theta}, \overline{\theta}\right]$, x_s^j lies below, in, or above $[\underline{x}, \overline{x}]$, and x_s^i lies below, in, or above $[\underline{x}, \overline{x}]$.

Accordingly, consider

$$\frac{1}{\sum_{s=1}^{t-1} g^i\left(x_s^i - x\right)} \sum_{\substack{s=1,\ldots,t-1 \\ \theta_s < \underline{\theta}}} u\left(\theta_s, a_s^j\right) g^i\left(x_s^i - x\right).$$

For $\underline{\theta}$ small enough, taking $a_s^j = 0$ for all s gives an upper bound on the absolute value of this expression. Furthermore, letting

$$\begin{aligned} g^{\min}\left(x^i\right) &= \inf_{x \in [b,c]} g^i\left(x^i - x\right) \\ \text{and } g^{\max}\left(x^i\right) &= \sup_{x \in [b,c]} g^i\left(x^i - x\right), \end{aligned}$$

we have

$$\left| \frac{1}{\sum_{s=1}^{t-1} g^i\left(x_s^i - x\right)} \sum_{\substack{s=1,\ldots,t-1 \\ \theta_s < \underline{\theta}}} u\left(\theta_s, a_s^j\right) g^i\left(x_s^i - x\right) \right| \leq \left| \frac{1}{\sum_{s=1}^{t-1} g^{\min}\left(x_s^i\right)} \sum_{\substack{s=1,\ldots,t-1 \\ \theta_s < \underline{\theta}}} u\left(\theta_s, 0\right) g^{\max}\left(x_s^i\right) \right|$$

for all $x \in [b, c]$. By the strong law of large numbers, the expression on the right-hand side a.s. approaches

$$L\left(\underline{\theta}\right) := \left| \frac{\int_{\theta \leq \underline{\theta}} \int_{X^i} u\left(\theta, 0\right) g^{\max}\left(x^i\right) p\left(\theta, x^i\right) dx^i d\theta}{\int_{\Theta} \int_{X^i} g^{\min}\left(x^i\right) p\left(\theta, x^i\right) dx^i d\theta} \right|,$$

where $p\left(\theta, x^i\right) = \phi\left(\theta\right) f^i\left(\frac{x^i - \theta}{\sigma}\right)$ represents the density associated with the combination

[15] It is important that the estimates in the proof for the compact state space are applied only to the types in the interval $[b, c]$, and not in the larger interval $[\underline{x}, \overline{x}]$. Since the additional error term can be made arbitrarily small independent of the initial interval $[b, c]$, the value of π' does not depend on the choice of interval $[\underline{x}, \overline{x}]$ here.

$\left(\theta, x^i\right)$ in any period. Since g^i is bounded by M^i and $u(\theta, 0) < 0$ for $\theta \leq \underline{\theta}$, we have

$$L(\underline{\theta}) \leq -M^i \frac{\int_{\theta \leq \underline{\theta}} u(\theta, 0) \phi(\theta) \, d\theta}{\int_\Theta \int_{X^i} g^{\min}(x^i) p(\theta, x^i) \, dx^i d\theta}.$$

Note that the denominator is positive since $g^i(\cdot)$ is continuous and has full support. Since $u(\cdot, a)$ is integrable with respect to the distribution $\Phi(\cdot)$ for each a, the numerator can be made arbitrarily small by choosing $\underline{\theta}$ small, and the denominator is unaffected by this choice.

A similar argument applies for $\theta_s \geq \overline{\theta}$, and for $x_s^j \notin [\underline{x}, \overline{x}]$, except with $a(\theta) := \arg\max_a |u(\theta, a)|$ instead of $a = 0$.

Finally, we consider the part of the sum in (19) where $x_s^i \notin [\underline{x}, \overline{x}]$ and $\theta \in \left[\underline{\theta}, \overline{\theta}\right]$. Since $\lim_{x \to \infty} g(x) = \lim_{x \to -\infty} g(x) = 0$, we can use a similar bound, except again with $a(\theta)$ instead of $a = 0$. It follows that for each $x \in [b, c]$, the contribution to the estimated payoff arising from $\left(\theta, x^i\right) \in \left[\underline{\theta}, \overline{\theta}\right] \times (-\infty, \underline{x})$ is bounded in absolute value by

$$\frac{\int_{\theta=\underline{\theta}}^{\overline{\theta}} \int_{x^i \leq \underline{x}} u(\theta, a(\theta)) \, g^{\max}\left(x^i\right) p\left(\theta, x^i\right) dx^i d\theta}{\int_\Theta \int_{X^i} g^{\min}(x^i) p(\theta, x^i) \, dx^i d\theta}.$$

Given any $\varepsilon > 0$, we can choose $\underline{x}$ small enough so that $g^{\max}\left(x^i\right) < \varepsilon$ for all $x^i \leq \underline{x}$, and therefore the numerator can be made arbitrarily small while the denominator remains constant. The bound for $\left(\theta, x^i\right) \in \left[\underline{\theta}, \overline{\theta}\right] \times (\overline{x}, \infty)$ is similar. □

Proof of Proposition 4.2. For the first statement, it suffices to show that $m_{\sigma,\tau}(x, x)$ converges to $m_\sigma(x, x)$ uniformly on compact subsets of X as $\tau \to 0$. Recall that

$$m_\sigma(x, k) = \frac{\int_\Theta \phi(\theta) f\left(\frac{x-\theta}{\sigma}\right) u\left(\theta, 1 - F\left(\frac{k-\theta}{\sigma}\right)\right) d\theta}{\int_\Theta \phi(\theta) f\left(\frac{x-\theta}{\sigma}\right) d\theta}.$$

For $\tau > 0$,

$$m_{\sigma,\tau}(x, k) = \frac{\int_\Theta \int_{X^i} \phi(\theta) f\left(\frac{x'-\theta}{\sigma}\right) \frac{1}{\tau} g\left(\frac{x'-x}{\tau}\right) u\left(\theta, 1 - F\left(\frac{k-\theta}{\sigma}\right)\right) dx' d\theta}{\int_\Theta \int_{X^i} \phi(\theta) f\left(\frac{x'-\theta}{\sigma}\right) \frac{1}{\tau} g\left(\frac{x'-x}{\tau}\right) dx' d\theta}$$

First consider the denominator of the last expression:

$$\int_{X^i}\int_{\Theta}\phi(\theta)\,f\left(\frac{x'-\theta}{\sigma}\right)d\theta\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx' = \int_{X^i}p_x(x')\,\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx'.$$

Fix a compact subset $Z \subset X$. Given $\delta > 0$ and $\varepsilon > 0$, there exists some $\tau > 0$ such that $\int_{x-\delta}^{x+\delta}\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx' > 1-\varepsilon$. Since p_x is continuous on X, it is uniformly continuous on Z; hence given $\eta > 0$, there exists some $\delta(\eta) > 0$ not depending on x' such that $p_x(x'') \in (p_x(x')-\eta, p_x(x')+\eta)$ whenever $x'' \in (x'-\delta(\eta), x'+\delta(\eta))$. Therefore, given $\varepsilon > 0$, there exists some $\tau > 0$ such that

$$\begin{aligned}&\int_{X^i}p_x(x')\,\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx'\\&\quad\in\left(\int_{x-\delta(\varepsilon)}^{x+\delta(\varepsilon)}p_x(x')\,\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx', \int_{x-\delta(\varepsilon)}^{x+\delta(\varepsilon)}p_x(x')\,\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx'+\varepsilon\sup p_x(x')\right)\end{aligned}$$

and

$$\int_{x-\delta(\varepsilon)}^{x+\delta(\varepsilon)}p_x(x')\,\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx' \in ((1-\varepsilon)(p_x(x)-\varepsilon), p_x(x)+\varepsilon).$$

Together these imply that

$$\int_{X^i}p_x(x')\,\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx' \in \left((1-\varepsilon)(p_x(x)-\varepsilon), p_x(x)+\varepsilon\left(1+\sup p_x(x')\right)\right),$$

and therefore $\int_{X^i}p_x(x')\,\frac{1}{\tau}g\left(\frac{x'-x}{\tau}\right)dx'$ is within $\varepsilon(1+\sup p_x(x'))$ of $p_x(x)$ regardless of x, as needed.

The argument for the numerator is the same except that $p_x(x')$ is replaced by

$$U(x';x) = \int_{\Theta}\phi(\theta)\,f\left(\frac{x'-\theta}{\sigma}\right)u\left(\theta, 1-F\left(\frac{x-\theta}{\sigma}\right)\right)d\theta.$$

All that is needed is to verify that $U(x';x)$ is bounded and continuous in x'. Boundedness is immediate from the boundedness of $u(\cdot)$ and $p_x(\cdot)$.

Let A and $B \in \mathbb{R}$ be upper bounds on $f(\cdot)$ and $\phi(\cdot)$ respectively. Given $\varepsilon > 0$, there exists some compact subset $\Theta(\varepsilon) \subset \Theta$ such that

$$U(x';x) \in \left(\int_{\Theta(\varepsilon)} \phi(\theta) f\left(\frac{x'-\theta}{\sigma}\right) u\left(\theta, 1-F\left(\frac{x-\theta}{\sigma}\right)\right) d\theta - \varepsilon AV, \right.$$
$$\left. \int_{\Theta(\varepsilon)} \phi(\theta) f\left(\frac{x'-\theta}{\sigma}\right) u\left(\theta, 1-F\left(\frac{x-\theta}{\sigma}\right)\right) d\theta + \varepsilon AV \right)$$

for all x and x'. Let d be the number of discontinuities of f on $\Theta(\varepsilon)$, which is finite. Since f is bounded, it is uniformly continuous on $\Theta(\varepsilon)$ wherever it is continuous. Accordingly, let $\delta(\varepsilon)$ be such that $|f\left(\frac{x'-\theta}{\sigma}\right) - f\left(\frac{x''-\theta}{\sigma}\right)| < \varepsilon$ whenever $|\frac{x'-\theta}{\sigma} - \frac{x''-\theta}{\sigma}| < \delta(\varepsilon)$ and $\frac{x'-\theta}{\sigma}$ lies at a distance of at least $\delta(\varepsilon)$ from any discontinuity of f. Then changing x' by at most $\delta(\varepsilon)$ changes $\int_{\Theta(\varepsilon)} \phi(\theta) f\left(\frac{x'-\theta}{\sigma}\right) u\left(\theta, 1-F\left(\frac{x-\theta}{\sigma}\right)\right) d\theta$ by at most $\varepsilon V + 2d\delta(\varepsilon) ABV$, which can be made arbitrarily small, as needed.

The argument for the second statement is similar, except that $\frac{1}{\sigma} f\left(\frac{x'-\theta}{\sigma}\right)$ takes on the role of $\frac{1}{\tau} g\left(\frac{x'-x}{\tau}\right)$. □

References

Argenziano, R. and I. Gilboa (2005). History as a coordination device. http://privatewww.essex.ac.uk/~rargenz/Argenziano_Gilboa_Revolutions_Sept2005.pdf.

Beggs, A. (2005). Learning in bayesian games with binary actions. http://www.economics.ox.ac.uk/Research/wp/pdf/paper232.pdf.

Billot, A., I. Gilboa, D. Samet, and D. Schmeidler (2005). Probabilities as similarity-weighted frequencies. *Econometrica 73*, 1125–1136.

Carlsson, H. (2004). Rational and adaptive play in two-person global games. Lund University.

Carlsson, H. and E. van Damme (1993). Global games and equilibrium selection. *Econometrica 61*, 989–1018.

Germano, F. (2004). Stochastic evolution of rules for playing normal form games. Universitat Pompeu Fabra working paper No. 761, http://www.econ.upf.edu/eng/research/onepaper.php?id=761.

Gilboa, I. and D. Schmeidler (2001). *A Theory of Case-Based Decisions*. Cambridge: Cambridge University Press.

Heinemann, F., R. Nagel, and P. Ockenfels (2004). The theory of global games on test: experimental analysis of coordination games with public and private information. *Econometrica 72*, 1583–1599.

Jehiel, P. and F. Koessler (2006). Revisiting games of incomplete information with analogy-based expectations. http://www.enpc.fr/ceras/jehiel/ABEEbayes.pdf.

Jehiel, P. and D. Samet (2004). Valuation equilibria. http://www.enpc.fr/ceras/jehiel/valuation.pdf.

LiCalzi, M. (1995). Fictitious play by cases. *Games and Economic Behavior 11*, 64–89.

Milgrom, P. and J. Roberts (1990). Rationalizability, learning and equilibrium in games with strategic complementarities. *Econometrica 58*, 1255–1277.

Morris, S. (1997). Interaction games: a unified analysis of incomplete information, local interaction, and random matching games. Santa Fe Institute working paper No. 97-08-072E, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=290880.

Morris, S. (2000). Contagion. *Review of Economic Studies 67*, 57–78.

Morris, S. and H. S. Shin (2003). Global games: theory and applications. In M. Dewatripont, L. Hansen, and S. Turnovsky (Eds.), *Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society)*, Cambridge, pp. 56–114. Cambridge University Press.

Nachbar, J. H. (1990). Evolutionary selection dynamics in games: convergence and limit properties. *International Journal of Game Theory 19*, 59–89.

Samuelson, L. and J. Zhang (1992). Evolutionary stability in asymmetric games. *Journal of Economic Theory 57*, 363–391.

Stahl, D. O. and J. Van Huyck (2002). Learning conditional behavior in similar stag hunt games. http://erl.tamu.edu/JVH_gtee/RL1.htm.

Individual researchers, as well as the on-line and printed versions of the CERGE-EI Working Papers (including their dissemination) were supported from the following institutional grants:

- Center of Advanced Political Economy Research [Centrum pro pokročilá politicko-ekonomická studia], No. LC542, (2005-2009),
- Economic Aspects of EU and EMU Entry [Ekonomické aspekty vstupu do Evropské unie a Evropské měnové unie], No. AVOZ70850503, (2005-2010);
- Economic Impact of European Integration on the Czech Republic [Ekonomické dopady evropské integrace na ČR], No. MSM0021620846, (2005-2011);

Specific research support and/or other grants the researchers/publications benefited from are acknowledged at the beginning of the Paper.

Published by
Charles University in Prague, Center for Economic Research and Graduate Education (CERGE)
and
Economics Institute ASCR, v. v. i. (EI)
CERGE-EI, Politických vězňů 7, 111 21 Prague 1, tel.: +420 224 005 153, Czech Republic.
Printed by CERGE-EI, Prague
Subscription: CERGE-EI homepage: http://www.cerge-ei.cz

ISSN 1211-3298
ISBN 978-80-7343-123-5 (Univerzita Karlova. Centrum pro ekonomický výzkum a doktorské studium)
ISBN 978-80-7344-112-8 (Národohospodářský ústav AV ČR, v. v. i.)

CERGE-EI
P.O.BOX 882
Politických vězňů 7
111 21 Praha 1
Czech Republic
http://www.cerge-ei.cz