



Edinburgh School of Economics
Discussion Paper Series
Number 151

Contagion through Learning

Jakub Steiner (University of Edinburgh)
Colin Stewart (University of Toronto)

Date
August 2007

Published by

School of Economics
University of Edinburgh
30 -31 Buccleuch Place
Edinburgh EH8 9JT
+44 (0)131 650 8361

<http://www.ed.ac.uk/schools-departments/economics>



THE UNIVERSITY *of* EDINBURGH

Contagion through Learning*

Jakub Steiner[†] Colin Stewart[‡]
University of Edinburgh University of Toronto

August 9, 2007

Abstract

We study learning in a large class of complete information normal form games. Players continually face new strategic situations and must form beliefs by extrapolation from similar past situations. We characterize the long-run outcomes of learning in terms of iterated dominance in a related incomplete information game with subjective priors. The use of extrapolations in learning may generate contagion of actions across games even if players learn only from games with payoffs very close to the current ones. Contagion may lead to unique long-run outcomes where multiplicity would occur if players learned through repeatedly playing the same game. The process of contagion through learning is formally related to contagion in global games, although the outcomes generally differ.

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

*An earlier version of this paper was distributed under the title “Learning by Similarity in Coordination Problems.” We are grateful to Eddie Dekel, Eduardo Faingold, Philippe Jehiel, George Mailath, Ben Polak, Marzena Rostek, and Larry Samuelson for their comments, and especially to Stephen Morris and Avner Shaked for many valuable discussions throughout the project. We also wish to thank the organizers of the VI Trento Summer School in Adaptive Economic Dynamics, and seminar participants at Cambridge, Edinburgh, MIT, NYU, Penn State, PSE, Queen’s, Stanford, Toronto, UBC, Washington U in St. Louis, Western Ontario, Wisconsin, Yale, the Econometric Society meetings in Minneapolis and Vienna, and the Global Games workshop at Stony Brook. Jakub Steiner benefited from the grant “Stability of the Global Financial System: Regulation and Policy Response” during his research stay at LSE.

[†]email: jakub.steiner@ed.ac.uk

[‡]email: colin.stewart@utoronto.ca

1 Introduction

In standard models of learning, players repeatedly interact in the *same* game, and use their experience from the history of play to decide which action to choose 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. Thus, a tacit assumption of standard learning models is that players extrapolate their experience from 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 state θ . The action *Invest* is strategically risky, as its payoff depends on the action of the opponent. The safe action, *Not Invest*, 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 not investing 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

	Invest	Not Invest
Invest	θ, θ	$\theta - 1, 0$
Not Invest	$0, \theta - 1$	$0, 0$

Table 1: Payoffs in the Example of Section 2.

players estimate payoffs for the game $\Gamma(\theta)$ on the basis of past experience with states 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.

Since investing is dominant for all sufficiently high states, there is some $\bar{\theta}$ above which players eventually learn to invest. Now consider a state just below $\bar{\theta}$, say $\bar{\theta} - \varepsilon$. At $\bar{\theta} - \varepsilon$, investing may not be dominant, but players view some games with values of θ above $\bar{\theta}$ 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 $\bar{\theta}$, where the opponent may not invest. Thus players learn to invest in games with states below, but close to $\bar{\theta}$, giving a new threshold $\bar{\theta}'$ above which both players invest.

Repeating the argument with $\bar{\theta}$ replaced by $\bar{\theta}'$, investment continues to spread to games with smaller states, even though these are not directly similar to games in the dominance region. The process continues until a threshold state 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 states by a symmetric process. 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 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

¹In other words, players place much more weight on states very close to the present one when forming their payoff estimates.

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 contagion is driven solely by a natural solution to the problem of learning the payoffs to one's actions 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.

The main tool for understanding the result of contagion through learning is a formal parallel to equilibrium 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 noisy, while correctly believing that other players perfectly observe the true state. More precisely, players learn not to play strategies that would be serially dominated in the modified version of the game (see Theorem 4.1).

The relationship between the long-run outcomes of similarity-based learning and serially undominated strategies in the modified game is quite robust. The result holds for a broad class of games and a large class of learning processes which vary in the knowledge players have of the environment. In addition, very little structure is imposed on the similarity functions used by the players in the learning process. Roughly speaking, the modified game result holds as long as payoffs and similarity are sufficiently continuous in the state.


The modified game result enables us to identify the long-run outcomes of learning by solving the modified game through an extension of the techniques of Carlsson and van Damme (1993), further developed by Morris and Shin (2003). The original

game has a continuum of equilibria, but contagion leads to a unique learning outcome when similarity is concentrated on nearby states. This outcome involves symmetric, monotone strategies, but depends on the shape of the similarity function. The outcomes of contagion through learning generally differ from those of contagion through incomplete information in global games. Since the process of contagion through learning is exactly parallel to the process of contagion through incomplete information in the modified game, the difference between the global games and learning outcomes can be understood as a product of the heterogeneity of the priors in the modified game.


2 Example

Before introducing the general model in Section 3, we elaborate on the example from the Introduction to illustrate in more detail the process of contagion through learning. 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 in state θ when the opponent chooses action a^{-i} . To simplify notation, we will refer to investing as Action 1 and not investing as Action 0.

The game is played repeatedly in periods $t \in \mathbb{N}$, with the state 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 estimate their payoffs as a weighted average of historical payoffs in which the weights are determined by the similarity between the current and past states. Strategic considerations play no role in these estimates: players treat the past actions of their opponents as given. Thus following any history $h_t = \{\theta_s, a_s^1, a_s^2\}_{s < t}$, the estimated payoff to Player i from



(a) Example similarity function $g(\cdot)$.



(b) Example payoff estimates. Dots represent observed past payoffs.

Figure 1: Example similarity function with corresponding payoff estimates according to Equation (1).

choosing Action a^i given the state θ_t is

$$r(\theta_t; h_t) = \frac{\sum_{s < t} g(\theta_s - \theta_t) U(a^i, a_s^{-i}, \theta_s)}{\sum_{s < t} g(\theta_s - \theta_t)}, \quad (1)$$

where $g(\cdot) \geq 0$ is the *similarity function* determining the relative weight assigned to past cases. Each Player chooses the action giving the highest estimated payoff. Beliefs may be chosen arbitrarily if the history contains no state similar to θ_t , that is, if $\sum_{s < t} g(\theta_s - \theta_t) = 0$.

For this example, suppose that $g(\cdot)$ is the piecewise-linear function illustrated in Figure 1(a). Figure 1(b) illustrates the estimated payoffs as a function of θ for a particular history of observed payoffs using this similarity function.

The learning process is stochastic, but suppose that the empirical distribution of realized cases may be approximated by the true distribution over θ (this idea is formalized in Section 3 below). If the opponent plays according to a fixed strategy $s^{-i}(\cdot)$,

Player i 's expected estimated return to investing in state $\theta \in [-b + \tau, 1 + b - \tau]$ is given by

$$\int_{\Theta} U(1, s^{-i}(\theta'), \theta') g(\theta' - \theta) d\theta'. \quad (2)$$

The expression (2) is formally equivalent to the conditional expectation $E[U(1, s^{-i}(\theta'), \theta') | \theta]$ when θ is an imprecise signal of θ' , with noise $\theta' - \theta$ distributed according to density $g(\cdot)$. Thus, in the long-run, the similarity-based learner behaves *as if* she observes only a noisy signal of the true state. Theorem 4.1 makes this connection precise by showing that players learn to play strategies that would be serially undominated in a *modified* game of incomplete information in which each holds these (subjective) beliefs about the information structure.

The long-run outcome of the learning process may be identified by solving this modified game. Suppose that both players follow cut-off strategies with threshold θ^* ; that is, both choose Action 0 at signals below θ^* and Action 1 at signals above θ^* . Each player assigns probability $\frac{1}{2}$ to the true state being greater than her own signal. Since each believes that the other player observes the true state, a player who receives exactly the threshold signal θ^* believes that the other player chooses Action 1 with probability $\frac{1}{2}$. In order for this strategy profile to be a Bayesian Nash equilibrium, players must be indifferent between their two actions at the threshold signal. Therefore, θ^* is the unique solution to

$$\frac{1}{2}U(1, 0, \theta) + \frac{1}{2}U(1, 1, \theta) = 0. \quad (3)$$

This equilibrium turns out to be the unique serially undominated strategy profile in the modified game, and therefore the unique long-run outcome of the learning process in the original game.

The condition defining the threshold in Equation 3 is identical to the condition defining the unique equilibrium threshold in global games with the same payoff structure (but with common priors over the structure of the noise). This agreement does not generally hold if similarity functions are asymmetric or there are more than two players. We leave a detailed discussion of this comparison to Section 5 below.

3 The Learning Process

The learning process is comprised of an estimation procedure by which players estimate payoffs of actions together with a strategic environment in which they repeatedly interact. We are interested in environments in which players continually face new strategic problems, and each time are uncertain of the payoffs they will receive from each of their available actions. This uncertainty may result from imperfect knowledge of the payoff function, or simply from not knowing opponents' strategies. Since the strategic problem is new, players draw on their experience with similar past problems to estimate their current payoffs.

A fixed set of $I \geq 2$ players interact in a one-parameter family of games $\Gamma(\theta)$. The state θ is drawn from a compact, convex set $\Theta \subset \mathbb{R}^N$. The action set A^i available to each player i is the same across all states in Θ . Each set A^i is assumed to be finite. As usual, we will write $A = \times_{i=1}^I A^i$ for the set of action profiles.

The payoff function of each player i will be denoted by $U^i(\theta, a)$ for $a \in A$. In order to capture varying degrees of initial knowledge of payoff functions, we write the payoff function as

$$U^i(\theta, a^i, a^{-i}) = u^i(\theta, a^i, v^i(\theta, a^i, a^{-i})),$$

with $v^i : \Theta \times A \rightarrow V^i$ for some arbitrary set V^i , and $u^i : \Theta \times A^i \times V^i \rightarrow \mathbb{R}$. The interpretation of these functions is that player i knows the functional form of $u^i(\cdot, \cdot, \cdot)$, but cannot predict the value of the function $v^i(\cdot, \cdot)$ based on the values of θ and a^i alone. Each player estimates the value of $v^i(\cdot, \cdot)$ using values observed in similar past situations.

The learning process is comprised of three elements:

1. An *initial strategy* $s_0^i : \Theta \rightarrow A^i$ for each player i .
2. A *similarity function* $g^i : \Theta \times \Theta \rightarrow \mathbb{R}_+$ for each player i , where, for each θ , $g(\theta, \cdot)$ is integrable.
3. A continuous distribution Φ over states Θ with full support and continuous density

$\phi(\cdot)$.²

In each period $t = 0, 1, \dots$, Nature draws a state θ_t according to the distribution Φ . Draws are independent across periods. Each player perfectly observes θ_t , then chooses an action $a^i \in A^i$ based on her experience in similar states θ_s , weighted according to the degree of similarity $g^i(\theta_s, \theta_t)$.

If $\sum_{s < t} g^i(\theta_s, \theta_t) = 0$, then Player i simply follows her initial strategy and chooses action $a_t^i = s_0^i(\theta_t)$. The interpretation of this case is that Player i perceives all past data to be irrelevant to the problem in state θ_t , and hence ignores it. All of our results below are independent of the initial strategies used by players in the learning process.³

If, on the other hand, $\sum_{s < t} g^i(\theta_s, \theta_t) > 0$, then Player i chooses action a_t^i according to

$$a_t^i \in \operatorname{argmax}_{a^i \in A^i} \frac{\sum_{s < t} u^i(\theta_t, a^i, v^i(\theta_s, a^i, a_s^{-i})) g^i(\theta_s, \theta_t)}{\sum_{s < t} g^i(\theta_s, \theta_t)}. \quad (4)$$

This action choice may be interpreted as if each player assigns particular beliefs to the unknown value $V_t^i(a^i) = v^i(\theta_t, a^i, a_t^{-i})$ for each a^i ; namely,

$$Pr(V_t^i(a^i) = v) = \frac{\sum_{s < t} g^i(\theta_s, \theta_t) \mathbb{1}(v(\theta_s, a^i, a_s^{-i}) = v)}{\sum_{s < t} g^i(\theta_s, \theta_t)}.$$

These are precisely the beliefs that would arise if players used kernel estimators to estimate the distribution of $V_t^i(a^i)$ conditional on the state θ_t using the kernel function $g^i(\cdot, \theta_t)$.

Following any history $h_t = (\theta_s, a_s)_{s=0}^{t-1}$, the learning process defines a strategy $s_t^i : \Theta \rightarrow A^i$ for each player describing the action that would be chosen at each possible state θ if the realized state was $\theta_t = \theta$. The process therefore gives rise to a probability distribution over sequences of strategy profiles (s_0, s_1, \dots) . All probabilistic results below are with respect to this distribution.

In order to form the payoff estimates in (4), Player i must only observe values

²The modified game result, Theorem 4.1, also holds for discrete distributions over Θ ; in fact, the proof for the discrete case is much simpler. Our main focus, however, will be on the continuous case, which better captures the idea that players cannot learn from repeated interaction in the same game.

³All of our results hold without modification if, instead of an initial strategy, players begin with an arbitrary finite history of play that is sufficiently rich to prevent the case $\sum_{s < t} g^i(\theta_s, \theta_t) = 0$ from occurring.

of $v^i(\cdot, \cdot, \cdot)$ at the end of each period s given the state θ_s and the actions a_s^{-i} of the opponents in that period. However, Player i must observe the value of $v^i(\theta_s, a^i, a_s^{-i})$ for every action $a^i \in A^i$, regardless of the action she actually chose in Period s . In some instances of the learning process, the value of $v^i(\theta_s, a^i, a_s^{-i})$ does not depend on a^i . In other cases, however, players must observe certain counterfactual values of v^i . The role of these counterfactuals is to prevent players from failing to learn simply because they never take a particular action. The model may thus be viewed as an approximation of one in which players choose according to the preceding rules with high probability in each period, but experiment with some small independent probability by choosing a random action from A^i .

The learning process may be viewed as an extension of fictitious play to an environment in which the game is changing. If the state θ_t is fixed over all periods t , then the learning process reduces to standard fictitious play.

The general learning model encompasses two procedures of particular interest:

- **Payoff-based learning:** A particular example of the model arises when players initially know nothing about the payoff structure. Let $u^i(\theta, a^i, v^i) \equiv v^i$ and let $v^i(\theta, a^i, a^{-i}) \equiv U^i(\theta, a^i, a^{-i})$. Under this specification, at the end of each round s , the players observe only the payoffs they received or would have received for each action $a^i \in A^i$. Each player then chooses her action to maximize a weighted average of these observed payoffs, where the weights are determined by the similarity between the current state and each previous one. Players in this learning process are strategically naïve in the sense that they do not reason about the actions of other players; indeed, they treat the problem simply as a single-person decision problem with unknown payoffs.
- **Strategy-based learning:** Another instance of the general learning model is the case of players who fully understand the payoff function $U(\theta, a^i, a^{-i})$ and need only estimate the opponents' action profile at θ_t . Formally, take $v^i(\theta, a^i, a^{-i}) \equiv a^{-i}$ and $u^i(\theta, a^i, v^i) \equiv U(\theta, a^i, v^i)$. The informational feedback required for this process is minimal: each player observes only her opponents' actions a_s^{-i} at the

end of each period s .

The players in the learning process are case-based decision makers (Gilboa and Schmeidler (2001)) in a strategic environment. Alternatively, our players may be viewed as statisticians satisfying the axioms of Billot et al. (2005). These statisticians seek to predict the outcome of an action based on a database of outcomes in similar cases. The authors prove that if the order of cases in the database is irrelevant, and if a combination of two databases leads to a convex combination of beliefs generated by the original databases, then the agents form beliefs using similarity-weighted averages.⁴

The irrelevance axiom of Billot et al. (2005) does not preclude players from putting more weight on more recent cases. We rule out time-dependent similarity functions in order to simplify the analysis. More generally, one could suppose that observations are discounted over time according to a nonincreasing sequence $\delta(\tau) \in (0, 1]$ by modifying equation (4) to include an additional factor of $\delta(t-s)$ in both sums. In the undiscounted model, the convergence results presented below rely on the property that changes in payoff estimates in a single round 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 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 round. However, as long as memory is “sufficiently long,” we expect this agreement to occur in a large fraction of rounds. 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.

We impose the following technical assumptions on the learning process:

A1. Bounded payoffs: there exist upper and lower bounds on $u^i(\theta, a^i, v^i)$ uniformly

⁴This connection is subject to the caveat that Billot et al. assume a finite outcome space, whereas the outcome space is infinite here.

over all $(\theta, a^i, v^i) \in \Theta \times A^i \times V^i$.

A2. Each similarity function g^i is bounded.

The following assumption ensures that players eventually obtain relevant data for every state:

A3. For each θ , $\int_{\Theta} \phi(\theta') g^i(\theta', \theta) d\theta' > 0$.

We require the following continuity assumption:

A4. For every a^i and a^{-i} , the expression $u^i(\theta, a^i, v^i(\theta', a^i, a^{-i})) g^i(\theta', \theta)$ is continuous in θ uniformly over all θ' .⁵

Note that the continuity in Assumption A4 is uniform over $(\theta, \theta', a^i, a^{-i}) \in \Theta \times \Theta \times A^i \times A^{-i}$ because Θ is compact and the actions sets are finite. Also note that in the case of Payoff-Based Learning described above, Assumption A4 holds if $g^i(\theta', \theta)$ is continuous.

4 Long-run Characterization

In this section, we characterize the long-run outcomes of the learning process from Section 3 in terms of the equilibria of a particular game, which we call the modified game. We begin by informally outlining an observation that lies at the core of this connection, before the formal statement in Theorem 4.1.

Suppose that the learning process converges to a time-invariant strategy profile $s(\theta)$. By the law of large numbers, Player i 's long-run estimated payoff for action a^i in state θ_t against the profile $s^{-i}(t)$ approaches

$$\frac{\int_{\Theta} u^i(\theta_t, a^i, v^i(\theta, a^i, s^{-i}(\theta))) \phi(\theta) g^i(\theta, \theta_t) d\theta}{\int_{\Theta} \phi(\theta') g^i(\theta', \theta_t) d\theta'} = \int_{\Theta} u^i(\theta_t, a^i, v^i(\theta, a^i, s^{-i}(\theta))) q^i(\theta|\theta_t) d\theta,$$

where

$$q^i(\theta|\theta_t) = \frac{\phi(\theta) g^i(\theta, \theta_t)}{\int_{\Theta} \phi(\theta') g^i(\theta', \theta_t) d\theta'}. \quad (5)$$

⁵Formally, for each a^i and a^{-i} , given any $\varepsilon > 0$, there exists some $\delta > 0$ (independent of θ and θ') such that $|u^i(\theta, a^i, v^i(\theta', a^i, a^{-i})) g^i(\theta', \theta) - u^i(\theta'', a^i, v^i(\theta', a^i, a^{-i})) g^i(\theta', \theta'')| < \varepsilon$ whenever $|\theta - \theta''| < \delta$.

That is, Player i 's expected estimated payoff at state θ_t against the strategy profile $s^{-i}(\theta)$ coincides with the expected payoff against the same strategy profile of a player with payoffs $u^i(\theta_t, a^i, v^i(\theta, a^i, s^{-i}(\theta)))$ and beliefs $q^i(\theta|\theta_t)$ over the state θ .

The virtual conditional belief $q^i(\theta|\theta_t)$ has a convenient interpretation. Suppose that the state θ is drawn according to the distribution $\Phi(\cdot)$, and Player i observes only a noisy signal θ_t of θ , where the signal is conditionally distributed according to density $\frac{g^i(\theta, \cdot)}{\int_{\Theta} g^i(\theta, \theta') d\theta'}$. Then $q^i(\theta|\theta_t)$ is precisely the density describing Player i 's posterior beliefs over θ after observing the signal θ_t . This interpretation motivates the following definition:

Definition 4.1. *The modified game is a Bayesian game with heterogenous priors. The players $i \in \{1, \dots, I\}$ simultaneously choose actions $a^i \in A^i$, where each A^i is finite. The state space is given by $\Omega = \Theta^{I+1}$, with typical member $(\theta, \theta^1, \dots, \theta^I)$, where each θ^i denotes the type of Player i , and θ is a common payoff parameter. Each player i has payoff function $u^i(\theta^i, a^i, v^i(\theta, a^i, a^{-i}))$. Player i assigns probability 1 to the event that $\theta^j = \theta$ for all $j \neq i$, and has prior beliefs over (θ, θ^i) given by the density $\phi(\theta)g^i(\theta, \theta^i)$.*

Whereas the class of games $\Gamma(\theta)$ describes the actual environment in which the players interact, the modified game describes a virtual setting with no direct interpretation in terms of reasoning or beliefs in the learning process. The modified game merely provides a useful tool for studying the learning outcomes because the learning process converges, in a sense that will be made precise below, to the set of strategies that are serially undominated in the modified game.

For any game with subjective priors, we may define (interim) dominated strategies in the same way as for Bayesian games with common priors.⁶ In fact, we will also use a stronger form of dominance in which the payoff difference exceeds some fixed $\pi \geq 0$. Consider any function $\mathbf{a}^i : \Theta \rightarrow 2^{A^i}$. We interpret $\mathbf{a}^i(\theta)$ as the set of admissible actions for Player i at type θ . The profiles $(\mathbf{a}^i)_i$ and $(\mathbf{a}^j)_{j \neq i}$ will be denoted, as usual, as \mathbf{a} and \mathbf{a}^{-i} respectively.

⁶Since no other notion of domination will be employed here, we henceforth drop the term ‘‘interim’’ and refer simply to ‘‘dominated strategies.’’

Definition 4.2. A strategy $s^i(\cdot)$ is said to be consistent with $\mathbf{a}^i(\cdot)$ if $s^i(\theta) \in \mathbf{a}^i(\theta)$ for all $\theta \in \Theta$. A strategy profile $s^{-i}(\cdot)$ is said to be consistent with $\mathbf{a}^{-i}(\cdot)$ if each component of $s^{-i}(\cdot)$ is consistent with the corresponding component of $\mathbf{a}^{-i}(\cdot)$.

For any profile $\mathbf{a}(\cdot)$, action $a^i \in \mathbf{a}^i(\theta)$ is said to be π -dominated⁷ at θ under the profile $\mathbf{a}(\cdot)$ if there exists $a^{i'} \in \mathbf{a}^i(\theta)$ such that for all $s^{-i}(\cdot)$ consistent with $\mathbf{a}^{-i}(\cdot)$,

$$E_{q(\theta'|\theta)} \left[u^i \left(\theta, a^{i'}, v^i \left(\theta', a^{i'}, s^{-i}(\theta') \right) \right) - u^i \left(\theta, a^i, v^i \left(\theta', a^i, s^{-i}(\theta') \right) \right) \middle| \theta \right] > \pi.$$

We define iterated elimination of π -dominated strategies in the usual way. For each i and $\pi > 0$, let $\mathbf{a}_{0,\pi}^i(\theta) \equiv A^i$. For $k = 1, 2, \dots$, define $\mathbf{a}_{k,\pi}^i(\theta)$ to be the set of actions that are not π -dominated for Type θ of Player i under the profile $\mathbf{a}_{k-1,\pi}(\cdot)$. The set of *serially π -undominated* actions for Type θ of Player i is given by $\mathbf{a}_{\infty,\pi}^i(\theta) = \bigcap_k \mathbf{a}_{k,\pi}^i(\theta)$. Since π -domination agrees with the usual notion of strict dominance when $\pi = 0$, we will drop the prefix π in that case.

The need to consider π -domination instead of ordinary strict dominance arises because of the difference between estimated payoffs following finite histories and their long-run expectations. In the proof of Theorem 4.1 below, we show that for any $\pi > 0$, estimated payoffs under the learning process almost surely eventually lie within π of the corresponding expected 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 dominance.

Lemma 4.1. Fix any type θ of Player i in the modified game and any $k \in \mathbb{N}$. If $a^i \notin \mathbf{a}_{k,0}^i(\theta)$, then there exists some $\pi > 0$ such that $a^i \notin \mathbf{a}_{k,\pi}^i(\theta)$.

Proof. See appendix. □

The main result of this section, given in the following theorem, shows that, in the long-run, players will not play strategies that are serially dominated in the modified

⁷The 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.

game. Note that strategies in each period of the learning process are defined over states Θ , which is identical to the type space for each player in the modified game. Strategies $s_t^i(\cdot)$ under the learning process may therefore be identified with strategies $s^i(\cdot)$ in the modified game; to keep notation simple, we will not distinguish between the two.

Theorem 4.1. *1. For any $k \in \mathbb{N}$ and any $\pi > 0$, the strategy profiles $s_t(\cdot)$ under the learning process are almost surely eventually consistent with $\mathbf{a}_{k,\pi}(\cdot)$.⁸*

2. The probability that the action profile in period t under the learning dynamics is consistent with the set of serially undominated strategies at θ_t in the modified game approaches one as time tends to infinity; that is,

$$\Pr (s_t^i(\theta_t) \text{ is consistent with } \mathbf{a}_{\infty,0}^i(\theta_t) \forall i) \rightarrow 1$$

as $t \rightarrow \infty$.

Proof. See appendix. □

Using the Strong Law of Large Numbers, it is relatively straightforward to show that in a given state against a fixed strategy, the long-run payoff estimate in the learning process approaches 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 iterated elimination of dominated strategies to occur under the learning dynamics, players must learn not to play dominated actions in finite time at an uncountable set of states. 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.

The concept of Bayesian Nash equilibrium with subjective priors is difficult to justify based on learning in a fixed game (see Dekel, Fudenberg, and Levine (2004)). When players learn by similarity, long-run behaviour corresponds naturally to equilibrium behaviour with subjective priors. Even when similarity is very narrowly concentrated,

⁸Recall that a property holds *eventually* if there exists some T such the property holds for all $t \geq T$.

so that, in a sense, differences in the corresponding priors are “small”, the consequences for behaviour can be significant, as demonstrated in the following section.

5 Coordination Games

We now focus on learning by similarity in a class of symmetric coordination games $\Gamma(\theta)$, where the distribution $\Phi(\theta)$ has support $\Theta = [\underline{\theta}, \bar{\theta}]$. Each of I players chooses between two actions, 0 and 1. The players share an identical payoff function. We normalize the payoff from Action 0 to be 0 in every state θ against every action profile. We denote by $U(\theta, l)$ the payoff from choosing Action 1 in state θ when $l \in \{0, \dots, I-1\}$ opponents choose action 1.

The similarity function is identical across players, and depends only on the difference $\theta' - \theta$ between states, and a parameter $\tau > 0$ according to

$$g^i(\theta', \theta) \equiv \frac{1}{\tau} g\left(\frac{\theta' - \theta}{\tau}\right),$$

where $g : \mathbb{R} \rightarrow \mathbb{R}_+$. We will focus on outcomes in the limit as τ tends to 0, where similarity is narrowly concentrated on nearby states. Away from this limit, similarity-based learning generally leads to inconsistent payoff estimates even if the opponents' strategies are fixed. When τ is small, if play converges, these inconsistencies become small except possibly in states close to discontinuities in payoffs.

As before, the learning process may take different forms depending on the feedback players receive over time. To capture this, we write the payoff to Action 1 as

$$U(\theta, l) = u(\theta, v(\theta, l)).$$

Whenever $\sum_{s < t} g\left(\frac{\theta_s - \theta_t}{\tau}\right) > 0$, the estimated payoff for Action 0 is simply 0, and that for Action 1 is given by

$$\frac{\sum_{s < t} u(\theta_t, v(\theta_s, l_s)) \frac{1}{\tau} g\left(\frac{\theta_s - \theta_t}{\tau}\right)}{\sum_{s < t} \frac{1}{\tau} g\left(\frac{\theta_s - \theta_t}{\tau}\right)}.$$

In addition to the general assumptions from Section 3, we assume the following:

- A5. **State Monotonicity:** The payoffs $u(\theta, 0)$ and $u(\theta, I - 1)$ are strictly increasing in θ .
- A6. **Extremal Payoffs at Extremal Profiles:** For all $l = 0, \dots, I - 1$ and all $\theta \in \Theta$, $u(\theta, 0) \leq u(\theta, l) \leq u(\theta, I - 1)$.⁹
- A7. **Dominance Regions:** There exist some $\underline{\theta}', \bar{\theta}' \in (\underline{\theta}, \bar{\theta})$ such that Action 0 is dominant at every state below $\underline{\theta}'$ and Action 1 is dominant at every state above $\bar{\theta}'$.
- A8. **Continuity:** The payoffs $u(\theta, 0)$ and $u(\theta, I - 1)$ are continuous in θ .

Note that, since $g^i(\theta, \cdot)$ is a probability density function, so is $g(\cdot)$. Let $G(\cdot)$ be the distribution function corresponding to the density $g(\cdot)$. Define the threshold θ^* to be the (unique) solution to

$$G(0)u(\theta, 0) + (1 - G(0))u(\theta, I - 1) = 0. \quad (6)$$

The existence of this solution is guaranteed by the existence of dominance regions (Assumption A7), and its uniqueness by state monotonicity (Assumption A5).

Proposition 5.1. *For any $\delta > 0$, there exists $\bar{\tau} > 0$ such that for any $\tau \in (0, \bar{\tau})$, in the learning process with parameter τ , the probability that all players choose Action 0 conditional on $\theta_t < \theta^* - \delta$ and Action 1 conditional on $\theta_t > \theta^* + \delta$ tends to 1 as $t \rightarrow \infty$.*

This proposition provides a stark contrast to learning in a fixed game. If instead of varying in each period, the state θ was fixed over all periods, then the learning process reduces to standard fictitious play (as long as $g^i(\theta, \theta) > 0$). For any θ outside of the dominance regions, there would be multiple long-run learning outcomes that depend on the initial strategies used by players in the game. For instance, if all players are initially coordinated on one of the two actions, then they will continue to choose this

⁹The extremal payoff assumption is essentially a weakened form of strategic complementarity, which differs from full strategic complementarity in not requiring that $u(\theta, l') \geq u(\theta, l)$ if $0 < l < l' < I - 1$.

action in every period. In contrast, Proposition 5.1 indicates that extrapolation from different past states may lead to a unique long-run outcome in many of these states θ , independent of the initial strategies players use in the learning process.

The following proof draws on techniques from the proofs of Propositions 2.1 and 2.2 in Morris and Shin (2003) for the corresponding result in global games.

Proof of Proposition 5.1. Define $m_\tau(\theta, k)$ to be the expected payoff to Action 1 for type θ in the modified game when the opponents play a threshold strategy with threshold k . That is

$$m_\tau(\theta, k) \equiv \frac{\int_{\underline{\theta}}^k \phi(\theta') \frac{1}{\tau} g\left(\frac{\theta' - \theta}{\tau}\right) u(\theta, v(\theta', 0)) d\theta' + \int_k^{\bar{\theta}} \phi(\theta') \frac{1}{\tau} g\left(\frac{\theta' - \theta}{\tau}\right) u(\theta, v(\theta', I - 1)) d\theta'}{\int_{\Theta} \phi(\tilde{\theta}) \frac{1}{\tau} g\left(\frac{\tilde{\theta} - \theta}{\tau}\right) d\tilde{\theta}}. \quad (7)$$

First, we prove that Action 0 is serially dominated for $\theta > \bar{\theta}^*$ and Action 1 is serially dominated for $\theta < \underline{\theta}^*$, where $\bar{\theta}^*$ and $\underline{\theta}^*$ are, respectively, the maximal and minimal roots of $m_\tau(\theta, \theta) = 0$.¹⁰ Note that the function $m_\tau(\theta, k)$ is continuous and decreasing in k . Moreover, for sufficiently small τ , the existence of dominance regions (Assumption A7) implies that $m_\tau(\theta, k)$ is negative for small enough values of θ and positive for large enough values.

Let $\bar{\theta}_0 = \bar{\theta}$, and for $k = 1, 2, \dots$, recursively define $\bar{\theta}_k$ to be the maximal solution to the equation

$$m_\tau(\theta, \bar{\theta}_{k-1}) = 0.$$

Let S_k denote the set of strategies remaining for each player after k rounds of deletion of dominated strategies. We will prove by induction that Action 0 is dominated for all types of each player above $\bar{\theta}_k$ against profiles of strategies from the set S_{k-1} . Suppose that the claim holds for $k - 1$. By Assumption A6, if opponents play strategies in S_{k-1} , then the payoff to Action 1 for any type θ is at least as large as if every opponent

¹⁰One could alternatively prove this statement by applying Theorem 5 of Milgrom and Roberts' (1990) to the *ex ante* game (in which heterogeneous priors are no longer an issue). However, the direct argument given here better illustrates the process of contagion, and avoids technical issues that arise in moving to the *ex ante* game.

played a cut-off strategy with threshold $\bar{\theta}_{k-1}$ (i.e. a strategy choosing Action 0 at any type below $\bar{\theta}_{k-1}$ and Action 1 at any type above $\bar{\theta}_{k-1}$). Hence the expected payoff for Action 1 at θ is at least $m_\tau(\theta, \bar{\theta}_{k-1})$ regardless of which strategies from S_{k-1} are chosen by the opponents. This expected payoff must be positive above the maximal root $\bar{\theta}_k$ because $m_\tau(\theta, \cdot)$ is continuous everywhere and positive for sufficiently large θ . Therefore, Action 0 is dominated above $\bar{\theta}_k$, as claimed.

Next, we show by induction that $(\bar{\theta}_k)_{k=1}^\infty$ is a nonincreasing sequence. Note first that $\bar{\theta}_1 \leq \bar{\theta}_0$ trivially because $\bar{\theta}_0$ lies at the upper boundary of Θ . Suppose that $\bar{\theta}_{k-1} \leq \bar{\theta}_{k-2}$. Then $m_\tau(\theta, \bar{\theta}_{k-1}) \geq m_\tau(\theta, \bar{\theta}_{k-2})$ because $m_\tau(\theta, k)$ decreases in k , and hence the maximal root of $m_\tau(\theta, \bar{\theta}_{k-1}) = 0$ must be weakly smaller than that of $m_\tau(\theta, \bar{\theta}_{k-2}) = 0$, which establishes the induction step.

The nonincreasing sequence $(\bar{\theta}_k)_{k=1}^\infty$ converges to some $\bar{\theta}^*$ which, from the continuity of m_τ , must be a solution to $m_\tau(\theta, \bar{\theta}^*) = 0$. Therefore, Action 0 is indeed serially dominated at every type above $\bar{\theta}^*$. The symmetric argument from below establishes that Action 1 is serially dominated below the minimal solution $\underline{\theta}^*$ of $m_\tau(\theta, \theta) = 0$.

Note that since $\int_{\theta-\varepsilon}^{\theta+\varepsilon} \frac{1}{\tau} g\left(\frac{\theta'-\theta}{\tau}\right) d\theta' = \int_{-\frac{\varepsilon}{\tau}}^{\frac{\varepsilon}{\tau}} g(z) dz$, given any $\delta > 0$ and $\varepsilon > 0$, there exists some $\bar{\tau} > 0$ such that for all $\tau \in (0, \bar{\tau})$, $\int_{\theta-\varepsilon}^{\theta+\varepsilon} \frac{1}{\tau} g\left(\frac{\theta'-\theta}{\tau}\right) d\theta' > 1 - \delta$. In particular, for any function $\psi(\cdot)$ that is continuous at θ , we have

$$\lim_{\tau \rightarrow 0} \int_{\Theta} \psi(\theta') \frac{1}{\tau} g\left(\frac{\theta'-\theta}{\tau}\right) d\theta' = \psi(\theta), \quad (8)$$

and similarly

$$\lim_{\tau \rightarrow 0} \int_{-\infty}^{\theta} \psi(\theta') \frac{1}{\tau} g\left(\frac{\theta'-\theta}{\tau}\right) d\theta' = \psi(\theta) G(0), \quad (9)$$

$$\text{and} \quad \lim_{\tau \rightarrow 0} \int_{\theta}^{+\infty} \psi(\theta') \frac{1}{\tau} g\left(\frac{\theta'-\theta}{\tau}\right) d\theta' = \psi(\theta) (1 - G(0)). \quad (10)$$

Moreover the convergence of the limits in Equations (8) through (10) is uniform over θ in some set X as long as the function $\psi(\theta)$ is uniformly continuous on X .

Applying Equations (8) through (10) to the definition of $m_\tau(\theta, \theta)$ from Equation

(7) gives

$$\begin{aligned} \lim_{\tau \rightarrow 0} m_\tau(\theta, \theta) &= G(0)u(\theta, v(\theta, 0)) + (1 - G(0))u(\theta, v(\theta, I - 1)) \\ &= G(0)U(\theta, 0) + (1 - G(0))U(\theta, I - 1) \end{aligned}$$

on the open interval $(\underline{\theta}, \bar{\theta})$. The convergence is uniform on any compact subinterval of Θ since $\phi(\theta)$, $u(\theta, 0)$ and $u(\theta, I - 1)$ are uniformly continuous on compact sets. We can choose such a compact subinterval $\bar{\Theta}$ of $(\underline{\theta}, \bar{\theta})$ to intersect with both dominance regions, so that all roots of $m_\tau(\theta, \theta) = 0$ must lie in $\bar{\Theta}$. Define $m(\theta) \equiv \lim_{\tau \rightarrow 0} m_\tau(\theta, \theta)$ for $\theta \in \bar{\Theta}$. Given any neighbourhood N of the unique root θ^* of $m(\theta)$, there exists some $\varepsilon > 0$ such that $m(\theta)$ is uniformly bounded away from 0 by ε outside of N . Choosing $\bar{\tau} > 0$ small enough so that whenever $\tau < \bar{\tau}$, $m_\tau(\theta, \theta)$ is within ε of $m(\theta)$ everywhere on $\bar{\Theta}$ guarantees that $m_\tau(\theta, \theta)$ has no root in $\bar{\Theta} \setminus N$. \square

Theorem 4.1 identifies a formal parallel between contagion through learning and contagion through incomplete information in the modified game. This connection explains in part why many features of the two kinds of contagion appear similar. However, the information structure of the modified game is inherently different from that of global games with a common prior. Consequently, important differences arise between the outcomes of contagion through learning and those of contagion in global games.

First, the equilibrium threshold in the standard binary action global game model is independent of the noise distribution, while the threshold in the similarity learning model depends on the similarity function $g(\cdot)$ (which determines the noise distribution in the modified game). The noise-independence result in global games is driven by the common prior, which generates, in equilibrium, uniform beliefs over l at the threshold type regardless of the noise distribution. With learning by similarity, beliefs over l at the threshold in the modified game depend on $g(\cdot)$ because of the heterogeneity of the priors. In fact, Equation (6) indicates that the threshold θ^* falls if we replace the similarity function $g(\cdot)$ with some function $\tilde{g}(\cdot)$ that first-order stochastically dominates

$g(\cdot)$. The similarity function $\tilde{g}(\cdot)$ may be interpreted as more optimistic than $g(\cdot)$ since it places more weight on better past states. Under this interpretation, optimism leads to more efficient coordination. Izmalkov and Yildiz (2007) obtain a similar result in a game with subjective priors closely related to the modified game studied here.

Moreover, without further restricting the similarity function, any threshold θ^* outside the dominance regions can arise as an outcome of learning for an appropriately chosen similarity function. This indeterminacy is consistent with the result of Weinstein and Yildiz (2007) that any equilibrium of any complete information game can be uniquely selected by an appropriate choice of perturbation in the universal type space. In the environment considered here, the particular perturbations that arise in the modified game as $g(\cdot)$ varies suffice to uniquely select any equilibrium in each complete information game $\Gamma(\theta)$.

Second, the outcome of contagion through learning generally differs from that in global games even if we restrict attention to symmetric similarity functions, for which $G(0) = \frac{1}{2}$. This restriction must hold if similarity weights are symmetric in the sense that θ receives the same weight at θ' that θ' receives at θ . Under this restriction, the threshold θ^* solves

$$\frac{u(\theta^*, 0) + u(\theta^*, I - 1)}{2} = 0, \quad (11)$$

while the standard global game threshold solves

$$\sum_{l=0}^{I-1} \frac{u(\theta^*, l)}{I} = 0. \quad (12)$$

These thresholds generally differ if payoffs are not linear in l . Moreover, the outcome threshold of the learning model is independent of the payoffs $u(\cdot, l)$ for values of l other than 0 and $I - 1$. In the modified game, the threshold player places zero probability on these intermediate values of l , and thus the model does not require full monotonicity of $u(\cdot, l)$ with respect to l . In global games with common priors, the threshold player has uniform beliefs over l , and full strategic complementarity is required for equilibrium uniqueness.

6 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 (Germano (2004), Katz (1996), LiCalzi (1995), Mengel (2007)). Jehiel and Koessler (2006) study an equilibrium concept in Bayesian games in which players partition states into analogy classes. Each player best responds to the strategy obtained by averaging opponents' strategies within each analogy class. These analogy classes can be viewed as arising from particular similarity functions in our learning model; however, the motivation is quite different. Whereas we propose similarity as a means of making inferences from different past situations, Jehiel and Koessler's analogies are based on players' inability to distinguish information obtained in various states. Their focus therefore centers on interesting deviations from standard equilibrium behavior arising from persistent errors in beliefs. We, on the other hand, focus on the case in which these errors are small, leading to a selection among equilibria. Furthermore, the contagion process that forms the focus of the present paper is precluded by their partition model of analogies since states cannot influence one another across elements of the partition.

A related departure from Bayesian Nash equilibrium is Eyster and Rabin's (2005) concept of cursed equilibrium. Players beliefs assign some probability to their opponents' true strategy profile and the remaining probability to the simplified, constant action profile, which is required to be correct "on average". Like Jehiel and Koessler, Eyster and Rabin use their model to capture large deviations from rational play.

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 varied under two different experimental treatments. The observed long-run behavior in a particular game contained in both sets varied across the treatments, indicating that subjects were influenced by their experience playing different games.

Carlsson (2004) proposes an evolutionary justification of global games equilibrium

using strategy-based learning by similarity. Carlsson offers an informal argument to suggest that the learning process can be approximated by the best-response dynamics of a modified game. Theorem 4.1 above formalizes this result in terms of long-run outcomes. The outcome of the learning process coincides with the global game prediction in Carlsson’s two-player model. With more than two players, Proposition 5.1 above indicates that the learning outcome generally shares only the qualitative features of the global game solution, while quantitatively they differ.

Argenziano and Gilboa (2005) study multiplicity of similarity-based learning outcomes in coordination problems. With finitely many states, the long-run outcome depends on historical accidents when games with dominant actions are sufficiently rare.

Milgrom and Roberts (1990) study supermodular games, of which the coordination environment studied here is 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 state spaces, where play of the game (at most) reveals the actions $s(\theta_t)$ assigned by strategies $s(\cdot)$ to the particular states θ_t 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 4.1 below.¹¹

An alternative approach to learning in certain binary action supermodular games is offered by Beggs (2005). Play almost surely converges to an equilibrium of the game if players follow learning rules that adapt threshold strategies based on payoffs from similar types.

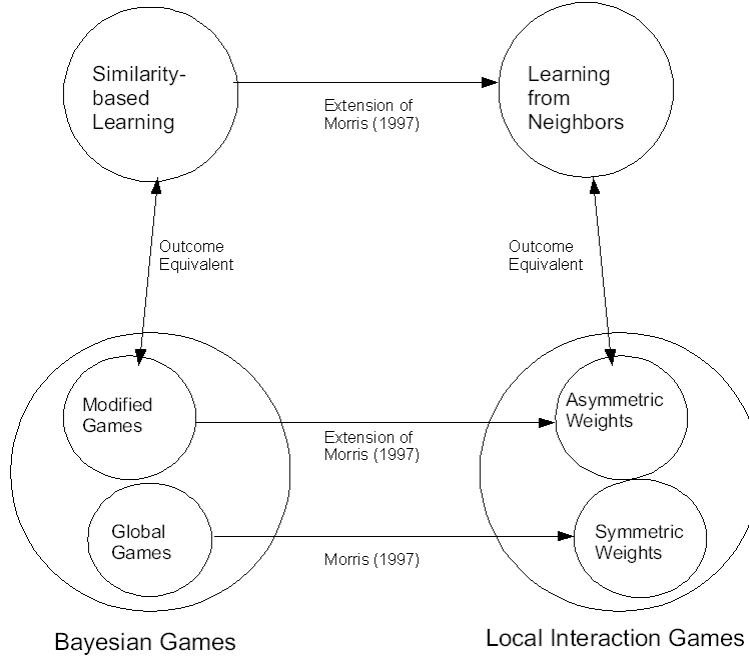


Figure 2: Atlas of Contagion

7 Discussion and Conclusion

Morris (1997) identified a formal relationship between contagion across types in incomplete information games and contagion across players in local interaction games. Starting with some incomplete information game, one may reinterpret the types in the game as players situated in various locations. Each of these players interacts with some subset of the population, her *neighbours*, and must choose the same action against all opponents. Payoffs in the local interaction game are obtained by a weighted sum of payoffs from interactions with all neighbours, where the weights correspond to the posterior beliefs over opponents' types in the incomplete information game (see Morris (1997) for details).

The model of learning by similarity discussed here may be reinterpreted in the same way. Each player may instead be viewed as a population of players with locations $\theta \in \Theta$.

¹¹In addition, Samuelson and Zhang (1992), Nachbar (1990) and Heifetz, Shannon, and Spiegel (2007) identify classes of learning processes under which players learn not to play serially dominated strategies. However, all three papers assume finite or real-valued strategy spaces. The strategy space in our environment, consisting of functions $s : \Theta \rightarrow \{0, 1\}$, is larger.

In every period, players at a randomly drawn location are matched to play a game $\Gamma(\theta)$. Players estimate payoffs based on other players' experience at nearby locations. Thus learning by similarity corresponds to learning from neighbours. Since this is merely a formal reinterpretation, the modified game result also holds in this setting. As a game of incomplete information, the modified game may be reinterpreted according to Morris (1997) as a local interaction game. The only difference from the usual case is that the heterogeneous priors in the modified game correspond to *asymmetric* weighting of payoffs in the corresponding local interaction game; thus, for instance, Player i 's payoff may depend on Player j 's action even if Player j 's payoff does not depend on Player i 's action.

The formal connections among the three sources of contagion described here are summarized in Figure 7. Contagion through learning is related to contagion in Bayesian games through the equivalence of outcomes with the modified game (Theorem 4.1). The modified games that arise in this way differ from global games because of heterogeneous priors. Each of these may be reinterpreted according to (Morris 1997) as local interaction games, with heterogeneous priors corresponding to asymmetric weights and common priors to symmetric weights. Learning by similarity may also be reinterpreted directly as learning from neighbours in local interaction, where the modified game result describes an equivalence of outcomes with certain local interaction games with asymmetric payoff weights.

An earlier version of this paper (Steiner and Stewart 2006) considers learning by similarity in games with incomplete information. The environment is close to that of Section 5, except that each player receives only a noisy signal $x_t^i = \theta_t + \sigma \varepsilon_t^i$ of the state θ_t in each period t . Players estimate payoffs based on payoffs from similar past types. The modified game result of Theorem 4.1 extends naturally to this setting, and may be used to demonstrate that there is a unique outcome of learning when both σ and τ are small. This outcome depends on the ratio $\frac{\sigma}{\tau}$. If σ is small relative to τ , then we recover the complete information learning outcome of Proposition 5.1 above. If σ is large relative to τ , then we recover the usual global game solution.

The theory of global games has shown that relaxing the common knowledge assumption in games in a particular way can lead to a unique selection among multiple equilibria. This paper identifies a similar effect that arises under learning if we relax the assumption that players learn from repeated play in exactly the same game. Moreover, the learning outcome is formally related to the equilibria of a global game with subjective priors, which we have called the modified game. While the connection to the modified game is very general, the set of learning outcomes may be difficult to identify in games outside of the coordination environment studied here. The unique outcome of learning in this environment relies on the dominance solvability of the modified game. In more general settings, learning outcomes correspond to rationalizable profiles of the game when beliefs are perturbed in a particular way that depends on the similarity function. Weinstein and Yildiz (2007) have shown that for any game, given any rationalizable strategy profile, there exists a perturbation of beliefs in the universal type space for which this profile is uniquely rationalizable. A natural question, then, is whether the corresponding result holds under learning by similarity; in other words, whether the class of perturbations that may be obtained in the modified game by varying the similarity function is large enough to uniquely select each equilibrium.

A Appendix

Lemma A.1. *For any $\varepsilon > 0$, there exists some $\delta > 0$ such that changing the opponents' strategies on a set of type profiles of Lebesgue measure at most δ changes the expected payoff of every type of Player i from each action a^i by at most ε .*

Proof of Lemma A.1. Denote i 's expected payoff from Action a^i at Type θ against the profile $s^{-i}(\cdot)$ by

$$\tilde{U}^i(\theta, a^i, s^{-i}(\cdot)) = \frac{\int_{\Theta} u^i(\theta, a^i, v^i(\theta', a^i, s^{-i}(\theta'))) \phi(\theta') g^i(\theta', \theta) d\theta'}{\int_{\Theta} \phi(\tilde{\theta}) g^i(\tilde{\theta}, \theta) d\tilde{\theta}}.$$

The denominator is bounded above zero because it is continuous and positive by As-

sumption A3, and hence attains a positive minimum on the compact set Θ . Recall that the functions $u^i(\cdot)$, $g^i(\cdot)$, and $\phi(\cdot)$ are bounded by assumption. Hence there exists a constant K such that if $s^{-i}(\cdot)$ changes only on a set of measure δ , then the numerator changes by at most $K\delta$. \square

Lemma A.2. *Fix a profile $\mathbf{a}(\cdot)$ and an arbitrary action a^i . For any $\delta > 0$, there exists some $\pi > 0$ such that the set of types of Player i for which action a^i is dominated but not π -dominated under the profile $\mathbf{a}(\cdot)$ has measure at most δ .*

Proof of Lemma A.2. Consider any decreasing sequence π_1, π_2, \dots such that $\lim_{n \rightarrow \infty} \pi_n = 0$. Let $\Theta(n)$ denote the set of types for which Action a^i is π_n -dominated under $\mathbf{a}(\cdot)$, and let $\bar{\Theta}$ denote the set of types for which a^i is dominated under $\mathbf{a}(\cdot)$. Then $\Theta(n)$ is a monotone sequence of sets, and it suffices to show that $\lim_{n \rightarrow \infty} \Theta(n) = \bar{\Theta}$.

Suppose for contradiction that $\bar{\Theta} \setminus \lim_{n \rightarrow \infty} \Theta(n)$ contains some type θ . Then there exists some action $a^{i'}$ that dominates a^i at θ under the profile $\mathbf{a}(\cdot)$, but does not π -dominate a^i at θ under $\mathbf{a}(\cdot)$ for any $\pi > 0$. Hence we have

$$\inf_{s^{-i}(\cdot) \in \mathbf{a}^{-i}(\cdot)} \tilde{U}^i(\theta, a^{i'}, s^{-i}(\cdot)) - \tilde{U}^i(\theta, a^i, s^{-i}(\cdot)) = 0,$$

where we abuse notation by writing $s^{-i}(\cdot) \in \mathbf{a}^{-i}(\cdot)$ to mean that $s^{-i}(\cdot)$ is consistent with $\mathbf{a}^{-i}(\cdot)$. Define a strategy profile $\bar{s}^{-i}(\cdot)$ by choosing

$$\bar{s}^{-i}(\theta') \in \operatorname{argmin}_{a^{-i} \in \mathbf{a}^{-i}(\theta')} u^i(\theta, a^{i'}, v^i(\theta', a^{i'}, a^{-i})) - u^i(\theta, a^i, v^i(\theta', a^i, a^{-i}))$$

for each θ' . The profile $\bar{s}^{-i}(\cdot)$ is consistent with $\mathbf{a}^{-i}(\cdot)$, and satisfies

$$\tilde{U}^i(\theta, a^{i'}, \bar{s}^{-i}(\cdot)) - \tilde{U}^i(\theta, a^i, \bar{s}^{-i}(\cdot)) = \inf_{s^{-i}(\cdot) \in \mathbf{a}^{-i}(\cdot)} \tilde{U}^i(\theta, a^{i'}, s^{-i}(\cdot)) - \tilde{U}^i(\theta, a^i, s^{-i}(\cdot)) = 0,$$

contradicting that $a^{i'}$ dominates a^i at θ under $\mathbf{a}(\cdot)$. \square

Lemma A.3. *For any k and θ , we have $\mathbf{a}_{k,\pi}^i(\theta) \subseteq \mathbf{a}_{k,\pi'}^i(\theta)$ whenever $\pi \leq \pi'$.*

Proof of Lemma A.3. Note that the statement is trivial for $k = 0$. Suppose for induction that the statement holds for k (for all θ). We need to show that if a^i is π' -dominated at θ under $\mathbf{a}_{k,\pi'}$ then a^i is π -dominated at θ under $\mathbf{a}_{k,\pi}$. Accordingly, suppose that a^i is π' -dominated at θ under $\mathbf{a}_{k,\pi'}$; that is, there exists $a^{i'} \in \mathbf{a}_{k,\pi'}^i$ for which

$$\tilde{U}(\theta, a^{i'}, s^{-i}(\cdot)) - \tilde{U}(\theta, a^i, s^{-i}(\cdot)) > \pi' \text{ for all } s^{-i}(\cdot) \text{ consistent with } \mathbf{a}_{k,\pi'}^{-i}. \quad (13)$$

Since, by the inductive hypothesis, we have $\mathbf{a}_{k,\pi}^{-i} \subseteq \mathbf{a}_{k,\pi'}^{-i}$, Inequality (13) implies

$$\tilde{U}(\theta, a^{i'}, s^{-i}(\cdot)) - \tilde{U}(\theta, a^i, s^{-i}(\cdot)) > \pi \text{ for all } s^{-i}(\cdot) \text{ consistent with } \mathbf{a}_{k,\pi}^{-i}. \quad (14)$$

If $a^{i'} \in \mathbf{a}_{k,\pi}^i(\theta)$, then we are done. Otherwise, there exists some $a^{i''} \in \mathbf{a}_{k,\pi}^i(\theta)$ such that $a^{i''}$ dominates $a^{i'}$ at θ under the profile $\mathbf{a}_{k,\pi}$. Thus we have

$$\tilde{U}(\theta, a^{i''}, s^{-i}(\cdot)) - \tilde{U}(\theta, a^{i'}, s^{-i}(\cdot)) > 0 \text{ for all } s^{-i}(\cdot) \text{ consistent with } \mathbf{a}_{k,\pi}^{-i}.$$

Combining this with Inequality (14) gives the result. \square

Proof of Lemma 4.1. Given any type θ for which $a^i \in \mathbf{a}_{k-1,0}^i \setminus \mathbf{a}_{k,0}^i$, there exists some $\pi(\theta) > 0$ such that, against any profile $s^{-i}(\cdot)$ consistent with $\mathbf{a}_{k-1,0}^{-i}(\cdot)$, the expected payoff for some action $a^{i'} \in \mathbf{a}_{k,0}^i(\theta)$ is at least $\pi(\theta)$ greater than that for Action a^i . By Lemma A.3, we have $a^{i'} \in \mathbf{a}_{k,\pi}^i(\theta)$ for all $\pi > 0$, and hence it suffices to show that $a^{i'}$ dominates a^i under the profile $\mathbf{a}_{k-1,\pi}(\cdot)$ for some $\pi > 0$. By Lemma A.1, it suffices to show that given any $\delta > 0$, there exists some $\pi > 0$ small enough so that, for any player i , $\mathbf{a}_{k-1,\pi}^i(\cdot)$ differs from $\mathbf{a}_{k-1,0}^i(\cdot)$ on a set of measure at most δ .

We proceed by induction. The result is trivial for $k = 1$. For $k > 1$, assume for induction that the result is true for $k - 1$; that is, assume that for any $\delta > 0$, there exists some $\pi > 0$ for which $\mathbf{a}_{k-2,\pi}^i(\cdot)$ differs from $\mathbf{a}_{k-2,0}^i(\cdot)$ on a set of measure at most δ .

By Lemma A.2, given $\delta > 0$, we can choose $\pi' > 0$ small enough so that the set

of types of Player i for which a Action a^i is dominated but not π' -dominated under $\mathbf{a}_{k-2,0}(\cdot)$ has measure at most δ . By Lemma A.1, starting from any strategy profile, there exists some $\delta' > 0$ such that changing the actions of at most a measure of δ' of the opponents' types changes the expected payoff of each type of Player i by at most $\frac{\pi'}{4}$. By the inductive hypothesis, we can choose $\pi'' > 0$ such that $\mathbf{a}_{k-2,\pi''}^{-i}(\cdot)$ differs from $\mathbf{a}_{k-2,0}^{-i}(\cdot)$ on a set of types of measure at most δ' . Consider $\pi = \min\left\{\frac{\pi'}{2}, \pi''\right\}$. We need to show that $\mathbf{a}_{k-1,\pi}^i(\cdot)$ differs from $\mathbf{a}_{k-1,0}^i(\cdot)$ on a set of types of measure at most δ .

Consider any type θ and actions $a^i, a^{i'} \in \mathbf{a}_{k-2,0}^{-i}(\theta)$. By Lemma A.3, a^i and $a^{i'}$ also belong to $\mathbf{a}_{k-2,\pi}^{-i}(\theta)$ for all $\pi > 0$. Also by Lemma A.3, $\mathbf{a}_{k-2,\pi}^{-i}(\theta') \subseteq \mathbf{a}_{k-2,\pi''}^{-i}(\theta')$ for all θ' . Therefore, we have

$$\sup_{s^{-i} \in \mathbf{a}_{k-2,\pi}^{-i}(\cdot)} \left[\tilde{U}(\theta, a^i, s^{-i}) - \tilde{U}(\theta, a^{i'}, s^{-i}) \right] \leq \sup_{s^{-i} \in \mathbf{a}_{k-2,\pi''}^{-i}(\cdot)} \left[\tilde{U}(\theta, a^i, s^{-i}) - \tilde{U}(\theta, a^{i'}, s^{-i}) \right], \quad (15)$$

where, as above, we write $s^{-i} \in \mathbf{a}_{k,\pi}^{-i}(\cdot)$ to mean that the strategy profile s^{-i} is consistent with $\mathbf{a}_{k,\pi}^{-i}(\cdot)$. By the definition of π'' , we have

$$\sup_{s^{-i} \in \mathbf{a}_{k-2,\pi''}^{-i}(\cdot)} \left[\tilde{U}(\theta, a^i, s^{-i}) - \tilde{U}(\theta, a^{i'}, s^{-i}) \right] \leq \sup_{s^{-i} \in \mathbf{a}_{k-2,0}^{-i}(\cdot)} \left[\tilde{U}(\theta, a^i, s^{-i}) - \tilde{U}(\theta, a^{i'}, s^{-i}) + \frac{\pi'}{2} \right]. \quad (16)$$

If Action a^i is π' -dominated by $a^{i'}$ for Type θ of Player i under $\mathbf{a}_{k-2,0}(\cdot)$, then

$$\sup_{s^{-i} \in \mathbf{a}_{k-2,0}^{-i}(\cdot)} \left[\tilde{U}(\theta, a^i, s^{-i}) - \tilde{U}(\theta, a^{i'}, s^{-i}) + \frac{\pi'}{2} \right] < -\frac{\pi'}{2}.$$

Combining Inequalities (15) and (16) gives the following: if action a^i is π' -dominated for i at Type θ under $\mathbf{a}_{k-2,0}(\cdot)$, then a^i must be $\frac{\pi'}{2}$ -dominated, and hence also π -dominated under $\mathbf{a}_{k-2,\pi}(\cdot)$. Therefore, if, at some θ , a^i is dominated under $\mathbf{a}_{k-2,0}(\cdot)$ but not π -dominated under $\mathbf{a}_{k-2,\pi}(\cdot)$, then a^i is dominated under $\mathbf{a}_{k-2,0}(\cdot)$ but not π' -dominated under $\mathbf{a}_{k-2,0}(\cdot)$. The latter can happen only on a set of types of measure δ .

We have shown that the set of types θ for which $a^i \in \mathbf{a}_{k-2,\pi}^i(\theta)$ but $a^i \notin \mathbf{a}_{k-2,0}^i(\theta)$ has measure at most δ . The result now follows since the number of players and the number of actions are both finite. \square

Proof of Theorem 4.1. Part (1). Assume for induction that there almost surely exists some period after which the strategies $s_t^i(\cdot)$ are consistent with $\mathbf{a}_{k-1,\pi}^i(\cdot)$.

In the first step, we consider payoff estimates at a fixed state θ^* . Suppose that action a^i is π -dominated by some action $a^{i'}$ for Type θ^* of Player i in the modified game under the profile $\mathbf{a}_{k-1,\pi}(\cdot)$. Let

$$\pi'(\theta^*) = \pi \int_{\Theta} g^i(\theta, \theta^*) d\Phi(\theta).$$

We will show that, under the learning process, there almost surely exists some period after which the estimated payoff to Action $a^{i'}$ at θ^* exceeds that to Action a^i by at least π . This will be the case if

$$\frac{1}{t} \sum_{s < t} [u^i(\theta^*, a^{i'}, v^i(\theta_s, a^{i'}, a_s^{-i})) - u^i(\theta^*, a^i, v^i(\theta_s, a^i, a_s^{-i}))] g^i(\theta_s, \theta^*) > \pi'(\theta^*). \quad (17)$$

For each θ, θ' and a^{-i} , let

$$\Delta(\theta, \theta', a^{-i}) = u^i(\theta, a^{i'}, v^i(\theta', a^{i'}, a^{-i})) - u^i(\theta, a^i, v^i(\theta', a^i, a^{-i})).$$

Keeping θ^* fixed, choose the strategy profile $s_{\min}^{-i}(\theta)$ to minimize the payoff advantage of $a^{i'}$ over a^i at θ^* ; that is,

$$s_{\min}^{-i}(\theta) \in \operatorname{argmin}_{a^{-i} \in \mathbf{a}_{k-1,\pi}^{-i}(\theta)} \Delta(\theta^*, \theta, a^{-i}).$$

Define a random variable

$$X = \Delta(\theta^*, \theta, s_{\min}^{-i}(\theta)) g^i(\theta, \theta^*),$$

where the distribution of X is induced from the distribution Φ of θ .

By the inductive hypothesis, opponents play actions in $\mathbf{a}_{k-1,\pi}^{-i}(\theta_s)$ in every period $s \geq t_0$ for some t_0 . For large enough $t > t_0$, periods up to t_0 receive an arbitrarily small weight in each player's payoff estimates. Thus assuming that $a_s^{-i} \in \mathbf{a}_{k-1,\pi}^{-i}(\theta_s)$ for all s introduces only an arbitrarily small error in the payoff estimates. Note that for any history $(\theta_s, a_s^i, a_s^{-i})_{s=1}^{t-1}$ in which $a_s^{-i} \in \mathbf{a}_{k-1,\pi}^{-i}(\theta_s)$, we have

$$\frac{1}{t} \sum_{s < t} \Delta(\theta^*, \theta_s, a_s^{-i}) g^i(\theta_s, \theta^*) \geq \frac{1}{t} \sum_{s < t} \Delta(\theta^*, \theta_s, s_{\min}^{-i}(\theta_s)) g^i(\theta_s, \theta^*),$$

so it suffices to prove that (17) holds when $a_s^{-i} = s_{\min}^{-i}(\theta_s)$ for every s .

By the Strong Law of Large Numbers, the weighted payoff difference

$$\frac{1}{t} \sum_{s < t} \Delta(\theta^*, \theta_s, s_{\min}^{-i}(\theta_s)) g^i(\theta_s, \theta^*)$$

almost surely tends to the expectation of X , that is, to

$$\int_{\Theta} \Delta(\theta^*, \theta, s_{\min}^{-i}(\theta)) g^i(\theta, \theta^*) d\Phi(\theta).$$

By the assumption that $a^{i'}$ π -dominates a^i , the last expression is greater than $\pi'(\theta^*)$. Therefore, there almost surely exists some round T such that Inequality (17) holds for every $t > T$, as desired.

In the second step, we will show that there exists some $\delta > 0$ such that if Inequality (17) holds at θ^* , then $s_t^i(\theta) \neq a^i$ for all $\theta \in (\theta^* - \delta, \theta^* + \delta)$. Let

$$\pi' = \inf_{\theta \in \Theta} \pi'(\theta).$$

By Assumption A3, $\pi'(\theta)$ is positive everywhere, and since it is continuous, the compactness of Θ guarantees that π' is bounded away from zero.

Define the function

$$\begin{aligned} k(h_t, \theta) &= \frac{1}{t} \sum_{s < t} [u^i(\theta, a^i, v^i(\theta_s, a^i, a_s^{-i})) - u^i(\theta, a^i, v^i(\theta_s, a^i, a_s^{-i}))] g^i(\theta_s, \theta) \\ &= \frac{1}{t} \sum_{s < t} \Delta(\theta, \theta_s, a_s^{-i}) g^i(\theta_s, \theta) \end{aligned}$$

for finite histories $h_t = (\theta_s, a_s^i, a_s^{-i})_{s < t}$ and states θ . Player i will not choose action a^i at state θ following history h_t if $k(h_t, \theta) > 0$. Thus the second step will be complete if we show that, for some $\delta > 0$, after any history h_t , we have

$$|k(h_t, \theta) - k(h_t, \theta')| < \pi'$$

whenever $|\theta - \theta'| < \delta$.

By Assumption A4, $\Delta(\theta, \theta_s, a_s^{-i}) g^i(\theta_s, \theta)$ is uniformly continuous in θ over all θ_s and a_s^{-i} . Hence the average $k(h_t, \theta)$ is also continuous in θ uniformly over all values of θ and all histories h_t , as needed.

Finally, partition the set Θ into a finite number of subsets $\Theta_1, \dots, \Theta_m$, each of diameter less than δ (where δ is chosen given π' as in the second step above). Consider any of these subsets Θ_l . If Θ_l contains some θ^* at which a^i is π -dominated by $a^{i'}$ (under some profile), then by the first step, $k(h_t, \theta^*)$ is eventually larger than π' . By the second step, $k(h_t, \theta)$ is therefore positive for *all* types in Θ_l . Hence there almost surely exists some period T_l after which Player i never plays action a^i at *any* state in Θ_l . Since there are only finitely many sets Θ_l , there almost surely exists some T after which Player i never plays action a^i at any state θ for which it is serially π -dominated for the corresponding type in the modified game. The result now follows from the finiteness of the action and player sets.

Part (2). First we claim that for any k , the probability that play is consistent with k rounds of IEDS in the modified game approaches one as time tends to infinity. In the proof of Lemma 4.1, we showed that for any $\delta > 0$, there exists some $\pi > 0$ such that the set of types θ at which k rounds of IEDS differ from k rounds of IE π DS has

measure at most δ . Thus the probability that θ_t lies in this set can be made arbitrarily small by choosing π to be sufficiently small. Note that outside of this set, play under the learning process is almost surely eventually consistent with k rounds of IEDS by Statement 1 of the theorem. This proves the claim.

For each $k = 1, 2, \dots$, denote by Θ_k^i the set of types of Player i for which all serially dominated actions are eliminated within the first k rounds of IEDS. Let $\Theta_k = \bigcap_{i=1, \dots, I} \Theta_k^i$. The sequence of sets Θ_k is nondecreasing in k , and converges to the set Θ . Hence the measure of the set $\Theta \setminus \Theta_k$ converges to zero as k tends to infinity. From the previous paragraph, the probability that play is consistent with IEDS on Θ_k tends to one over time, and the probability that the state θ_t lies in $\Theta \setminus \Theta_k$ can be made arbitrarily small by choosing k to be sufficiently large.

□

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.
- Dekel, E., D. Fudenberg, and D. K. Levine (2004). Learning to play bayesian games. *Games and Economic Behavior* 46, 282–303.
- 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.
- Heifetz, A., C. Shannon, and Y. Spiegel (2007). The dynamic evolution of preferences. *Economic Theory* 32, 251–286.
- Izmalkov, S. and M. Yildiz (2007). Investors' sentiments.
- Jehiel, P. and F. Koessler (2006). Revisiting games of incomplete information with analogy-based expectations. <http://www.enpc.fr/ceras/jehiel/ABEEbayes.pdf>.
- Katz, K. (1996). *Similar games*. Ph. D. thesis, University of Pennsylvania.
- LiCalzi, M. (1995). Fictitious play by cases. *Games and Economic Behavior* 11, 64–89.
- Mengel, F. (2007). Learning across games. IVIE Working Paper Serie AD 2007-05.
- 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.

Steiner, J. and C. Stewart (2006). Learning by similarity in coordination problems.
Edinburgh School of Economics discussion paper 151.

Weinstein, J. and M. Yildiz (2007). A structure theorem for rationalizability with
application to robust predictions of refinements. *Econometrica* 75, 365–400.