

Coordination in Dynamic Environments

A Dissertation

Presented to the Faculty of the Graduate School

of

Yale University

in Candidacy for the Degree of

Doctor of Philosophy

by

Colin Stewart

Dissertation Directors: Professor Stephen Morris, Professor Benjamin Polak, and

Professor Dirk Bergemann

December 2007

UMI Number: 3293389

INFORMATION TO USERS

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

UMI[®]

UMI Microform 3293389

Copyright 2008 by ProQuest Information and Learning Company.

All rights reserved. This microform edition is protected against unauthorized copying under Title 17, United States Code.

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

© 2007 by Colin Stewart
All rights reserved.

Abstract

Coordination in Dynamic Environments

Colin Stewart

2007

The first chapter, joint with Jakub Steiner, studies a learning process in which subjects extrapolate their experience from 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 equilibria of static 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.

The second chapter considers the equilibrium selection problem in coordination games when players interact on an arbitrary social network. I examine the impact of the network structure on the robustness of the usual risk dominance prediction as mutation rates vary. For any given network, a sufficiently large bias in mutation probabilities favoring the non-risk dominant action overturns the risk dominance prediction; bounds are obtained on the size of this bias depending on the network structure. As the size of the population grows large, the risk dominant equilibrium is highly robust in some networks. This is true in particular if the risk dominant action spreads contagiously in the network and there does not exist a sufficiently cohesive finite group of players. Examples demonstrate that robustness does not coincide with fast convergence.

The third chapter, joint with Amil Dasgupta and Jakub Steiner, studies how the presence of multiple participation opportunities coupled with private learning about payoffs affects the ability of agents to coordinate efficiently in global coordination games. Two

players face the option to invest irreversibly in a project in one of many rounds. The project succeeds if some underlying state variable θ is positive and both players invest, possibly asynchronously. In each round they receive informative private signals about θ , and asymptotically learn the true value of θ . Players choose in each period whether to invest or to wait for more precise information about θ . We show that with sufficiently many rounds, both players invest with arbitrarily high probability whenever investment is socially efficient. This result stands in sharp contrast to the usual static global game outcome in which players coordinate on the risk-dominant action. We provide a foundation for these results in terms of higher order beliefs.

Acknowledgements

I am very grateful to Stephen Morris for his guidance throughout my time at Yale. Working with Stephen has been a great privilege for me. I am also deeply indebted to Ben Polak, from whom I learned many invaluable lessons about research, teaching, and mentorship. My coauthors Jakub Steiner and Amil Dasgupta have also earned my gratitude. I particularly thank Jakub for countless stimulating and challenging discussions.

The three chapters of this dissertation have benefited from the comments of many people, received either directly or through my coauthors. In particular, I would like to thank Dirk Bergemann, Martin Cripps, Eduardo Faingold, Dino Gerardi, Philippe Jehiel, George Mailath, John Moore, Emre Ozdenoren, Marzena Rostek, Larry Samuelson, and Avner Shaked. I am also grateful for the comments of many seminar and conference participants.

Contents

1	Learning by Similarity in Coordination Problems	10
1.1	Introduction	10
1.2	Example	14
1.3	The Learning Model	16
1.3.1	Long-run Characterization	21
1.4	Limit Results and Comparative Statics	25
1.4.1	Narrow Similarity and Small Noise	25
1.4.2	Comparative Statics	31
1.4.3	The Environmental Multiplier	34
1.5	Related Literature	36
1.6	Conclusion	39
1.7	Appendix	40
2	Robust Conventions and the Structure of Social Networks	54
2.1	Introduction	54
2.2	Literature review	56
2.3	The model	57
2.4	Ellison's radius and coradius	61
2.5	Fixed populations	63
2.6	Large population games	72

2.7	Waiting times	77
2.8	Discussion and conclusion	80
3	Efficient Dynamic Coordination with Private Learning	85
3.1	Introduction	85
3.1.1	Literature Review	90
3.2	Model	92
3.3	Analysis	94
3.3.1	The failure of coordination in the synchronous game	96
3.3.2	The success of coordination in the asynchronous game	97
3.4	Higher Order Beliefs	101
3.4.1	The synchronous case	101
3.4.2	The asynchronous case	102
3.5	Conclusion	110

List of Figures

2.1	$\frac{2}{3}$ -cohesive and $\frac{1}{2}$ -cohesive sets of nodes in a finite lattice.	64
2.2	Nearest neighbor interaction on the circle.	69
2.3	Uniform interaction.	79
2.4	Regions of size m for $m = 4$ (Morris, 2000).	79

List of Tables

1.1	Payoffs in the Example of Section 1.2.	11
-----	--	----

Chapter 1

Learning by Similarity in Coordination Problems

1.1 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 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

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

Table 1.1: Payoffs in the Example of Section 1.2.

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 1.1 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 1.3. 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.

Since investing is dominant for all sufficiently high fundamentals, there is some $\bar{\theta}$ above which players eventually learn to invest. Now consider a fundamental 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 fundamentals 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 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. 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 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 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

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

to that from incomplete information, players in the general model of Section 1.3 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 more noisy 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 1.1).

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 1.3 describes the general learning model, and characterizes its long-run behavior in terms of the modified game. Section 1.4 fully identifies the long-run state in the the limit of small noise and narrow similarity distributions, and examines comparative statics. Section 1.5 reviews the related literature.

1.2 Example

This section presents an example to illustrate in more detail the process of contagion through learning before describing the general model in Section 1.3.

The underlying family of coordination problems consists of the 2-player games in Table 1.1. 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 $\{\theta_s, a_s^1, a_s^2\}_{s < t}$, the estimated payoff to player i from choosing action a^i given the fundamental θ_t is

$$\frac{\sum_{s < t} g(\theta_s - \theta_t) u(a^i, a_s^{-i}, \theta_s)}{\sum_{s < t} g(\theta_s - \theta_t)}, \quad (1.1)$$

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 1.3 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 $l : \Theta \rightarrow \{0, 1\}$ 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

$$\int_{\Theta} (\theta' + l(\theta') - 1) g(\theta' - \theta) d\theta' \quad (1.2)$$

from choosing to invest, and 0 from the safe action. The expression (1.2) 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 long-run, the similarity-based learner behaves *as if* her observation of θ is 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 $\tilde{\theta}(\theta) > 1 - \bar{\alpha}$, for some $\bar{\alpha} > 0$.

Following this second learning stage, we may apply the same argument for fundamentals θ with $\tilde{\theta}(\theta) = 1 - \bar{\alpha} - \alpha$ 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 $\bar{\theta}$ is reached at which the gain in the estimated return to investing due to the

opponent investing above $\bar{\theta}$ is exactly offset by the loss in this return if the opponent chooses the safe action below $\bar{\theta}$.

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 $\bar{\theta}$. 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 $\bar{\theta} = \underline{\theta} = \theta^*$. At this threshold, players are indifferent between their two actions given their long-run payoff estimates. Thus we have

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

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 $\tilde{\theta}(\theta^*) + \frac{c-1}{2} = 0$.

In Section 1.3, 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.

1.3 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 $p(\theta, x^1, x^2)$ denote the probability density associated with the combination (θ, x^1, x^2) when $\sigma > 0$, we have²

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

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

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

³The reason for defining payoffs so as to allow a nonlinear 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.

Assumptions A2 and A3 are quite standard in global games (see, e.g., Morris and Shin (2003)). 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 $t = 1, 2, \dots$, 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 $u(\theta_t, a_t^i)$ 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 $g^i : \mathbb{R} \rightarrow \mathbb{R}$ 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 nonnegative everywhere and integrable, and we will normalize it to be a probability density function. Following a history $h_t = (\theta_s, x_s^1, x_s^2, a_s^1, a_s^2)_{s=1, \dots, t}$, type x^i of player i forms the *estimated return* to

action 1 by

$$r(x^i; h_t) := \frac{\sum_{s=1}^t u(\theta_s, a_s^{-i}) g^i(x_s^i - x^i)}{\sum_{s=1}^t g^i(x_s^i - x^i)} \quad (1.3)$$

whenever the denominator on the right-hand side is nonzero. 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 $\sum_{s=1}^t g^i(x_s^i - x^i) = 0$,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, Gilboa, Samet and Schmeidler (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. (2005), 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 model, 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 $r(x^i; h^t)$ place equal weight on all past

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

⁵This connection is subject to the caveat that Billot et al. (2005) assume a finite outcome space, whereas the outcome space is infinite here.

observations regardless of how much time has elapsed. More generally, one could suppose that observations are discounted over time according to a nonincreasing sequence $\delta(\tau) \in (0, 1]$ by modifying equation (1.3) 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 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, $\lim_{x \rightarrow \infty} g^i(x) = \lim_{x \rightarrow -\infty} g^i(x) = 0$. 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 $p_x^i(\cdot)$ denote the marginal density corresponding to the distribution of the signal of player i . With incomplete information, we have $p_x^i(x) = \int_{\Theta} \phi(\theta) f^i\left(\frac{x-\theta}{\sigma}\right) d\theta$, and with complete information, $p_x^i(x) = \phi(x)$.

⁶All of the results hold if instead we suppose that for each i , there exists some x^i below which g^i is nondecreasing, and above which it is nonincreasing.

A7. The marginal densities $p_x^i(\cdot)$ 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.

1.3.1 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 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 long-run).

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 (1.3). Against a fixed strategy $a^j(x')$ of the opponent, the expected value of this numerator is proportional to

$$\frac{\int_{X^i} \int_{X^j} \int_{\Theta} u(\theta, a^j(x')) g^i(x - x^i) p(\theta, x, x') d\theta dx' dx}{\int_{X^i} \int_{X^j} \int_{\Theta} p(\theta, x, x') g^i(x - x^i) d\theta dx' dx} = \int_{X^j} \int_{\Theta} u(\theta, a^j(x')) q^i(\theta, x'|x^i) d\theta dx', \quad (1.4)$$

where

$$q^i(\theta, x^j|x^i) = \frac{\int_{X^i} p(\theta, x, x^j) g^i(x - x^i) dx}{\int_{\Theta} \int_{X^j} \int_{X^i} p(\theta, x, x') g^i(x - x^i) dx dx' d\theta}. \quad (1.5)$$

Note that the right-hand side of (1.4) is precisely the expected payoff to player i from playing action 1 given the posterior beliefs $q^i(\theta, x^j|x^i)$, 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 $q^i(\theta, x^j|x^i)$. 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

$$q^i(\theta|x^i) = \frac{\phi(\theta)g^i(\theta - x^i)}{\int_{\Theta} \phi(\theta)g^i(\theta - x^i) d\theta}. \quad (1.6)$$

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 $u^i(\theta, a^i, a^{-i})$ denote the payoff to player i from the action profile (a^i, a^{-i}) 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

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

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 1.1 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 domination.

Lemma 1.1. *Suppose that action $a \in \{0, 1\}$ is serially dominated for type \bar{x} of player i . Then there exists some $\pi > 0$ such that action a is serially π -dominated for type \bar{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

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

⁸The 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.

strategies (IE π DS). 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 IE π DS differs from IEDS in the first $N - 1$ rounds can be made to have 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 long-run, players will not play strategies that are serially dominated in the modified game.

Theorem 1.1. *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.

1.4 Limit Results and Comparative Statics

1.4.1 Narrow Similarity and Small Noise

By Theorem 1.1, 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 ($\sigma \rightarrow 0$) and narrow similarity ($\tau \rightarrow 0$), while holding the ratio $\frac{\sigma}{\tau}$ fixed.

Proposition 1.1. *For any $\delta > 0$, there exists $\bar{\gamma} > 0$ such that for any $\gamma \in (0, \bar{\gamma})$, 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

$$\int_0^1 u(\theta, l) dH(l) = 0 \quad (1.7)$$

and $H(\cdot)$ is defined by

$$H(l) = \int_{\Xi} g(\xi) \left(1 - F \left(\frac{\tilde{\tau}}{\tilde{\sigma}} \xi + F^{-1}(1-l) \right) \right) d\xi, \quad (1.8)$$

and for $\sigma = 0$, θ^* solves

$$G(0)u(\theta, 0) + (1 - G(0))u(\theta, 1) = 0, \quad (1.9)$$

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 (1.8) 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

$$\tilde{m}_{\sigma, \tau}(x, k) \equiv \int_{\Theta} q_\theta(\theta|x) u \left(\theta, 1 - F \left(\frac{k - \theta}{\sigma} \right) \right) d\theta, \quad (1.10)$$

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 $x > \bar{\theta}^*$ and action 1 is serially dominated for $x < \underline{\theta}^*$, where $\bar{\theta}^*$ 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

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

Substituting this expression and interchanging the order of integrals in (1.10), 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 $\lim_{\tau \rightarrow 0} q_{x'}(x'|x)$ and $\lim_{\sigma \rightarrow 0} m_{\sigma}(x, k)$. 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

⁹The basic idea is that if action 1 has been eliminated for all types below $k < \theta^*$, 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 monotonicity of the sequences used in Morris and Shin's (2003) proof.

continuous at x , we have

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

It follows that

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

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 $\sigma \rightarrow 0$,

$$\Psi_\sigma(l; x + \sigma\xi, x) \rightarrow 1 - F(\xi + F^{-1}(1 - l)) \quad (1.11)$$

uniformly.

Step 4 consists of taking the limit as $\gamma \rightarrow 0$ and combining the limits from the step 3.

Accordingly, we have

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

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 (1.11) gives

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

Thus defining

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

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 \bar{x} such that action 1 is dominated for $x < \underline{x}$ and action 0 is dominated for $x > \bar{x}$. Thus we may restrict attention to signals in some compact set \bar{X} , on which $\tilde{m}_{\tilde{\sigma}\gamma, \tilde{\tau}\gamma}(x, x)$ converges to $\tilde{m}(x, x)$ uniformly. Given any neighbourhood 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 $\bar{\gamma} > 0$ small enough so that whenever $\gamma < \bar{\gamma}$, $\tilde{m}_{\tilde{\sigma}\gamma, \tilde{\tau}\gamma}(x, x)$ is within ε of $\tilde{m}(x, x)$ everywhere on \bar{X} , guarantees that $\tilde{m}_{\tilde{\sigma}\gamma, \tilde{\tau}\gamma}(x, x)$ has no root in $\bar{X} \setminus N$. \square

Proposition 1.1 can be generalized from 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

$$u(\theta, l) \equiv \sum_{k=1}^N \binom{N-1}{k-1} l^{k-1} (1-l)^{N-k} \tilde{u}(\theta, k) \quad (1.12)$$

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

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

non-decreasing in l . With this definition of payoffs, Proposition 1.1 applies to the N player case unchanged, and the limit long-run threshold is characterized by equation (1.7) with the distribution $H(l)$ satisfying (1.8). Moreover, let us rewrite the pure strategy profile payoff as $\tilde{u}(\theta, k) = v\left(\theta, \frac{k}{N}\right)$, where $v : \mathbb{R} \times [0, 1] \rightarrow \mathbb{R}$. Then $\lim_{N \rightarrow \infty} u(\theta, l) = v(\theta, l)$, as all the weight in the summation in (1.12) becomes concentrated at $k = l$. In particular, as N grows large, the long-run threshold converges to the solution of

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

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 (1.7) when ρ is finite, and by (1.9) 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 1.1. *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 1.1 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 1.1 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 1.2. *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, $\lim_{\rho \rightarrow 0} \theta_{\rho}^* = \theta_{\text{gg}}^*$*
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, $\lim_{\rho \rightarrow \infty} \theta_{\rho}^* = \theta_{\infty}^*$.*

Proof. See Appendix. □

1.4.2 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 $\gamma \rightarrow 0$, 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 1.3. *If the sign of $\frac{\partial u}{\partial z}(\theta, l; z)$ is the same for all θ and l , then*

$$\text{sign} \left(\frac{\partial \theta_{\tilde{\tau}}^*}{\partial z} \right) = \text{sign} \left(\frac{\partial \theta_{g\tilde{\sigma}}^*}{\partial z} \right) = -\text{sign} \left(\frac{\partial u}{\partial z} \right),$$

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

Proof. By Proposition 1.1, the long-run threshold $\theta_{\tilde{\tau}}^*$ is the solution to

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

By the implicit function theorem,

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

The denominator is positive by Assumption A1, that returns are increasing in θ . The sign of numerator is equal to $\text{sign} \left(\frac{\partial u}{\partial z} \right)$. 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 $[0, 1]$. \square

Proposition 1.3 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 value¹¹, and decreases with the

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

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 (1.8) in Proposition 1.1 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, 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 1.4. *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 (1.8) in Proposition 1.1 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

¹²Consider 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 $\tilde{u}(\theta, l) = u(\theta, lz)$.

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 $\theta^* \leq \tilde{\theta}^*$. □

1.4.3 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 (2003) to that via learning.

The underlying game Γ_σ of this example is characterized by the payoff function in Table 1.1, 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 1.1, the long-run behavior is consistent with IEDS in the modified game $\Gamma_{\sigma, \tau}^m$. We use the procedure utilized in the proof of Proposition 1.1 according to which the solution of $\Gamma_{\sigma, \tau}^m$ 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 ,

$$\Theta|x \sim N\left(\frac{y(\sigma^2 + \tau^2) + x\omega^2}{\omega^2 + \sigma^2 + \tau^2}, \frac{\omega^2(\sigma^2 + \tau^2)}{\omega^2 + \sigma^2 + \tau^2}\right). \quad (1.13)$$

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 (1.13) 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 1.1, $\tilde{m}_{\sigma, \tau}(x, k) = A(x) - F\left(\frac{k - A(x)}{B}\right)$, and the threshold x^* is the root of

$$\tilde{m}_{\sigma, \tau}(x, x) = A(x) - F\left(\frac{x - A(x)}{B}\right) = 0. \quad (1.14)$$

Keeping σ and τ fixed, the left hand side of (1.14) 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 $\tau \rightarrow 0$. Applying the implicit function theorem to equation (1.14) gives

$$\frac{\partial x^*}{\partial y} = -\frac{\sigma^2 + f((x - y)D)D(\omega^2 + \sigma^2)}{\omega^2 - f((x - y)D)D(\omega^2 + \sigma^2)}, \quad (1.15)$$

where $D = \frac{1}{\sqrt{(\omega^2 + \sigma^2)\left(\frac{2\omega^2}{\sigma^2}\right) + 1}}$. This result is equivalent to the result in Morris and Shin (2003, Section 3.1). The coincidence is a consequence of Proposition 1.2, 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 (1.15) is larger due to the strategic behavior in the global game model. Although (1.15) also holds in the learning model (in the limit $\tau \rightarrow 0$), 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 $\sigma \rightarrow 0$ holding $\tau > 0$ fixed. In this case

$$\frac{\partial x^*}{\partial y} = -\frac{\tau^2 + f((x-y)E)E(\omega^2 + \sigma^2)}{\omega^2 - f((x-y)E)E(\omega^2 + \sigma^2)}, \quad (1.16)$$

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 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.

1.5 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 long-run. 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 (2007) 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 (2007) 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 chose 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 (2007) 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

¹³In a related paper, Jehiel and Samet (2007) 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 (2004) offers an informal argument to suggest that the learning process can be approximated by the best-response dynamics of a modified game. Theorem 1.1 above formalizes the corresponding result for binary action global games when similarity is instead used to estimate payoffs directly. Carlsson's (2004) 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 1.1 above.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.

¹⁴In 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.

1.6 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.

1.7 Appendix

Proof of Lemma 1.1. First note that the uniform continuity of $g^i(\cdot)$ implies that the beliefs $q^i(\cdot|x^i)$ are uniformly continuous in x^i for every (θ, x^j) , 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 \bar{x} such that action 1 is π -dominated for $x < \underline{x}$ and action 0 is π -dominated for $x > \bar{x}$. 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 $\mu(s, s') \in \mathbb{R}_+ \cup \{\infty\}$ denote the Lebesgue measure of the set of types on which s and s' differ. Let $U(s, x)$ 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 N th 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 π_1, π_2, \dots such that $\lim_{n \rightarrow \infty} \pi_n = 0$. 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 \bar{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 \bar{X} in the limit, for otherwise $\bar{X} \setminus \lim_{n \rightarrow \infty} X(n)$ is nonempty, and any type contained in it cannot receive a positive payoff when playing action 1 against s_0 , contradicting the definition of \bar{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 nondecreasing 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. \square

Proof of Theorem 1.1. 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 noncompact spaces if the similarity function has full support.

By the lemma, it suffices to prove the result for IE π DS 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 IE π DS.

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

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

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

$$\frac{\int_{\Theta} \int_{X^j} \int_{X^i} u(\theta, l(x')) p(\theta, x, x') g^i(x - \tilde{x}^i) dx dx' d\theta}{\int_{\Theta} \int_{X^j} \int_{X^i} p(\theta, x, x') g^i(x - \tilde{x}^i) dx dx' d\theta} < -\pi \quad (1.17)$$

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

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

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

for every θ , inequality (1.17) implies that

$$\int_{\Theta} \int_{X^j} \int_{X^i} u(\theta, \bar{l}(x')) p(\theta, x, x') g^i(x - \tilde{x}^i) dx dx' d\theta < -\pi', \quad (1.18)$$

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 $P_{\delta}(X^i)$ and $P_{\delta}(\Theta)$ respectively. To simplify the notation below, we assume that the partition $P_{\delta}(X^j)$ may be chosen so that $\bar{l}(\cdot)$ is $P_{\delta}(X^j)$ -measurable, and for $\mu^j \in P_{\delta}(X^j)$, we will write $\bar{l}(\mu^j)$ for $\bar{l}(x^j)$ 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 $P_{\delta}(X^j)$ 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 $(\rho, \mu^1, \mu^2) \in P_{\delta}(\Theta) \times P_{\delta}(X^1) \times P_{\delta}(X^2)$, 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 $(\theta_t, x_t^1, x_t^2) \in (\rho, \mu^1, \mu^2)$ is within η of the probability associated with the event (ρ, μ^1, μ^2) . Since the number of such events is finite, there will almost surely come a time after which this is true for all $(\rho, \mu^1, \mu^2) \in P_{\delta}(\Theta) \times P_{\delta}(X^1) \times P_{\delta}(X^2)$.

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,

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

since the denominator on the right-hand side of (1.3) is (eventually) positive. Accordingly, we have

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

where η is chosen to be sufficiently small so that each term $\Pr(\rho, \mu^1, \mu^2) - \eta$ 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

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

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

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 $P_\delta(X^i)$ with δ small enough so that $g^i(x - \tilde{x}^i) - g^i(x' - \tilde{x}^i) \in (-\varepsilon, \varepsilon)$ whenever $x, x' \in \mu^i$ for some $\mu^i \in P_\delta(X^i)$. That is, letting

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

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

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

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

and therefore

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

which, by (1.18), 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 $\Pr(x^i \in [b, c]) > 1 - \delta$, and choose any $x \in [b, c]$. We want to show that for any $\varepsilon > 0$ there exist $\underline{x}, \bar{x}, \underline{\theta}, \bar{\theta}$ such that there will almost surely be some period T for which

$$\left| \frac{1}{\sum_{s=1}^{t-1} g^i(x_s^i - x)} \sum_{\substack{s=1, \dots, t-1 \\ (\theta_s, x_s^i, x_s^j) \notin [\underline{\theta}, \bar{\theta}] \times [\underline{x}, \bar{x}] \times [\underline{x}, \bar{x}]}} u(\theta_s, a_s^j) g^i(x_s^i - x) \right| < \varepsilon \quad (1.19)$$

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

¹⁵It is important that the estimates in proof for the compact state space are applied only to the types in the interval $[b, c]$, and not in the larger interval $[\underline{x}, \bar{x}]$. 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 $[\underline{x}, \bar{x}]$ here.

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

Accordingly, consider

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

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

$$g^{\min}(x^i) = \inf_{x \in [b, c]} g^i(x^i - x)$$

and $g^{\max}(x^i) = \sup_{x \in [b, c]} g^i(x^i - x),$

we have

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

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

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

where $p(\theta, x^i) = \phi(\theta) f^i\left(\frac{x^i - \theta}{\sigma}\right)$ represents the density associated with the combination (θ, x^i) 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 $\theta_s \geq \bar{\theta}$, and for $x_s^j \notin [\underline{x}, \bar{x}]$, except with $a(\theta) := \arg \max_a |u(\theta, a)|$ instead of $a = 0$.

Finally, we consider the part of the sum in (1.19) where $x_s^i \notin [\underline{x}, \bar{x}]$ and $\theta \in [\underline{\theta}, \bar{\theta}]$. Since $\lim_{x \rightarrow \infty} g(x) = \lim_{x \rightarrow -\infty} g(x) = 0$, 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 $(\theta, x^i) \in [\underline{\theta}, \bar{\theta}] \times (-\infty, \underline{x})$ is bounded in absolute value by

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

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

Proof of Proposition 1.2. 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 $\tau \rightarrow 0$. 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

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

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

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

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

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

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 $\frac{1}{\tau} g\left(\frac{x' - x}{\tau}\right)$. □

Bibliography

- Argenziano, R. and Gilboa, I.: 2005, History as a coordination device. Mimeo.
- Beggs, A.: 2005, Learning in bayesian games with binary actions. Mimeo.
- Billot, A., Gilboa, I., Samet, D. and Schmeidler, D.: 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 van Damme, E.: 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.
- Gilboa, I. and Schmeidler, D.: 2001, *A Theory of Case-Based Decisions*, Cambridge University Press, Cambridge.
- Heinemann, F., Nagel, R. and Ockenfels, P.: 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 Koessler, F.: 2007, Revisiting games of incomplete information with analogy-based expectations. Forthcoming in *Games and Economic Behavior*.

- Jehiel, P. and Samet, D.: 2007, Valuation equilibrium, *Theoretical Economics* **2**, 163–185.
- LiCalzi, M.: 1995, Fictitious play by cases, *Games and Economic Behavior* **11**, 64–89.
- Milgrom, P. and Roberts, J.: 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.
- Morris, S.: 2000, Contagion, *Review of Economic Studies* **67**, 57–78.
- Morris, S. and Shin, H. S.: 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 University Press, Cambridge, pp. 56–114.
- Nachbar, J. H.: 1990, Evolutionary selection dynamics in games: convergence and limit properties, *International Journal of Game Theory* **19**, 59–89.
- Samuelson, L. and Zhang, J.: 1992, Evolutionary stability in asymmetric games, *Journal of Economic Theory* **57**, 363–391.
- Stahl, D. O. and Van Huyck, J.: 2002, Learning conditional behavior in similar stag hunt games. Mimeo.

Chapter 2

Robust Conventions and the Structure of Social Networks

2.1 Introduction

In large population coordination games, social conventions are often thought to provide common expectations of behavior, thereby allowing coordination on a particular equilibrium. A fundamental problem in the literature on conventions has been to understand which properties lead to the selection of a particular convention in the presence of multiple equilibria. Unique outcomes generally emerge in the long-run under evolutionary dynamics in which agents play myopic best responses except for a small probability of mutation. These outcomes are called *stochastically stable*. The basic question posed in this paper is the following: how robust is the stochastically stable equilibrium selection to changes in mutation rates?

Foster and Young (1990) and Kandori, Mailath and Rob (1993) introduced the criterion of stochastic stability, and showed that when players are randomly matched to play a 2×2 coordination game, coordination on the risk dominant action is the unique stochastically stable outcome. Bergin and Lipman (1996) criticize this approach on the grounds that the results are sensitive to the formulation of mutation rates that are freely chosen by the

modeller. They show, in particular, that any equilibrium of the dynamic process without mutations is selected for some specification of mutation probabilities. When the matching process selects from neighbors in a social network, the risk dominance prediction holds for the specific formulations of mutation rates that have appeared in the literature, regardless of the structure of the network (Peşki (2004), Young (1998)). However, the robustness of this prediction as the mutation rates vary depends heavily on this structure. The main results of this paper provide sufficient conditions on the structure of the interaction network to guarantee robustness.

First we consider a fixed population interacting on a given network. The risk dominance prediction is found to be robust when all players interact with roughly the same number of other players, and there do not exist small, highly cohesive clusters in the network. If there exist some players who interact with many others, then it is possible that mutation by these players alone could influence enough of the population to move away from the risk dominant action to a different equilibrium. If the number of required mutations is small, then the risk dominant equilibrium will tend not to be robust. Similarly, if there exists a small, highly cohesive set of players in the network, then mutation by these players alone suffices to move away from the equilibrium coordinated on the risk dominant action.

In general, a sufficiently large bias in mutation probabilities always suffices to overturn the risk dominance prediction in a population of a given size. However, as the size of the population grows, there exist networks for which this prediction is robust to arbitrarily large mutation biases. As in the case of a fixed population, the nonexistence of a finite highly-cohesive cluster in the network is a necessary condition for the risk dominance prediction to be robust. When the risk dominant action spreads contagiously in the network, this nonexistence condition is also sufficient. Intuitively, contagion makes it possible to reach coordination on the risk dominant action from any initial strategy profile with a relatively small number of mutations. As the size of the population grows, a much larger number of mutations is required to move away from this coordinated equilibrium. When mutation probabilities are small, this implies that transitions to the risk dominant equilibrium occur

much more frequently than transitions away from it, regardless of the specification of mutation rates.

Ellison (1993) argues that when mutation probabilities are small, the expected time to convergence to a stochastically stable outcome may be unreasonably long for practical applications. Furthermore, he demonstrates that the structure of the interaction network has a strong influence on the speed of convergence. Intuitively, fast convergence occurs if “not many” simultaneous mutations are required for best response dynamics to lead play to the predicted outcome from any initial strategy profile. It may therefore be tempting to believe that the structural properties generating fast convergence should coincide with those leading to robustness to varying mutation rates. Section 2.7 presents examples demonstrating that this intuition is false: neither of these properties is sufficient to guarantee the other. This suggests that in order to evaluate the relevance of stochastically stable outcomes in a particular game, it is necessary to examine both the speed of convergence and the robustness to varying mutation rates.

2.2 Literature review

The model studied in this paper is based on that of Kandori et al. (1993) (henceforth KMR), in which a large population of agents are matched in each of an infinite sequence of periods to play a 2×2 coordination game. Matches are drawn according to a uniform distribution over the entire population. Each player chooses a best response to the distribution of actions in the preceding period, except for a small probability of mutation, in which the action is chosen randomly according to a uniform distribution. Mutation probabilities are independent across players and periods, and constant across players and strategy profiles. If there is a strictly risk dominant action, A , in the 2×2 game, KMR show that as the mutation probabilities tend to zero, the probability that the population will be coordinated on A in a given period tends to one in the long-run. In other words, coordination on A is the unique stochastically stable state.

Young (1998) and Peşki (2004) consider variants of the KMR model in which each player is matched according to a uniform distribution over a particular subset of the population, namely that player's neighbors in a fixed social network. Under the uniform mutation rates considered by Peşki and the payoff-dependent mutation rates used by Young, the risk dominance result continues to hold subject to mild regularity conditions. Goyal and Vega-Redondo (2005) and Hojman and Szeidl (2006) find similar results when players interact in endogenously formed networks as long as the cost of forming links is small.

A more general model of evolutionary processes is employed by Bergin and Lipman (1996), who show that any distribution over states that is stable in the process without mutations can be obtained as the (unique) long-run distribution as mutation probabilities tend to zero for some specification of mutation rates. Blume (2003) addresses this critique of stochastic stability results by considering a class of payoff-dependent mutation probabilities, and identifying conditions on these probabilities under which the usual results are preserved.

Lee, Szeidl and Valentinyi (2003) study the robustness of the risk dominance prediction to varying mutation rates when players lie on a 2-dimensional torus, and each interacts with her four nearest neighbors. They show that, for given mutation rates, as the size of the torus grows large, the risk dominant equilibrium will eventually be stochastically stable. In contrast, with the general interaction structures considered here, the analogous result need not hold.

2.3 The model

A population of N agents forms the nodes of a (social) network $\Gamma = (V, L)$, where V is the set of nodes, and L is a set of unordered pairs of distinct elements of V . The elements of L are called the *links* of the network, and nodes $i, j \in V$ are said to be *neighbors* in Γ if $\{i, j\} \in L$. We will say that i *interacts with* j if i and j are neighbors in Γ . Note that the neighbor relation is symmetric, so that i interacts with j whenever j interacts with i .

Each agent plays one of two actions, A or B . Payoffs from each interaction are given by a function $u(\cdot, \cdot)$ corresponding to the matrix

	A	B	
A	a, a	c, d	
B	d, c	b, b	

The 2×2 game with these payoffs will be referred to as the *underlying game*.

The following restrictions are imposed on the payoffs:

1. (A, A) and (B, B) are Nash equilibria; that is, $a > d$ and $b > c$.
2. A is strictly risk dominant; that is, $a + c > b + d$.

The first of these conditions ensures that the equilibrium selection problem is nontrivial. The case that has received the most attention in the literature is when the risk dominant and payoff dominant equilibria differ; under the best response dynamics considered here, however, payoff dominance plays no role, so it is not necessary to identify the payoff dominant equilibrium. As will become clear below, the analysis is trivial when risk dominance is not strict (i.e. when $a + c = b + d$), and the choice of A as the strictly risk dominant action is therefore without loss of generality.

Player i 's payoff $U_i(s_i, s_{-i})$ from playing s_i when the remaining agents play the profile s_{-i} is given by adding the payoffs $u(s_i, s_j)$ over all neighbors j of i . Formally, payoffs are given by

$$U_i(s_i, s_{-i}) = \sum_{\{i,j\} \in L} u(s_i, s_j).$$

In order to distinguish it from the underlying game, the N player game with these payoffs will be referred to as the *population game*. Assume for simplicity that each agent's best response correspondence is single-valued. In other words, letting $\delta(i) = \#\{j \in V \mid \{i, j\} \in L\}$, assume that $\frac{b-c}{a-d+b-c}\delta(i)$ is not an integer for any i .

In the *unperturbed best response dynamics*, the population game is played over infinitely

many periods $t = 0, 1, \dots$. Starting from some strategy profile in period 0, in each period $t \geq 1$, each player updates her action with independent probability $\pi \in (0, 1)$, and otherwise plays the same action as in period $t - 1$. When updating, players myopically choose the best response to the strategy profile played in the preceding period. The reason for introducing randomness into the updating process is technical: it guarantees that the stable states of the unperturbed process are precisely the Nash equilibria of the game. This is formalized in Lemma 2.1.

Note that action A is a best response for player i in the population game if and only if the fraction of her neighbors choosing action A is at least

$$p = \frac{b - c}{a - d + b - c}.$$

Therefore, for any initial strategy profile, the unperturbed dynamics depend only on the values of π and p .

The *perturbed best response dynamics* agree with the unperturbed dynamics except that players may “mutate” by switching to an action that is not a best response. Fix $\alpha > 0$ and $\varepsilon \in (0, 1)$. In each period in which a given player i is called upon to update her strategy, i mutates to B with probability ε^α if A is the best response to the strategy profile of the previous period, and mutates to A with probability ε if B is the best response. In both cases, player i plays her best response otherwise. Note that ε and α depend neither on the player nor the state. Random draws are independent across players and time.

The parameter α captures the bias in mutations toward action A . When α is small, players are much more likely to mutate to action B when A is a best response than they are to mutate to A when B is a best response. If mutations are interpreted as experimentation by boundedly rational players, such a bias may result, for example, from a tendency to try to attain the payoff-dominant outcome. As usual in models of this type, only the orders of magnitude of the mutation probabilities are relevant in determining the stochastically stable outcomes. If the probability ε^α were to be replaced by $\alpha\varepsilon$ in the above formulation,

then α may affect the stationary distribution over outcomes, but not the set of stochastically stable states.

Both the unperturbed and the perturbed best response dynamics define finite Markov chains whose states are the strategy profiles of the population game. Recall that two states σ, σ' are said to *communicate* in a Markov chain if, beginning from σ , there is a positive probability that σ' will occur within a finite number of periods, and vice versa. A *recurrent class* is a set of states within which each pair of states communicate, and from which no other state occurs with positive probability in finite time. A Markov chain is *irreducible* if the entire state space forms a recurrent class; otherwise, it is *reducible*. A finite Markov chain possesses a unique stationary distribution if and only if it is irreducible (see, e.g., Young (1998)).

Whereas the unperturbed best response dynamics form a reducible Markov chain, the Markov chain defined by the perturbed dynamics is irreducible. For each ε and α , let $\mu^{\varepsilon, \alpha}(\cdot)$ denote the stationary distribution of the perturbed process.

Definition 2.1. *Given α , the state σ is stochastically stable if $\lim_{\varepsilon \rightarrow 0} \mu^{\varepsilon, \alpha}(\sigma) > 0$.*

The main question to be addressed here concerns the extent to which, depending on the structure of the interaction network, mutations must be biased in favor of action B in order to overturn the risk dominance prediction. Accordingly, define the *mutation robustness threshold* $\bar{\alpha}$ to be

$$\bar{\alpha} := \inf \{ \alpha \mid \sigma_A \text{ is stochastically stable} \},$$

where σ_A denotes the state in which all players play action A . If $\alpha < \bar{\alpha}$, then the stochastically stable states may contain the equilibrium σ_B coordinated on B , or may contain only *coexistent conventions*, equilibria in which the population is not coordinated on a single action.

2.4 Ellison's radius and coradius

This section introduces some Markov chain terminology and describes Ellison's radius-coradius method (Ellison 2000), which will be used to compute bounds on the threshold value $\bar{\alpha}$. The reader who is familiar with this material may wish to jump ahead to the next section.

Fix $\alpha > 0$, and let $P_\varepsilon(\sigma, \sigma')$ denote the transition probability from σ to σ' in the Markov chain describing the perturbed dynamics. For any states σ, σ' , define the *transition cost* $c(\sigma, \sigma')$ to be the unique real number satisfying

$$\lim_{\varepsilon \rightarrow 0} \frac{P_\varepsilon(\sigma, \sigma')}{\varepsilon^{c(\sigma, \sigma')}} \in (0, \infty).$$

Note that since best response updating is random, there may be different ways to transition from σ to σ' in a single period depending on which players update and which mutate. The cost $c(\sigma, \sigma')$ is the minimum value of the sum $n_A + \alpha n_B$ over all such single-period transitions, where n_A is the number of mutations required from B to A , and n_B is the number required from A to B .

A *path* from σ to σ' is a finite sequence $(\sigma_0, \sigma_1, \dots, \sigma_n)$ of distinct states such that $\sigma_0 = \sigma$ and $\sigma_n = \sigma'$. Let $\Pi(\sigma, \sigma')$ denote the set of all such paths. Define the *cost* $c(\bar{\sigma})$ of the path $\bar{\sigma} = (\sigma_0, \sigma_1, \dots, \sigma_n)$ to be

$$c(\bar{\sigma}) = c(\sigma_0, \sigma_1) + c(\sigma_1, \sigma_2) + \dots + c(\sigma_{n-1}, \sigma_n).$$

Paths that minimize the transition cost between states play a special role in identifying the stochastically stable states, as these are the transitions that occur most frequently in the limit as mutation probabilities vanish. Accordingly, for each pair of states σ, σ' , define the *minimal cost* $m(\sigma, \sigma')$ by

$$m(\sigma, \sigma') = \min_{\bar{\sigma} \in \Pi(\sigma, \sigma')} c(\bar{\sigma}).$$

Let Ω be a union of recurrent classes of the unperturbed dynamics. The *basin of attraction* $B(\Omega)$ of Ω is the set of states from which some state in Ω is almost surely reached in finite time under the unperturbed dynamics. Equivalently, $B(\Omega)$ consists of those states from which there exists a zero-cost path to some state in Ω , but there exists no such path to any recurrent class not in Ω .

Ellison (2000) defines the *radius* $R(\Omega)$ of the basin of attraction of Ω by

$$R(\Omega) := \min_{\sigma \in \Omega, \sigma' \notin B(\Omega)} m(\sigma, \sigma').$$

Thus the radius $R(\Omega)$ is the lowest cost associated with any transition that does not almost surely return to Ω under the unperturbed dynamics. Similarly, the *coradius* $C(\Omega)$ of the basin of attraction of Ω is defined by

$$C(\Omega) := \max_{\sigma' \notin B(\Omega), \sigma \in \Omega} m(\sigma', \sigma).$$

Thus, starting from any initial state, the coradius of Ω is the greatest cost that could be necessary in order to reach Ω . Ellison shows that if $R(\Omega) > C(\Omega)$ then Ω contains the set of stochastically stable states. Intuitively, when this is the case, transitions to Ω occur more frequently than transitions away from it, so as the mutation probabilities approach zero, much more time is spent at states in Ω than at any other state.

Since the radius-coradius condition is sufficient but not necessary, it can be used to identify upper bounds on the threshold $\bar{\alpha}$ by taking $\Omega = \{\sigma_A\}$. If, for some β , $R(\{\sigma_A\}) > C(\{\sigma_A\})$ whenever $\alpha > \beta$, then we have $\bar{\alpha} \leq \beta$. Similarly, by taking Ω to be the union over all recurrent classes except $\{\sigma_A\}$, then we obtain a lower bound on $\bar{\alpha}$ by identifying β such that $R(\Omega) > C(\Omega)$ whenever $\alpha < \beta$.

2.5 Fixed populations

Given a fixed network, it is possible to derive bounds on the mutation robustness threshold $\bar{\alpha}$ based on the network structure and the value of the payoff parameter p . Since both of the methods described in the previous section involve transition costs among recurrent classes of the unperturbed dynamics, we begin by identifying these classes.

Lemma 2.1. *The recurrent classes of the unperturbed best response dynamics are precisely the singleton sets containing the Nash equilibria of the population game.*

Proof. Clearly each Nash equilibrium forms a recurrent class.

For the converse, we must show that, beginning from any state σ_0 , a Nash equilibrium will be reached with positive probability in finite time. It suffices to construct a finite sequence of states $\sigma_0, \sigma_1, \dots, \sigma_m$ such that for each $k = 1, \dots, m$, σ_k differs from σ_{k-1} only through best response updating by a single player. Without loss of generality, suppose that action A is a best response for some player who plays action B under σ_0 . Choose any such player i , and define σ_1 to be equal to σ_0 except that player i plays action A . Repeat this step until a state σ_r is reached at which no such player remains. Now repeat this process beginning from σ_r , except with actions A and B reversed.

I claim that the final state σ_m attained under this process is a Nash equilibrium. Suppose for contradiction that player i plays an action σ_m^i that is not a best response under σ_m . It is clear by construction that $\sigma_m^i = B$ and player i 's best response under σ_m is A . Since the number of i 's neighbors playing action A is nonincreasing along the path $\sigma_r, \dots, \sigma_m$, A must also be a best response for player i under each σ_k for $k = r, \dots, m$. Therefore, player i must choose action B under σ_r , contradicting the construction of σ_r . □

Pęski (2004) considers the special case of the present model in which $\alpha = 1$. He shows that the state σ_A is stochastically stable regardless of the structure of the network, which, in our terminology, immediately implies the upper bound $\bar{\alpha} \leq 1$ on the mutation robustness

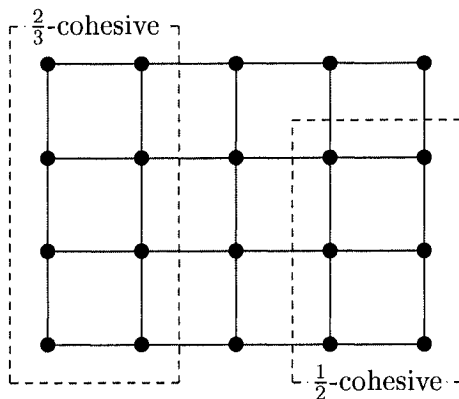


Figure 2.1: $\frac{2}{3}$ -cohesive and $\frac{1}{2}$ -cohesive sets of nodes in a finite lattice.

threshold. To obtain tighter bounds, some definitions are required concerning structural properties of networks.

Definition 2.2. The degree $\delta(i)$ of node i is the number $\#\{j \in V \mid \{i, j\} \in L\}$ of its neighbors in the network.

$$\text{Let } \delta_{\min} = \min_{i \in V} \delta(i) \text{ and } \delta_{\max} = \max_{i \in V} \delta(i).$$

Definition 2.3. Given $r \in [0, 1]$, a subset S of the set of nodes V of the network $\Gamma = (V, L)$ is r -cohesive in Γ if for every $i \in S$,

$$\#\{j \in S \mid \{i, j\} \in L\} \geq r \#\{j \in V \mid \{i, j\} \in L\}.$$

In words, each node in S has a fraction of at least r of its neighbors in S .

Figure 4 exhibits two sets of r -cohesive nodes in a finite lattice. For each set, the given value of r is the largest for which the set is r -cohesive.

The cohesiveness of sets of nodes in the network is directly related to the best response dynamics of the interaction game. If i lies in a p -cohesive set of nodes S , then A is a best response for i whenever all other players in S play A . Similarly, if S is $(1 - p)$ -cohesive, then B is a best response for i whenever all other players in S play B . Identifying each

strategy profile with the set S of agents playing A , S is a Nash equilibrium if and only if it is p -cohesive and its complement $V \setminus S$ in V is $(1 - p)$ -cohesive.

Consider the $(1 - p)$ -cohesive sets in Γ . These are partially ordered by inclusion and include the sets \emptyset and V . A *chain* of $(1 - p)$ -cohesive sets of length l is an increasing sequence of $l + 1$ distinct sets $V_0 \subsetneq V_1 \subsetneq \dots \subsetneq V_l$. Such a chain is *maximal* if $V_0 = \emptyset$, $V_l = V$, and for each $k = 1, \dots, l$, there does not exist any $(1 - p)$ -cohesive set U such that $V_{k-1} \subsetneq U \subsetneq V_k$.

Lemma 2.2. *When $\alpha = 1$, the coradius of the basin of attraction of σ_A is at most the length l of the shortest maximal chain of $(1 - p)$ -cohesive sets.*

Proof. Let σ_B denote the state in which all players choose action B , and identify each state with the set of agents choosing action B . Thus, in particular, $\sigma_A = \emptyset$ and $\sigma_B = V$. First we show that $C(\sigma_A) = m(\sigma_B, \sigma_A)$. For each state σ let $B(\sigma)$ denote the set of agents for which B is a best response to σ in the population game. Given any path $\bar{\sigma} = (\sigma_0, \dots, \sigma_k) \in \Pi(\sigma_B, \sigma_A)$, define the set

$$M_A(\bar{\sigma}) := \{v \in V \mid \exists j \in \{1, \dots, k\} \text{ such that } v \in \sigma_{j-1} \setminus \sigma_j \text{ and } v \in B(\sigma_{j-1})\}.$$

Thus the set $M_A(\bar{\sigma})$ consists of all nodes that mutate to action A at some point along the path $(\bar{\sigma})$.

I claim that for any state σ , there exists a zero-cost path from $\sigma \setminus M_A(\bar{\sigma})$ to σ_A ; that is, beginning from σ , if all agents in $M_A(\bar{\sigma})$ switch to (or remain at) action A , then it is possible to reach σ_A through best responses alone. If not, then let $j < k$ be the largest index for which there exists a state $\sigma' \subset (\sigma_j \setminus M_A(\bar{\sigma}))$ such that $m(\sigma \setminus M_A(\bar{\sigma}), \sigma') = 0$. Then there is some $v \in \sigma' \setminus \sigma_{j+1}$ for which $v \in B(\sigma')$. Since $\sigma' \subset \sigma_j$, $v \in B(\sigma')$ implies that $v \in B(\sigma_j)$, and hence $v \in M_A(\bar{\sigma})$, contradicting that $v \in \sigma'$ and $\sigma' \cap M_A(\bar{\sigma}) = \emptyset$.

The claim implies that, for any state σ and any path $\bar{\sigma} \in \Pi(\sigma_B, \sigma_A)$, the minimal cost $m(\sigma, \sigma_A)$ is at most $|M_A(\bar{\sigma})|$. If $\bar{\sigma}$ is a cost-minimizing path from σ_B to σ_A , then it cannot

involve any mutations to action B . Thus we have

$$m(\sigma_B, \sigma_A) = |M_A(\bar{\sigma})| = \max_{\sigma} m(\sigma, \sigma_A),$$

and therefore, $C(\sigma_A) = m(\sigma_B, \sigma_A)$.

Let $\emptyset = V_0 \subset V_1 \subset \dots \subset V_l = V$ be a shortest maximal chain of $(1-p)$ -cohesive sets. For $i \in \{1, \dots, l\}$, consider the initial state $\sigma_0 = V_i$ in which precisely the agents in V_i choose action B . Suppose some agent $v_0 \in V_i \setminus V_{i-1}$ mutates to action A , so that the state becomes $\sigma_1 = V_i \setminus \{v_0\}$. If $V_i \setminus \{v_0\} \neq V_{i-1}$, then $V_i \setminus \{v_0\}$ cannot be $(1-p)$ -cohesive by the maximality of the chain $V_0 \subset \dots \subset V_l$. Since V_{i-1} is $(1-p)$ -cohesive, there must exist some $v_1 \in (V_i \setminus \{v_0\}) \setminus V_{i-1}$ such that $v_1 \notin B(V_i \setminus \{v_0\})$. Repeating this argument with $\sigma_2 = V_i \setminus \{v_0, v_1\}$ in place of $V_i \setminus \{v_0\}$, and continuing recursively in this fashion gives rise to a path $\bar{\sigma} = (\sigma_0, \sigma_1, \dots, \sigma_k)$ of states such that $\sigma_k = V_{i-1}$ and $c(\bar{\sigma}) = 1$. Connecting these paths in sequence over all $i \in \{1, \dots, l\}$ gives a path in $\Pi(\sigma_B, \sigma_A)$ having cost l , as needed. \square

Definition 2.4. *An r -cohesive set of nodes S is a minimal r -cohesive set if it contains no nonempty r -cohesive proper subset.*

Theorem 2.1. *Let k be the size of the smallest nonempty $(1-p)$ -cohesive set whose complement in V is p -cohesive (or empty), and let K be the size of the largest minimal $(1-p)$ -cohesive set in V . We have the following upper bound:*

$$\bar{\alpha} \leq \frac{(\delta_{\max} + \delta_{\min} - 2 \lfloor p\delta_{\min} \rfloor)(N - K + 1)}{(\delta_{\min} - 2 \lfloor p\delta_{\min} \rfloor)k}.$$

Proof. Let R and C respectively denote the radius and coradius of σ_A when $\alpha = 1$. For general α , the coradius of σ_A is the same, whereas its radius is equal to αR . The state σ_A is therefore stochastically stable if $\alpha R > C$, indicating that the ratio $\frac{C}{R}$ is an upper bound on the threshold $\bar{\alpha}$. The proof proceeds in two steps. The first is to demonstrate that $R \geq \frac{\delta_{\min} - 2 \lfloor p\delta_{\min} \rfloor}{\delta_{\max} + \delta_{\min} - 2 \lfloor p\delta_{\min} \rfloor} k$. Showing that $C \leq N - K + 1$ then gives the result.

Suppose that the initial state is σ_A . If the radius is R , then there exists a set S_0 of R agents and a sequence i_1, \dots, i_m of agents such that B is a best response for i_l if every agent in $S_0 \cup \{i_1, \dots, i_{l-1}\}$ chooses B (that is, if a fraction of at least $(1 - p)$ of i_l 's neighbors lie in $S_0 \cup \{i_1, \dots, i_{l-1}\}$), and there is a coexistent convention in which at least one agent in $S_0 \cup \{i_1, \dots, i_r\}$ chooses B .

Recall that any coexistent convention corresponds to a partition of the nodes into a p -cohesive set and a $(1 - p)$ -cohesive set. Let S_B denote the set of agents who play B in the final equilibrium. Note that we may restrict ourselves to the subgraph containing only those links involving at least one node in S_B . Let $\delta_R(S_B)$ be the sum of the degrees of the R nodes in this subgraph having the largest degrees. That is, denoting the subgraph by $\Gamma(S_B)$, and the degree of node j in $\Gamma(S_B)$ by $\delta_{\Gamma(S_B)}(j)$, define

$$\delta_R(S_B) := \max_{\{j_1, \dots, j_R\}} \sum_{l=1}^R \delta_{\Gamma(S_B)}(j_l).$$

In order for the sequential best response condition mentioned above to hold, it must be the case that for each i_l , the number of edges connecting i_l to nodes in $S_0 \cup \{i_1, \dots, i_{l-1}\}$ is at least $\lceil (1 - p) \delta(i_l) \rceil$. After adding node i_1 , that leaves at most

$$\delta_R(S_B) - \lceil (1 - p) \delta(i_1) \rceil + \delta(i_1) - \lceil (1 - p) \delta(i_1) \rceil = \delta_R(S_B) - (\delta(i_1) - 2 \lfloor p \delta(i_1) \rfloor)$$

edges to be connected to other i_l 's. Continuing recursively in this fashion, we obtain the following inequality:

$$\lceil (1 - p) \delta(i_r) \rceil \leq \delta_R(S_B) - \sum_{j=1}^{r-1} (\delta(i_j) - 2 \lfloor p \delta(i_j) \rfloor).$$

This implies the weaker condition that

$$\delta_R(S_B) \geq \sum_{j=1}^r (\delta(i_j) - 2 \lfloor p \delta(i_j) \rfloor).$$

Since $\delta(i_j) - 2 \lfloor p\delta(i_j) \rfloor \geq \delta_{\min} - 2 \lfloor p\delta_{\min} \rfloor$ for each i_j , this implies that

$$R\delta_{\max} \geq r(\delta_{\min} - 2 \lfloor p\delta_{\min} \rfloor).$$

By definition of k , we must have $R + r \geq k$, which, when substituted for r in the last expression, gives

$$R \geq \frac{\delta_{\min} - 2 \lfloor p\delta_{\min} \rfloor}{\delta_{\max} + \delta_{\min} - 2 \lfloor p\delta_{\min} \rfloor} k,$$

as desired.

For the coradius, C , note that the existence of a minimal $(1-p)$ -cohesive set U of size K implies that there is a maximal chain of $(1-p)$ -cohesive sets $\emptyset = V_0 \subset V_1 \subset \dots \subset V_l = V$ such that $V_1 = U$. By Lemma 2.2, the coradius of σ_A is at most l , which cannot exceed $N - K + 1$. \square

In the special case in which there are no coexistent conventions in the population game, the structural conditions of Theorem 2.1 follow from a simpler property of the network, namely, the existence of a small p -cohesive set. This observation is formalized in the following result.

Corollary 2.1. *Suppose that there are no coexistent conventions in the population game. If the network Γ contains a p -cohesive set of size m , then we have the bound*

$$\bar{\alpha} \leq \frac{(\delta_{\max} + \delta_{\min} - 2 \lfloor p\delta_{\min} \rfloor)(m - 1)}{(\delta_{\min} - 2 \lfloor p\delta_{\min} \rfloor)N}.$$

Proof. If we show that in the absence of any coexistent convention, the existence of a p -cohesive group S of size m implies the existence of a minimal $(1-p)$ -cohesive group of size at least $N - m + 2$, then we are done. To see this, it suffices to show that there is no $(1-p)$ -cohesive group of nodes containing at most one element of S . Note first that since S is p -cohesive, any set S' of nodes containing exactly one element of S cannot be $(1-p)$ -cohesive since the node in both sets cannot have enough neighbors in S' . Thus we




Figure 2.2: Nearest neighbor interaction on the circle.

need only consider sets S' disjoint from S .

Suppose for contradiction that S' is a $(1 - p)$ -cohesive set disjoint from S . Now apply the following recursive procedure to the remaining nodes in $V \setminus (S \cup S')$: (i) take all nodes having a fraction of at least $(1 - p)$ of their neighbors in S' and assign them to S' ; (ii) repeat step (i) until there are no more such nodes, and assign all remaining nodes to the set S . I claim that the resulting partition of the nodes describes a coexistent convention. By abuse of notation, let S and S' denote the resulting sets after all other nodes have been assigned, so that $S \cup S' = V$. It is clear by construction that S' is $(1 - p)$ -cohesive, and that each element of S has a fraction of at most $(1 - p)$ of its neighbors in S' . But then since $S \cup S' = V$, each element of S must have a fraction of at least p of its neighbors in S , proving the claim. \square

The p -cohesiveness condition of the corollary may be interpreted as the existence of a small clique. In a regular network that does not support any coexistent convention for the given value of p , the existence of a single clique that is small relative to the size of the population is sufficient to guarantee the robustness of the risk dominance prediction. The regularity assumption precludes the existence of a leader who, by single-handedly changing her action, could affect the incentives of a large number of players.

For the given structural properties, Theorem 2.1 is tight, as the following example demonstrates.

Example 2.2 (Nearest-neighbor interaction on the circle). *An even number, N , of*

players lie at distinct points on a circle. Each agent interacts with the immediate neighbors lying to each side (see Figure 2). Formally, let $\{1, \dots, N\}$ be the set of agents. The nodes i and j are neighbors in the network if and only if $i - j \equiv \pm 1 \pmod{N}$. For any $p \in (0, \frac{1}{2})$, it suffices for one agent to choose action A in order for A to be a best response for both of her neighbors. Therefore, there are no coexistent conventions in the population game, as any such convention must involve at least two neighboring agents who play different actions. In order to apply Corollary 2.1, note that the set $\{1, 2\}$ is p -cohesive, so we may take $m = 2$. Further, we have $\delta_{\max} = \delta_{\min} = 2$, and $\lfloor p\delta_{\min} \rfloor = 0$. Corollary 2.1 therefore gives the bound

$$\bar{\alpha} \leq \frac{2}{N}.$$

As N grows large, the threshold $\bar{\alpha}$ tends to zero. The risk-dominance prediction is therefore strongly robust for this network when the population is large.

To check that this bound is tight, we may compute the precise value of the threshold $\bar{\alpha}$. Since there are only two recurrent classes, σ_A and σ_B , Ellison's radius-coradius method provides both necessary and sufficient conditions for stochastic stability whenever the inequality $R(\sigma) > C(\sigma)$ is strict. Note that, beginning from σ_B , it suffices for a single agent to mutate to action A in order for there to exist a zero-cost path to σ_A . Thus we have $C(\sigma_A) = r(\sigma_B, \sigma_A) = 1$. To compute the radius of σ_A , note that if two neighboring players choose action A , then A will be a best response for both no matter what actions the other players take. Thus any path from σ_A to σ_B must involve a mutation to action B by at least one from every pair of neighboring agents, which implies that $R(\sigma_A) \geq \frac{N}{2}\alpha$. Conversely, there exists a zero-cost path to σ_B from the state σ in which precisely the even-numbered players choose action A . Since σ can be reached from σ_A by $\frac{N}{2}$ mutations to action B , we have $R(\sigma_A) = \frac{N}{2}\alpha$. Combining these results, σ_A is stochastically stable precisely when $\frac{N}{2}\alpha \geq 1$, and therefore

$$\bar{\alpha} = \frac{2}{N},$$

demonstrating that the bound of Theorem 2.1 is tight.

The preceding upper bounds exploit Ellison's radius-coradius theorem to identify conditions under which coordination on the risk dominant action is stochastically stable. Applying the same theorem to the collection of all recurrent classes except for σ_A gives rise to a lower bound.

Theorem 2.3. *Suppose that there exists a $(1 - p)$ -cohesive set of size r in Γ . Then the we have the lower bound $\bar{\alpha} \geq \frac{1}{r-1}$.*

Proof. Let Ω be the set of all Nash equilibria of the population game except for σ_A . Beginning from any state outside of Ω , either there exists a zero-cost path to some state in Ω , or there exists a zero-cost path to σ_A . Thus for computing the radius and coradius of Ω it suffices to consider paths to and from σ_A .

Let S be a $(1 - p)$ -cohesive set of size r , and let σ denote the state in which all agents in S play action B , and all other agents play action A . Let σ' be identical to σ except that one of the agents x in S plays A . Since S is $(1 - p)$ -cohesive, B is the best response for x to the profile σ' . Hence we have $c(\sigma', \sigma) = 0$, and $m(\sigma_A, \sigma) \leq \alpha(r - 1)$. In order to reach σ_A from σ , at least one agent in S must mutate to action A . Therefore, σ lies in the basin of attraction of Ω , and we have $C(\Omega) \leq \alpha(r - 1)$. Since σ_A cannot be reached from any state in Ω without at least one mutation to action A , the radius $R(\Omega)$ is at least one. Therefore, the condition $R(\Omega) > C(\Omega)$ holds whenever $1 > \alpha(r - 1)$, in which case every stochastically stable state lies in Ω . \square

The lower bound of Theorem 2.3 is also tight. In the trivial example of two interacting agents, the unperturbed dynamics are symmetric with respect to the two actions. The state σ_B is therefore the unique stochastically stable state whenever $\alpha < 1$, and the threshold $\bar{\alpha}$ is equal to one.

2.6 Large population games

One difficulty in interpreting the preceding finite network results is that, without a compelling justification for any particular mutation probabilities, the question of how small the value of the threshold $\bar{\alpha}$ must be in order to accept the risk dominance prediction remains open. In general, this value is bounded below by the reciprocal of the population size. Thus it is natural to consider conditions under which $\bar{\alpha}$ approaches zero as the population grows large, ensuring that coordination on A is stochastically stable for a wide range of mutation probabilities.

Consider a network $\Gamma = (V, L)$ on a countably infinite set of nodes V . Assume that there is a uniform upper bound $\Delta \in \mathbb{N}$ on the number of neighbors of any node; that is, assume that for all $i \in V$,

$$\#\{j \in V \mid \{i, j\} \in L\} \leq \Delta.$$

The approach taken here to understand stochastic stability in large population games will be to consider an increasing sequence of finite networks that approach the infinite network in the limit. Such a sequence may be obtained from a *labelling* of the set of nodes V , that is, from a bijection $\iota : \mathbb{N} \rightarrow V$. Given any labelling ι , define for each $n \in \mathbb{N}$ the subnetwork $\Gamma_\iota(n)$ of Γ whose nodes are given by the set $V_\iota(n) := \iota(\{1, \dots, n\})$, and whose links $L_\iota(n)$ consist of all links in L between any two nodes in $V_\iota(n)$; thus

$$L_\iota(n) := \{\{i, j\} \in L \mid \iota^{-1}(\{i, j\}) \subset \{1, \dots, n\}\}.$$

Fixing the payoffs in the underlying 2×2 game, define for each $n \in \mathbb{N}$ the mutation robustness threshold $\bar{\alpha}_\iota(n)$ to be the value of $\bar{\alpha}$ for the network $\Gamma_\iota(n)$.

Definition 2.5. Mutation robustness *holds in the infinite network* $\Gamma = (V, L)$ if there exists a labelling ι of V such that $\lim_{n \rightarrow \infty} \bar{\alpha}_\iota(n) = 0$.

If mutation robustness holds, then the range of mutation probabilities giving rise to σ_A as a stochastically stable outcome can be made arbitrarily large by taking a sufficiently

large population. Note that the set of limit points of the sequence $\bar{\alpha}_\iota(n)$ depends in general on the choice of labelling ι . For example, it is always possible to choose a labelling such that, for each n , the network $\Gamma(n)$ contains at least one isolated node that has no neighbors. In this case, since such an isolated node forms a $(1 - p)$ -cohesive set of size 1, it follows from Theorem 2.3 that $\bar{\alpha}_\iota(n) = 1$ for all n .

Morris (2000) characterizes conditions under which the risk dominant action spreads contagiously in a similar model, which differs only in that the dynamics are deterministic: in each period, every player chooses a best response to the strategy profile of the previous period. *Contagion* occurs in the infinite network Γ if, starting from some initial strategy profile in which only a finite number of agents play action A , every member of the population plays A in the limit as time tends to infinity. Morris shows in particular that contagion occurs in Γ if and only if there exists a labelling ι of the nodes of Γ such that for some sufficiently large $N \in \mathbb{N}$, A is a best response for $\iota(n)$ whenever each node $\iota(1), \dots, \iota(n - 1)$ plays A and $n \geq N$.

Although we have not defined analogues of Ellison's radius and coradius for games played on infinite networks, the occurrence of contagion corresponds intuitively to the coradius of σ_A being finite. Thus mutation robustness should hold as long as the radius of σ_A is infinite. Morris (2000) shows that it is impossible for action B to spread to an infinite set of agents from an initial strategy profile in which only a finite set of agents play B , suggesting that mutation robustness should hold as long as there is no coexistent convention in which a finite set of agents plays B . This intuition is formalized in the following theorem.

Theorem 2.4. *If there exists a finite $(1 - p)$ -cohesive set of nodes in Γ then mutation robustness does not hold. Conversely, if contagion occurs in Γ and there does not exist a $(1 - p)$ -cohesive set of nodes, then mutation robustness holds.*

Proof. For the first part, let S be a $(1 - p)$ -cohesive set of nodes in Γ of finite size m . It suffices to note that given any labelling ι , there exists some N such that ι assigns a label

of at most N to every node in S . By Theorem 2.3, $\bar{\alpha}_\iota(n) \geq \frac{1}{m}$ for all $n > N$.

For the converse, first consider, for each n , the coradius of σ_A in the game played on the network $\Gamma_\iota(n)$. Since contagion occurs, there exists some labelling ι for which there is some absolute bound M_{CR} and some sufficiently large N such that whenever $n \geq N$, this coradius is at most M_{CR} . To prove this, choose a finite set S of nodes from which, if all of these choose A , best response dynamics lead to all agents choosing A . Consider best response dynamics where, in period 0, only members of S choose action A . For each node v , there is some earliest period $k(v)$ after which A is always a best response for v as the best response dynamics are iterated. The desired labelling is any for which $k \circ \iota$ is nondecreasing (that is, the lowest labels are assigned to the nodes that switch to A earliest). Let $M_{CR} = |S|$. It is clear by construction that, beginning from σ_B , mutation of all M_{CR} nodes in S is sufficient to lead to σ_A .

All that remains is to show that for some labelling ι satisfying the requirement of the preceding paragraph, the radius of σ_A in $\Gamma_\iota(n)$ tends to infinity as n grows large. Note that any $(1-p)$ -cohesive group in $\Gamma_\iota(n)$ must contain some member of S by the way in which the labelling ι was chosen. Recall that the degrees of the nodes of Γ are uniformly bounded by some number Δ . For each n , and each $d \in \mathbb{N}$, let $g^d(n) \in \mathbb{N}$ be the smallest number for which all nodes within distance d of any node in $\iota(\{1, \dots, n\})$ are in $\Gamma_\iota(g^d(n))$.

Suppose that the radius of σ_A in $\Gamma_\iota(n)$ does not tend to infinity with n . Then there exists some $M_R \in \mathbb{N}$ such that for each N , there exists some $n > N$ for which the radius of σ_A in $\Gamma_\iota(n)$ is at most M_R . I claim that there exists a number K , depending only on p and Δ , such that beginning from σ_A , any number m of mutations to B can lead, through best response dynamics, to at most Km players choosing B . Assuming for now that the claim is true, let $N = g^{KM_R+1}(|S|)$. For some $n > N$, there exists a $(1-p)$ -cohesive set S_{1-p} in $\Gamma_\iota(n)$ of size at most KM_R , for otherwise the radius of σ_A would be greater than M_R for all $n > N$. As noted above, this set S_{1-p} must contain some element of S . Since we may assume without loss of generality that this set is connected (otherwise take some component), it follows that every node in S_{1-p} lies within distance KM_R of some

member of S , and therefore that every neighbor of every node in S_{1-p} lies in $\Gamma_l(n)$. But then S_{1-p} is $(1-p)$ -cohesive in Γ , contradicting the assumption that Γ contains no finite $(1-p)$ -cohesive group.

All that remains is to prove the claim of the preceding paragraph. Accordingly, let S_{1-p} be any (finite) $(1-p)$ -cohesive group of size M in an arbitrary network Ω , and suppose that m mutations suffice for the members of S_{1-p} to switch to playing B . Then there exists a labelling $\kappa : \{1, \dots, M\} \rightarrow S_{1-p}$ such that for each $n > m$, a fraction of at least $(1-p)$ of $\kappa(n)$'s neighbors in Ω lie in the set $\{1, \dots, n-1\}$. For each $l = 1, \dots, M$, let δ_l be the degree of $\kappa(l)$. For each $l > m$, there must be at least $(1-p)\delta_l$ links connecting $\kappa(l)$ to nodes with smaller labels, and hence at most $p\delta_l$ links connecting $\kappa(l)$ to nodes with higher labels. Thus we have

$$\sum_{l=1}^m \delta_l + p \sum_{j=m+1}^M \delta_j \geq (1-p) \sum_{k=m+1}^M \delta_k.$$

Assuming a uniform upper bound of Δ on the degrees of the nodes in Ω , this implies that

$$m\Delta \geq (1-2p) \sum_{k=m+1}^M \delta_k.$$

Assuming that Ω contains no solitary nodes, so that $\delta_k \geq 1$ for all k , this gives

$$m\Delta \geq (1-2p)(M-m),$$

and therefore,

$$m \geq \frac{(1-2p)}{\Delta + (1-2p)} M.$$

Taking $K = \frac{\Delta + (1-2p)}{(1-2p)}$ therefore gives the desired result. \square

In a similar model, Lee et al. (2003) consider the interaction structure formed by a 2-dimensional torus, and find that an analogue of mutation robustness holds for all values of p , that is, given any model of the mutation probabilities, the risk dominant equilibrium

is stochastically stable when the population on the torus is sufficiently large. As they explain, this result is driven by the existence of small p -cohesive sets of nodes that cover the entire network combined with a stochastic form of contagion. Significantly, however, the result depends on the fact that, for each $p \in (0, \frac{1}{2})$, the size of the smallest $(1 - p)$ -cohesive set of nodes on the torus grows without bound as the size of the network grows large. On the other hand, the stochastic contagion underlying their robustness result is a weaker property than the deterministic contagion of Theorem 2.4, suggesting that it may be possible to generalize this result.

In networks possessing enough symmetry, contagion cannot occur if there exists a finite $(1 - p)$ cohesive set of nodes. To be precise about the relevant notion of symmetry, we require the following definition:

Definition 2.6. *an automorphism ϕ of Γ is a bijection $\phi : V \rightarrow V$ such that x and y are neighbors in Γ if and only if $\phi(x)$ and $\phi(y)$ are neighbors in Γ .*

Thus an automorphism of a network is a permutation of its nodes that preserves the link structure, and each nontrivial automorphism corresponds to a symmetry of the network.

Proposition 2.1. *Suppose that for each $x \in V$ there exist infinitely many $y \in V$ such that there is some automorphism ϕ of Γ satisfying $\phi(x) = y$. Then contagion cannot occur if there exists a finite $(1 - p)$ -cohesive set in Γ .*

Proof. Suppose that contagion occurs in Γ , and that there exists a finite $(1 - p)$ -cohesive set of nodes C . Then there exists a finite set of nodes S whose complement does not contain a $(1 - p)$ -cohesive set (see Morris (2000)). Let d_C be the diameter of the set C ; that is, d_C is the greatest distance between any two nodes in C . Since S is finite and the degrees of the nodes of Γ are uniformly bounded, given any $d \in \mathbb{N}$, there are only finitely many nodes y for which there exists a node in S within distance d of y . Therefore there exists some y lying at a distance strictly greater than d from all nodes of S such that y is the image of some $x \in C$ under some automorphism ϕ of Γ . By construction, the set $\phi(C)$ is disjoint from S , and $\phi(C)$ is $(1 - p)$ -cohesive since C is, contradicting the choice of S . \square

The following corollary is immediate from Theorem 2.4 and Proposition 2.1.

Corollary 2.2. *Under the symmetry assumption of Proposition 2.1, the occurrence of contagion is a sufficient condition for mutation robustness.*

2.7 Waiting times

Ellison (1993) argues that the relevance of stochastically stable outcomes depends on the expected waiting time to convergence, which in turn depends on the interaction structure. Young (1998) extends an argument due to Ellison (1993) to bound the expected waiting time in local interaction games when each node in the network lies in a sufficiently close-knit group. Close-knittedness is a clustering property similar to, but stronger than, the r -cohesiveness used above. Fast convergence occurs under Young's conditions because the required mutations can take place in small steps, each of which is much more likely to occur than are many simultaneous mutations.

In a more general setting, Ellison (2000) bounds the expected waiting time using only the coradius of the set of stochastically stable states, showing that a small coradius is sufficient to ensure fast convergence. Since a small coradius of σ_A tends to favor mutation robustness, one might expect fast convergence and mutation robustness to be closely related. The examples below demonstrate that this intuition is false in general. First, however, we must give a more precise definition of fast convergence for large networks.

Consider a sequence $\bar{\Gamma} = (\Gamma_1, \Gamma_2, \dots)$ of networks such that $|V(\Gamma_{N'})| > |V(\Gamma_N)|$ whenever $N' > N$; that is, the size of the population is strictly increasing along the sequence. Any labelling ι of the nodes of an infinite network Γ naturally gives rise to such a sequence $\bar{\Gamma}_\iota$ by taking $\Gamma_N = \Gamma_\iota(N)$ for all N .

Definition 2.7. *Fast convergence occurs in $\bar{\Gamma}$ if there exists some T not depending on N such that, for each N , from any initial state, the expected time until σ_A is reached in Γ_N is $O(\varepsilon^{-T})$ when $\alpha = 1$.*

This definition captures the idea that, for convergence to be fast, the expected waiting time should not grow by orders of magnitude as the network becomes large. Since best response dynamics typically require more periods to adjust following mutations in a large network compared to a small one, the constant implied by the big- O will generally depend on N . However, the order of magnitude, as measured by the exponent T , must remain bounded as the network grows large.

We say that fast convergence occurs in an infinite network Γ if there exists a labelling ι such that fast convergence occurs in the sequence $\bar{\Gamma}_\iota$. If contagion occurs in Γ , then there exists a labelling ι and a set of nodes S of size K such that, for large enough N , any state in which all members of S choose action A lies in the basin of attraction of σ_A in the network $\Gamma_\iota(N)$. In particular, fast convergence holds in Γ with $T = K$ since, from any initial state, mutation of all members of S to action A is sufficient for the unperturbed dynamics to lead to σ_A .

Example 2.5 (Nearest-neighbor interaction on the circle). *In this case, mutation robustness and fast convergence both hold. For each N , let Γ_N be the network corresponding to nearest-neighbor interaction on a circle of size N , as in Example 2.2. Since σ_A can be reached through the unperturbed dynamics whenever a single agent chooses action A , fast convergence holds with $T = 1$. From Example 2.2, the threshold $\bar{\alpha}$ tends to zero as N grows large.*

The preceding example captures the intuition that mutation robustness and fast convergence coincide if σ_A can be reached from any initial state by a small number of mutations. This coincidence, however, does not extend more generally, as the following examples demonstrate.

Example 2.6 (Uniform interaction). *In this case, fast convergence fails for any p , but the threshold $\bar{\alpha}$ is small when p is small. Let Γ_N be the complete network on N nodes; that is, every player interacts with every other player (see Figure 3). Ellison (1993) shows that fast convergence fails in this network. The mutation robustness threshold, however,*




Figure 2.3: Uniform interaction.




Figure 2.4: Regions of size m for $m = 4$ (Morris, 2000).

is approximately constant as the network grows large. Since there are only two equilibria of the population game, the radius and coradius of σ_A may be used to compute the precise value of $\bar{\alpha}$. Accordingly, we have $C(\sigma_A) = \lceil p(N - 1) \rceil$ and $R(\sigma_A) = \lceil (1 - p)(N - 1) \rceil$, and hence

$$\bar{\alpha} = \frac{\lceil p(N - 1) \rceil}{\lceil (1 - p)(N - 1) \rceil}.$$

In particular, the threshold $\bar{\alpha}$ is small for large N when the payoff parameter p is small.

Example 2.7 (Regions of size m). In this case, for some values of p , fast convergence holds but mutation robustness fails. Consider an infinite network Γ in which the nodes correspond to elements of $\mathbb{Z} \times \{1, \dots, m\}$. Each node (i, j) interacts with all $m - 1$ other nodes having coordinates (i, \cdot) , as well as to the two nodes $(i - 1, j)$ and $(i + 1, j)$ (see Figure 4). The sets $\{\{i, 1\}, \dots, \{i, m\}\}$ are the regions of the network, within which interaction is uniform, and between which links are relatively rare. For each m , contagion occurs

if and only if $p < \frac{1}{m+1}$. If $p > \frac{1}{m+1}$, then any two adjacent regions together form a $(1-p)$ -cohesive set, so by Theorem 2.4, mutation robustness holds in Γ if and only if $p < \frac{1}{m+1}$. On the other hand, fast convergence holds for all $p < \frac{1}{2}$.¹ Intuitively, when $p > \frac{1}{m+1}$, convergence is fast because transitions to σ_A can occur in many small steps through a sequence of coexistent conventions. This rich structure of conventions, however, also ensures that mutation robustness fails because, no matter how large the network, some coexistent convention may be reached from σ_A by only a fixed number of mutations.

2.8 Discussion and conclusion

The structure of social networks has been widely studied in the sociology literature, and a number of regularities have been empirically observed in a variety of settings (see Newman (2003) for a survey). We may consider, then, how these properties relate to the structural conditions described above that are relevant for mutation robustness in order to assess the relevance of the risk dominance prediction in real-world networks when mutation probabilities are unmodelled. This discussion must, however, necessarily remain vague since definitions of observed network properties vary, and quantification is difficult in general.

- *Small-world networks.* A number of real-world networks have been found to have significant local clustering, but at the same time a small global diameter relative to certain highly structured networks; that is, the distance between any two nodes is “small” given the size of the population. Such networks are said to possess the *small-world property*. Clustering alone, if sufficiently dense, can correspond to the existence of small highly cohesive groups of nodes, thereby placing a lower bound on the threshold parameter $\bar{\alpha}$ that depends on the size of the smallest such group. In large populations, the existence of such a group may preclude mutation robustness.

¹The proof of this result relies on a strengthening, also due to Ellison, of the radius-coradius method, in which the coradius is replaced by the (smaller) *modified coradius* (Ellison 2000). One can show that, for an appropriate labelling of Γ , the modified coradius of σ_A in $\Gamma_i(N)$ is at most $2m$ regardless of N . Theorem 2 of Ellison (2000) then implies that fast convergence holds for $T = 2m$. The details of the modified coradius calculation are somewhat involved, and are omitted here.

Furthermore, Morris (2000) provides sufficient conditions for contagion that include “low neighbor growth,” which is inconsistent with a small diameter. Thus insofar as contagion may contribute to mutation robustness, the small diameter property of small-world networks also appears to be contrary to the structural properties required for a robust risk dominance prediction, although one must be cautious here since this is based only on sufficient conditions for mutation robustness.

- *Community structure.* The nodes in networks sometimes form identifiable groups or communities in such a way that the density of links is much higher within groups than between them. As in the regions example of the preceding section, if these communities are sufficiently strong in the sense that a sufficiently large proportion of links in the network are within groups, then highly cohesive sets of nodes will exist, some of which will be small relative to the size of the population if many communities exist in the network. Thus community structure may also prevent mutation robustness from occurring.
- *Scale-free networks.* The distribution of degrees of nodes in a purely random network, in which there is a fixed independent probability that a link exists between any two nodes of the network, is binomial, approaching a Poisson distribution as the population grows large (Newman 2003). In real-world networks, the degree distribution typically features a heavier tail than that for random networks, corresponding to a greater number of high-degree nodes. Recall that the upper bound on the threshold value $\bar{\alpha}$ given in Theorem 2.1 is strongest when all nodes have the same degree, and becomes weaker as the distribution becomes more dispersed. The presence of high degree nodes can reduce the radius of coordination on the risk dominant action, lessening the bias in mutations necessary to overturn the stability of this equilibrium.
- *Tie strength.* The model employed here assumes for simplicity that all links are given equal weight in each player’s payoffs. In general, however, these weights may differ, for example because of non-uniform probabilities of matching. The analysis

extends naturally to this more general setting, with the role of p -cohesiveness replaced by *weighted* p -cohesiveness: action A is a best response for player i if the share of weights associated with those of i 's neighbors who play action A is at least p . Strong links, corresponding to those which are assigned higher payoff weights, tend to exhibit greater clustering than weak links (Granovetter 1973). Weighting links will therefore increase the likelihood that a small highly (weighted) cohesive set will exist in the network, which again limits the size of the bias necessary to overturn the risk dominance prediction.

To summarize, for each of the network properties that have been most prominent in the empirical literature, none contributes to mutation robustness. This suggests the need to be careful when modelling mutation rates in local interaction environments, as large biases in mutation probabilities may not be necessary to alter the set of stochastically stable outcomes.

Bibliography

- Bergin, J. and Lipman, B. L.: 1996, Evolution with state-dependent mutations, *Econometrica* **64**, 943–956.
- Blume, L. E.: 2003, How noise matters, *Games and Economic Behavior* **44**, 251–271.
- Ellison, G.: 1993, Learning, local interaction, and coordination, *Econometrica* **61**, 1047–1071.
- Ellison, G.: 2000, Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution, *Review of Economic Studies* **67**, 17–45.
- Foster, D. P. and Young, H. P.: 1990, Stochastic evolutionary game dynamics, *Theoretical Population Biology* **38**, 219–232.
- Goyal, S. and Vega-Redondo, F.: 2005, Network formation and social coordination, *Games and Economic Behavior* **50**, 178–207.
- Granovetter, M.: 1973, The strength of weak ties, *American Journal of Sociology* **78**, 1360–1380.
- Hojman, D. and Szeidl, A.: 2006, Endogenous networks, social games and evolution, *Games and Economic Behavior* **55**, 112–130.
- Kandori, M., Mailath, G. J. and Rob, R.: 1993, Learning, mutation, and long run equilibria in games, *Econometrica* **61**, 29–56.

- Lee, I. H., Szeidl, A. and Valentinyi, A.: 2003, Contagion and state dependent mutations, *Advances in Theoretical Economics* **3**(1).
- Morris, S.: 2000, Contagion, *Review of Economic Studies* **67**, 57–78.
- Newman, M.: 2003, The structure and function of complex networks, *SIAM Review* **45**, 167–256.
- Pęski, M.: 2004, State-independent evolution on networks. Mimeo.
- Young, H. P.: 1998, *Individual Strategy and Social Structure*, Princeton University Press, Princeton, NJ.

Chapter 3

Efficient Dynamic Coordination with Private Learning

3.1 Introduction

Coordination problems arise in a wide variety of economic situations. A typical example is of a setting where the successful implementation of some socially beneficial project depends on whether enough agents participate. Such settings may lead to coordination failure, which arises when a given group of agents fails to participate in the project despite the fact that it is in their collective interest to do so.

The traditional theoretical analysis of coordination problems, where payoffs are typically assumed to be commonly known, has been plagued by the existence of multiple equilibria. For a given payoff rule, there exists at least one equilibrium with coordination failure, and one without. Such analysis is unable, therefore, to quantify the extent and relevance of coordination failure, since it is not possible to assign probabilities across equilibria.

The recent literature on global games (Carlsson and van Damme 1993, Morris and Shin 2003) has made substantial progress in resolving the problem of multiplicity in the analysis of coordination problems. This literature has identified an important class of coordination

games, in which underlying payoffs are observed with small amounts of idiosyncratic noise, where the multiplicity of equilibria is eliminated. The “refinement” thus achieved allows us to quantify the extent of coordination failure, and indeed coordination failures do occur in the unique equilibrium of the canonical global game. Whether coordination failure arises depends on the payoffs of the underlying complete information game. Roughly speaking, agents are only able to coordinate on some risky action in the unique equilibrium of a global game if that action is *risk dominant*, i.e., it is optimal for each agent to choose that action in the underlying complete information game even when there is only a “low” probability that his fellow players will choose that action.¹ This can only happen if the benefits that arise from the action conditional upon success are high relative to the cost of undertaking it. Thus, the global games literature has negative implications for the ability of agents to coordinate on socially beneficial actions: only projects that involve “little strategic risk” will be implemented in equilibrium. In all other cases, coordination failure will arise.

In this paper we evaluate how the incidence of coordination failure in global games is affected by the presence of multiple opportunities to participate between which players privately learn about the fundamental. The canonical global game requires that all agents choose their actions simultaneously. To what extent would the incidence of coordination failure change if we allowed for some asynchronicity in the actions of potential participants in a coordination problem?

To be specific, consider a setting in which the success of a socially beneficial investment project depends on the total number of agents who invest over the course of T distinct periods. Two players choose at which period (if any) to invest irreversibly, while observing noisy private signals (where the standard deviation of the noise is denoted σ_t) about the underlying state variable (θ). At each period t , the information structure is that of a canonical global game. We assume that agents privately learn the fundamental θ asymp-

¹More precisely, equilibria of the underlying complete information game survive in the induced global game only if they are p -dominant (Morris, Rob and Shin 1995) for “low” p . Exactly how low p must be depends on the structure of the game. In two player games, p -dominant equilibria for $p < \frac{1}{2}$ survive. See Kajii and Morris (1997) for a generalization of this idea.

totically, i.e., $\sigma_t \rightarrow 0$ as $t \rightarrow \infty$. The project succeeds if the fundamental is good and each player invests in some period. Note that it is *not* necessary for success that both players invest in the same period. Thus, this is an asynchronous investment game. The choice between early vs late investment is driven by a trade-off: early investment generates higher payoffs if the project succeeds, while late investors have more accurate private information about payoffs. As in a standard global game, we assume that there exist values of θ that make investment dominant ($\theta \geq 1$) or dominated ($\theta \leq 0$). To fix ideas, imagine that the fundamental is such that investing is *not* risk-dominant. This means that if agents had to choose their actions simultaneously in some period, say T , and thus play a static global game, then, in the limit as noise vanishes, coordinated investment could not be supported as an equilibrium outcome, and coordination failure arises. To what extent will the possibility of choosing actions asynchronously affect the incidence of coordination failure? We report the following results:

1. Coordination failure almost never arises in a sufficiently long asynchronous investment game.

For any $\underline{\theta} > 0$ and $\varepsilon > 0$ there exists some \underline{T} such that for any $T \geq \underline{T}$, investment succeeds with probability at least $1 - \varepsilon$ in the asynchronous game with T periods whenever $\theta > \underline{\theta}$. Thus, in the limit as $T \rightarrow \infty$, the project succeeds whenever $\theta > 0$. In addition, as noise in observation vanishes (i.e., $\sigma_t \rightarrow 0$ for all t) there is also no delay in investment: players successfully coordinate on implementing the project immediately, thus achieving the social optimum.

2. The forces driving our results can be cleanly characterized in terms of higher order beliefs in the asynchronous coordination game.

Building on the standard belief operators of Monderer and Samet (1989), we construct an *asynchronous p-belief* operator which is suitable for characterizing behaviour in our asynchronous investment game. Using this operator, we show that by choosing sufficiently long asynchronous investment games, it is possible to generate adequate

levels of generalized approximate common knowledge (i.e., generalized common p -belief for arbitrarily high p) in order to support asynchronous coordination. The generalization lies in allowing the required beliefs to be attained at different times.

If *synchronous* participation at the last round T was necessary for the success of the project, then players invest only if they commonly p -believe in the standard sense of Monderer and Samet (1989), for sufficiently high p , that the fundamental allows success. It is now well understood (see, for example, Morris and Shin (2003)) that, however small the private errors are, the global games information structure does not generate common p -belief for $p > \frac{1}{2}$, and thus coordination fails whenever investment is not risk-dominant. In our setting, players do not have to participate synchronously at T , but both players must participate eventually by period T . In such a situation, only a relaxed version of common beliefs is necessary for coordination. Fix a probability of success $p \in (0, 1)$ sufficient to induce players to invest in period t . Both players will invest by period T , if they both believe with probability p by period T that the fundamental is good, they both believe with probability p by period T that they both believe with probability p by period T that the fundamental is good...etc. We refer to such an event as asynchronous common p -belief of event $\theta > 0$. This variation of standard common belief turns out not to be very demanding in our setting.

To obtain some intuition for why asynchronous approximate common knowledge is attained in long games, consider a game with infinitely many rounds in which each player asymptotically privately learns the fundamental. Then if the fundamental lies in some open set G , all players will eventually p -believe G almost surely for any $p < 1$. This makes event G p -evident in an asynchronous sense, which in turn implies asynchronous common p -belief of G . The shortcoming of this line of argument is that it relies on the assumption that the fundamental is asymptotically perfectly revealed to players. It is thus not clear whether the argument extends to long but

finite games in which some information about the fundamental remains uncovered. Existing results on static global games show that there is an important discontinuity in the structure of standard common beliefs as information about the fundamental becomes infinitely precise. We find that such a discontinuity does *not* arise when the infinite asynchronous game is approximated by a sequence of finite asynchronous games. Approximate asynchronous common knowledge *is* attained even if small uncertainty about the fundamental remains.

Our results suggest that allowing asynchronicity and private learning in coordination problems may substantially reduce the extent of coordination failure in global games. In addition to being of theoretical interest, our results are potentially widely applicable. For example, consider the problem of foreign direct investment (FDI) into a newly liberalizing emerging market. Payoffs from FDI depend on whether the emerging economy “takes off”, which in turn depends on the amount of FDI. Thus, this is a coordination problem. In addition, it may not matter precisely that all FDI takes place at the same time, but simply that it occurs during the first several months to several years of the liberalization programme. It is not uncommon for liberalization to be accompanied by government subsidies to early investors. Yet, it is also likely that late investors will have better information about the state of the underlying emerging markets. Finally, it seems unlikely that a great deal of reliable public information is available about the prospects of emerging economies – relevant information is garnered via individual research, and is thus at least partially private. Thus, the class of stylized games outlined above represents trade-offs that are not dissimilar to those outlined in this applied context. The FDI example is not unique. Indeed, it may be reasonable to argue that several of the applications studied to date using global games (e.g. currency crises, bank runs, financial contagion etc.) may well have an element of asynchronicity to them.

The learning process is exogenous in our model, and thus we are silent about its source. This process may be viewed as a reduced form of social learning. It is possible to model this

social learning more explicitly by using the methods developed in Dasgupta (2007), who endogenizes the late private signals as noisy observations of early investment levels. Since our aim is to study the consequences of private learning, and not its source, we employ the exogenous information structure.

The rest of the paper is organized as follows. In section 3.2 we outline the model. Section 3.3 states our main result, while section 3.4 explains the efficiency result in terms of asynchronous common p -belief. Section 3.5 concludes. Before proceeding to the main model, however, we first outline the related literature.

3.1.1 Literature Review

Our analysis originated in the work of Dasgupta (2007). Dasgupta outlines conditions under which the provision of the option to delay combined with private learning improves the ability of agents to coordinate efficiently in two-stage global games. We use Dasgupta's modeling framework for an analysis of a different but related question. We do not compare the coordination outcomes in different finite games², rather, studying long finite games, we approximate a limit case in which players learn the true value of the fundamental asymptotically, and we find that players coordinate efficiently in the limit. Another example of a dynamic global game with private learning can be found in Heidhues and Melissas (2006).

Our explanation of the efficiency result in terms of higher order beliefs builds on the work of Monderer and Samet (1989) and Morris and Shin (2007). For the purposes of explicating the higher order beliefs foundations of our result, we map our dynamic game to a static global game. We then use the methodology of (Morris and Shin 2007) which enables a characterization of rationalizable actions in terms of generalized common beliefs for a broad class of static games. The generalized common belief operators relevant for our game have a natural interpretation: they differ from the standard common beliefs of Monderer and

²Dasgupta's main contribution is the finding that an option of delay increases efficiency; investment in the early stage decreases, but late investments more than compensate the decline in the early stage. We do not study such trade-off.

Samet (1989) “only” in that players do not care about the time at which opponents attain the required beliefs. Our analysis also shows that private learning generates the relevant generalized common beliefs in our game. Indeed, such common beliefs are very easy to attain in our setting: they essentially coincide with first order beliefs.

The current analysis bears a general connection to models of information dynamics in multi-stage global games (e.g., Chamley (2003), Angeletos, Hellwig and Pavan (2007)). In contrast to our work, papers in this strand of the literature focus on learning from endogenously generated *public* signals, and focus on the robustness of equilibrium uniqueness in global games. We restrict attention to pure private information settings with a unique long-run outcome, and focus on characterizing the (lack of) incidence of coordination failure.

Beyond the literature on global games, our analysis is related to the work of Cripps, Ely, Mailath and Samuelson (2006). These authors delineate general conditions under which agents asymptotically attain approximate common knowledge via private learning. The analysis of Cripps *et al* has important implications for long-run outcomes in situations which can be divided into two distinct phases: agents learn privately in the first phase, and attempt to coordinate *synchronously* in the second. We study situations in which those two phases are merged together: players attempt to coordinate *asynchronously* while they privately learn about payoffs. Both papers study whether private learning leads to approximate common knowledge. However, *different* concepts of approximate common knowledge are relevant for ensuring successful coordination in synchronous and asynchronous coordination games, because the payoff-relevant events differ in these two types of games. Cripps *et al* study standard common beliefs as defined in (Monderer and Samet 1989), while we study an asynchronous form of common beliefs. The two concepts turn out to have very different properties. In our model, private learning fails to deliver common knowledge in the standard sense as studied by Cripps *et al*, but succeeds in delivering asynchronous common beliefs. This explains why coordination failure arises in the synchronous coordination

game but does not arise in our asynchronous game.³

Gale (1995) provides an elegant analysis that bounds the extent of inefficient delay in dynamic coordination games with complete information. Based on the observation that players can induce other players to invest by investing early, Gale establishes an upper bound on inefficient delay in coordinated investment. For a small number of players his result implies nearly efficient coordination. However, the bound increases linearly with the number of players. Though we formulate our model only for two players, this is only for the sake of exposition, and our results on the elimination of coordination failure generalize immediately to settings with arbitrarily many players. In addition, the reasoning behind the results is very different. While Gale uses backward induction, and hence relies on perfect information, we use elimination of serially dominated strategies under incomplete information, in a setting in which players do not observe each others' past choices. Hörner (2004) studies a model similar to that of Gale (1995), and finds that patient players coordinate efficiently when they receive a single noisy signal of payoffs prior to the play of the game.

3.2 Model

Two players $i \in \{1, 2\}$ play a joint investment game Γ_T , with $T \in \mathbb{N}$.⁴ The game consists of T rounds, all of which may take place within a finite, possibly short time window. In each round $t \in \{1, \dots, T\}$, each player chooses one of the two actions $a_t^i \in \{0, 1\}$; we interpret Action 1 as “invest”, and Action 0 as “wait”. Each player may invest in at most one round. The payoffs depend on the action profiles and the value of a fundamental parameter $\theta \in \mathbb{R}$ describing the characteristics of the project. The fundamental θ is drawn before the first round according to an improper uniform distribution on \mathbb{R} , and remains fixed over all

³In an earlier paper, Ely (2003) informally discusses the notion of asynchronous common belief, but only to contrast it to the standard common belief which is the relevant concept for the types of problems he considers.

⁴All of our results generalize easily to any finite number of players. Indeed, if the learning process is viewed as a reduced form model of social learning, then the assumption of a large number of players becomes natural.

rounds.

The players do not observe the true value of the fundamental θ ; instead, they receive private noisy signals of the value of θ in every round. Specifically, each player i receives a signal $\tilde{x}_t^i = \theta + \tilde{\sigma}_t \epsilon_t^i$ in round t , where the errors ϵ_t^i are drawn from $N(0, 1)$ and are independent across players and rounds. The standard errors $\tilde{\sigma}_t$ are strictly positive for all t . Player i does not observe the choices of player $-i$ before the end of the game.

Players form their beliefs in each period about the true value of the fundamental through Bayesian updating given their received signals. The resulting beliefs over θ conditional on a sequence $(\tilde{x}_{t'}^i)_{t'=1}^t$ are distributed as $N(x_t^i, \sigma_t^2)$, where

$$x_t^i = \frac{\sum_{t'=1}^t \tilde{x}_{t'}^i \frac{1}{\tilde{\sigma}_{t'}^2}}{\sum_{t'=1}^t \frac{1}{\tilde{\sigma}_{t'}^2}},$$

and

$$\frac{1}{\sigma_t^2} = \sum_{t'=1}^t \frac{1}{\tilde{\sigma}_{t'}^2}.$$

We will refer to x_t^i as the *cumulative signal*, and to σ_t as the *cumulative standard error*.

We assume that players asymptotically *privately* learn the true fundamental; that is,

$$\lim_{t \rightarrow \infty} \sigma_t = 0. \tag{3.1}$$

Note that, since each standard error $\tilde{\sigma}_t$ is strictly positive, each cumulative standard error σ_t is also strictly positive. Thus even though players learn the true fundamental in the limit over all periods, some uncertainty remains in each round.

The success of the project is determined at the end of the game, based on the fundamental θ and the actions of the players:

- For $\theta \leq 0$, the project fails regardless of the players' actions.
- For $\theta \geq 1$, the project succeeds regardless of the players' actions.
- For $0 < \theta < 1$, the project succeeds if and only if both players invest in *some* round,

possibly asynchronously.

Each player's payoff in the game depends on whether and in which round the player invested, and whether the project succeeded. The payoffs are

- 0 if the player never invests,
- $\delta^t b$ if the player invests in round t and the project succeeds, and
- $-\delta^t c$ if the player invests in round t and the project fails,

where the parameters b and c are both strictly positive, and $\delta \in (0, 1)$.

The payoffs in this game are consistent with a wide variety of applied settings. For example, they can be easily understood in the context of the FDI example discussed in the introduction. Future payoffs from FDI are positive only if enough foreign firms participate, and the state of the domestic economy (θ) is not too weak. Net benefits from successful FDI participation decline for later participants due to declining subsidies from the emerging market government. Net costs in the event of failure decline for late participants as well, due to a smaller lock-in period for valuable resources.

We now proceed to analyze this game, and show the crucial role of asynchronicity and asymptotic learning in eliminating coordination failure.

3.3 Analysis

The payoffs outlined above imply two simple properties of the best response correspondence, which we describe below in Lemmas 3.1 and 3.2. Our main results, in turn, can be fully stated in terms of these two properties.

Lemma 3.1. *There exists some $\underline{p} \in (0, 1)$ such that, in any round t , waiting is the unique best response for any type that believes the project will succeed with probability less than \underline{p} .*

Proof. The payoff to investing immediately is

$$\delta^t (pb + (1 - p)(-c)).$$

The minimum value to waiting is 0. Thus, \underline{p} is defined by

$$\delta^t (\underline{p}b + (1 - \underline{p})(-c)) = 0,$$

or equivalently,

$$\underline{p} = \frac{c}{b + c},$$

as needed. ■

Lemma 3.2. *There exists some $\bar{p} \in (0, 1)$ such that, in any round t , investing is the unique best response for any type that believes the project will succeed will probability greater than \bar{p} .*

Proof. The payoff to investing immediately is

$$\delta^t (pb + (1 - p)(-c)).$$

The maximum value to waiting is $\delta^{t+1}b$. Thus, \bar{p} is defined by

$$\delta^t (\bar{p}b + (1 - \bar{p})(-c)) = \delta^{t+1}b$$

or equivalently,

$$\bar{p} = \frac{\delta b + c}{b + c},$$

as needed. ■

We note that $1 > \bar{p} > \underline{p} > 0$. Finally, we observe that the existence of $\bar{p} < 1$ implies that however great the amount future information, any player will choose to invest immediately if she is sufficiently optimistic.

We are now in a position to state our main results, which demonstrate the stark difference between synchronous and asynchronous coordination games. We begin with the benchmark synchronous case.

3.3.1 The failure of coordination in the synchronous game

As a benchmark to compare our results to the existing literature on static global games, consider the following static, synchronous version of the game Γ_T . In the synchronous version, which we label by Γ_T^S , for $0 < \theta < 1$, the project succeeds if and only if both players invest synchronously at round T . All other features remain unchanged. We show that for any $\theta < 1$ coordinated investment fails with arbitrarily high probability as long as T is big enough, whenever $\underline{p} > \frac{1}{2}$.

Proposition 3.1. *Fix any $\underline{p} \in (\frac{1}{2}, 1)$. For any $\bar{\theta} < 1$ and $\varepsilon > 0$ there exists some \underline{T} such that for any $T \geq \underline{T}$, the project fails with probability at least $1 - \varepsilon$ in Γ_T^S whenever $\theta \leq \bar{\theta}$.*

This result is a consequence of results from the extant literature on static global games (see Morris and Shin (2003)), and so we only discuss the argument informally. The game played at round T is a canonical static global game with signals x_T^1 and x_T^2 with precision σ_T . The unique equilibrium of this game is characterized by a threshold, x_T^* , such that players invest if and only if their signals satisfy $x_T^i \geq x_T^*$. A player observing the threshold signal $x_T^i = x_T^*$ assigns probability $\frac{1}{2} < \underline{p}$ to the event that her opponent received a signal above x_T^* . This is a consequence of the Laplacian beliefs property of the global games information structure discussed in (Morris and Shin 2003). Unless the threshold player assigns probability almost \underline{p} to the event $\theta \geq 1$, she strictly prefers to wait. Thus, the distance of the indifference point x_T^* from 1 must be on the order of σ_T , and hence, as T becomes large and σ_T small, x_T^* approaches 1.

In order for coordination failure to occur in the synchronous game, it is not essential that agents can invest only at round T . In fact, a similar result would hold in an alternative benchmark game where agents are free to choose in which round to invest, but the project succeeds only if they both end up investing in the same round.

In sharp contrast to these synchronous settings, we now show that coordination almost never fails in the asynchronous game for $\theta > 0$.

3.3.2 The success of coordination in the asynchronous game

The following proposition establishes that, in the game with many rounds, both players are likely to invest whenever the fundamental allows for success of the project ($\theta > 0$).

Proposition 3.2. *Fix any $\bar{p} \in (0, 1)$. Suppose that both players play serially interim undominated strategies. For any $\underline{\theta} > 0$ and $\varepsilon > 0$ there exists some \underline{T} such that for any $T \geq \underline{T}$, the project succeeds with probability at least $1 - \varepsilon$ in Γ_T whenever $\theta > \underline{\theta}$.*

Proof. Fix $q \in (\bar{p}, 1)$. Denote the event that player i q -believes event E at t by $B_q^{(i,t)}(E)$; that is, $B_q^{(i,t)}(E) = \{x_t^i \mid \Pr(E|x_t^i) \geq q\}$. Denote by $l_t^{\theta^*,q}(\theta)$ the probability that the player has q -believed that $\theta > \theta^*$ at least once up to and including round t :

$$l_t^{\theta^*,q}(\theta) = \Pr \left(\bigcup_{t'=1,\dots,t} B_q^{(i,t')}(\theta \geq \theta^*) \mid \theta \right).$$

Let $l^{\theta^*,q}(\theta)$ denote $\lim_{t \rightarrow \infty} l_t^{\theta^*,q}(\theta)$. This limit exists because $l_t^{\theta^*,q}(\theta)$ is non-decreasing in t .

Lemma 3.3. *For all $0 < q < 1$ and all $\theta^* \in \mathbb{R}$: $l^{\theta^*,q}(\theta^*) = 1$.*

The main idea of the proof of Lemma 3.3 is the following: conditional on θ^* , the probability that a player q -believes $\theta > \theta^*$ is $1 - q$ in each round, but with the complication that the posterior probabilities $p^{(i,t)} = \Pr(\theta > \theta^* | x^{(i,t)})$ are correlated across rounds. We will show, roughly, that beliefs across sufficiently distant rounds t and t' are approximately independent. The intuition for this is that if the amount of information that a player receives between t and t' is large relative to what she knew at t , then the information at t has only a negligible impact at t' . For long games, we can choose a long subsequence of rounds such that all rounds in the subsequence are sufficiently distant. Hence the probability of q -believing $\theta > \theta^*$ in at least one of these rounds approaches one as the number of rounds grows large.

Proof of Lemma 3.3. For any sequence $\tau = (t_1, t_2, \dots)$, let

$$l^{\theta^*, q}(\theta^*; \tau) = \Pr \left(\bigcup_{t'=t_1, t_2, \dots} B_q^{(i, t')}(\theta > \theta^*) \mid \theta^* \right).$$

We have $l^{\theta^*, q}(\theta^*; \tau) \leq l^{\theta^*, q}(\theta^*)$, so it suffices to show that, for any $\varepsilon > 0$, there exists some sequence τ for which $l^{\theta^*, q}(\theta^*; \tau) \geq 1 - \varepsilon$.

Given $\varepsilon > 0$, let $\underline{x} < \theta^*$ be such that $\Pr(x_1 < \underline{x} | \theta^*) < \varepsilon(1 - \varepsilon)$. Let $t_1 = 1$ and choose each subsequent t_k so that

$$\Pr(x_{t_k} < \underline{x} | \theta^*) < \varepsilon^k(1 - \varepsilon).$$

For this sequence t_1, t_2, \dots , we have

$$\Pr(x_{t_k} < \underline{x} \text{ for some } k | \theta^*) < \varepsilon.$$

Thus it suffices to show that, as long as $x_{t_k} \geq \underline{x}$ for all k , there almost surely exists some period t_k in this sequence at which the player q -believes that $\theta > \theta^*$. By the Borel-Cantelli Lemma, it suffices to show that for some $\delta > 0$ and some subsequence τ of this sequence, the player q -believes that $\theta > \theta^*$ with independent probability δ in each period in τ .

A player q -believes in period t that $\theta > \theta^*$ as long as $\frac{x_t - \theta^*}{\sigma_t} > \Phi^{-1}(q)$, where $\Phi(\cdot)$ denotes the standard normal distribution function. Given that $x_{t_k} \geq \underline{x}$, for $t > t_k$, we have

$$\begin{aligned} x_t &= \frac{\sigma_t^2}{\sigma_{t_k}^2} x_{t_k} + \sigma_t^2 \sum_{s=t_k+1}^t \frac{\tilde{x}_s}{\tilde{\sigma}_s^2} \\ &\geq \frac{\sigma_t^2}{\sigma_{t_k}^2} \underline{x} + \sigma_t^2 \sum_{s=t_k+1}^t \frac{\tilde{x}_s}{\tilde{\sigma}_s^2}. \end{aligned}$$

Hence $\frac{x_t - \theta^*}{\sigma_t} > \Phi^{-1}(q)$ whenever

$$\frac{\frac{\sigma_t^2}{\sigma_{t_k}^2} \underline{x} + \sigma_t^2 \sum_{s=t_k+1}^t \frac{\tilde{x}_s}{\tilde{\sigma}_s^2} - \theta^*}{\sigma_t} > \Phi^{-1}(q). \quad (3.2)$$

Since $\frac{1}{\sigma_t^2} = \frac{1}{\sigma_{t_k}^2} + \sum_{s=t_k+1}^t \frac{1}{\tilde{\sigma}_s^2}$, Inequality (3.2) is equivalent to

$$(\underline{x} - \theta) \frac{\sigma_t}{\sigma_{t_k}^2} + \left(\sigma_t \sqrt{\frac{1}{\sigma_t^2} - \frac{1}{\sigma_{t_k}^2}} \right) \left(\sqrt{\sum_{s=t_k+1}^t \frac{1}{\tilde{\sigma}_s^2}} \left(\frac{\sum_{s=t_k+1}^t \frac{\tilde{x}_s}{\tilde{\sigma}_s^2}}{\sum_{s=t_k+1}^t \frac{1}{\tilde{\sigma}_s^2}} - \theta^* \right) \right) > \Phi^{-1}(q).$$

The first term on the left-hand side of this inequality tends to zero as t grows large. The second term is a product of two factors, the first of which tends to one as t grows large, and the second of which is a standard normal random variable independent of the realizations of all signals up to period t_k . Therefore, for small enough $\delta > 0$, the inequality holds with probability at least δ when t is sufficiently large given δ , q , and \underline{x} . \blacksquare

Lemma 3.4. *Suppose that both players play serially interim undominated strategies. For any $\underline{\theta} > 0$ and $q \in (\bar{p}, 1)$, there exists some \underline{T} such that for any $T > \underline{T}$, player i invests in round t of game Γ_T if she believes at t with probability at least q that $\theta \geq \underline{\theta}$.*

Proof. Let θ^{**} be the infimum of those $\underline{\theta}$ for which the statement holds. We must show that $\theta^{**} = 0$. We proceed by a contagion argument. The statement clearly holds for $\underline{\theta} \geq 1$. The proof consists of showing that if the statement holds for all $\underline{\theta} > \theta^*$ for some $\theta^* > 0$, then there exists $\varepsilon > 0$ such that it holds for all $\underline{\theta} > \theta^* - \varepsilon$. Thus we must have $\theta^{**} = 0$, for otherwise taking $\theta^* = \theta^{**}$ would give a contradiction.

Assume that the statement is true for all $\underline{\theta} > \theta^*$ for some $\theta^* > 0$. Then there exists some T'' such that in any game Γ_T with $T > T''$, player $-i$ invests at t if she q -believes that $\theta > \theta^*$. Thus, whenever $\bigcup_{t'=1, \dots, T} B_q^{(-i, t')}(\theta > \theta^*)$ is true, player $-i$ will invest at some t in the game Γ_T .

Fix some $r \in \left(\frac{\bar{p}}{q}, 1\right)$. Lemma 3.3 implies that there exists T' such that

$$l_{T'}^{\theta^*, q}(\theta^*) > r.$$

The function $l_{T'}^{\theta^*, q}(\cdot)$ is continuous. Therefore, there exists some $\varepsilon \in (0, \theta^*)$ such that

$$l_{T'}^{\theta^*, q}(\theta^* - \varepsilon) \geq r.$$

Since $l_T^{\theta^*,q}(\theta)$ is non-decreasing in T and θ , we have

$$l_T^{\theta^*,q}(\theta) > r$$

for all $T > T'$ and $\theta > \theta^* - \varepsilon$.

Now consider a game Γ_T with $T > \max(T'', T')$. Suppose player i q -believes at t that $\theta > \theta^* - \varepsilon$. Since $T > T'$, $\theta > \theta^* - \varepsilon$ implies that $l_{-i,T}^{\theta^*,q}(\theta) \geq r$. Since $T > T''$, by hypothesis, player $-i$ invests at t if she q -believes that $\theta > \theta^*$ at t . Thus, conditional on $\theta > \theta^* - \varepsilon$ the probability that player $-i$ invests is no less than r . Therefore, at t , player i attaches probability at least rq to the event that the project succeeds. Since $rq > \bar{p}$, this implies that player i invests at t . ■

The proposition follows from Lemmas 3.3 and 3.4. Fix $\underline{\theta} > 0$ and $\varepsilon > 0$. By Lemma 3.4, there exists some T' such that each player invests in the game Γ_T with $T > T'$ if she q -believes that $\theta \geq \underline{\theta}$. By Lemma 3.3, there exists some T'' such that for $T > T''$, when the fundamental is at least $\underline{\theta}$, the probability that both players q -believe that $\theta \geq \underline{\theta}$ in some round in Γ_T is greater than $1 - \varepsilon$. Taking $\underline{T} = \max\{T', T''\}$ gives the result. ■

Thus, in sharp contrast the synchronous case, coordination failure arises with vanishing probability in the asynchronous case as the number of rounds grows large. In addition, if, as in the synchronous case, we let observation noise vanish, we get the even stronger implication that there is no delay in successful coordination. This is a corollary of Proposition 3.2. To make this idea precise, consider a family of sequences $(\sigma\sigma_t)_{t=1}^{\infty}$, where $\sigma > 0$ is a scaling factor, and $(\sigma_t)_{t=1}^{\infty}$ is some fixed sequence with strictly positive members converging to 0. We will denote by $\Gamma_T(\sigma)$ game with T rounds and noise parameters $(\sigma\sigma_t)_{t=1}^T$.

Corollary 3.1. *For any $\underline{\theta} > 0$ and $\varepsilon > 0$ there exists some $\bar{\sigma} > 0$ and \underline{T} such that for any $\sigma < \bar{\sigma}$, in any equilibrium of $\Gamma_T(\sigma)$ with $T > \underline{T}$, both players invest in round 1 with probability at least $1 - \varepsilon$ whenever $\theta \geq \underline{\theta}$.*

What explains the stark difference in outcomes in the synchronous and asynchronous coordination games? One instructive way to interpret this difference arises out of character-

izing the higher order beliefs of players in these games. We turn to such a characterization in the next section.

3.4 Higher Order Beliefs

It is well-known that the coordination failure arising in static global games can be explained by the lack of approximate common knowledge. The finding that coordination failure does not arise in our asynchronous global game indicates that some aspects of higher order beliefs differ between synchronous and asynchronous global games. The current section is devoted to examining this difference.

First, we introduce notation for payoff-relevant sets of fundamentals: $G = (0, +\infty)$ and $U = [1, +\infty)$. If $\theta \in G$, the project may succeed. If $\theta \in U$, the project must succeed.

3.4.1 The synchronous case

We first *informally* review the well-known result for the static global game. Consider the simple static game obtained when the dynamic game described in Section 3.2 has only one round; that is, when $T = 1$. The following discussion is based on Morris and Shin (2003).

Let $B_p^i(E)$ denote the set of i 's types that assign probability at least p to the event E ; for types $B_p^i(E)$ we say that i p -believes E . Let $B_p(E)$ denote the profiles at which both players p -believe E .

To simplify the exposition, assume (for this subsection only) that investment of both players is necessary for the project to succeed whenever $\theta > 0$, so that there is no upper dominance region. Then the best response of each player is to invest if and only if she \underline{p} -believes both that the fundamental θ is in G , and that $-i$ invests. Therefore, investment is rationalizable only for types that \underline{p} -believe the following list:

- G ,
- that her opponent \underline{p} -believes G ,

- that her opponent p -believes that she herself \underline{p} -believes G ,
- etc.

Hence both players will invest only on the intersection

$$\bigcap_{k \geq 1} [B_{\underline{p}}]^k(G),$$

which is denoted by $C_{\underline{p}}(G)$ and called common \underline{p} -belief of G .

However, common \underline{p} -belief of G is difficult to achieve in static global games. Suppose $\underline{p} > \frac{1}{2}$. In that case, player i \underline{p} -believes G only if $x^i \geq x^{(1)} = 0 + \sigma F^{-1}(\underline{p})$. But for common \underline{p} -belief of G , player i must also believe that the opponent's signal exceeds $x^{(1)}$. This belief occurs only if $x^i \geq x^{(2)} = x^{(1)} + \sqrt{2}\sigma F^{-1}(\underline{p})$. Continuing to higher orders of beliefs, we get conditions $x^i \geq x^{(k)}$ where $x^{(k)} = x^{(k-1)} + \sqrt{2}\sigma F^{-1}(\underline{p})$ for all $k > 1$. Since $F^{-1}(\underline{p}) > 0$ for $\underline{p} > \frac{1}{2}$, the sequence $x^{(k)}$ diverges, and hence there is no state at which G is common \underline{p} -belief. Note that this argument holds for arbitrarily small σ .

If we take a snapshot of our dynamic game at any round t , the information structure is identical to that of the static global game with $\sigma = \sigma_t$. Hence the common \underline{p} -belief in the above static sense is *not* achieved in any of the rounds of the dynamic game. This explains the coordination failure described in Proposition 3.1 — the game studied there is essentially a static game with $\sigma = \sigma_T$, and the fact that σ_T decreases in T is irrelevant as long as $\sigma_T > 0$.

3.4.2 The asynchronous case

The discussion so far indicates that the difference between the asynchronous and synchronous games does not lie in the ability to generate standard common p -belief. In this respect private learning does not help. The difference must lie in the relevant concept of common beliefs which characterize the set of types for which investment is rationalizable. The less restrictive conditions under which the project succeeds in the asynchronous game lead to

less demanding belief operators and to a concept of common belief which is satisfied at a large set of states.

Definitions

In what follows, for convenience, we refer to the beliefs and actions of player i at date t as the beliefs and actions of *agent* (i, t) . Let Θ denote the set of possible fundamentals, and $X^{(i,t)}$ the set of types of agent (i, t) for $i \in \mathcal{I} = \{1, 2\}$ and $t \in \{1, \dots, T\}$.⁵

We now define relevant events:

- A Θ -event F_Θ is a subset of Θ . Such events describe the fundamental, θ .
- An (i, t) -event $F^{(i,t)}$ is a subset of $X^{(i,t)}$. Such events describe the type of agent (i, t) .
- An i -event $F^i = \times_{t \leq T} F^{(i,t)}$ is a list of (i, t) -events, with each member of the list describing the type of agent (i, t) .
- A *compound* event $F = F_\Theta \times (\times_{i \in \mathcal{I}} F^i)$ is a list containing (i, t) -events for each $i \in \{1, 2\}$ and $t \in \{1, \dots, T\}$, together with one Θ -event.

We will abuse notation by identifying each Θ -event F_Θ with the compound event $F_\Theta \times (\times_{i,t} X^{t,i})$, $\times_{i \in \mathcal{I}} F^i$ with the compound event $\Theta \times (\times_{i \in \mathcal{I}} F^i)$, and so on.

We say that an i -event F^i is *eventually true* (or *holds eventually*) if $\cup_{t \leq T} F^{(i,t)}$ is true, that is, if the true state lies in $\cup_{t \leq T} F^{(i,t)}$.

For each player i we define an operator $\alpha^{T,i}(\cdot)$ that maps each compound event $F = F_\Theta \times (\times_{(j,t)} F^{(j,t)})$ to

$$\alpha^{T,i}(F) \equiv F_\Theta \cap \left(\bigcap_{j \in \mathcal{I} \setminus \{i\}} \left(\bigcup_{t \leq T} F^{(j,t)} \right) \right). \quad (3.3)$$

The operator $\alpha^{T,i}(\cdot)$ has a useful interpretation. Suppose that the project succeeds only if the fundamental lies in F_Θ and all players invest by round T . Suppose that each agent

⁵In our simple setup $\Theta = X^{(i,t)} = \mathbb{R}$ for all (i, t) .

(j, t) invests only on the event $F^{(j,t)}$. Then $\alpha^{T,i}(F)$ is the event that the project succeeds by round T , conditional on player i investing. Hence $\alpha^{T,i}(F)$ is the payoff-relevant event for player i .

Next we define relevant belief operators:

- The belief operator $A_p^{T,(i,t)}(\cdot)$ of agent (i, t) maps each compound event F to the set of types of (i, t) that assign probability at least p to $\alpha^{T,i}(F)$; that is,

$$A_p^{T,(i,t)}(F) = B_p^{(i,t)}(\alpha^{T,i}(F)).$$

We refer to $A_p^{T,(i,t)}(F)$ by saying that agent (i, t) asynchronously p -believes F . Note that $A_p^{T,(i,t)}(F)$ is an (i, t) -event.

- The belief operator $A_p^{T,i}(\cdot)$ of player i maps each compound event to a list of (i, t) -events, with each members describing, for some t , the types of (i, t) that asynchronously p -believe F ; that is,

$$A_p^{T,i}(F) \equiv \times_{t \leq T} A_p^{T,(i,t)}(F).$$

- The belief operator $A_p^T(\cdot)$ maps each compound event to a list of (i, t) -events, with each member describing, for some (i, t) , the types of (i, t) that asynchronously p -believe F ; that is,

$$A_p^T(F) \equiv \times_{i \in \mathcal{I}} A_p^{T,i}(F).$$

- The asynchronous common belief operator $D_p^T(\cdot)$ is defined (on compound events F) by $D_p^T(F) \equiv \bigcap_{k=1}^{\infty} [A_p^T]^k(F)$.

The interpretation of asynchronous common belief $D_p^T(G)$ of a compound event G resembles the interpretation of the usual static concept of common belief. The event $D_p^T(G)$ is a list of events, with each member describing for some agent (i, t) the types of (i, t) that p -believe the following list:

- G ,
- that her opponent eventually p -believes G ,
- that her opponent eventually p -believes that her she herself eventually p -believe G ,
- etc.

The interpretation of asynchronous common belief differs from the interpretation of standard common belief on page 101 only in the insertion of the qualifier “eventually”. We now proceed to utilize this concept of asynchronous common belief to delineate the set of types for which investment is rationalizable in the asynchronous game.

Rationalizability

In the first step, we formulate a sufficient condition for rationalizability of investment in the *asynchronous* game in terms of rationalizability in a related *simultaneous move* game. This will allow us to use the results on higher order beliefs in simultaneous move games from Morris and Shin (2007). We refer to this associated game as the *characteristic game*, and define it as follows:

Definition 3.1. *The **characteristic game** $\tilde{\Gamma}^T$ is a simultaneous move game with $2T$ players denoted by (i, t) for $i \in \{1, 2\}$ and $t \in \{1, \dots, T\}$. The information structure is as in the asynchronous game: the fundamental θ is drawn according to an improper uniform distribution on \mathbb{R} , and each player (i, t) observes a signal $x^{(i,t)} \sim N(\theta, \sigma_t^2)$. Each player chooses an action from $\{0, 1\}$, which we interpret as Not-Invest and Invest respectively. We say that the project succeeds either if $\theta \geq 1$, or if $\theta > 0$ and for each $i \in \{1, 2\}$, at least one of the players $\{(i, 1), \dots, (i, T)\}$ invests. The payoff for player (i, t) for not investing is 0, and for investing is \tilde{b} if the project succeeds, and $-\tilde{c}$ if the project does not succeed. The parameters \tilde{b} and \tilde{c} satisfy*

$$\underline{p}(\tilde{c}, \tilde{b}) = \bar{p}(c, b).^6$$

⁶Note that we require the probability of success that was *sufficient* for immediate investment in the

Note that players in the characteristic game are analogues of agents in the asynchronous game. However, we continue to refer to player (i, t) in the characteristic game as agent (i, t) , and the collection $\{(i, t)\}_t$ as player i , as would be appropriate in the asynchronous game.

The asynchronous and characteristic games have the same number of agents. In the characteristic game, investment is a best response for agent (i, t) *if and only if* she $\bar{p}(c, b)$ -believes that project succeeds. In the asynchronous game, investment is a best response for agent (i, t) *if* she $\bar{p}(c, b)$ -believes that project succeeds. Hence, rationalizability of investment in the characteristic game is a sufficient condition for rationalizability of investment in the dynamic game. We proceed to characterize the rationalizability of investment in the characteristic game.

One technical complication we face is that our game is of *common* and not *private* values. Players are sure of their own payoff parameters in private value games, and hence they suffer only from strategic uncertainty; this makes common beliefs directly applicable.⁷ In common value games, players suffer also from uncertainty over the fundamental, which requires a slight modification in the relevant belief operators. We introduce these modified operators below and use them to characterize the set of types for which investment is rationalizable. The introduction of the modified operators is only a technical step; later on, we identify sufficient conditions for rationalizability of investment in terms of the unmodified operators defined in Section 3.4.2 above, and thereafter the modified ones will not be needed.

Define the following operators:

- $R_p^{T,(i,t)}(F) \equiv A_p^{T,(i,t)}((F \cap G) \cup U)$.
- $R_p^{T,i}(F) \equiv \times_{t \leq T} R_p^{T,(i,t)}(F)$.

asynchronous game to be *necessary* in the characteristic game. It must then be the case that payoffs in the characteristic game are less favorable to investing than in the original asynchronous game. It is easy to check that $\tilde{b} = b - \delta b$ and $\tilde{c} = c + \delta b$ works. Finally, note that since the characteristic game is static, $\underline{p}(\tilde{c}, \tilde{b}) = \bar{p}(c, b)$ is necessary and sufficient for investment.

⁷This is the reason why most of the higher order beliefs literature deals with private value games. In our case the common value setup is dictated by the motive of private learning in our model.

- $R_p^T(F) \equiv \times_{i \in \mathcal{I}} R_p^{T,i}(F)$.
- $Q_p^T(F) \equiv \bigcap_{k=1}^{\infty} [R_p^T]^k(F)$.

The motivation for the operator $R_p^{T,(i,t)}(\cdot)$ is as follows: suppose agents (i, t) invest at types $F^{(i,t)}$, and consider the compound event $F = \Theta \times (\times_{i,t} F^{(i,t)})$. Then $A_p^{T,(i,t)}(F \cap G) \cup U$ is the event that (i, t) asynchronously p -believes that the project succeeds since success occurs when the fundamental is good and all players eventually invest ($F \cap G$), or when the fundamental is in the upper dominance region (U).

Proposition 3.3. (*Morris and Shin 2007*) *Investment is rationalizable in the characteristic game at type $x^{(i,t)}$ if and only if $x^{(i,t)}$ is an element of*

$$R_{\bar{p}}^{T,(i,t)}(Q_{\bar{p}}^T(G)).$$

Proof. See Morris and Shin (2007). ■

To obtain some intuition for Proposition 3.3, consider iterated deletion of actions which are never best responses. After the first round of deletion, investment survives for types $A_{\bar{p}}^{T,(i,t)}(G) = R_{\bar{p}}^{T,(i,t)}(G)$. After the second round, investment survives for types

$$A_{\bar{p}}^{T,(i,t)}((R_{\bar{p}}^T(G) \cap G) \cup U) = R_{\bar{p}}^{T,(i,t)}(R_{\bar{p}}^T(G)).$$

After the third round, investment survives for types

$$A_{\bar{p}}^{T,(i,t)}\left(\left([R_{\bar{p}}^T]^2(G) \cap G\right) \cup U\right) = R_{\bar{p}}^{T,(i,t)}\left([R_{\bar{p}}^T]^2(G)\right),$$

and so on.

The following lemma specifies sufficient conditions for the events $R_p^T(F)$ and $Q_p^T(F)$ to occur in terms of the events $A_q^T(F)$ and $D_q^T(F)$ for sufficiently high q .

Lemma 3.5. Suppose $q \geq \frac{p+1}{2}$ and $A_q^{T,(i,t)}(F) \subseteq A_q^{T,(i,t)}(G)$. Then

$$(i) \quad A_q^{T,(i,t)}(F) \subseteq R_p^{T,(i,t)}(F),$$

and $(ii) \quad D_q^T(F) \subseteq Q_p^T(F).$

Proof. Since $A_q^{T,(i,t)}(F) \subseteq A_q^{T,(i,t)}(G)$, we have

$$A_q^{T,(i,t)}(F) \subseteq A_q^{T,(i,t)}(F) \cap A_q^{T,(i,t)}(G).$$

The right-hand side of this last expression is contained in $A_{2q-1}^{T,(i,t)}(F \cap G)$, which is in turn contained in $R_{2q-1}^{T,(i,t)}(F)$. Since $q \geq \frac{p+1}{2}$, we have $2q - 1 \geq p$, and hence $R_{2q-1}^{T,(i,t)}(F) \subseteq R_p^{T,(i,t)}(F)$. This proves part (i).

We now use part (i) to prove part (ii) of the lemma. The event $D_q^T(F)$ is contained in $A_q^{T,(i,t)}(F)$. By part (i), $D_q^T(F)$ is contained in $R_p^{T,(i,t)}(F)$. Since this containment holds for all (i, t) , the event $D_q^T(F)$ is contained in $R_p^T(F)$. Furthermore, $D_q^T(F)$ is contained in $A_q^{T,(i,t)}(D_q^T(F))$, and hence also in $A_q^{T,(i,t)}(R_p^T(F))$. Applying part (i) again gives containment in $R_p^T(R_p^T(F)) = [R_p^T]^2(F)$. Continuing in this fashion, we obtain containment in $[R_p^T]^k(F)$ for any order k . ■

We are now ready to state sufficient conditions for rationalizability of investment in terms of the operators $A_q^{T,(i,t)}(\cdot)$ and $D_q^T(\cdot)$.

Proposition 3.4. *Investment is rationalizable in the characteristic game for types of agent (i, t) in*

$$A_q^{T,(i,t)}(D_q^T(G)) \tag{3.4}$$

for $q \geq \frac{p+1}{2}$.

Proof. A sufficient condition for the event $R_p^{T,(i,t)}(Q_p^T(G))$ to occur is for $A_q^{T,(i,t)}(D_q^T(G))$ to occur with $q \geq \frac{p+1}{2}$. To see this, note first that $A_q^{T,(i,t)}(D_q^T(G)) \subseteq A_q^{T,(i,t)}(G)$, so the

conditions of Lemma 3.5 are satisfied with $F = D_q^T(G)$. Hence we have

$$A_q^{T,(i,t)}(D_q^T(G)) \subseteq R_p^{T,(i,t)}(D_q^T(G)) \subseteq R_p^{T,(i,t)}(Q_p^T(G)),$$

where the second containment follows from part (ii) of Lemma 3.5 with $F = G$. By Proposition 3.3, investment is rationalizable for types of agent (i, t) in $R_p^{T,(i,t)}(Q_p^T(G))$.

■

Characterization of asynchronous beliefs

Section 3.4.2 established sufficient conditions for rationalizability of investment in terms of asynchronous common beliefs. In this section, we show that asynchronous common belief is easily attained in sufficiently long games.

Following Monderer and Samet (1989), we say that E is an *asynchronous p -evident event* (for T rounds) if $E \subseteq A_p^T(E)$. The following proposition restates a result due to Monderer and Samet (1989), but in the asynchronous setting.

Proposition 3.5. *A state ω lies in $D_p^T(F)$ if and only if there exists an asynchronous p -evident event E containing ω such that $E \subseteq A_p^T(F)$.*

Proof. See Monderer and Samet (1989), Proposition 3. ■

We use the characterization of asynchronous common beliefs from Proposition 3.5 to prove the next result.

Proposition 3.6. *For all $r > q$, there exists some \underline{T} such that for all $T \geq \underline{T}$,*

$$A_r^{T,(i,t)}(G) \subseteq A_q^{T,(i,t)}(D_q^T(G)).$$

Proof. Let $\alpha^T(F) \equiv \bigcap_i \alpha^{T,i}(F)$. Recalling the definition of $\alpha^{T,i}(\cdot)$ from page 103, $\alpha^T(F)$ may be interpreted as the event that $F^{(i,t)}$ is eventually true for each player i and F_Θ holds.⁸

⁸We apply the operator $\alpha^T(\cdot)$ only to compound events for which $F_\Theta = \Theta$, and hence F_Θ holds trivially.

For $\theta^* = 0$, Lemma 3.3 states that for all $r < 1$ and all $\theta \in G$,

$$\lim_{T \rightarrow \infty} \Pr(\alpha^T(A_r^T(G)) | \theta) = 1.$$

Denoting $\Pr(\alpha(A_r^T(G)) | \theta = 0)$ by s_T , we have

$$A_r^{T,(i,t)}(G) \subseteq A_{r \cdot s_T}^{T,(i,t)}(A_r^T(G))$$

because any type of agent (i, t) that assigns probability r to G assigns probability at least $r \cdot s_T$ to the event that G holds and the opponent eventually r -believes G .

Since $r > q$ and $s_T \rightarrow 1$ as $T \rightarrow \infty$, the product $r \cdot s_T$ exceeds q for sufficiently large T . Hence we have

$$A_r^{T,(i,t)}(G) \subseteq A_q^{T,(i,t)}(A_r^T(G)), \quad (3.5)$$

and, since this holds for all agents (i, t) , $A_r^T(G)$ is an asynchronous q -evident event.

By Proposition 3.5, the event $A_r^T(G)$ must be contained in $D_q^T(G)$. Combining this with (3.5) gives

$$A_r^{T,(i,t)}(G) \subseteq A_q^{T,(i,t)}(A_r^T(G)) \subseteq A_q^{T,(i,t)}(D_q^T(G)),$$

as needed. ■

Proposition 3.6 indicates that the sufficient conditions for rationalizability of investment given above are not demanding when T is large. All that is needed is *first-order* r -belief of G with $r > \frac{\bar{v}+1}{2}$, which is achieved for signals exceeding $F^{-1}(r)\sigma_t$. Investment is therefore rationalizable for all positive signals except in a small neighborhood of 0 of size on the order of σ_t .

3.5 Conclusion

Static coordination games represent a useful abstraction for studying coordination problems in the real world. However, the associated requirement of synchronicity in participation

may be a strong restriction: the outcomes generated in such models may not be good representations for real-world coordination problems where agents are able to participate at different points of time and can learn about payoffs while deciding when to participate. We illustrate the radical difference between synchronous and asynchronous coordination problems within the framework of global games. In canonical synchronous (one-shot) global games, the risk-dominant equilibrium of the underlying complete information game is selected. Thus, coordination failure is endemic in static global games: there exist a wide class of payoffs for which players fail to efficiently coordinate in the unique equilibrium of the canonical global game despite the fact that it is in their collective interest to do so. At the other extreme, we introduce a class of enriched asynchronous global games where agents have an infinity of opportunities to participate, while they asymptotically and privately learn the true payoffs. In such games, we show that equilibrium play ensures Pareto dominant outcomes. Coordination failure is eliminated.

Irreversibility plays an important role in our analysis, and, more generally, in the analysis of dynamic coordination games. The tendency towards efficiency in our model is related to the fact that we chose the efficient rather than the inefficient action to be irreversible. This assumption is natural in the context of many applications, including the leading example of foreign direct investment which we used to motivate our stylized model. However, in other applications, alternative assumptions may be more appropriate. Had we chosen differently, that is, had we assumed that the project succeeds only if all players choose to invest in *all* rounds, the project would always fail except in the upper dominance region. The coordination outcome in dynamic coordination games is, therefore, sensitive to the details of the dynamic setup. A deeper understanding of dynamic coordination problems may pinpoint detailed changes in the design of coordination processes that could help to avoid coordination failures. Our results provide a benchmark for such design exercises.

While it is useful as a benchmark exercise to study the extreme cases in which players learn nothing or everything during the play of the game, or when investment is fully reversible vs irreversible, from an applied perspective it is of greatest interest to learn

about intermediate cases, i.e., about finite-rounds asynchronous global games with private learning during which players learn something but not everything. These intermediate cases remain interesting problems for future research.

Bibliography

- Angeletos, G.-M., Hellwig, C. and Pavan, A.: 2007, Dynamic global games of regime change: learning, multiplicity and timing of attacks, *Econometrica* **75**, 711–756.
- Carlsson, H. and van Damme, E.: 1993, Global games and equilibrium selection, *Econometrica* **61**, 989–1018.
- Chamley, C.: 2003, Dynamic speculative attacks, *American Economic Review* **93**, 603–621.
- Cripps, M., Ely, J., Mailath, G. and Samuelson, L.: 2006, Common learning. Mimeo.
- Dasgupta, A.: 2007, Coordination and delay in global games, *Journal of Economic Theory* **134**, 195–225.
- Ely, J.: 2003, Common learning. Mimeo.
- Gale, D.: 1995, Dynamic coordination games, *Economic Theory* **5**, 1–18.
- Heidhues, P. and Melissas, N.: 2006, Equilibria in a dynamic global game: the role of cohort effects, *Economic Theory* **28**, 531–557.
- Hörner, J.: 2004, Payoff-dominance in dynamic coordination games. Mimeo.
- Kajii, A. and Morris, S.: 1997, The robustness of equilibria to incomplete information, *Econometrica* **65**, 1283–1312.
- Monderer, D. and Samet, D.: 1989, Approximating common knowledge with common belief, *Games and Economic Behavior* **1**, 170–190.

Morris, S., Rob, R. and Shin, H.: 1995, P-dominance and belief potential, *Econometrica* **63**, 145–167.

Morris, S. and Shin, H.: 2007, Common belief foundations of global games. Mimeo.

Morris, S. and Shin, H. S.: 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 University Press, Cambridge, pp. 56–114.