



Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Games and Economic Behavior 48 (2004) 124–138

GAMES and
Economic
Behavior

www.elsevier.com/locate/geb

Adaptive play by idiosyncratic agents

David P. Myatt* and Chris C. Wallace

Department of Economics, Oxford University, Manor Road, Oxford OX1 3UQ, UK

Received 3 December 2002

Available online 6 November 2003

Abstract

Equilibrium selection in coordination games has generated a large literature. Kandori, Mailath and Rob [Econometrica 61 (1993) 29] and Young [Econometrica 61 (1993) 57] studied dynamic models of aggregate behaviour where agents best-respond to observations of population play. Crucially, infrequent mistakes (“mutations”) allow agents to take actions contrary to current trends and prevent initial configurations from determining long-run play. An alternative approach is offered here: Trembles are added to payoffs so that with some probability it is optimal to act against the flow of play. The long-run distribution of population behaviour is characterised—modes correspond to stable Bayesian Nash equilibria. Allowing the variance of payoff trembles to vanish (a purification process) a single equilibrium is played almost always in the long run. Kandori, Mailath, and Rob, and Young, show that the number of contrary actions required to escape an equilibrium determines selection; here, the likelihood that such actions are taken is equally important.

© 2003 Elsevier Inc. All rights reserved.

JEL classification: C72; C73

Keywords: Adaptive play; Payoff idiosyncrasy; Evolution; Equilibrium selection

1. Introduction

Many games have multiple Nash equilibria. Optimising agents might play a Nash equilibrium, but which one? Motivated in part by this question, researchers have modelled the adaptive evolution of play. Kandori, Mailath and Rob (KMR, 1993) and Young (1993) analysed the periodic strategy revision of boundedly rational players. Revising players

* Corresponding author.

E-mail addresses: david.myatt@economics.ox.ac.uk (D.P. Myatt),
christopher.wallace@economics.ox.ac.uk (C.C. Wallace).

observe the distribution of either current or historical strategy choices, and choose a myopic best response. This generates a Markov process that is path dependent: Play will become locked in to any pure strategy Nash equilibrium. To remove this feature, revising players are assumed to err with some probability—referred to as a “mutation.” Such mutations allow the process to move between equilibria, and permit analysis of the long-run distribution of play, independent from any initial conditions. As the probability of a mutation is allowed to vanish, this distribution places almost all weight on play corresponding to a single Nash equilibrium. In the class of 2×2 symmetric coordination games, such a selection process picks the Harsanyi–Selten (1988) risk-dominant equilibrium. This is because a relatively large number of mutations are required to escape from a risk-dominant equilibrium.

In the KMR–Young framework the probability of a mutation is independent of the state of play. This plays a crucial rôle in the selection process. Bergin and Lipman (1996) demonstrated that general state dependent mutations can result in the selection of any equilibrium. Since fully general results are unavailable, it is suggested that mutation specifications should derive from reasonable underlying justifications. This paper responds to this idea. The simple example of 2×2 symmetric coordination games is considered. Payoffs are subject to trembles, yielding a Bayesian game. Errant behaviour, therefore, is caused by the idiosyncrasy of preferences rather than mistaken choices. It follows that long-run play may be analysed while retaining the use of pure best responses. Furthermore, the long-run distribution may be partially characterised without the need for vanishing mutation probabilities: In fact, the modes of this distribution coincide with stable Bayesian Nash equilibria.

Equilibrium selection with “vanishing mutations” is still possible, via a purification process. Allowing the variances of payoff trembles to tend to zero at the same rate, the associated mutation probabilities vanish to zero at different rates. In the sense of Bergin and Lipman (1996), mutations are (endogenously) state dependent. In this environment, equilibrium selection depends upon the likelihood of observing a contrary action, as well as the number of contrary actions needed to escape from a particular Nash equilibrium. Moreover, the selection criterion applies when the population consists of only two players, and merely a single mutation is required to escape from an equilibrium.

The approach matches that of Blume (1999), where mutation probabilities depend upon the difference in expected payoffs. He describes a sufficient condition of “skew symmetry” for the selection of a risk-dominant equilibrium. In the context of a Bayesian game, this condition holds when payoff tremble variances are independent of the strategy profile. The analysis presented here moves further, by identifying precisely how skew-symmetric the noise process needs to be.

Others have studied alternative models of state dependent mutations: van Damme and Weibull (1998) analyse a model where mutations are endogenously derived from agents making costly attempts to control their “trembling hands.” Foster and Young (1990) have perhaps the earliest discussion of state-dependent mutations (although with no reference to risk-dominance and equilibrium selection). Binmore and Samuelson (1997) also present a model with state dependent mutations. They proceed directly to the specification of the elements of a tridiagonal Markov matrix in a single-revision dynamic and hence their approach lacks an explicit economic model, but is nonetheless related.

2. Adaptive play of a Bayesian game

The model is based upon a symmetric 2×2 coordination game with payoffs:

$$\begin{array}{cc} & \begin{array}{cc} 1 & 2 \end{array} \\ \begin{array}{c} 1 \\ 2 \end{array} & \begin{array}{|cc|} \hline a & c \\ \hline a & b \\ \hline c & d \\ \hline \end{array} \quad \text{or equivalently} \quad \begin{array}{cc} & \begin{array}{cc} 1 & 2 \end{array} \\ \begin{array}{c} 1 \\ 2 \end{array} & \begin{array}{|cc|} \hline a-c & 0 \\ \hline 0 & d-b \\ \hline 0 & d-b \\ \hline \end{array} \end{array}, \quad (1)$$

where $a > c$ and $d > b$ ensure that the game has two pure strategy Nash equilibria (1, 1) and (2, 2). For simplicity, assume that payoffs are generic in the following sense: $dn/(a+d) \notin \mathbb{Z}$ for $n \in \mathbb{Z}$. Without loss of generality, it is assumed that $a - c > d - b$, ensuring that the equilibrium (1, 1) is risk-dominant (Harsanyi and Selten, 1988). Players care only about the difference in expected payoffs when making a choice, and hence the coordination game is strategically equivalent to the pure coordination game on the right hand side of Eq. (1). It is further without loss of generality to set $b = c = 0$ throughout. In this formulation, the mixed strategy Nash equilibrium entails mixing probabilities of $[x^*, 1 - x^*]$ where $x^* = d/(a+d) < 1/2$ since $a > d$ by assumption.

The payoffs a and d represent mean utilities. To generate a Bayesian game, each player has idiosyncratic payoffs \tilde{a} and \tilde{d} , generated by the addition of normally and independently distributed payoff trembles:

$$\begin{aligned} \tilde{a} &= a + \sigma \varepsilon_a \\ \tilde{d} &= d + \sigma \varepsilon_d \end{aligned} \quad \text{where} \quad \begin{bmatrix} \varepsilon_a \\ \varepsilon_d \end{bmatrix} \sim N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \xi_a^2 & \rho \xi_a \xi_d \\ \rho \xi_a \xi_d & \xi_d^2 \end{bmatrix} \right).$$

The parameters ξ_a and ξ_d allow the variance of the trembles to be strategy profile specific.¹ Blume (1999, Section 6) offers a similar “random utility” approach. In his model any payoff noise is added directly to the expected payoff difference of the two pure strategies, rather than to the payoffs from particular strategy profiles. Such a specification is equivalent to setting $\rho = -1$ and $\xi_a = \xi_d$. The same observation can also be made of related papers by Brock and Durlauf (2001) and Blume and Durlauf (2001). The parameter σ is a common scaling factor which is allowed to vanish for the limiting results of Section 4. The normal distribution proves convenient for the subsequent analysis.² Its crucial property, however, is the unboundedness of the support, allowing either strategy to be dominant with some probability.

Play evolves adaptively among a finite population of n players. In a single period, each individual plays a fixed strategy against a randomly selected opponent from the remaining $n - 1$ players. The state of play is the number of players using strategy 1, denoted $z \in Z = \{0, \dots, n\}$. At the end of each period, a randomly selected player is replaced. The new entrant is equipped with newly trembled payoffs \tilde{a} and \tilde{d} . This player observes the

¹ A degree of freedom may be eliminated without loss of generality by setting either $\xi_a = 1$ or $\xi_d = 1$.

² For the limiting results as $\sigma \rightarrow 0$, the key features are the asymptotic properties of the densities and hazard rates of the disturbances. Thus any other distribution sharing these features should lead to similar results. It is also worth noting that full generality of trembles, particularly allowing trembles to vary by state, as Bergin and Lipman (1996) have shown, would lead to inconclusive results.

strategy distribution among $n - 1$ remaining incumbents and selects a best response. This is a “one step at a time” dynamic (a birth-death process), similar to the one proposed by Binmore and Samuelson (1997), but unlike the process of KMR (1993) in which all players revise together. Although the dynamic is set in discrete time, it would be equivalent to cast the model in continuous time, and allow individual players to be woken by independent Poisson alarm clocks and subsequently replaced.

An entrant’s choice depends only upon the current strategy distribution, and hence the adaptive play dynamic describes a homogeneous Markov chain on the state space Z . Beginning in state z , and following the exit of an incumbent, there will be either $i = z$ or $i = z - 1$ of the remaining incumbents using strategy 1. Use $x = i/(n - 1)$ to denote the fraction using strategy 1. Against this strategy distribution, the payoffs to pure strategies 1 and 2 are $x(a + \sigma \varepsilon_a)$ and $(1 - x)(d + \sigma \varepsilon_d)$ respectively. An entrant chooses strategy 1 whenever the first term is larger than the second. Rearranging, this occurs whenever: $(1 - x)\varepsilon_d - x\varepsilon_a < [xa - (1 - x)d]/\sigma$. The left hand side is normal with zero mean and variance $x^2\xi_a^2 + (1 - x)^2\xi_d^2 - 2x(1 - x)\rho\xi_a\xi_d$. It follows that the entrant will choose pure strategy 1 with probability:

$$\Pr[1 | x] = \Phi\left(\frac{xa - (1 - x)d}{\sigma\sqrt{x^2\xi_a^2 + (1 - x)^2\xi_d^2 - 2x(1 - x)\rho\xi_a\xi_d}}\right)$$

where Φ represents the cumulative distribution function of the standard normal. By inspection, it is clear that an entrant is more likely to choose strategy 1 when x (and hence z) is sufficiently high. This fact is reflected in Definition 1, where the notation $\lceil u \rceil$ indicates the smallest integer above u , and $\lfloor u \rfloor$ indicates the largest integer below u .³

Definition 1. The *basins of attraction* for strategies 1 and 2 are:

$$Z_1 = \{\lceil (n - 1)x^* + 1 \rceil, \dots, n\} \quad \text{and} \quad Z_2 = \{0, \dots, \lfloor (n - 1)x^* \rfloor\}.$$

The *basin depth* faced by an entrant is $\kappa(x)^2$ or equivalently κ_i^2 , where:

$$\kappa(x) = \frac{xa - (1 - x)d}{\sqrt{x^2\xi_a^2 + (1 - x)^2\xi_d^2 - 2x(1 - x)\rho\xi_a\xi_d}} \quad \text{and} \quad \kappa_i = \kappa\left(\frac{i}{n - 1}\right)$$

so that an entrant will choose pure strategy 1 with probability $\Phi(\kappa(x)/\sigma)$.

The “flow of play” is toward a strategy from states within its basin of attraction. Starting from $z \in Z_1$, an entrant will observe at least $i = z - 1$ incumbents using strategy 1. Following Definition 1, $i = z - 1 \geq \lceil (n - 1)x^* \rceil > (n - 1)x^*$,⁴ and hence $x = i/(n - 1) > x^*$: The expected payoff from strategy 1 is higher than that from strategy 2. Equivalently, $\kappa(x) = \kappa_i > 0$ and hence $\Pr[1 | x] > 1/2$. Similarly, from a state $z \in Z_2$, any entrant is more likely to choose strategy 2. Finally, state $z = \lceil (n - 1)x^* \rceil$ belongs to neither basin of attraction. The most likely entrant choice depends upon the identity of the exiting player.

Whereas the basins of attraction reflect the flow of play, the basin depth indexes the difficulty of moving against that flow. Consider once again a state $z \in Z_1$. An entrant is

³ More formally $\lceil u \rceil = \min\{k \in \mathbb{Z}: k \geq u\}$ and $\lfloor u \rfloor = \max\{k \in \mathbb{Z}: k \leq u\}$; $u \notin \mathbb{Z} \Rightarrow \lceil u \rceil > \lfloor u \rfloor$.

⁴ By assumption $(n - 1)x^*$ does not take an integer value.

most likely to play strategy 2 when a strategy 1 player exits. In this case, the probability of a contrary action is simply $1 - \Phi(\kappa_i/\sigma)$ where $i = z - 1$. For large κ_i^2 , it is highly unlikely the choice will be made against the flow of play.

The transition probabilities of adaptive play may now be calculated. Write $p_{zz'} = \Pr[z_{t+1} = z' | z_t = z]$ for the probability of a transition from state z to state z' . A single player is replaced each period, and hence $p_{zz'} = 0$ for $|z - z'| > 1$. For states $z < n$, the probability of a step up is

$$p_{z,z+1} = \frac{n-z}{n} \times \Phi\left(\frac{\kappa_z}{\sigma}\right).$$

A step up requires the exit of strategy 2 incumbent, which occurs with probability $(n-z)/n$. The entrant must choose strategy 1, which occurs with probability $\Phi(\kappa_z/\sigma)$. Similar procedures lead to expressions for other transitions:

$$p_{z,z-1} = \frac{z}{n} \times \left[1 - \Phi\left(\frac{\kappa_{z-1}}{\sigma}\right)\right] \quad \text{and}$$

$$p_{z,z} = \frac{n-z}{n} \times \left[1 - \Phi\left(\frac{\kappa_z}{\sigma}\right)\right] + \frac{z}{n} \Phi\left(\frac{\kappa_{z-1}}{\sigma}\right).$$

3. Long-run behaviour

Adaptive play exhibits path dependence in the short run. An entrant to state $z = n$, for instance, is most likely to choose strategy 1, and hence the process remains in the basin of attraction Z_1 . In the long run, however, contrary actions against the flow of play allow the process to escape from a basin of attraction. Formally, the Markov chain is *ergodic*, and its long-run behaviour is independent of any initial conditions. To see this, first observe that the process is *irreducible*: There is positive probability of moving between any two states in a finite number of steps, and hence permanent “lock in” to any state cannot occur. Second, all states are *aperiodic*: There is positive probability of remaining in a state, and this prevents the occurrence of established cycles. These two features are sufficient for ergodicity (Grimmett and Stirzaker, 2001).

By the Ergodic Theorem, any finite ergodic Markov chain has a unique stationary distribution. This is a probability vector $[\pi_z]_{z \in Z}$ satisfying $\pi_z = \lim_{t \rightarrow \infty} \Pr[z_t = z]$, independent of any initial conditions. No matter where the process starts, in the long run the process will be in state z with probability π_z . The stationary distribution is the unique solution to the detailed balance equations $\pi_z = \sum_{z' \in Z} \pi_{z'} p_{z'z}$, $\forall z \in Z$. The one-step-a-time nature of adaptive play ensures that these equations take a particularly simple form:

$$\pi_0 = \pi_0 p_{00} + \pi_1 p_{10}, \quad \pi_n = \pi_n p_{nn} + \pi_{n-1} p_{(n-1)n} \quad \text{and}$$

$$\pi_z = \pi_z p_{zz} + \pi_{z-1} p_{(z-1)z} + \pi_{z+1} p_{(z+1)z}.$$

Solving these equations yields the unique solution:

$$\pi_z = \frac{q_z}{\sum_{j \in Z} q_j} \quad \text{where} \quad q_z = \prod_{0 \leq j < z} p_{j(j+1)} \prod_{z < j \leq n} p_{j(j-1)}$$

$$\Rightarrow \frac{\pi_z}{\pi_{z+1}} = \frac{p_{(z+1)z}}{p_{z(z+1)}}. \quad (2)$$

This familiar form is common to birth-death processes. The final element of Eq. (2) has an easy interpretation: The long-run relative likelihood of neighbouring states is the relative probability of jumping backward and forward between them. Employing the explicit transition probabilities for the adaptive play dynamic yields the following.

Lemma 1. *Adaptive play has a unique stationary (ergodic) distribution π satisfying:*

$$\pi_z = \frac{q_z}{\sum_{j \in Z} q_j} \quad \text{where } q_z = \prod_{0 \leq j < z} \frac{n-j}{n} \Phi\left(\frac{\kappa_j}{\sigma}\right) \prod_{z < j \leq n} \frac{j}{n} \left(1 - \Phi\left(\frac{\kappa_{j-1}}{\sigma}\right)\right). \quad (3)$$

Proof. Substitute the transition probabilities into Eq. (2). \square

The distribution characterised by Lemma 1 is related to the Bayesian Nash equilibria of the underlying stage game. Suppose that a fraction $x < \Phi(\kappa(x)/\sigma)$ of incumbent players are using strategy 1. A strategy 1 player is less likely to exit (with probability x) than to enter (with approximate probability $\Phi(\kappa(x)/\sigma)$). Similarly, a strategy 2 player is more likely to exit than to enter. In expectation, the number of strategy 1 players is growing. Identical logic suggests that the number of strategy 1 players is likely to decrease when $x > \Phi(\kappa(x)/\sigma)$. Intuitively, the process is moving toward stable fixed points of $\Phi(\kappa(x)/\sigma)$. Such a fixed point $\tilde{x} = \Phi(\kappa(\tilde{x})/\sigma)$ is, of course, a Bayesian Nash equilibrium of the stage game. Lemma 2 describes such equilibria. An equilibrium (and hence fixed point) \tilde{x} will be denoted *stable* if it is a downcrossing. Formally, for small ϵ , $\Phi(\kappa(\tilde{x} - \epsilon)/\sigma) > \tilde{x} - \epsilon$ and $\Phi(\kappa(\tilde{x} + \epsilon)/\sigma) < \tilde{x} + \epsilon$. This means that a sequence of best responses will lead from any point in a neighbourhood of \tilde{x} toward \tilde{x} . An equilibrium will be denoted *unstable* if the opposite is true.

Lemma 2. *For σ sufficiently small, there are three Bayesian Nash equilibria: $\tilde{x}_L < \tilde{x}_M < \tilde{x}_H$. The central BNE \tilde{x}_M is unstable, the remaining equilibria are stable. Moreover, $\tilde{x}_L \rightarrow 0$, $\tilde{x}_M \rightarrow x^*$ and $\tilde{x}_H \rightarrow 1$ as $\sigma \rightarrow 0$. For σ sufficiently large, there is a unique and stable BNE \tilde{x} . For $a > d \Leftrightarrow x^* < 1/2$ this stable BNE satisfies $\tilde{x} > 1/2$.*

Proof. See Appendix A. \square

The correspondence between stable Bayesian Nash equilibria and the long-run behaviour of the adaptive play dynamic is stated formally in Proposition 1, and illustrated graphically in Fig. 1. Notice that for σ large enough ($\sigma = 2.25$) there is a single Bayesian Nash equilibrium above the halfway point ($\tilde{x} > 1/2$), and the corresponding ergodic distribution is unimodal. On the other hand, when σ is sufficient small (for instance, $\sigma = 1.25$), there are three Bayesian Nash equilibria. Two of these equilibria are stable (downcrossings), and are associated with the two maxima of the bimodal ergodic distribution.

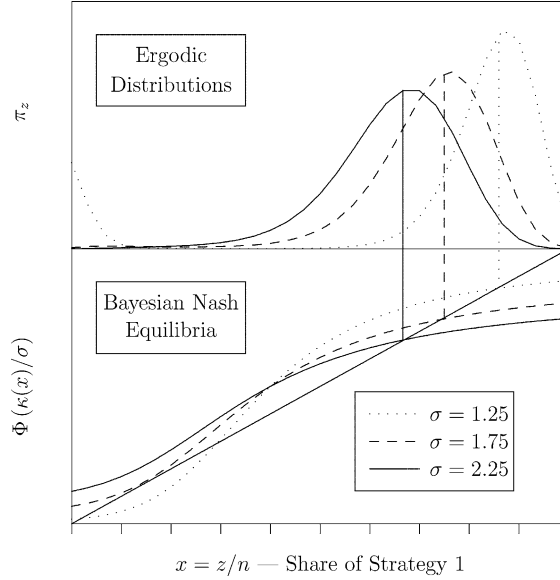


Fig. 1. Parameters are $a = 3$, $d = 2$, $\xi_a = \xi_d = 1$, $\rho = 0$ and $n = 30$.

Proposition 1. *The local maxima (modes) of the ergodic distribution coincide with stable Bayesian Nash equilibria of the stage game. The local minima coincide with unstable Bayesian Nash equilibria. Formally, for sufficiently large n :*

$$x < \Phi\left(\frac{\kappa(x)}{\sigma}\right) \Rightarrow \pi_{\lfloor xn \rfloor} < \pi_{\lceil xn \rceil} \quad \text{and} \quad x > \Phi\left(\frac{\kappa(x)}{\sigma}\right) \Rightarrow \pi_{\lfloor xn \rfloor} > \pi_{\lceil xn \rceil}.$$

As n grows large all weight focuses on a single stable Bayesian Nash equilibrium \tilde{x} :

$$\lim_{n \rightarrow \infty} \sum_{\lfloor (\tilde{x} - \epsilon)n \rfloor \leq z \leq \lceil (\tilde{x} + \epsilon)n \rceil} \pi_z = 1 \quad \text{for any } \epsilon > 0.$$

If $a > d$, $\xi_a = \xi_d$, and σ is sufficiently small for there to be three Bayesian Nash equilibria $\tilde{x}_L < \tilde{x}_M < \tilde{x}_H$ (see Lemma 2), then all weight focuses on \tilde{x}_H .

Proof. Take any x such that $x < \Phi(\kappa(x)/\sigma)$, so that $1 - x > 1 - \Phi(\kappa(x)/\sigma)$. Then:

$$\frac{\pi_{\lfloor xn \rfloor}}{\pi_{\lceil xn \rceil}} = \frac{p_{\lfloor xn \rfloor \lfloor xn \rfloor}}{p_{\lfloor xn \rfloor \lceil xn \rceil}} = \frac{\lfloor xn \rfloor [1 - \Phi(\kappa_{\lfloor xn \rfloor} / \sigma)]}{(n - \lfloor xn \rfloor) \Phi(\kappa_{\lfloor xn \rfloor} / \sigma)} \xrightarrow{n \rightarrow \infty} \frac{x [1 - \Phi(\kappa(x) / \sigma)]}{(1 - x) \Phi(\kappa(x) / \sigma)} < 1.$$

Analysis of further limiting behaviour ($n \rightarrow \infty$) is relegated to Appendix A. \square

4. Equilibrium selection

KMR (1993) and Young (1993) allow the probability of an error (a mutation) to vanish. In the limit, the ergodic distribution places all weight on a single state, and “selects” a pure strategy Nash equilibrium. The specification considered here involves no errors—players

best respond against the observed population frequency, contingent on their own payoffs. The analogue of a mutation is a contrary action taken against the flow of play. Suppose that a fraction $x > x^*$ of observed incumbents are using strategy 1, so that $\kappa(x) > 0$. Without payoff idiosyncrasy, an entrant would choose strategy 1. With payoff idiosyncrasy, an entrant will choose the contrary strategy 2 with probability $1 - \Phi(\kappa(x)/\sigma)$. Following a Harsanyi (1973) purification process ($\sigma \rightarrow 0$) ensures that the probability of such a “mutation” vanishes to zero.⁵ The use of purification to select equilibria generates the following observations.

The probability of a contrary action differs across states: For $\kappa(y) > \kappa(x) > 0$, such probabilities satisfy $1 - \Phi(\kappa(y)/\sigma) < 1 - \Phi(\kappa(x)/\sigma)$, and are in some sense state dependent. Bergin and Lipman (1996), however, employed a more stringent definition of state dependence. They demonstrated that the critical feature is the rate at which a mutation probability vanishes to zero in the selection process. Taking $\sigma \rightarrow 0$:

$$\begin{aligned} \frac{1 - \Phi(\kappa(x)/\sigma)}{1 - \Phi(\kappa(y)/\sigma)} &= \frac{\phi(\kappa(y)/\sigma)/(1 - \Phi(\kappa(y)/\sigma))}{\phi(\kappa(x)/\sigma)/(1 - \Phi(\kappa(x)/\sigma))} \times \frac{\phi(\kappa(x)/\sigma)}{\phi(\kappa(y)/\sigma)} \\ &\rightarrow \frac{\kappa(y)/\sigma}{\kappa(x)/\sigma} \exp\left(-\frac{\kappa(x)^2 - \kappa(y)^2}{2\sigma^2}\right) \\ &\rightarrow \infty. \end{aligned}$$

The first equality follows from multiplying and dividing both numerator and denominator by the normal density ϕ . This yields a ratio of hazard rates as the first term, and a ratio of densities as the second term. The second step follows from recognising that the hazard rate of the normal distribution is asymptotically linear: Formally $u - [\phi(u)/(1 - \Phi(u))] \rightarrow 0$ as $u \rightarrow \infty$, yielding the ratio on the left.⁶ The exponential term on the right follows from the combination of the two normal densities. The third step follows from dominance of exponential terms in the limit, and that $\kappa(x)^2 < \kappa(y)^2$ by assumption. Notice that the rate at which the mutation probability dies away to zero is determined by the basin depth $\kappa(x)^2$. Hence the endogenously generated mutations meet the definition of state dependence proposed by Bergin and Lipman (1996).

A further observation is that the model fits into the framework described by Blume (1999), where the probability of a mutation (the “noise” specification) depends upon the expected payoff difference between the strategies. Facing a proportion x of strategy 1 opponents, this difference is $ax - d(1 - x)$, which is a monotonic transformation of x . The adaptive play process described here is, therefore, a special case of Blume’s specification. He obtained the following results: When the noise process is “skew symmetric,” meaning that the mutation probability depends only on the absolute difference in payoffs, then the risk-dominant equilibrium is selected in the limit. Here, the noise process is skew symmetric whenever $\xi_a = \xi_d$. It remains to investigate selection when $\xi_a \neq \xi_d$. A first

⁵ Of course, this is only one example of a purification process. There are other notions, for example that of Aumann et al. (1983).

⁶ Apply l’Hôpital’s rule to obtain

$$\lim_{u \rightarrow \infty} \left\{ \frac{\phi(u)}{1 - \Phi(u)} - u \right\} = \lim_{u \rightarrow \infty} \frac{\phi(u) - u(1 - \Phi(u))}{1 - \Phi(u)} = \lim_{u \rightarrow \infty} \frac{\Phi(u) - 1}{-\phi(u)} = 0.$$

step is to establish that only extreme states (corresponding to pure Nash equilibria) matter in the limit.

Lemma 3. Interior states carry no weight in the limit:

$$z \notin \{0, n\} \Rightarrow \lim_{\sigma \rightarrow 0} \pi_z = 0.$$

Proof. Consider a state $z \in Z_1$ where $z < n$. Note that $\kappa_z > 0$. Allowing $\sigma \rightarrow \infty$:

$$\pi_z \leq \frac{\pi_z}{\pi_{z+1}} = \frac{p_{(z+1)z}}{p_{z(z+1)}} = \frac{z+1}{n} \times \frac{n}{n-z} \times \frac{1 - \Phi(\kappa_z/\sigma)}{\Phi(\kappa_z/\sigma)} \rightarrow 0,$$

which follows since $\Phi(\kappa_z/\sigma) \rightarrow 1$ as $\sigma \rightarrow 0$. \square

It follows that $\lim_{\sigma \rightarrow 0} (\pi_0 + \pi_n) = 1$. The relative likelihood of observing these extreme states is determined by the relative difficulty of moving between them. Two factors influence this. The first is the *basin width*. This is the number of steps (or mutations) required to escape from a basin of attraction. Since $x^* < 1/2$, the basin Z_1 is wider than Z_2 . In the models of KMR (1993) and Young (1993) each mutative step is taken with equal probability. It follows that the strategy with the widest basin of attraction (in this case, the risk-dominant strategy 1) is selected. Here, however, the *basin depth* is of importance. Both width and depth together influence the selection process. An extension of Definition 1 reflects these ideas.

Definition 2. Basin volumes are defined as $B_1 = \sum_{z \in Z_1} \kappa_{z-1}^2$ and $B_2 = \sum_{z \in Z_2} \kappa_z^2$.

The basin depths and volumes are both illustrated in Fig. 2. The solid line illustrates the basins of attraction when the payoff trembles have equal variances. The basin of attraction for strategy 1 is both wider and deeper. In contrast, the broken line illustrates the case where the payoff trembles have unequal variances. Whereas the basin Z_2 is narrower, it is also far deeper. This is because the payoff \tilde{a} is more volatile, the probability of observing a contrary action in this basin is much higher, and hence the basin volume is greater. The effect of basin volume is demonstrated in the central selection result of the paper.

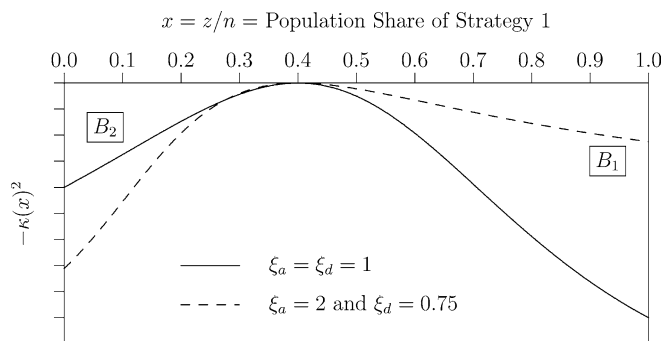


Fig. 2. Basin depths and volumes. Parameters are $a = 3$ and $d = 2$.

Proposition 2. *The strategy with the largest basin volume is selected as $\sigma \rightarrow 0$. Formally, if $B_1 > B_2$ then $\lim_{\sigma \rightarrow 0} \pi_n = 1$ and if $B_2 > B_1$ then $\lim_{\sigma \rightarrow 0} \pi_0 = 1$.*

Proof. Using Lemma 1 take the ratio of π_n and π_0 to obtain:

$$\frac{\pi_n}{\pi_0} = \frac{\prod_{z < n} \frac{n-z}{n} \Phi(\kappa_z/\sigma)}{\prod_{z > 0} \frac{z}{n} (1 - \Phi(\kappa_{z-1}/\sigma))} = \frac{\prod_{z \in Z_1} \Phi(\kappa_{z-1}/\sigma)}{\prod_{z \in Z_2} (1 - \Phi(\kappa_z/\sigma))} \frac{\prod_{z \in Z_2} \Phi(\kappa_z/\sigma)}{\prod_{z \in Z_1} (1 - \Phi(\kappa_{z-1}/\sigma))}.$$

The above follows from the cancellation of terms such as z/n and $(n - z)/n$, and re-indexing and separating the products as appropriate. The numerator and denominator of the first term on the right-hand side both tend to unity as $\sigma \rightarrow 0$. For instance, when $z \in Z_1$, reference to Definition 1 confirms that $\kappa_{z-1} > 0$. It follows strategy selection is determined by the second term, where both numerator and denominator tend to zero. Write this term as

$$\begin{aligned} \frac{\prod_{z \in Z_2} \Phi(\kappa_z/\sigma)}{\prod_{z \in Z_1} (1 - \Phi(\kappa_{z-1}/\sigma))} &= \frac{\prod_{z \in Z_2} \phi(\kappa_z/\sigma)}{\prod_{z \in Z_1} \phi(\kappa_{z-1}/\sigma)} \\ &\times \frac{\prod_{z \in Z_1} \phi(\kappa_{z-1}/\sigma)/(1 - \Phi(\kappa_{z-1}/\sigma))}{\prod_{z \in Z_2} \phi(\kappa_z/\sigma)/\Phi(\kappa_z/\sigma)}. \end{aligned}$$

The first term is explicitly:

$$\frac{\prod_{z \in Z_2} \phi(\kappa_z/\sigma)}{\prod_{z \in Z_1} \phi(\kappa_{z-1}/\sigma)} = (2\pi)^{(n-2\lceil(n-1)x^*\rceil)/2} \exp\left(-\frac{\sum_{z \in Z_2} \kappa_z^2 - \sum_{z \in Z_1} \kappa_{z-1}^2}{2\sigma^2}\right). \quad (4)$$

Consider a typical element in the numerator of the second term. Since $\kappa_{z-1} > 0$ for $z \in Z_1$, $\kappa_{z-1}/\sigma \rightarrow +\infty$ as $\sigma \rightarrow 0$. Since the hazard rate of the normal is asymptotically linear, it follows that $\lim_{\sigma \rightarrow 0} \phi(\kappa_{z-1}/\sigma)/(1 - \Phi(\kappa_{z-1}/\sigma)) = \lim_{\sigma \rightarrow 0} \kappa_{z-1}/\sigma$. A similar argument holds for the denominator, where $\kappa_z < 0$. Hence,

$$\frac{\prod_{z \in Z_1} \phi(\kappa_{z-1}/\sigma)/(1 - \Phi(\kappa_{z-1}/\sigma))}{\prod_{z \in Z_2} \phi(\kappa_z/\sigma)/\Phi(\kappa_z/\sigma)} \rightarrow \frac{\sigma^{\lceil(n-1)x^*\rceil} \prod_{z \in Z_1} \kappa_{z-1}}{\sigma^{n-\lceil(n-1)x^*\rceil} \prod_{z \in Z_2} (-\kappa_z)}$$

which is polynomial in σ . The first term, (4), is exponential in σ . The limit becomes

$$\begin{aligned} \lim_{\sigma \rightarrow 0} \frac{\pi_n}{\pi_0} &= \lim_{\sigma \rightarrow 0} \left(\frac{\sqrt{2\pi}}{\sigma}\right)^{n-2\lceil(n-1)x^*\rceil} \exp\left(-\frac{\sum_{z \in Z_2} \kappa_z^2 - \sum_{z \in Z_1} \kappa_{z-1}^2}{2\sigma^2}\right) \\ &\times \frac{\prod_{z \in Z_1} \kappa_{z-1}}{\prod_{z \in Z_2} (-\kappa_z)}. \end{aligned}$$

The exponential term dominates asymptotically, and hence the limit diverges whenever the numerator of the fractional term in the exponent is negative. In conclusion:

$$\sum_{z \in Z_1} \kappa_{z-1}^2 > \sum_{z \in Z_2} \kappa_z^2 \Rightarrow \lim_{\sigma \rightarrow 0} \frac{\pi_n}{\pi_0} = \infty \quad \text{and}$$

$$\sum_{z \in Z_1} \kappa_{z-1}^2 < \sum_{z \in Z_2} \kappa_z^2 \Rightarrow \lim_{\sigma \rightarrow 0} \frac{\pi_n}{\pi_0} = 0,$$

which is precisely the desired result. \square

Proposition 2 yields an immediate and simple corollary.

Corollary 1. For $n = 2$, strategy 1 is selected ($\lim_{\sigma \rightarrow 0} \pi_n = 1$) whenever $a/\xi_a > d/\xi_d$.

The $n = 2$ case is of interest precisely because it corresponds to a situation where the population size matches the number of players. In this case, the models of KMR (1993) and Young (1993) do not offer a selection result. Each basin of attraction has a single member, and hence escape from either requires exactly one mutation. Here, however, a selection result is obtained. This is because the probability of a contrary action depends upon the payoffs and tremble variances of the game. In the $n = 2$ case, selection is driven entirely by basin depth and not by basin width. From Corollary 1 it is clear that strategy 1 is selected whenever $\xi_a = \xi_d$, since its risk dominance corresponds to $a > d$. This is also true for larger populations.

Corollary 2. Suppose that $\xi_a = \xi_d$. Strategy 1 is risk-dominant and hence selected.

Proof. See Appendix A. \square

When payoffs are trembled in the same way, risk-dominant equilibria continue to be selected, confirming the analysis of Blume (1999). When trembles are payoff specific (equivalent to a skew asymmetric noise specification), then either equilibrium may be selected. The basin volume of Definition 2 captures the relative influence of these factors. The degree to which asymmetry in tremble variances can overcome the effect of relative payoff size is illustrated in Fig. 3. To the right of the vertical dotted line, strategy 1 is risk-dominant. When payoff variances are equal ($\xi_a^2/(\xi_a^2 - \xi_d^2) = 1/2$) the selection criterion coincides with risk dominance. As the trembles become more skew asymmetric (for instance, for relatively large ξ_a^2) the criterion diverges. The dashed line, for instance, shows the equilibrium selected for $n = 5$.

Acknowledgments

Thanks are due to Ben Cooper, Huw Dixon, Jon Dworak, Glenn Ellison, Justin Johnson, Paul Klemperer, Meg Meyer, Hyun Shin and Peter Sørensen, anonymous reviewers, and seminar participants at Nuffield College, Oxford IES, SUNY at Stony Brook, University College London and York University for helpful comments. This paper previously appeared as Nuffield Discussion Paper 130 under the title “Adaptive Dynamics with Payoff Heterogeneity.”

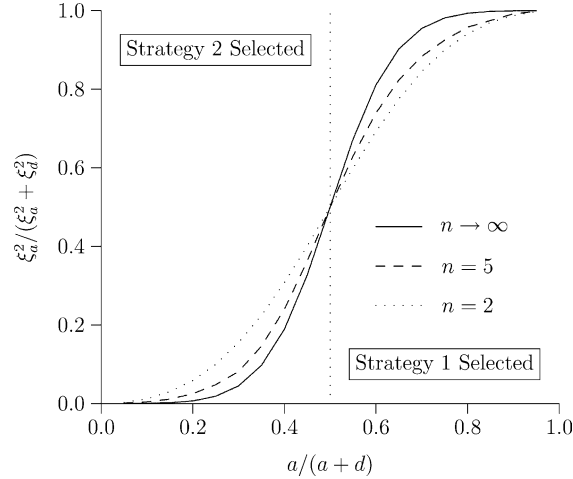


Fig. 3. Strategy selection regions ($\rho = 0$).

Appendix A. Omitted results

Proof of Lemma 2. Fixed points of $\Phi(\kappa(x)/\sigma)$ are roots of $f(x) = \Phi(\kappa(x)/\sigma) - x$. Notice that $f'(x) = \phi(\kappa(x)/\sigma)\kappa'(x)/\sigma - 1$. As $\sigma \rightarrow \infty$, $f'(x) \rightarrow -1$ uniformly for $x \in [0, 1]$. Thus, for sufficiently large σ , $f(x)$ is decreasing everywhere. Hence $f(x)$ only has one root in the neighbourhood of $x = \frac{1}{2}$. When $\sigma \rightarrow 0$, $f(x) \rightarrow 1 - x$ if $x > x^*$ and $f(x) \rightarrow -x$ if $x < x^*$, so there cannot be a fixed point unless it is local to $\{0, x^*, 1\}$. Consider the interval $[0, \varepsilon]$. For sufficiently small σ , $f(x)$ is decreasing in this interval. Moreover, $f(0) > 0$ and $f(\varepsilon) < 0$. Therefore there is exactly one root in this interval. A similar argument applies to $[1 - \varepsilon, 1]$. Now consider $[x^* - \varepsilon, x^* + \varepsilon]$. Then $f(x^* - \varepsilon) < 0$ and $f(x^* + \varepsilon) > 0$. Again there is at least one root in the interval. $\Phi(\kappa(x)/\sigma)$ is strictly increasing. A fixed point of $\Phi(\kappa(x)/\sigma)$ corresponds to a fixed point of its inverse. Local to x^* the derivative of the inverse is less than one. This locality expands as σ gets small. Within this region there can be only one fixed point of the inverse and hence in this interval the root of $f(x)$ is unique. \square

Proof of Proposition 1. Take a proportion x of strategy 1 players and a corresponding state $z = \lceil nx \rceil$ or $z = \lfloor nx \rfloor$. Take logs of Eq. (3) and divide by n :

$$\begin{aligned} \frac{\log q_z}{n} &= \frac{1}{n} \left\{ \sum_{j < z} \log \left[\frac{n-j}{n} \Phi \left(\frac{\kappa_j}{\sigma} \right) \right] + \sum_{j > z} \log \left[\frac{j}{n} \left(1 - \Phi \left(\frac{\kappa_{j-1}}{\sigma} \right) \right) \right] \right\} \\ &\rightarrow Q(x) \equiv \int_0^x \log(1-y) \Phi \left(\frac{\kappa(y)}{\sigma} \right) dy + \int_x^1 \log y \left(1 - \Phi \left(\frac{\kappa(y)}{\sigma} \right) \right) dy \end{aligned}$$

which is an integral approximation similar in spirit to that employed by Young (1998). Differentiation reveals that $Q'(x) = 0 \Leftrightarrow x = \Phi(\kappa(x)/\sigma)$. In other words, maxima and

minima \tilde{x} of $Q(x)$ correspond to Bayesian Nash equilibria. This reaffirms the first part of the proposition. For almost all payoff specifications, $Q(x)$ has a unique global maximiser \hat{x} . For arbitrarily small ϵ , define

$$\Delta_\epsilon = \{x \in [0, 1]: Q(x) \geq Q(\hat{x}) - \epsilon\}.$$

For sufficiently small ϵ , this is a convex neighbourhood of \hat{x} . $\Delta_{2\epsilon}$ may be defined similarly. Now consider the following:

$$\frac{\sum_{z: \frac{z}{n} \in \Delta_\epsilon} \pi_z}{\sum_{z: \frac{z}{n} \notin \Delta_{2\epsilon}} \pi_z} = \frac{\sum_{z: \frac{z}{n} \in \Delta_\epsilon} q_z}{\sum_{z: \frac{z}{n} \notin \Delta_{2\epsilon}} q_z} \geq \frac{\min_{z: \frac{z}{n} \in \Delta_\epsilon} q_z}{\max_{z: \frac{z}{n} \notin \Delta_{2\epsilon}} q_z} \times \frac{\#\{z: \frac{z}{n} \in \Delta_\epsilon\}}{\#\{z: \frac{z}{n} \notin \Delta_{2\epsilon}\}}.$$

The right-hand term converges to a non-zero number. Taking the left hand term:

$$\begin{aligned} \log \left[\frac{\min_{z: \frac{z}{n} \in \Delta_\epsilon} q_z}{\max_{z: \frac{z}{n} \notin \Delta_{2\epsilon}} q_z} \right] &= n \left[\min_{z: \frac{z}{n} \in \Delta_\epsilon} \frac{\log q_z}{n} - \max_{z: \frac{z}{n} \notin \Delta_{2\epsilon}} \frac{\log q_z}{n} \right] \\ &\xrightarrow{n \rightarrow \infty} n[(Q(\hat{x}) - \epsilon) - (Q(\hat{x}) - 2\epsilon)] = n\epsilon \\ &\xrightarrow{n \rightarrow \infty} \infty. \end{aligned}$$

It follows that all weight in the ergodic distribution must accrue in a neighbourhood of \hat{x} . It remains to show that $\hat{x} = \tilde{x}_H$ when $\xi_a = \xi_d$. For sufficiently small σ there are two stable Bayesian Nash equilibria satisfying $\tilde{x}_L < 1/2 < \tilde{x}_H$ (Lemma 2). It is sufficient, therefore, to show that $\hat{x} > 1/2$. Suppose not, so that $\hat{x} < 1/2$. Since \hat{x} is the unique global maximiser of $Q(x)$ it must be the case that $Q(\hat{x}) > Q(1 - \hat{x})$. Computing the difference of these two terms:

$$\begin{aligned} Q(\hat{x}) - Q(1 - \hat{x}) &= \int_{\hat{x}}^{1-\hat{x}} \log y \left(1 - \Phi \left(\frac{\kappa(y)}{\sigma} \right) \right) dy - \int_{\hat{x}}^{1-\hat{x}} \log(1-y) \Phi \left(\frac{\kappa(y)}{\sigma} \right) dy \\ &= \int_{\hat{x}}^{1-\hat{x}} \log y \left(1 - \Phi \left(\frac{\kappa(y)}{\sigma} \right) \right) dy - \int_{\hat{x}}^{1-\hat{x}} \log y \Phi \left(\frac{\kappa(1-y)}{\sigma} \right) dy \\ &= \int_{\hat{x}}^{1-\hat{x}} \log \left[\frac{1 - \Phi(\kappa(y)/\sigma)}{\Phi(\kappa(1-y)/\sigma)} \right] dy \\ &= \int_{\hat{x}}^{1/2} \log \left[\frac{1 - \Phi(\kappa(y)/\sigma)}{\Phi(\kappa(1-y)/\sigma)} \right] dy + \int_{1/2}^{1-\hat{x}} \log \left[\frac{1 - \Phi(\kappa(y)/\sigma)}{\Phi(\kappa(1-y)/\sigma)} \right] dy \\ &= \int_{\hat{x}}^{1/2} \log \left[\frac{1 - \Phi(\kappa(y)/\sigma)}{\Phi(\kappa(1-y)/\sigma)} \right] dy + \int_{\hat{x}}^{1/2} \log \left[\frac{1 - \Phi(\kappa(1-y)/\sigma)}{\Phi(\kappa(y)/\sigma)} \right] dy \\ &= \int_{\hat{x}}^{1/2} \log \left[\frac{1 - \Phi(\kappa(y)/\sigma)}{\Phi(\kappa(1-y)/\sigma)} \right] \left[\frac{1 - \Phi(\kappa(1-y)/\sigma)}{\Phi(\kappa(y)/\sigma)} \right] dy. \end{aligned}$$

For this to be positive, it must be true that for some $y < 1/2$ the integrand is positive. This reduces to:

$$\begin{aligned} 1 - \Phi(\kappa(y)/\sigma) > \Phi(\kappa(1-y)/\sigma) &\Leftrightarrow -\kappa(y) > \kappa(1-y) \\ &\Leftrightarrow -ya + (1-y)d > (1-y)a - yd \\ &\Leftrightarrow d > a. \end{aligned}$$

But this is a contradiction, since by assumption $a > d$. \square

Proof of Corollary 2. The basin volume condition is

$$\sum_{j=0}^{\lfloor x^*(n-1) \rfloor} \frac{(ja - (n-j-1)d)^2}{j^2 + (n-j-1)^2} < \sum_{j=\lceil x^*(n-1) \rceil}^{n-1} \frac{(ja - (n-j-1)d)^2}{j^2 + (n-j-1)^2}.$$

Consider the right hand term. Make a change of variable $k = n - j - 1$ to yield

$$\sum_{j=0}^{\lfloor x^*(n-1) \rfloor} \frac{(ja - (n-j-1)d)^2}{j^2 + (n-j-1)^2} < \sum_{k=0}^{n-1-\lceil x^*(n-1) \rceil} \frac{((n-k-1)a - kd)^2}{k^2 + (n-k-1)^2}.$$

If strategy 1 is risk-dominant, then $a > d$ and $\lceil x^*(n-1) \rceil < n - 1 - \lceil x^*(n-1) \rceil$. The condition becomes

$$\begin{aligned} &\sum_{j=0}^{\lfloor x^*(n-1) \rfloor} \frac{(ja - (n-j-1)d)^2 - ((n-j-1)a - jd)^2}{j^2 + (n-j-1)^2} \\ &< \sum_{k=\lceil x^*(n-1) \rceil}^{n-1-\lceil x^*(n-1) \rceil} \frac{((n-k-1)a - kd)^2}{k^2 + (n-k-1)^2}. \end{aligned}$$

The right-hand side of this expression is positive. Multiply out the numerator in each of the left-hand terms:

$$(ja - (n-j-1)d)^2 - ((n-j-1)a - jd)^2 = (j^2 - (n-j-1)^2)(a^2 - d^2).$$

This is negative since $a^2 > d^2$ and $j^2 < (n-j-1)^2$ for $j \leq \lfloor x^*(n-1) \rfloor$. The inequality holds, yielding sufficiency. Necessity follows since if $a < d$ the same procedure establishes the dominance of strategy 2. \square

Calculation of basin volume for large n . Figure 3 displays the basin volume selection criterion for large n . Using $\lambda = a/(a+d)$, $\psi = \xi_a^2/(\xi_a^2 + \xi_d^2)$ and setting $\rho = 0$, the basin depth is proportional to:

$$\tilde{\kappa}(x)^2 = \frac{(\lambda x - (1-\lambda)(1-x))^2}{\psi x^2 + (1-\psi)(1-x)^2}. \tag{A.1}$$

Dividing the basin volume by n , obtain

$$\frac{B_1}{n} = \frac{1}{n} \sum_{z \in Z_1} \kappa_{z-1}^2 \xrightarrow{n \rightarrow \infty} \int_{x^*}^1 \kappa(x)^2 dx \propto \int_{x^*}^1 \tilde{\kappa}(x)^2 dx$$

and similarly for B_2 . An explicit form for the last integral is available:

$$\int \tilde{\kappa}(x)^2 dx = x + \frac{\lambda^2 - \psi - 2\lambda\psi + 2\psi^2}{\sqrt{\psi(1-\psi)}} \arctan \left\{ \frac{x - (1-\psi)}{\sqrt{\psi(1-\psi)}} \right\} \\ + (\lambda - \psi) \log(\psi x^2 + (1-\psi)(1-x)^2) + \text{const.}$$

Differentiation of the above yields the expression from Eq. (A.1). This enables the calculation of the line for $n \rightarrow \infty$ in Fig. 3. \square

References

- Aumann, R.J., Katznelson, Y., Radner, R., Rosenthal, R.W., Weiss, B., 1983. Approximate purification of mixed strategies. *Math. Operations Res.* 8, 3277–3341.
- Bergin, J., Lipman, B., 1996. Evolution with state-dependent mutations. *Econometrica* 64 (4), 943–956.
- Binmore, K.G., Samuelson, L., 1997. Muddling through: noisy equilibrium selection. *J. Econ. Theory* 74, 235–265.
- Blume, L.E., 1999. How noise matters. SFI working paper. Mimeo.
- Blume, L.E., Durlauf, S.N., 2001. The interactions-based approach to socioeconomic behaviour. In: Durlauf, S.N., Young, H.P. (Eds.), *Social Dynamics*. MIT Press, Cambridge, MA, pp. 15–44.
- Brock, W.A., Durlauf, S.N., 2001. Discrete choice with social interactions. *Rev. Econ. Stud.* 68 (2), 235–260.
- Foster, D., Young, H.P., 1990. Stochastic evolutionary game dynamics. *Theor. Population Biol.* 38, 219–232.
- Grimmett, G.R., Stirzaker, D.R., 2001. *Probability and Random Processes*, 3rd Edition. Oxford Univ. Press, Oxford.
- Harsanyi, J.C., 1973. Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *Int. J. Game Theory* 2, 1–23.
- Harsanyi, J.C., Selten, R., 1988. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA.
- Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation and long-run equilibria in games. *Econometrica* 61 (1), 29–56.
- van Damme, E., Weibull, J.W., 1998. Evolution with mutations driven by control costs. Mimeo. CentER, Tilburg University and The Stockholm School of Economics.
- Young, H.P., 1993. The evolution of conventions. *Econometrica* 61 (1), 57–84.
- Young, H.P., 1998. *Individual Strategy and Social Structure*. Princeton Univ. Press, Princeton, NJ.