ADAPTIVE DYNAMICS WITH PAYOFF HETEROGENEITY

DAVID P MYATT AND CHRIS WALLACE

May 1997

Abstract. A finite population of agents playing a 2×2 symmetric game evolves by adaptive best response. The assumption that players make *mistakes* is dropped in favour of one where players *differ*, via payoff heterogeneity. Arbitrary mutations are thus replaced with an economically justified specification. The *depth* as well as the *width* of basins of attraction is important when determining long-run behaviour. With vanishing noise and *balanced* payoff variances, the risk dominant equilibrium is selected. *Unbalanced* variances may result in the selection of other equilibria, including the payoff dominant. The ergodic extrema correspond exactly to the Bayesian Nash equilibria of the underlying trembled stage game. This enables an analysis of the ergodic distribution for non-vanishing noise and larger populations.

1. Introduction

> "That was excellently observed," say I when I read a passage in another where his opinion agrees with mine. When we differ, then I pronounce him to be mistaken.
>
> Jonathan Swift: *Thoughts on Various Subjects.*

The Nash equilibrium concept in strategic form games requires players to choose payoff maximising actions, given their beliefs, and further imposes a consistency requirement on the beliefs of all players. This concept, however, places strong rationality and epistemic requirements on players.[1] Furthermore, many games have multiple Nash equilibria, leading to an equilibrium selection problem if the play of a particular game is to be predicted.

Motivated in part by dissatisfaction with the requirements of Nash equilibrium and with a desire to select among a multiplicity, a research programme has developed aiming to model the dynamics of boundedly rational agents. Influenced by the evolutionary literature, pioneering papers such as Kandori, Mailath and Rob (KMR, [8]) and Young [14] specify adaptive dynamics in which finite populations of players

Date: This version (4.0) May 6, 1997.

Keywords. Adaptive dynamics, payoff heterogeneity, evolution, generalised risk dominance.

We gratefully acknowledge financial support from the ESRC, awards R00429624089 and R00429534346 respectively. Thanks are due to Ben Cooper, Huw Dixon, Jon Dworak, Justin Johnson, Paul Klemperer, Meg Meyer, Hyun Shin and Peter Sørensen for helpful comments.

[1] Aumann and Brandenburger [1] present a detailed analysis of these requirements. Not only must players maximise given their beliefs, but there must also be mutual knowledge of payoffs, beliefs and conjectures.

revise their strategies periodically. They observe the distribution of strategies in the population and proceed to play a best response.[2] A purely deterministic dynamic remains path-dependent, and the search for sharper conclusions leads these authors to study perturbed Markov processes. Following the specification of a base non-ergodic dynamic, perturbations are introduced, yielding irreducible Markov chains. The associated ergodic distribution is then characterised as the perturbations die away. For generic games with strict Nash equilibria, the limiting process places all weight on a single equilibrium. Within the class of 2×2 symmetric games, these models focus on the Harsanyi-Selten [6] risk-dominant equilibrium.

These models involve naive belief formation and also drop the maximisation hypothesis. The base dynamic specifies payoff maximisation given beliefs. Beliefs are based merely on observed frequencies, and players are thus required only to know their own payoffs. Such a process is path-dependent. For example, consider a 2×2 coordination game. An entrant to a population uniformly playing one strategy will respond with the same, resulting in the critical nature of the initial population state. To obtain an ergodic process, KMR [8] and Young [14], among others, introduce *mutations*, hence ensuring the irreducibility of the Markov chain. The interpretation is that players make *mistakes* when updating their strategies. Long-run selection results are obtained when these mutations are allowed to vanish in the limit. These results are therefore based upon a weakening of payoff maximisation which is integral to the notion of Nash equilibrium.

The rôle of such perturbations is critical but, absent a reasonable model, the specification of state-independent mutations is arbitrary. State-independent mistakes imply that an agent fails to play an optimal response with the same probability irrespective of strategy frequencies, an unnatural feature in a coordination scenario. This observation is central to an elegant critique by Bergin and Lipman [2]. They demonstrate that mutations may be chosen such that any stationary distribution of the base process is the limiting ergodic distribution of a perturbed process. In particular, any strict Nash equilibrium of a strategic form game is selected for a suitable mutation model. In the light of this result, they argue (with good reason) that any specific model of mutations must be justified; an arbitrary model will not suffice.

Intuitively, the basin of attraction for each equilibrium drives long-run selection. Uniform mutation rates ensure that such a basin has constant *depth*. Hence the *width* of the basin governs long-run behaviour. Mutation rates that differ by state result in varying basin depth. It is then the overall *volume* of the attraction basin that determines the long-run outcome.

The essence of the irreducibility of the adaptive process is that for a given state, different entering or revising players may take different actions. KMR [8] and Young [14] generate this by assuming that entrants may fail to optimise through occasional error. The weakening of payoff maximisation via mistakes is premature. It is more plausible to maintain the maximisation hypothesis and instead allow players to *differ* in their preferences over outcomes. The best response of two distinct agents in an identical situation may therefore be different.

[2]In fact, KMR [8] allow a wider class of underlying deterministic dynamics, and Young [14] allows new or revising players to sample from a truncated history of play.

This paper presents an adaptive response model with payoff heterogeneity. Trembles are added to a mean payoff matrix, and allowed to vary across strategy profiles. This captures the idiosyncratic preferences of individuals. Notions of *balanced* and *unbalanced* trembles are introduced, corresponding to the relative variability of payoffs. Facing a strategy frequency, an agent plays a best response *without error*. To the observer, however, the action may appear mistaken relative to mean payoffs. Apparent "mutations" arise as a natural consequence of this procedure. The approach has a strong economic justification and moreover generates state-dependent mutations.[3] These mutations have attractive properties, in particular they are small for uniform populations. Payoff trembles are fully parametrised using the normal, yielding precise results. Furthermore, this distributional assumption is intuitively appealing.

In the adaptive response dynamic employed, a finite population is subject to turnover. Each period, a randomly selected agent is replaced. The new entrant, equipped with a trembled payoff matrix, plays a best response to the existing strategy frequency. Implicit to this schema, an underlying renewal process is envisioned in which players periodically revise their strategies or are replaced in continuous time. Binmore and Samuelson [3] develop this approach more formally. Note that in a continuous time framework, simultaneous revisions (such as those used in KMR [8]) would be a measure-zero event.

The analysis uses the graph-theoretic techniques of Freidlin and Wentzell [5], so profitably employed in earlier work, to characterise the stationary distribution of the associated Markov process. The limit is taken as tremble variances fall to zero. Although these converge to zero at the *same* rate, they endogenously give rise to mutations converging to zero at *different* rates. Hence the observations of Bergin and Lipman [2] are pertinent. The parametric specification of the model allows a convenient decomposition of transition probabilities into densities and hazards. Asymptotic properties of hazard rates enable the derivation of an exact criterion for equilibrium selection. This condition relies solely on mean payoffs and the balance of the trembles. Heuristically, the mean payoffs determine the width of attraction basins whilst the trembles determine the depth.

For balanced trembles in the 2×2 coordination game, the dynamic selects the risk-dominant equilibrium as heterogeneity vanishes. This reinforces the conclusions of previous authors for this class of games. If trembles are unbalanced, however, the risk-dominated equilibrium may be selected. A concept of *generalised-risk-dominance* is introduced that better predicts the long-run outcome for small trembles. This concept captures the rôle of payoff variability in determining the riskiness of equilibria. If an equilibrium is both risk-dominant and generalised-risk-dominant, then it is selected.

A further issue to address is the relevance of the ergodic distribution. As noted by Ellison [4] and others, for small mutations and large populations, the transition times between long-run equilibria are large. Ellison markedly reduces transition times by introducing local interaction of players, which here is analogous to a small population. With this in mind, results for vanishing heterogeneity are particularly applicable when

[3] Binmore and Samuelson [3] 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. Hence their approach lacks an explicit economic justification.

the population is small. For larger populations, the limiting results continue to hold, but are of less interest. For such populations the ergodic distribution is only of interest for non-vanishing noise.

Following this observation, the ergodic distribution is examined for fixed positive heterogeneity. The judicious choice of a parametric specification allows such an approach. This type of analysis has been absent from previous research. With positive noise, the distribution places positive weight on all states and no equilibrium is "selected". Hence, the extrema of the invariant distribution are determined. A convenient integral approximation to rooted tree weights is available, which becomes exact for a larger population — precisely the case of interest. This approximation is employed to show that the extrema of the ergodic distribution correspond to the Bayesian Nash equilibria of the underlying trembled stage game. Therefore, a strong relationship between adaptive behaviour and rational play is revealed. This relationship is further explored in a companion paper, Myatt and Wallace [11].

Typically, the trembled stage game has either one or three Bayesian Nash equilibria. For sufficiently small heterogeneity, there are three. As heterogeneity vanishes, they converge to the two pure and one mixed Nash equilibria of the unperturbed stage game.[4] The novel approach taken here allows an inspection of the close relationship between the Bayesian Nash equilibria and the ergodic distribution of the adaptive response dynamic. The two Bayesian Nash equilibria that correspond to the pure equilibria lie at the modal points of the ergodic distribution. The "mixed" Bayesian Nash equilibrium corresponds to a local minimum. As heterogeneity grows, all but one of these Bayesian Nash equilibria eventually vanish. The remaining equilibrium corresponds to the risk-dominant pure Nash equilibrium. Thus a bimodal ergodic distribution becomes unimodal for sufficiently large heterogeneity, and the surviving mode lies close to the risk-dominant strategy.

The argument proceeds as follows. In Section 2, the model with payoff perturbations is presented, together with a motivating example of PC adoption. This example is in a similar vein to that given by Kandori and Rob [9]. The analysis takes place in Section 3. The stationary distributions for vanishing heterogeneity are characterised in Section 3.3. Using the payoff assumptions the long-run selection criterion is obtained. In Section 3.4 the results for non-vanishing noise are presented and the ergodic distribution for this case characterised. Returning to the example in Section 4, the results are discussed and illustrated. These elements are drawn together along with some concluding remarks in Section 5. For convenience, omitted results are collected in an appendix.

[4]This is analogous to the purification argument of Harsanyi [7].

2. The Model

Section 2.1 describes the trembled stage game played by the agents. The dynamic via which the population evolves is outlined in Section 2.2. To illustrate this model, an example is presented in Section 2.3.

2.1. The Trembled Stage Game.

The starting point for the analysis is the familiar symmetric 2×2 strategic form game with generic payoffs:

	1	2
1	a, a	b, c
2	c, b	d, d

Notice that this game can be represented by a 2×2 matrix.

Definition 1. *The* mean payoff matrix *is defined as:*

$$\Lambda = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

A player equipped with payoffs Λ, entering a population is a mean payoff entrant*.*

Coordination games will be of particular interest, and will be the focus of subsequent analysis. This is when $a > c$ and $d > b$. Such a game has two pure Nash equilibria, $(1,1)$ and $(2,2)$ with associated security payoffs b and c. The symmetric mixed equilibrium entails mixing probabilities of $[x^*, 1 - x^*]$ where:

$$x^* = \frac{(d-b)}{(d-b)+(a-c)}$$

Note that equilibrium $(1,1)$ risk-dominates $(2,2)$ if and only if $a - c > d - b$, corresponding exactly to $x^* < \frac{1}{2}$.

The payoffs Λ may be viewed as the mean payoffs for any entering player. Any particular agent has heterogeneous preferences which are generated by the addition of payoff trembles. Each payoff is subject to an independent Gaussian disturbance.[5] The variances of these disturbances may be strategy profile specific, with a common scaling factor which is allowed to vanish for limiting results. Clearly these disturbances have a fully-parametric form. This, however, is a natural representation of differing payoffs across players. In particular, one might view differences over a particular payoff to be the resulting sum of many individual idiosyncratic factors, yielding the normal distribution as a natural specification.[6] Furthermore, this formulation allows clear closed-form results to be obtained. In the light of Bergin and Lipman [2], full generality of trembles, particularly allowing trembles to vary by state, leads to inconclusive results, and hence the approach is justified.

[5]The independence assumption is further discussed in Appendix A.1. The results are unaffected by tremble correlation.

[6]The later analysis demonstrates that the key features are the asymptotic properties of the densities and hazard rates of the disturbances. Thus any other distribution sharing these features will lead to similar results.

Definition 2. *Define the* payoff heterogeneity matrix Ψ *as:*

$$\Psi = \begin{bmatrix} \sigma_a & \sigma_b \\ \sigma_c & \sigma_d \end{bmatrix}$$

An entrant has trembled payoff matrix $\widetilde{\Lambda}$ *where:*

$$\widetilde{\Lambda} = \begin{bmatrix} \tilde{a} & \tilde{b} \\ \tilde{c} & \tilde{d} \end{bmatrix} = \begin{bmatrix} a + \sigma\varepsilon_a & b + \sigma\varepsilon_b \\ c + \sigma\varepsilon_c & d + \sigma\varepsilon_d \end{bmatrix}$$

where $\varepsilon_i \sim N\left(0, \sigma_i^2\right)$, *with* $E\left[\varepsilon_i\varepsilon_j\right] = 0$ *for* $i \neq j$ *and* σ *is a common scaling factor.*

Note that the payoff heterogeneity matrix Ψ determines the *relative* size of the payoff trembles. The overall size is determined by the scaling factor σ. For later convenience, the following definitions are introduced.

Definition 3. *The* payoff balance *of* Λ *and the* tremble balance *of* Ψ *are respectively:*

$$\lambda = 1 - x^* = \frac{(a-c)}{(a-c)+(d-b)} \quad \text{and} \quad \psi = \frac{\left(\sigma_a^2 + \sigma_c^2\right)}{\left(\sigma_a^2 + \sigma_c^2\right) + \left(\sigma_b^2 + \sigma_d^2\right)}$$

A game has balanced *trembles if* $\psi = \frac{1}{2}$. *Otherwise it has* unbalanced *trembles.*

Definition 4. *Strategy* 1 generalised-risk-dominates *strategy* 2 *whenever:*

$$\frac{a-c}{\sqrt{\sigma_a^2 + \sigma_c^2}} > \frac{d-b}{\sqrt{\sigma_b^2 + \sigma_d^2}}$$

Using the notions of balance, this is equivalent to $\lambda/(1-\lambda) > \sqrt{\psi/(1-\psi)}$. Assembling these components yields the trembled stage game:

Definition 5. *Define the* trembled stage game $\mathcal{G}$ *as the triple:*

$$\mathcal{G} = \langle \Lambda, \Psi, \sigma \rangle$$

Notice that $\mathcal{G}$ also represents a Bayesian game of incomplete information. This observation will prove useful in Section 3.4.

2.2. **The Adaptive Response Dynamic.** Take a finite population of n players. During a period each player repeatedly plays randomly selected opponents from the remaining $n-1$ players. Their strategies are fixed during each period. Denote the number of players using strategy 1 in a particular period as z, a member of the finite state space $Z = \{0, \ldots, n\}$. At the end of each period, a randomly selected member of the population leaves, and is replaced by another player with a newly trembled payoff matrix $\widetilde{\Lambda}$. This player observes the strategy distribution among the incumbents and selects a best response to this frequency.

Although the dynamic is described via the entry of new players, an equivalent scenario is one where members of the population periodically revise their strategy. In this interpretation it is assumed that the preferences of an updating individual will have changed since the last revision. This is represented as a fresh draw of the payoff matrix $\widetilde{\Lambda}$. For small noise this is a reasonable assumption. In the results, behaviour in the limit as $\sigma \to 0$ is examined, corresponding exactly to this case. However, the analysis also considers non-vanishing σ. Here, the procedure can be justified by noting that players are more likely to revise their strategy whenever their preferences change. In actual fact, this is a reasonable interpretation for any size of σ, and so is the preferred one.

A second issue is the assumption that only a single player revises their strategy each period. In KMR [8], however, all players in the population revise simultaneously. Lone revisions are more realistic. To see this, embed the model in a continuous time framework. Consider individual players revising periodically according to an underlying Poisson process. In this case, during any small period of time, at most one revision will be observed with high probability. This is the approach of Binmore and Samuelson [3].

2.3. **An Illustrative Example.** A simple example is outlined which will later illustrate the main results. In spirit, this follows a leading example of Kandori and Rob [9]. Consider a population of n academics in a research institution. All members use personal computers (PCs) to conduct their work, and may adopt either the IBM or Apple Macintosh (Mac) standards. Institution members interact during the course of their work, and receive payoffs according to their PC and that adopted by their colleagues. The strategic form game for mean payoff agents is:

	IBM	Mac
IBM	5, 5	4, 2
Mac	2, 4	6, 6

The payoffs are chosen as a stylised representation of the following criteria: Players benefit from compatibility; given compatibility, Mac adoption results in higher productivity than IBM adoption; the loss from incompatibility is less severe for IBM users than Mac users due to wider outside support for the IBM standard. This game has two pure strategy Nash equilibria corresponding to the two standards. Note that although Mac is payoff dominant, IBM is risk-dominant ($x^* = \frac{2}{5} < \frac{1}{2}$). The game is thus a Rousseau [12] stag-hunt.

Using the formal notation of Section 2.1, this becomes the trembled stage game $\mathcal{G}_{PC} = \langle \Lambda_{PC}, \Psi_{PC}, \sigma \rangle$. Of particular interest are the cases $\Psi_{PC} \in \{\Psi_B, \Psi_U\}$ where:

$$\Psi_B = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \text{ and } \Psi_U = \begin{bmatrix} 0 & \frac{3}{4} \\ 2 & 0 \end{bmatrix}$$

corresponding to balanced and unbalanced respectively. The second configuration is specified to reflect the increased risk of being a lone Mac user.

Each period a randomly selected member leaves the institution and a replacement chooses their computer system. As an alternative interpretation, researchers periodically replace their PCs. At replacement time, however, it is assumed that the requirements of the researcher are different from those at the time of the original purchase, and hence payoffs are a fresh draw from the payoff distribution.

3. Analysis

The analysis begins by examining the play of a new entrant, and concludes that the adaptive response dynamic is an irreducible Markov chain on the state space Z. Using familiar graph-theoretic techniques, the invariant distribution is characterised. Limits are taken as heterogeneity falls. Finally the case of non-vanishing heterogeneity is analysed.

3.1. **Entrant Response and Markov Properties.**

Consider a new entrant. Represent the fraction of the $n-1$ incumbents playing strategy 1 as x. The payoffs from strategies 1 and 2 are respectively:

$$
\begin{aligned}
y_1 &= x\tilde{a} + (1-x)\tilde{b} = xa + (1-x)b + x\varepsilon_a + (1-x)\varepsilon_b \\
y_2 &= x\tilde{c} + (1-x)\tilde{d} = xc + (1-x)d + x\varepsilon_c + (1-x)\varepsilon_d
\end{aligned}
$$

so that:

$$
y_1 - y_2 = x(a-c) - (1-x)(d-b) + x(\varepsilon_a - \varepsilon_c) + (1-x)(\varepsilon_b - \varepsilon_d) \tag{1}
$$

An entrant chooses strategy 1 whenever $y_1 > y_2$, or equivalently $y_1 - y_2 > 0$. Rearranging Equation (1) this occurs whenever:

$$
x(\varepsilon_c - \varepsilon_a) + (1-x)(\varepsilon_d - \varepsilon_b) < x(a-c) - (1-x)(d-b) \tag{2}
$$

The left hand side of Equation (2) is normally distributed with zero mean and variance $\sigma^2\left(x^2\left(\sigma_a^2+\sigma_c^2\right)+(1-x)^2\left(\sigma_b^2+\sigma_d^2\right)\right)$. Facing frequency x, the probability that the entrant responds with strategy 1 is:

$$
\Pr[1|x] = \Phi\left(\frac{[x(a-c)-(1-x)(d-b)]}{\sigma\sqrt{x^2\left(\sigma_a^2+\sigma_c^2\right)+(1-x)^2\left(\sigma_b^2+\sigma_d^2\right)}}\right)
$$

where Φ represents the standard Gaussian distribution. If i of the incumbents play strategy 1, then $x = i/(n-1)$. The notation is simplified by the following definition:

Definition 6. *Define the* basin depth *as κ_i^2 where κ_i satisfies:*

$$
\kappa_i = \frac{[i(a-c)-(n-i-1)(d-b)]}{\sqrt{i^2\left(\sigma_a^2+\sigma_c^2\right)+(n-i-1)^2\left(\sigma_b^2+\sigma_d^2\right)}} \tag{3}
$$

Using this notation, strategies 1 and 2 are chosen with probabilities $\Phi(\kappa_i/\sigma)$ and $\Phi(-\kappa_i/\sigma)$ respectively.

The best response probabilities depend only on the strategy frequency among incumbents. By assumption, the player leaving each period is chosen at random. The new strategy frequency thus depends only on the previous state. Since the normal distribution has full support, either strategy may chosen. Hence any state may be reached in a finite number of steps with positive probability. Summarising:

Proposition 1. *The adaptive response dynamic is a homogeneous Markov process on the finite state space Z. Moreover, it has a unique invariant (ergodic) distribution.*

Proof. No states are transient. Standard results show that the process has a unique invariant distribution — see for instance Theorem 11.2 of Stokey and Lucas [13]. ∎

The transition matrix P corresponds to the Markov chain on Z. The following lemma gives a convenient characterisation.

Lemma 1. *The transition probabilities p_{ij} of the Markov matrix P satisfy:*

$$p_{ij} = \left\{ \begin{array}{ll} [i/n]\,\Phi\left(-\kappa_{i-1}/\sigma\right) & j = i-1 \\ [i/n]\,\Phi\left(\kappa_{i-1}/\sigma\right) + \left[(n-i)/n\right]\Phi\left(-\kappa_i/\sigma\right) & j = i \\ \left[(n-i)/n\right]\Phi\left(\kappa_i/\sigma\right) & j = i+1 \end{array} \right\} \tag{4}$$

and are zero elsewhere.

Proof. Start in state i. The process cannot move to $j < i-1$ or $j > i+1$, since there is a lone replacement. A move to state $i+1$ requires the loss of a strategy 2 player and the gain of a strategy 1 player. The former occurs with probability $(n-i)/n$. The best response of an entrant to the frequency $i/(n-1)$ is strategy 1 with probability $\Phi\left(\kappa_i/\sigma\right)$. A similar argument establishes the cases $j = i$ and $j = i-1$. ∎

Example 1. *The complete Markov matrix for the case $n = 3$ and $\sigma = 1$ is:*

$$P = \frac{1}{3}\left[\begin{array}{llll} 3\Phi\left(-\kappa_0\right) & 3\Phi\left(\kappa_0\right) & 0 & 0 \\ \Phi\left(-\kappa_0\right) & \Phi\left(\kappa_0\right) + 2\Phi\left(-\kappa_1\right) & 2\Phi\left(\kappa_1\right) & 0 \\ 0 & 2\Phi\left(-\kappa_1\right) & 2\Phi\left(\kappa_1\right) + \Phi\left(-\kappa_2\right) & \Phi\left(\kappa_2\right) \\ 0 & 0 & 3\Phi\left(-\kappa_2\right) & 3\Phi\left(\kappa_2\right) \end{array} \right]$$

3.2. **The Ergodic Distribution.** The Markov matrix is characterised by its main, sub and super diagonals. Given this convenient tridiagonal form, the stationarity equation $\mu P = \mu$ may be solved for the ergodic distribution. Following previous work, however, the graph-theoretic approach of Freidlin and Wentzell [5] is employed. These techniques were used by KMR [8], Young [14] and Ellison [4] among others, and allow an immediate and intuitive closed solution.

The Freidlin and Wentzell [5] approach constructs a directed graph on the state space Z with edge weights corresponding to Markov transition probabilities. The directed edge set $E \subseteq Z \times Z$, has weights $p : E \mapsto R^+$, where the first and second coordinates represent source and target nodes respectively. A *tree* rooted at z is a set of edges $h \subseteq Z \times Z$ such that each node $i \neq z$ has a unique successor. All sequences of edges lead to z, which has no successor. The collection of trees rooted at z is H_z. The weight of such a tree h is then:

$$w_h = \prod_{(i,j)\in h} p_{ij}$$

Following a standard notation, sum over all trees rooted at z to obtain:

$$q_z = \sum_{h \in H_z} w_h \tag{5}$$

At each step of the Markov chain, a route opens from each node to another. This yields a directed edge set on the state space. Restricting to rooted trees gives route sets which eventually lead to a specified node z. Appendix B gives a derivation and explanation of the following lemma, due to Freidlin and Wentzell [5, Chapter 6, Lemma 3.1] and briefly discussed in KMR [8].

Lemma 2. *The invariant distribution μ satisfies:*

$$\mu_z = \frac{q_z}{\sum_{z' \in Z} q_{z'}} = \frac{\sum_{h \in H_z} \prod_{(i,j) \in h} p_{ij}}{\sum_{z' \in Z} \sum_{h \in H_{z'}} \prod_{(i,j) \in h} p_{ij}}$$

This lemma provides an immediate closed form for the invariant distribution. The relative weights of any two states in this distribution may be assessed by considering the ratio $q_z/q_{z'}$. Appendix B also contains more detailed discussion of this technique.

In the context of the adaptive response dynamic, Equation (5) takes a simple form. Since the transition matrix P is tridiagonal, observe that associated to each state there is a unique positively weighted rooted tree.

Lemma 3. *For the adaptive response dynamic:*

$$q_z = \prod_{0 \le i < z} p_{i(i+1)} \prod_{z < i \le n} p_{i(i-1)}$$

Proof. Begin at state 0. With positive probability the directed graph can remain at 0 (creating a loop) or proceed to state 1. From state 1, returning to state 0 again creates a loop, and hence the only positively weighted path entails proceeding to node 2. Node z is reached eventually. A symmetric argument holds beginning at node n. ∎

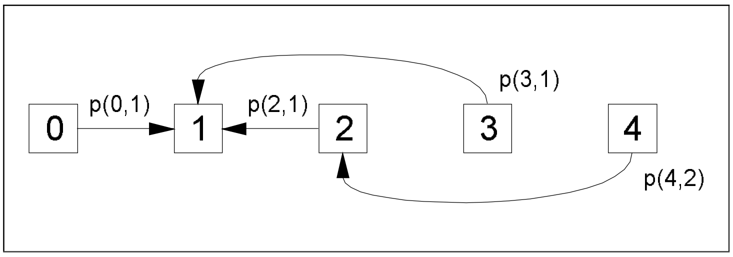

FIGURE 1. Tree Rooted at $z = 1$ with Zero Weight

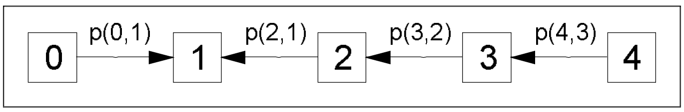

FIGURE 2. Tree Rooted at $z = 1$ with Positive Weight

Figures 1 and 2 illustrate this result graphically. Figure 1 depicts a tree for the case $n = 4$ rooted at $z = 1$ with zero weight. Two of the edges leapfrog nodes, resulting in a zero element in the product of the transition probabilities. Figure 2 shows the unique positively weighted rooted tree for this node.

Combine Lemma 3 with the transition probabilities from Equation (4) to obtain:

$$q_z = \frac{1}{n^n} \prod_{0 \le i < z} (n-i)\Phi\left(\frac{\kappa_i}{\sigma}\right) \prod_{z < i \le n} i\Phi\left(-\frac{\kappa_{(i-1)}}{\sigma}\right) \tag{6}$$

This characterisation simplifies further for the extreme states n and 0. These correspond to the products of super and sub-diagonal elements respectively.

Lemma 4. *At the extreme states $z \in \{0, n\}$, the invariant distribution satisfies:*

$$\frac{q_n}{q_0} = \prod_{0 \le i \le n-1} \frac{\Phi(\kappa_i/\sigma)}{\Phi(-\kappa_i/\sigma)} = \prod_{0 \le i \le n-1} \frac{\Phi(\kappa_i/\sigma)}{1 - \Phi(\kappa_i/\sigma)} \tag{7}$$

Proof. The power and factorial terms cancel. Re-index the equation to yield (7). ∎

3.3. Long-Run Equilibrium with Vanishing Heterogeneity.

Although Equations (6) and (7) give convenient forms for the invariant distribution, tighter results are available in the limit. Previous authors have considered the limit of the ergodic distribution as the probability of a mistake falls to zero. The analogue in this model is vanishing heterogeneity, taking $\sigma \to 0$. Exact conditions are established for the selection of an equilibrium with vanishing trembles.

Focus throughout will be on the case of two pure strategy Nash equilibria. Recall that the optimal response for a mean payoff entrant is to play strategy 1 when facing a frequency $x > x^*$. For convenience define $i^* = \lceil x^*(n-1) \rceil$, the least integer i such that $i \ge x^*(n-1)$. A mean payoff entrant plays strategy 1 when $i \ge i^*$ incumbents play strategy 1. As $\sigma \to 0$ the system places more weight on certain states. The following terminology formalises this notion:

Definition 7. *For vanishing heterogeneity,* z dominates z' *whenever* $\lim_{\sigma \to 0} q_z / q_{z'} = \infty$.

The following lemma will also be of use:

Lemma 5. $\{\kappa_i\}_{i=0}^{n-1}$ *satisfy a single crossing property. Generically* $\kappa_i > 0 \Leftrightarrow i \ge i^*$.[7]

Proof. The sign of κ_i is determined by the numerator of (3), and the result holds. ∎

Consider the recurrent classes of an unperturbed Markov chain. It is standard that for suitably perturbed processes the ergodic distribution focuses all weight on these classes as perturbations go to zero. The recurrent classes correspond to the extreme states $z \in \{0, n\}$. Hence a mixture state is dominated for vanishing heterogeneity:

Proposition 2. $z \in \{0, n\}$ *dominates any mixture state for vanishing heterogeneity.*

Proof. For $z \ge i^*$ use Equation (6) to form q_n/q_z. Re-indexing the denominator:

$$\lim_{\sigma \to 0} \frac{q_n}{q_z} = \lim_{\sigma \to 0} \frac{\prod_{z \le i < n} (n-i) \Phi(\kappa_i/\sigma)}{\prod_{z < i \le n} i \Phi\left(-\kappa_{(i-1)}/\sigma\right)} = \lim_{\sigma \to 0} \prod_{z \le i < n} \frac{(n-i)}{(i+1)} \frac{\Phi(\kappa_i/\sigma)}{1 - \Phi(\kappa_i/\sigma)}$$

Now $\kappa_i > 0$ and hence $\Phi(\kappa_i/\sigma) \to 1$ as $\sigma \to 0$. The denominator in each term tends to zero, and hence $\lim_{\sigma \to 0} q_n/q_z = \infty$. For $z < i^*$, compare q_z to q_0. ∎

The main result determines which of the extreme states is dominant. Former models have shown that the strategy with the *widest* basin of attraction is selected. In the current model the *depth* of the basin, κ_i^2, also plays a key rôle. Therefore, the basin *volume* is of critical importance:

[7] All the results hold for non-generic cases, in which $\kappa_{i^*} = 0$.

Definition 8. *Define the* basin volume *for strategies* 1 *and* 2 *respectively as:*

$$B_1^n = \frac{1}{n} \sum_{i^* \le i \le n-1} \kappa_i^2 \quad \textit{and} \quad B_2^n = \frac{1}{n} \sum_{0 \le i < i^*} \kappa_i^2$$

The equilibrium selected is determined by the relative weight of the extreme rooted trees. The weight of such a tree is the product of one step transition probabilities. Each of these probabilities is a cumulative normal term Φ or $1 - \Phi$. These may be rewritten as the ratio of a normal density and a hazard rate. Employing asymptotic properties of these elements, the main result obtains:

Proposition 3. *For vanishing heterogeneity, the strategy with the largest basin volume is selected. That is, strategy* 1 *dominates strategy* 2 *if and only if* $B_1^n > B_2^n$.

Proof. Separate the product of Equation (7) to obtain:

$$\frac{q_n}{q_0} = \frac{\prod_{i \ge i^*} \Phi\left(\kappa_i/\sigma\right)}{\prod_{i<i^*} \left(1 - \Phi\left(\kappa_i/\sigma\right)\right)} \frac{\prod_{i<i^*} \Phi\left(\kappa_i/\sigma\right)}{\prod_{i \ge i^*} \left(1 - \Phi\left(\kappa_i/\sigma\right)\right)}$$

Using the single crossing property from Lemma 5, the numerator and denominator of the left hand term tend to unity. Strategy selection is thus determined by the limit of the right hand term:

$$\lim_{\sigma \to 0} \frac{\prod_{i<i^*} \Phi\left(\kappa_i/\sigma\right)}{\prod_{i \ge i^*} \left(1 - \Phi\left(\kappa_i/\sigma\right)\right)}$$

Notice that both numerator and denominator tend to zero. Re-write this ratio as:

$$\frac{\prod_{i<i^*} \Phi(\kappa_i/\sigma)}{\prod_{i \ge i^*} \left(1 - \Phi(\kappa_i/\sigma)\right)} = \frac{\prod_{i<i^*} \phi(\kappa_i/\sigma)}{\prod_{i \ge i^*} \phi(\kappa_i/\sigma)} \times \frac{\prod_{i \ge i^*} \phi(\kappa_i/\sigma)/\left(1 - \Phi(\kappa_i/\sigma)\right)}{\prod_{i<i^*} \phi(\kappa_i/\sigma)/\Phi(\kappa_i/\sigma)}$$

The first term is explicitly:

$$\frac{\prod_{i<i^*} \phi(\kappa_i/\sigma)}{\prod_{i \ge i^*} \phi(\kappa_i/\sigma)} = (2\pi)^{(n-2i^*)/2} \exp\left(-\frac{\sum_{i<i^*} \kappa_i^2 - \sum_{i \ge i^*} \kappa_i^2}{2\sigma^2}\right) \tag{8}$$

Consider a typical element in the numerator of the second term:

$$\lim_{\sigma \to 0} \frac{\phi(\kappa_i/\sigma)}{1 - \Phi(\kappa_i/\sigma)} = \lim_{\sigma \to 0} \frac{\kappa_i}{\sigma} \tag{9}$$

Recall the hazard rate $\phi(u)/(1 - \Phi(u))$ of the normal distribution. By a standard result this is asymptotically linear as $u \to \infty$.[8] Since $\kappa_i > 0$ for $i \ge i^*$, (9) holds. An identical argument holds for the denominator, where $\kappa_i < 0$. Hence:

$$\lim_{\sigma \to 0} \frac{\prod_{i<i^*} \phi(\kappa_i/\sigma)/\Phi(\kappa_i/\sigma)}{\prod_{i \ge i^*} \phi(\kappa_i/\sigma)/\left(1 - \Phi(\kappa_i/\sigma)\right)} = \lim_{\sigma \to 0} \frac{\sigma^{i^*}}{\sigma^{n-i^*}} \frac{\prod_{i<i^*} \kappa_i}{\prod_{i \ge i^*} (-\kappa_i)}$$

which is polynomial in σ. The first term (8) is exponential in σ. The limit becomes:

$$\lim_{\sigma \to 0} \frac{q_n}{q_0} = \lim_{\sigma \to 0} \left(\frac{\sqrt{2\pi}}{\sigma}\right)^{n-2i^*} \exp\left(-\frac{\sum_{i<i^*} \kappa_i^2 - \sum_{i \ge i^*} \kappa_i^2}{2\sigma^2}\right) \frac{\prod_{i<i^*} \kappa_i}{\prod_{i \ge i^*} (-\kappa_i)} \tag{10}$$

[8] This result is reproduced in Appendix A.2 as Lemma 6.

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

$$\lim_{\sigma\to 0}\frac{q_n(\sigma)}{q_0(\sigma)} = +\infty \Leftrightarrow \sum_{i\geq i^*}\kappa_i^2 > \sum_{i<i^*}\kappa_i^2$$

which is exactly the required inequality. ∎

Equation (10) involves polynomial and exponential terms in σ. If the selected equilibrium (for example, strategy 1) is risk-dominant, then both terms diverge to $+\infty$, reinforcing selection. If the selected equilibrium is not risk-dominant ($n < 2i^*$), however, then these terms compete. Therefore, a smaller σ is required to approximate the limit. For the two player case the result is particularly clear.

Corollary 1. *For $n = 2$ strategy 1 dominates for vanishing heterogeneity whenever it is generalised-risk-dominant. Equivalently:*

$$\frac{d-b}{\sqrt{\sigma_b^2+\sigma_d^2}} < \frac{a-c}{\sqrt{\sigma_a^2+\sigma_c^2}} \tag{11}$$

Proof. If $n = 2$ then $i^* = 1$. The basin volumes are then:

$$B_1^n = \frac{\kappa_1^2}{2} = \frac{1}{2}\frac{(a-c)^2}{\sigma_a^2+\sigma_c^2} \quad \text{and} \quad B_2^n = \frac{\kappa_0^2}{2} = \frac{1}{2}\frac{(d-b)^2}{\sigma_b^2+\sigma_d^2}$$

which upon inspection yield the desired inequality. ∎

From Equation (11) it is clear that in a game with balanced trembles ($\psi = \frac{1}{2}$), equilibrium $(1,1)$ is selected if and only if $x^* < \frac{1}{2}$. This corresponds exactly to risk-dominance. The following corollary shows that this holds for the general n player population. The proof is relegated to Appendix A.3.

Corollary 2. *Consider a coordination game with balanced trembles. Strategy 1 dominates for vanishing heterogeneity if and only if it is risk-dominant.*

For balanced trembles this reaffirms the results of previous authors. However, the selection of a risk-dominant equilibrium may fail via the introduction of asymmetry in payoff heterogeneity. In Section 4 an illustration of this property is provided.

These results are particularly applicable to small populations, where transition times are low. Intuitively, a small number of contrarian entrants is sufficient to tip the process toward the opposing basin of attraction. With larger populations transition times can escalate. Formally, however, the results of this section continue to hold. Moreover, the basin volume condition of Proposition 3 is well approximated by an integral for sufficiently large n. The following definition is useful:

Definition 9. *The* asymptotic basin depth *is $\kappa(x)^2$, where:*

$$\kappa(x) = \frac{[x(a-c)-(1-x)(d-b)]}{\sqrt{x^2\left(\sigma_a^2+\sigma_c^2\right)+(1-x)^2\left(\sigma_b^2+\sigma_d^2\right)}}$$

Hence $\kappa_i = \kappa(i/(n-1))$. Similarly:

Definition 10. *The* asymptotic basin volumes *for strategies 1 and 2 respectively are:*

$$B_1^\infty = \int_{x^*}^{1} \kappa(x)^2 dx \quad \textit{and} \quad B_2^\infty = \int_0^{x^*} \kappa(x)^2 dx$$

Notice that $\lim_{n\to\infty} B_1^n = B_1^\infty$ and $\lim_{n\to\infty} B_2^n = B_2^\infty$. For sufficiently large populations, strategy 1 is selected whenever $B_1^\infty > B_2^\infty$. The explicit solutions to these integrals are given in Appendix A.4. They depend only on the payoff balance λ and the tremble balance ψ.

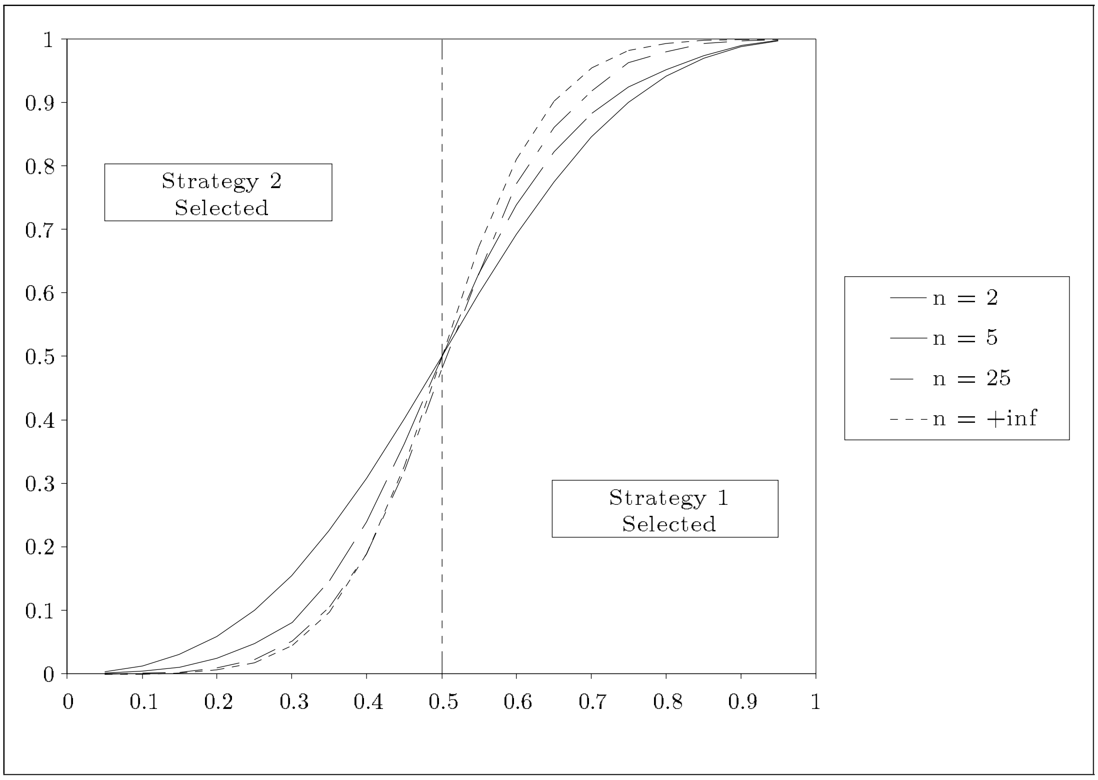

FIGURE 3. Strategy Selection Regions and Boundaries for Varying n.

Figure 3 illustrates the selection criterion for various population sizes. A trembled stage game $\mathcal{G}$ is characterised by its payoff balance λ and its tremble balance ψ. The line plotted for each n represents $B_1^n = B_2^n$. To the right of this line strategy 1 is selected. Strategy 1 is risk-dominant to the right of the vertical line. It is generalised-risk-dominant to the right of the line for $n = 2$. Notice that if a strategy is both risk-dominant and generalised-risk-dominant then it is selected. Risk-dominance continues to provide a reasonable basis for selection, but the generalised dominance property is needed for sufficiency. There are regions of the balance space for which a risk-dominant strategy is not selected, irrespective of population size. Selection depends upon n in the region between the lines for $n = 2$ and $n = \infty$.

3.4. **Long-Run Equilibrium with Non-Vanishing Heterogeneity.** As the Ellison [4] critique makes clear, taking mutations (or in this case, heterogeneity) to the limit greatly amplifies the transition times between the long-run equilibria. This is particularly acute for large populations. A large number of idiosyncratic entrants is required to tip the population out of a basin of attraction. For small noise and large populations, initial conditions are a more appropriate focus of attention. Hence, in this section heterogeneity is fixed at some $\sigma > 0$. The ergodic distribution is characterised for larger n, as this allows a convenient integral approximation to the tree weights. Following this method, the modes of the distribution are calculated and the shape of the distribution is examined.

Recall that the ergodic distribution is determined by the relative weights of rooted trees. Define $x_i = i/n$. The weight q_z of the unique rooted tree at node z is given by Equation (6). Denoting $q_z = q(x_z, n)$:

$$\begin{aligned} q(x,n) &= \prod_{0 \le i < xn} (1-x_i)\Phi\left(\frac{\kappa(x_i)}{\sigma}\right) \prod_{nx < i \le n} x_i \Phi\left(-\frac{\kappa(x_i)}{\sigma}\right) \\ &= \exp\left(\sum_{0 \le i < xn} \log(1-x_i)\Phi\left(\frac{\kappa(x_i)}{\sigma}\right) + \sum_{xn < i \le n} \log x_i \Phi\left(-\frac{\kappa(x_i)}{\sigma}\right)\right) \end{aligned}$$

The summations in the exponent are well approximated by integrals for larger n. Hence, fixing n, $q(x,n)$ is approximately $q(x)$:

$$q(x) = \exp\left(n\left[\int_0^x \log(1-y)\Phi\left(\frac{\kappa(y)}{\sigma}\right)dy + \int_x^1 \log y \Phi\left(-\frac{\kappa(y)}{\sigma}\right)dy\right]\right) \tag{12}$$

Differentiating to find the extrema of the ergodic distribution obtains:

Proposition 4. *For larger n, the extrema of the ergodic distribution correspond exactly to the Bayesian Nash equilibria of the trembled stage game $\mathcal{G}$.*

Proof. Facing a frequency x, a player's optimal response is strategy 1 with probability $\Phi(\kappa(x)/\sigma)$. Bayesian Nash equilibria correspond to fixed points of $x \to \Phi(\kappa(x)/\sigma)$. Consider now the extrema of Equation (12). The first derivative is zero when:

$$\log(1-x)\Phi\left(\frac{\kappa(x)}{\sigma}\right) = \log x\left(1 - \Phi\left(\frac{\kappa(x)}{\sigma}\right)\right) \Leftrightarrow x = \Phi(\kappa(x)/\sigma)$$

Thus the first order condition is satisfied at fixed points of $\Phi(\kappa(x)/\sigma)$. ∎

Proposition 5. *Downcrossings of $\Phi(\kappa(x)/\sigma)$ correspond to local maxima, upcrossings to local minima.*

Proof. Evaluating the second derivative of $q(x)$ at a fixed point gives:

$$q''(x) \gtrless 0 \Leftrightarrow \left(\frac{1}{x} + \frac{1}{1-x}\right)\left(\frac{\kappa'(x)}{\sigma}\phi\left(\frac{\kappa(x)}{\sigma}\right) - 1\right) \gtrless 0$$

The sign is determined by the second term in this expression. A downcrossing (upcrossing) of $\Phi(\kappa(x)/\sigma)$ occurs exactly when this term is negative (positive). ∎

Hence, to examine the modes of the ergodic distribution it is sufficient to analyse the Bayesian Nash Equilibria of $\mathcal{G}$. The following proposition is taken from Myatt and Wallace [11], and establishes some of the properties of these equilibria. The proof is contained in Appendix A.5.

Proposition 6. *For σ sufficiently small there are three Bayesian Nash equilibria. For σ sufficiently large there is a single equilibrium.*

Heuristically, suppose that strategy 1 is risk-dominant, so that $x^* < \frac{1}{2}$. Then there is always at least one fixed point above x^*. The other fixed points (if any) lie below x^*. To see this, notice that $\kappa(x^*) = 0$. Hence $\Phi(\kappa(x^*)/\sigma) = \frac{1}{2}$ for all σ. Clearly, $\Phi(\kappa(x^*)/\sigma) > x^*$, and thus there is a fixed point above x^*. This is apparent from Figure 4 which plots $\Phi(\kappa(x^*)/\sigma)$ and its fixed points for two different values of σ.

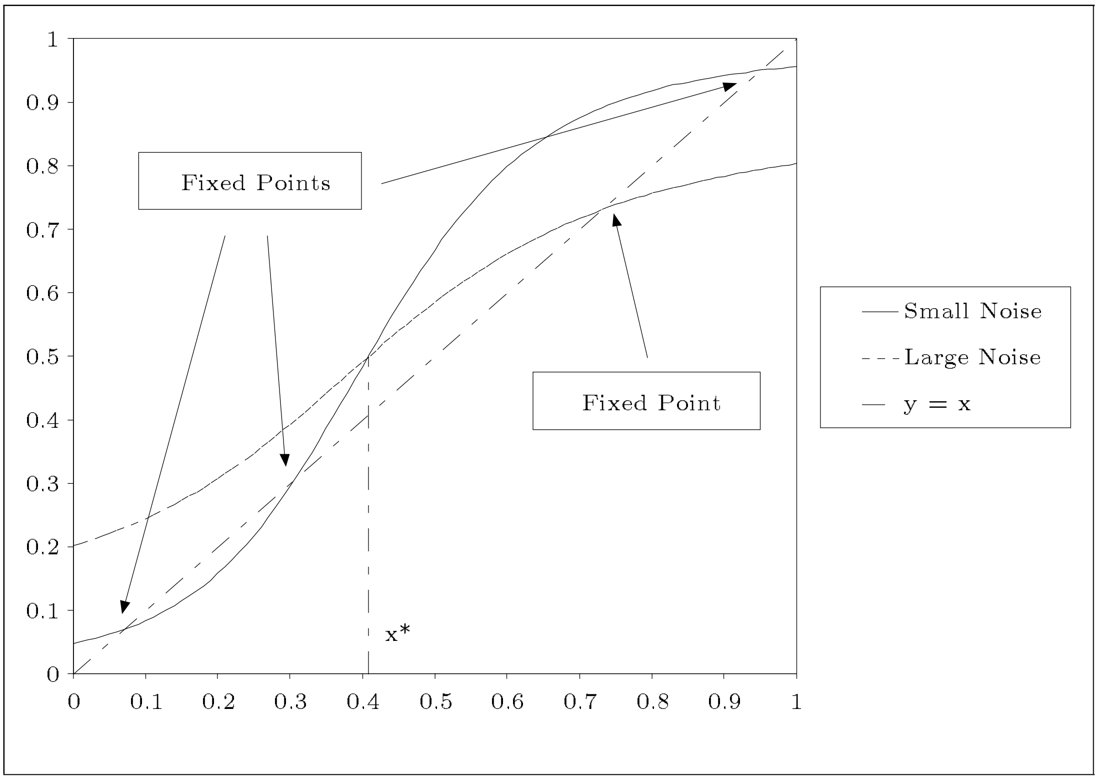

FIGURE 4. Fixed Points of $\Phi(\kappa(x)/\sigma)$.

The downcrossings of $\Phi(\kappa(x)/\sigma)$ correspond to the modes of $q(x)$. An upcrossing corresponds to a local minimum. When strategy 1 is risk-dominant there is a single mode to the right of x^* for large σ. For small σ there is also a local minimum and a mode to the left. Thus, for sufficiently small σ, the ergodic distribution is bimodal. These modes correspond to the two extreme Bayesian Nash equilibria, which in turn correspond to the pure strategy Nash equilibria as heterogeneity vanishes. There is a local minimum between the two modes, corresponding to the third Bayesian Nash equilibrium. As heterogeneity vanishes, this converges to the mixed strategy Nash equilibrium of the unperturbed stage game. As σ grows two of the Bayesian Nash equilibria eventually disappear. The ergodic distribution becomes unimodal. The only mode to survive is the one to the right of x^*. Recall that this is the mode that corresponds to the risk-dominant equilibrium as heterogeneity vanishes. These properties are best discussed with reference to an example. In Section 4 ergodic distributions are plotted for the examples introduced in Section 2.3.

4. Discussion

Here the results of the previous section are discussed and illustrated. Section 4.1 considers the case of vanishing heterogeneity. Basins of attraction are graphed for both balanced and unbalanced trembles, using the example $\mathcal{G}_{PC}$ of section 2.3. Section 4.2 considers this same example for non-vanishing heterogeneity. Graphical illustrations clearly demonstrate the connection between ergodic distributions and Bayesian Nash equilibria.

4.1. **Vanishing Heterogeneity.** Consider first the case of balanced trembles (Ψ_B). Following Corollary 2, the risk-dominant equilibrium (IBM) is selected for vanishing heterogeneity.

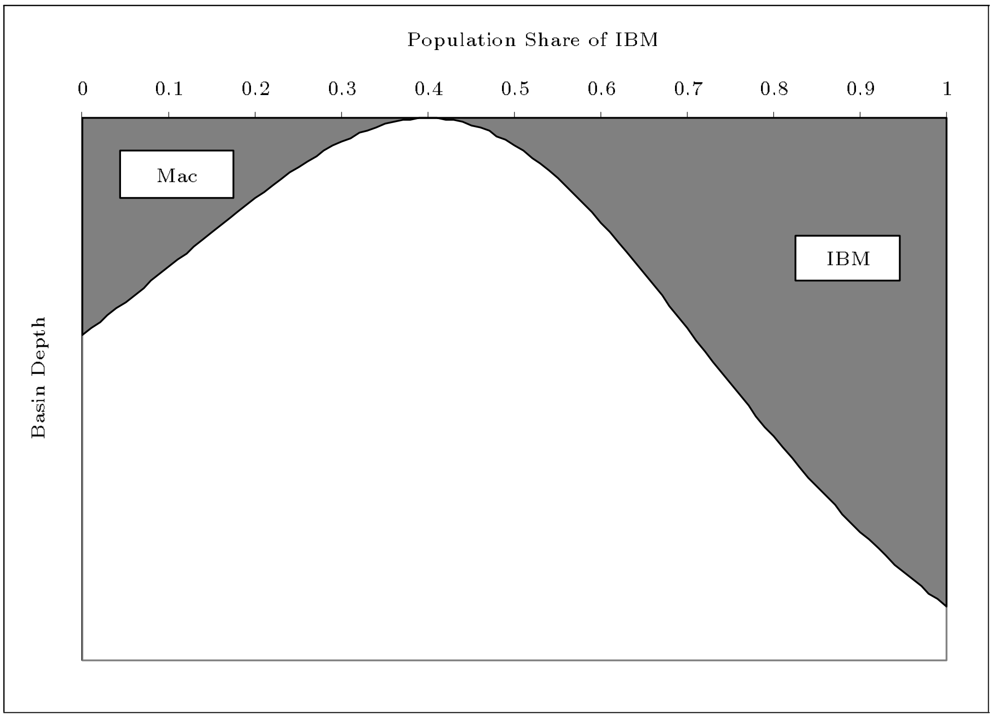

Figure 5. Asymptotic Basins of Attraction for Balanced Trembles Ψ_B

The asymptotic basins of attraction are illustrated in Figure 5, clearly showing that IBM has the larger basin and is thus dominant. Following the analysis so far, it is clear that the balanced tremble PC adoption game is equivalent to a pure coordination game with:

$$\Lambda'_{PC} = \begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix} \quad \text{and} \quad \Psi'_B = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

Using this formulation, an intuitive explanation for the result is available. First focus on a two player population. Consider a researcher entering an institution with an incumbent IBM user. For the entrant to adopt Mac, the researcher's idiosyncratic preference for lone Apple use must exceed 3. If the IBM user exits at the end of the period, the process tips into the Mac equilibrium. Beginning with an incumbent

Mac user, however, the idiosyncratic preference of the entrant has only to exceed 2 to enable the tip.

The conclusions are reversed for the case of unbalanced trembles (Ψ_U). The payoff-dominant equilibrium (Mac) is selected for vanishing heterogeneity, for all population sizes. Note that this is the generalised-risk-dominant equilibrium.

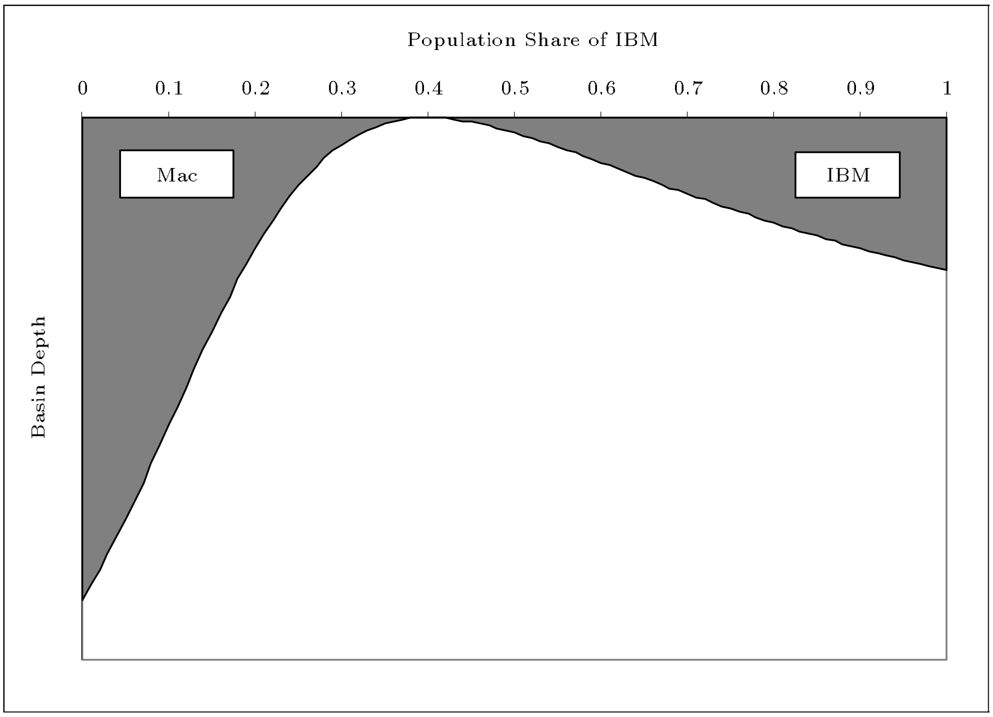

FIGURE 6. Asymptotic Basins of Attraction for Unbalanced Trembles Ψ_U.

As can be seen from Figure 6, although the IBM basin is wider, the Mac basin is far deeper, leading to Mac dominance. Once again, this result may be understood using a transformation to a pure coordination game:

$$\Lambda'_{PC} = \begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix} \quad \text{and} \quad \Psi'_U = \begin{bmatrix} 0 & 0.75 \\ 2 & 0 \end{bmatrix}$$

A new entrant facing an IBM incumbent still has a high hurdle to jump in order to adopt the Mac standard. However, lone Mac users are now more idiosyncratic, so that the probability of an entrant selecting against the incumbent standard is higher for IBM incumbents than for Mac. It is thus easier to tip out of an IBM equilibrium.

Heterogeneity may be reinterpreted in a mutation framework. Recall that in state $i > i^*$ a mean-payoff entrant adopts strategy 1. In this model a "mutation" occurs when the entrant plays strategy 2. This is not a mistake; rather the player has sufficiently idiosyncratic preferences. This happens with probability $1 - \Phi(\kappa_i/\sigma)$.

Proposition 7. *The model generates state-dependent "mutations".*

Proof. For states $j > i > i^*$, consider the ratio of mutation probabilities:

$$\frac{1-\Phi(\kappa_i/\sigma)}{1-\Phi(\kappa_j/\sigma)} = \frac{\phi(\kappa_j/\sigma)/(1-\Phi(\kappa_j/\sigma))}{\phi(\kappa_i/\sigma)/(1-\Phi(\kappa_i/\sigma))}\exp\left(-\frac{\kappa_i^2-\kappa_j^2}{2\sigma^2}\right)$$

Let $\sigma \to 0$. As in Proposition 3, the hazard terms are asymptotically linear, hence the exponential term dominates. Generically $\kappa_i^2 \neq \kappa_j^2$ and so the exponential term is 0 or ∞ in the limit. Thus "mutations" do not converge to zero at the same rate. ∎

Bergin and Lipman [2] have shown that the presence of state-dependent mutations can influence the equilibrium selected. In order to make useful predictions therefore, they suggest that a reasonable economic model should underlie the "mutation" process. This is such a model.

4.2. **Non-Vanishing Heterogeneity.** The selection results discussed in the previous section provide good predictions for small populations. Figure 7 plots various ergodic distributions for the PC adoption game with unbalanced trembles. For $n = 5$ as $\sigma \to 0$ the distribution quickly places large weight on the selected strategy.

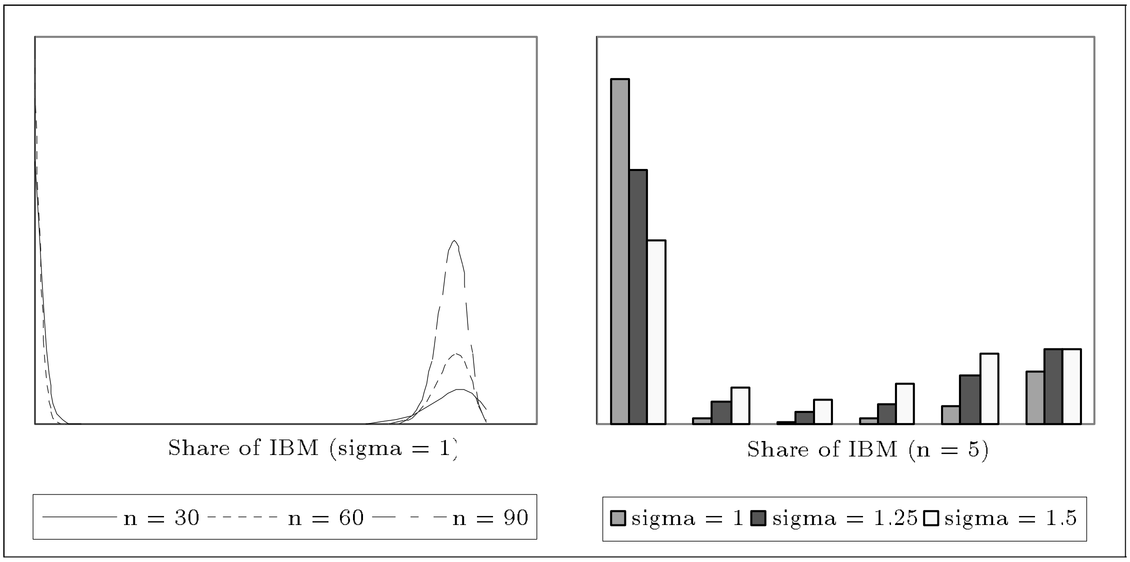

FIGURE 7. Ergodic Distributions for Ψ_U.

For larger populations, however, the limiting results are less useful. For $\sigma = 1$, allowing the population to grow shifts weight back towards the risk-dominant equilibrium. The results of Section 3.4 are more applicable under these circumstances.

Figure 8 fixes the population size at $n = 30$ in the PC adoption game with unbalanced trembles. The ergodic distribution is plotted for three values of σ. Below, the fixed points of $\Phi(\kappa_i/\sigma)$ are graphed. The close connection between the Bayesian Nash equilibria and the ergodic distribution as derived in Proposition 4 is illustrated. Upcrossings correspond to the modes and downcrossings to local minima. This has an intuitive explanation. Consider an incumbent population with frequency composition such that proportion x plays strategy 1, a Bayesian Nash equilibrium. An entrant has variable payoffs, and therefore plays strategy 1 with probability x. Thus the population composition is not expected to change. Furthermore, if the incumbent frequency is close to a downcrossing, then the response probability will be even closer. The population composition is thus expected to move toward the Bayesian Nash equilibrium. The opposite applies to an upcrossing. Hence the former are modes and the latter minima of the ergodic distribution.

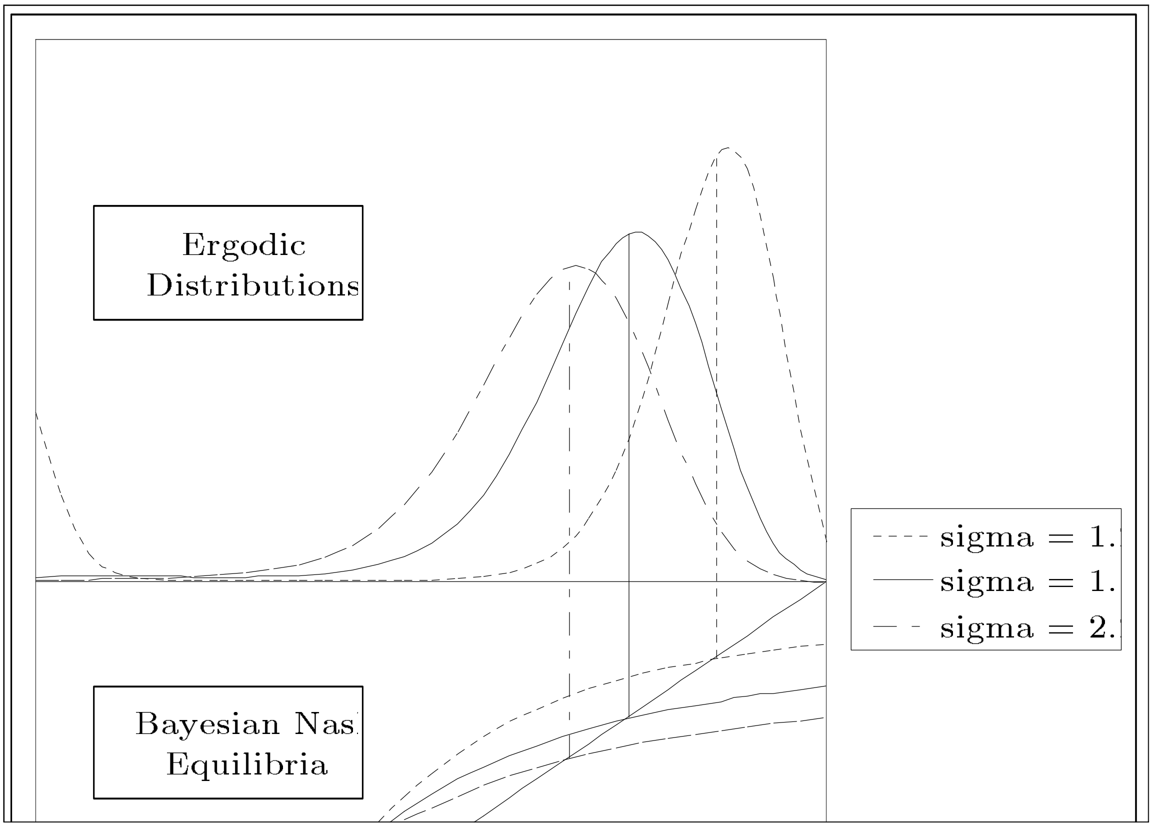

FIGURE 8. Ergodic Distributions and Bayesian Nash Equilibria ($n = 30$).

When risk-dominance coincides with generalised-risk-dominance (for example, when trembles are balanced), the selected equilibrium with vanishing heterogeneity coincides with that under non-vanishing heterogeneity. When risk-dominance and generalised-risk-dominance diverge, however, predictions are less clear-cut. Although the generalised-risk-dominant equilibrium may be selected for small σ or small n, when the population is large and noise is non-zero, risk-dominance is once again the criterion of choice. By inspection of Figures 7 and 8, the ergodic mode that lies closer to the risk-dominant equilibrium gains weight as the population increases. As σ grows, this is the only mode to survive.

5. Conclusion

This paper is motivated by the belief that individuals differ rather than err. Although players certainly make mistakes on occasion, different strategy selections in identical scenarios are more likely due to idiosyncratic preferences. Drawing on both noncooperative and evolutionary game theory, the adaptively rational approach taken by other authors is a profitable one. However, dropping the assumption that players maximise is premature in a modelling context.

A simple model of adaptive dynamics is posited. It captures the notion of differing preferences via the introduction of payoff trembles. By specifying reasonable parametric distributions for these trembles, a highly tractable model with sharp conclusions is obtained. In particular, the relative variability of payoffs can be critical to equilibrium selection in 2×2 symmetric coordination games. The model endogenously generates state-dependent "mutations" which converge to zero at different speeds in different states. Hence the critique of Bergin and Lipman [2] is applicable. The risk-dominant equilibrium need not be selected therefore, but in fact in the presence of reasonable economically derived assumptions, it continues to play an important rôle.

Previous models examine equilibrium selection as mutation rates tend to zero. Here, vanishing heterogeneity is analogous. For balanced trembles the results support those obtained by KMR [8] and Young [14], with the selection of the risk-dominant equilibrium. With unbalanced trembles, however, this may not be the case. An exact condition is derived for an equilibrium to be selected. It is not merely the *width* of attraction basins that is critical, but also the *depth* of such basins. In consequence, generalised-risk-dominance is a more appropriate criterion. It is best applied in scenarios with small populations where transition times are reasonable.

Ergodic distributions are examined for non-vanishing heterogeneity. For larger populations it is shown that the ergodic modes correspond to Bayesian Nash equilibria of the trembled stage game. Using this result for larger noise, risk-dominance becomes the key determinant of long-run behaviour, further strengthening the conclusions of earlier work.

The approach taken here maintains the assumption that agents maximise payoffs. Players are not fully rational however, as they play a best response to an incumbent frequency. This is reasonable, as players engage in repeated play in a slowly evolving population. KMR [8] may be interpreted similarly. If agents play only once following an observation of history (as in Young [14]), then this is less appropriate. In such an environment greater rationality on the part of the players might be expected. In related work, Myatt and Wallace [11] allow for greater sophistication. In that paper, players use history as a starting point for iterative reasoning. This process leads to the selection of a Bayesian Nash equilibrium contingent on the past. History, which provides the context for decision making, evolves as a result. The methods of Section 3 prove useful in the subsequent analysis.

In conclusion, for too long evolutionary game theory has followed its biological roots in economic application. Commonality of payoffs may be appropriate in a biological context. As Maynard Smith [10] argues: "...Darwinian fitness provides a natural and genuinely one-dimension scale." Uniform preferences are a less reasonable premise in an economic framework. State-independent mutation of strategies is also a sensible component of a biological model. Neither state independence nor "mutation" of strategies can be supported by empirical observation in economics. This paper constitutes an attempt to ground evolutionary results in economic assumption. Despite the change in modelling paradigm due to these alterations, the result is not so different. The consequence is a renewed emphasis on risk-dominance, and consensus is reached.

References

[1] R. Aumann and A. M. Brandenburger, *Epistemic conditions for Nash equilibrium*, Econometrica, 63 (1995), pp. 1161–80.

[2] J. Bergin and B. Lipman, *Evolution with state-dependent mutations*, Econometrica, 64 (1996), pp. 943–956.

[3] K. G. Binmore and L. Samuelson, *Muddling through: Noisy equilibrium selection*. May 1996.

[4] G. Ellison, *Learning, local interaction and coordination*, Econometrica, 61 (1993), pp. 1047–1071.

[5] M. I. Freidlin and A. D. Wentzell, *Random Perturbations of Dynamical Systems*, Springer-Verlag, Berlin/New York, 1984.

[6] J. Harsanyi and R. Selten, *A General Theory of Equilibrium Selection in Games*, MIT Press, Cambridge, 1988.

[7] J. C. Harsanyi, *Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points*, International Journal of Game Theory, 2 (1973), pp. 1–23.

[8] M. Kandori, G. J. Mailath, and R. Rob, *Learning, mutation and long-run equilibria in games*, Econometrica, 61 (1993), pp. 29–56.

[9] M. Kandori and R. Rob, *Evolution of equilibria in the long run: A general theory and applications*, Journal of Economic Theory, 65 (1995), pp. 383–414.

[10] J. Maynard Smith, *Evolution and the Theory of Games*, Cambridge University Press, 1982.

[11] D. P. Myatt and C. Wallace, *Sophisticated play with payoff heterogeneity*, research note, Nuffield College, February 1997.

[12] J.-J. Rousseau, *The inequality of man*, in Rousseau's Social Contract and Discourses, G. Cole, ed., J. M. Dent, London, 1913, pp. 157–246. First published 1755.

[13] N. L. Stokey, R. E. Lucas, and E. C. Prescott, *Recursive Methods in Economic Dynamics*, Harvard University Press, Cambridge, 1989.

[14] H. P. Young, *The evolution of conventions*, Econometrica, 61 (1993), pp. 57–84.

Appendix A. Omitted Results

This appendix provides proofs omitted from the main text.

A.1. **Correlated Trembles.** It is clear that the trembled stage game is equivalent to a pure coordination game with payoff and tremble matrices:

$$\Lambda = \begin{bmatrix} a & 0 \\ 0 & d \end{bmatrix} \quad \text{and} \quad \Psi = \begin{bmatrix} \sigma_a & 0 \\ 0 & \sigma_d \end{bmatrix}$$

Retain the earlier assumptions, except correlate ε_a and ε_d with coefficient ρ. Facing a frequency x, an entrant responds with strategy 1 whenever:

$$(1-x)\varepsilon_d - x\varepsilon_a < xa - (1-x)d$$

The left hand side is now distributed as:

$$N\left(0, \sigma^2\left(x^2\sigma_a^2 + (1-x)^2\sigma_d^2 - 2x(1-x)\rho\sigma_a\sigma_d\right)\right)$$

Basin depth κ_i^2 now satisfies $\kappa_i = \hat{\kappa}\left(i/(n-1)\right)$ where:

$$\hat{\kappa}(x) = \frac{xa-(1-x)d}{\sqrt{x^2\sigma_a^2 + (1-x)^2\sigma_d^2 - 2x(1-x)\rho\sigma_a\sigma_d}}$$

Using this generalisation, all the results continue to hold.

A.2. **Asymptotic Linearity of Normal Hazards.**

Lemma 6. *The hazard $\phi(x)/(1-\Phi(x))$ is asymptotically linear as $x \to \infty$.*

Proof. Applying l'Hôpital's rule:

$$\lim_{x\to\infty}\left\{\frac{\phi(x)}{1-\Phi(x)} - x\right\} = \lim_{x\to\infty}\frac{\phi(x)-x(1-\Phi(x))}{1-\Phi(x)} = \lim_{x\to\infty}\frac{\Phi(x)-1}{-\phi(x)} = 0$$

which gives the desired result. ∎

A.3. **Corollary 2 — Selection with Balanced Trembles.**

Proof. For simplicity of exposition, and without loss of generality, take $c = b = 0$. With balanced trembles, the basin volume condition becomes:

$$\sum_{0\le i<i^*}\frac{(ia-(n-i-1)d)^2}{i^2+(n-i-1)^2} < \sum_{i^*\le i\le n-1}\frac{(ia-(n-i-1)d)^2}{i^2+(n-i-1)^2}$$

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

$$\sum_{0\le i<i^*}\frac{(ia-(n-i-1)d)^2}{i^2+(n-i-1)^2} < \sum_{0\le j\le n-1-i^*}\frac{((n-j-1)a-jd)^2}{j^2+(n-j-1)^2}$$

If strategy 1 is risk-dominant, then $a > d$ and $i^* < n-1-i^*$. The condition becomes:

$$\sum_{0\le i<i^*}\frac{(ia-(n-i-1)d)^2-((n-i-1)a-id)^2}{i^2+(n-i-1)^2} < \sum_{i^*\le i\le n-1-i^*}\frac{(ia-(n-i-1)d)^2}{i^2+(n-i-1)^2}$$

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

$$(ia - (n-i-1)d)^2 - ((n-i-1)a - id)^2 = (i^2 - (n-i-1)^2)(a^2 - d^2)$$

This is negative since $a^2 > d^2$ and $i^2 < (n-i-1)^2$ for $i < i^*$. In conclusion then; the inequality holds, yielding sufficiency. Necessity follows since if $a < d$ the same procedure establishes the dominance of strategy 2. ∎

A.4. **Asymptotic Basin Volume.** In Section 3.3, the asymptotic basin volume is considered. Recall $B_1^\infty = \int_{x^*}^1 \kappa(x)^2 dx$ and $B_2^\infty = \int_0^{x^*} \kappa(x)^2 dx$. Using the terminology of Section 2, the asymptotic basin depth of Definition 9 is proportional to:

$$\widetilde{\kappa}(x)^2 = \frac{(\lambda x - (1-\lambda)(1-x))^2}{\psi x^2 + (1-\psi)(1-x)^2} \tag{13}$$

Using this formulation, an explicit form for the integral $\int \widetilde{\kappa}(x)^2 dx$ is available.

Proposition 8. *The basin volume satisfies:*

$$\int \widetilde{\kappa}(x)^2 dx = x + \frac{\lambda^2 - \psi - 2\lambda\psi + 2\psi^2}{\sqrt{\psi(1-\psi)}} \arctan\left\{\frac{x-(1-\psi)}{\sqrt{\psi(1-\psi)}}\right\} + (\lambda - \psi)\log\left(\psi x^2 + (1-\psi)(1-x)^2\right) + \textit{constant}$$

Proof. Differentiate the above to obtain the expression from Equation (13). ∎

Using this explicit form, the basin volumes B_1^∞ and B_2^∞ may be calculated. The plot of Figure 3 for $n = +\infty$ is obtained by solving $B_1^\infty = B_2^\infty$ for values λ and ψ.

A.5. **Proposition 6 — Fixed Points of $\Phi(\kappa(x)/\sigma)$.**

Proof. Fixed points of $\Phi(\kappa(x)/\sigma)$ correspond to roots of:

$$f(x) = \Phi\left(\frac{\kappa(x)}{\sigma}\right) - x$$

Notice that $f'(x) = \phi(\kappa(x)/\sigma)\,\kappa'(x)/\sigma - 1$. As $\sigma \to \infty$, $f'(x) \to -1$ uniformly for $x \in [0,1]$. Thus, for sufficiently large σ, $f(x)$ is decreasing everywhere. Hence $f(x)$ only has one root local to $x = \frac{1}{2}$.

When $\sigma \to 0$, $f(x) \to 1 - x$ if $x > x^*$ and $f(x) \to -x$ if $x < x^*$, so there cannot be a fixed point unless it is local to $\{0, x^*, 1\}$. Consider the interval $0 \le x \le \varepsilon$. 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 $1 - \varepsilon \le x \le 1$.

Now consider $x^* - \varepsilon \le x \le x^* + \varepsilon$. 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. ∎

Appendix B. Ergodic Distributions and Rooted Trees

Here a brief explanation of the graph-theoretic methods used in Section 3 is provided.

B.1. **Maps between Markov states.** Consider a homogeneous Markov chain on the state space $Z = \{0, 1, \ldots, n\}$ with generic member z. Represent the transition probabilities by the $(n+1) \times (n+1)$ Markov matrix $P = [p_{ij}]$, satisfying $p_{ij} \geq 0$ and $\sum_j p_{ij} = 1$ for all i, where as usual $p_{ij} = \Pr[z_{t+1} = j | z_t = i]$.

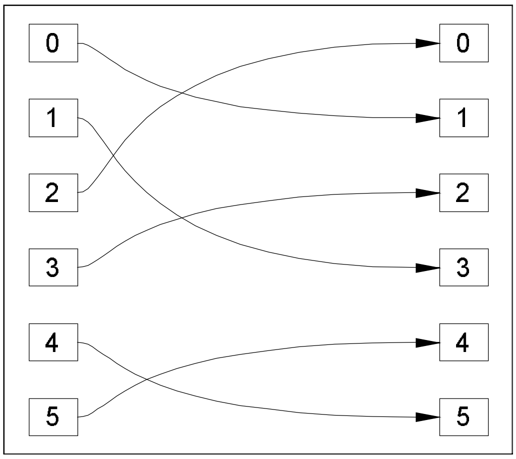

Figure 9. A function $g : Z \mapsto Z$ for $n = 5$

Construct a function $g : Z \mapsto Z$ from the state space into itself, and denote the set of all such functions as G. An event at time t may be viewed as a selection of a map from the set G, as illustrated in Figure 9. Taking the power set 2^G as the appropriate algebra, a finite measure may be constructed on this space:

$$\mu[g] = \prod_{i \in Z} p_{ig(i)}$$

A notion of weight is immediate.

Definition 11. *Define the weight of a function $g \in G$ by the measure $\mu[g]$.*

This extends to give $\mu_{ij} = \mu\{g : g(i) = j\}$, the total weight of all functions which map i to j. Using this:

$$\mu_{ij} = \sum_{g:g(i)=j} \prod_{k \in Z} p_{kg(k)} = p_{ij} \sum_{g:g(i)=j} \prod_{k \neq i} p_{kg(k)}$$

Clearly, the following lemma obtains.

Lemma 7. *The measure μ on $\{G, 2^G\}$ satisfies $p_{ij} = \mu_{ij} / \sum_k \mu_{ik}$.*

Proof. Since the terms of the summation do not depend on the restriction of $g(i)$:

$$\frac{\mu_{ij}}{\mu_{ik}} = \frac{p_{ij}}{p_{ik}} \frac{\sum_{g:g(i)=j} \prod_{k \neq i} p_{kg(k)}}{\sum_{g:g(i)=k} \prod_{j \neq i} p_{jg(j)}} = \frac{p_{ij}}{p_{ik}}$$

The result clearly follows. ∎

B.2. **Graph Representation.** The map g may be re-interpreted graph-theoretically. Construct a directed graph on the nodes formed by the Markov states. Then G represents the set of all directed graphs such that each node has a unique successor. This re-formulation is illustrated in Figure 10, where the graph corresponding to the function g of Figure 9 is illustrated. Similarly, $\mu[g]$ is the weight of the tree g.

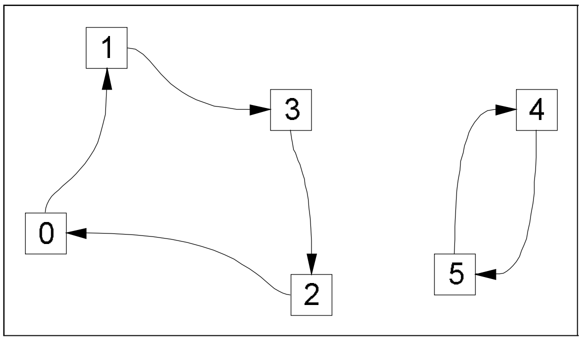

FIGURE 10. Graph representation of $g : Z \mapsto Z$

Notice, however, that the graph of Figure 10 contains two cycles. Hence the nodes present in the limit depend on the starting node. This is not true, however, if a restriction to directed paths with a unique closed loop is imposed.

Definition 12. *Define G_z as the set of directed graphs with unique successors for each node, such that the unique closed loop contains the state z.*

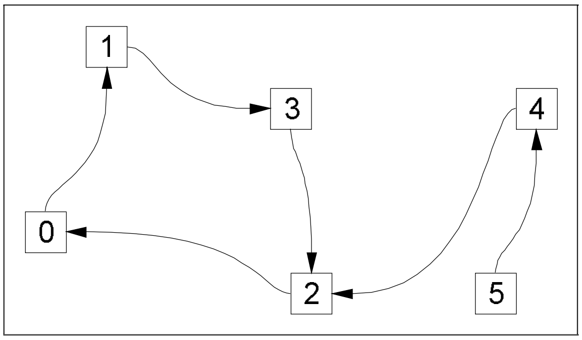

FIGURE 11. Graph with a unique closed loop containing states $0, 1, 2, 3$

In Figure 11 modifications are made to Figure 10, giving a graph with a single closed loop. Note that this closed loop contains the nodes 0, 1, 2 and 3 and hence $g \in G_z$ for z in this range. Intuitively, at each time coordinate in the process, an event occurs corresponding to $g \in G$. If $g \in G_z$ then the resulting path leads toward state z. Consider now the weight of $g \in G_z$:

$$\mu[g] = \prod_{i \in Z} p_{ig(i)} = p_{zg(z)} \prod_{i \neq z} p_{ig(i)}$$

Notice that the product in this second expression is the weight of the graph formed by dropping the edge $z \longrightarrow g(z)$. This yields a *tree rooted at* z; starting from any node the directed edges form a path leading to node z. Denote such a tree as:

$$h \in H_z \Leftrightarrow h = \{(i, g(i)) : i \neq z.g \in G_z\}$$

Thus H_z gives the set of trees rooted at z. This leads to:

$$\mu(h) = \prod_{i \neq z} p_{ig(i)}$$

giving the weight associated with the tree. Furthermore, the total weight of all trees rooted at z is:

$$q_z = \sum_{h \in H_z} \prod_{i \neq z} p_{ih(i)}$$

where the notation $h(i)$ is equivalent to $g(i)$ with the exception of node z, for which it is not defined. It is easy to see, then, that corresponding to any $g \in G_z$ there is a unique tree rooted at z such that:

$$\mu\,[g] = p_{zg(z)} \prod_{i \neq z} p_{ih(i)}$$

Moreover, any tree rooted at z may be extended to an appropriate member of G_z by adding an edge from z. This is illustrated in Figure 12. The solid edges depict a tree rooted at node $z = 2$, and the broken edges represent possible extensions to form a member of G_z.

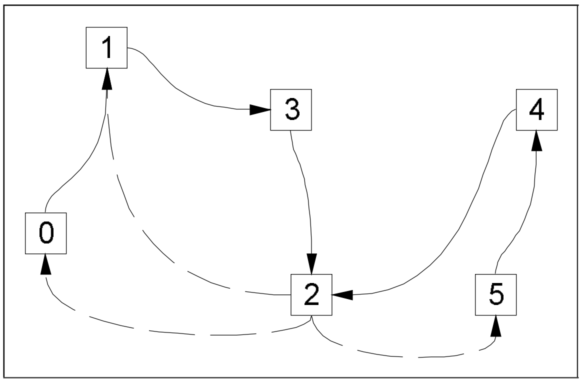

FIGURE 12. A tree $h \in H_2$ rooted at $z = 2$ and some extensions

Clearly, such an extension procedure may be undertaken for any tree rooted at node z. Hence:

$$G_z = \bigcup_{h \in H_z} \bigcup_{i \neq z} \{h \cup (z, i)\}$$

To obtain the weight of all graphs in G_z, sum to obtain:

$$\mu\,[G_z] = \sum_{g \in G_z} \mu\,[g] = \sum_{h \in H_z} \sum_{i \neq z} p_{zi} \prod_{i \neq z} p_{ih(i)} = \sum_{i \neq z} p_{zi} q_z$$

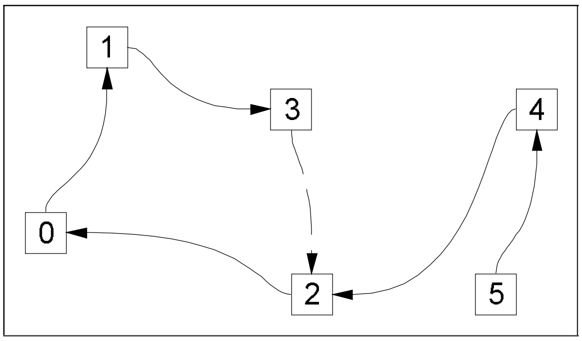

FIGURE 13. A tree $h \in H_3$ rooted at $z = 3$ and an extension to $g \in G_2$

A tree was obtained from $g \in G_z$ by deleting an edge leading *from* z. Notice than an alternative would be to remove an edge leading *to* z. For instance:

$$\mu[g] = \prod_{i \in Z} p_{ig(i)} = p_{ig(i)} \prod_{z \neq i} p_{zg(z)}$$

It is clear, then, that G_z may be constructed by considering sets of trees rooted at other nodes, and the adding an edge from the root node of each of these trees to the node z. This idea is illustrated in Figure 13. Thus:

$$G_z = \bigcup_{i \neq z} \bigcup_{h \in H_i} \{h \cup (i, z)\}$$

and:

$$\mu[G_z] = \sum_{g \in G_z} \mu[g] = \sum_{i \neq z} \sum_{h \in H_i} p_{iz} \prod_{k \neq i} p_{kh(k)} = \sum_{i \neq z} p_{iz} q_i$$

Notice then that $\sum_{i \neq z} p_{zi} q_z = \sum_{i \neq z} p_{iz} q_i$. This becomes:

$$q_z(1 - p_{zz}) = \sum_i p_{iz} q_i - p_{zz} q_z \Leftrightarrow q_z = \sum_i p_{iz} q_i$$

Forming the vector $q = [q_z]$ this expression is $q = Pq$. But then q is a scaled ergodic distribution. The following lemma (due to Freidlin and Wentzell [5]) is obtained.

Lemma 8. *The ergodic distribution of the Markov chain μ^* satisfies:*

$$\mu_z = \frac{q_z}{\sum_i q_i} = \frac{\sum_{h \in H_z} \prod_{k \neq z} p_{kh(k)}}{\sum_i \sum_{h \in H_i} \prod_{j \neq i} p_{jh(j)}}$$

Notice that an ergodic process is required, since otherwise all rooted trees would have zero weight, and the above expression would be ill-defined.

NUFFIELD COLLEGE, OXFORD OX1 1NF, UK
E-mail address: david.myatt@nuf.ox.ac.uk

NUFFIELD COLLEGE, OXFORD OX1 1NF, UK
E-mail address: christopher.wallace@nuf.ox.ac.uk