DOMINATED STRATEGIES AND EQUILIBRIUM SELECTION

DAVID P. MYATT AND CHRIS WALLACE*

January 2000

Abstract. A rational player will never play a strictly dominated strategy. It might be tempting therefore to eliminate such strategies from any subsequent analysis. However, if equilibrium selection is an issue it may be wrong to do so. In models of adaptive learning with state-independent mutations, (after KMR, 1993) global risk-dominance is a sufficient condition for selection in larger games. This paper presents an alternative perspective. State-dependent "mutations" are endogenously generated by a reasonable underlying economic framework where players are different (following Myatt and Wallace, 1997). Agents play best responses to frequency observations in m strategy, 2 player games. As the idiosyncrasy in the population vanishes, dominated strategies do not lose their significance. It is always possible to add a strategy to a coordination game which (a) is strictly dominated, (b) retains the global risk-dominance of the originally selected equilibrium but (c) results in the selection of an alternative equilibrium. Dominated strategies are important features of coordination games not to be discarded at once. They can provide a novel and intuitive solution to the equilibrium selection problem.

"I am glad of all details," remarked my friend,
"whether they seem to you to be relevant or not."

Sherlock Holmes in The Adventure of the Copper Beeches

1. Introduction

1.1. Dominated Strategies. A rational player will never play a strictly dominated strategy. Consequently, a strictly dominated strategy can play no part in a Nash equilibrium profile. Indeed, only mutual knowledge of rationality is required for the deletion of such strategies from a normal form game. Such strategies may thus be deemed irrelevant details of a game, and are neglected in subsequent analysis by many theorists.[1]

Date Printed. January 21, 2000 (Version 3.0).
JEL Subject Classification. C72 and C73.
Keywords. Dominated Strategies, Equilibrium Selection, Global Risk-Dominance, Idiosyncrasy.
*St Catherine's College, Oxford OX1 3UJ, United Kingdom, david.myatt@economics.ox.ac.uk and Nuffield College, Oxford OX1 1NF, United Kingdom, christopher.wallace@nuf.ox.ac.uk respectively. This paper appears in an amended form as Chapter 3 of Myatt's doctoral thesis, Oxford 1999. The authors would like to thank Peyton Young and Kevin Roberts for helpful comments and discussion.

[1] The survival of strictly dominated strategies in an evolutionary context is considered by Hofbauer and Weibull (1996) and Palomino (1995).

Common knowledge of rationality allows the iterated deletion of dominated strategies.[2] Such a procedure may still leave an embarassment of riches, however, with multiple Nash equilibria in the rationalisable set. An equilibrium selection problem is posed. The analyst might then restrict focus to the reduced game, constructing an appropriate mechanism to highlight a particular equilibrium as the most likely outcome. Such a restriction is premature. Whereas the deleted strategies can play no part in the equilibria themselves, might it not be possible for such strategies to determine selection between equilibria? This question provides the motivation for the current work. This paper re-evaluates the rôle played by dominated strategies in equilibrium selection. It concludes that dominated strategies may have far-reaching consequences for selection. In particular, the addition of a dominated strategy may preserve the global risk-dominance of an equilibrium while ensuring its deselection. The analyst should indeed be glad of all details when examining a game.

1.2. Equilibrium Selection and Mistaken Mutations. Recent research has been active in tackling the selection of equilibria. Key papers such as Kandori, Mailath and Rob (1993, KMR) and Young (1993) model the dynamics of boundedly rational agents. The construction of an evolutionary process is the central tool. A finite population is specified, in which players are either replaced or revise their strategies periodically. Entering players are boundedly rational, and adopt a best response to incumbent strategy frequencies. Such dynamics are path-dependent. To see this, consider a population commonly adopting a symmetric pure strategy Nash equilibrium profile. An entrant to the population will respond with the same strategy. To establish ergodicity, these authors perturb the associated Markov process to obtain an irreducible chain. These perturbations die away, and the ergodic distribution is characterised in the limit.

This research programme draws heavily on the evolutionary literature. The mutations of biological game-theoretic models are interpreted as mistakes by entering players, yielding the appropriate perturbations. Typically, an entrant fails to play a best response to an observed frequency with small probability, instead adopting one of the other strategies available. The requirement of payoff maximisation intrinsic to Nash play is thus weakened considerably. The approach does allow dominated strategies to play a part in the selection process, however. For instance, suppose that an entire population adopts a particular symmetric pure Nash profile. Whereas a strictly dominated strategy can never be a best response, it may be accidentally chosen by a mistaken player. If a sufficiently large number of errants also make such a choice, then the best response to the new frequency

[2] For more on what is necessary, see the recent book by Rubinstein (1998).

may lead to a third strategy. The mistaken choices may drag the population away from one equilibrium, with the best response structure pulling it in towards another.

This mechanism may be called into doubt via a deeper consideration of the rôle of mistakes. A player may mistakenly choose the incorrect response for a number of reasons. For instance, the player may have incorrectly observed the frequency to which a response must be made. Alternatively, the player may have observed the correct frequency, but calculated the best response incorrectly. Notice that the first reason may have consequences for selection among the rationalisable set. However, a dominated strategy can never be a best response, and hence incorrect observation of opposing frequencies can never lead to its selection. The second reason may still provide a justification for the strategy choice. An intelligent player, however, may be expected to recognise the dominated nature of the strategy, and hence would be most unlikely to err in this way. This conjecture is particularly acute if the strategy is dominated by another pure strategy rather than a mixture – it must surely be valid to expect that any intelligent player would immediately discount the use of such an action. In summary, it is unlikely that a sophisticated agent would select a dominated strategy, calling into question this approach.

The device of mistaken mutations is also subject to the critique of Myatt and Wallace (1997) and (1999). Those papers argue that the relaxation of the maximisation hypothesis is premature. The key requirement for ergodicity is that there be positive probability of contrarian behaviour by population entrants. A contrarian may be judged to have chosen the "wrong" strategy relative to the payoffs specified in the normal form. A normal form game, however, is an abstraction from reality and is better interpreted as a representation of the mean payoffs to an agent. Any particular player is likely to have idiosyncratic preferences. This leads to a static Bayesian game as the appropriate stage component. A player may then "mutate" not because of a mistake, but rather because the player is different. An appropriate best response may thus appear as a "mistake" to the modeller. A formal model is an idealisation of a strategic interaction. Players doubtless both differ and err; however, from a modelling perspective it is more appropriate to focus on the latter characteristic. The approach of Myatt and Wallace (1997) and (1999) is to introduce "trembles" to the payoffs of entrants.

The modelling paradigm of idiosyncratic agents immediately justifies the retention of dominated strategies. Adding payoff trembles with full support, a particular action can never be strictly dominated – there is always the possibility that an entering idiosyncrant will find an action optimal to play. Moreover, such payoff trembles lead immediately to endogenously generated state-dependent "mutations". A strategy may yield payoffs close to those of another strategy which dominates the former in mean payoffs. Only mild idiosyncrasy on the part of an entrant is thus required to adopt the latter. In such a case there is a strong and justified mechanism whereby the dominated strategy may drag a population away from a particular equilibrium. Once again, the best response structure of the game may then direct the population to another equilibrium.

1.3. The Model and Results. The model explored here builds upon that of Myatt and Wallace (1997), extending the ideas to the case of two player, m strategy symmetric games. The base component is a trembled normal form game, play of which evolves according to the Adaptive dynamic.

In the Adaptive Play dynamic, entrants to a population play a best response to the incumbent strategy frequency. The version considered here follows the previous work in allowing the population to update one player at a time. The graph theoretic methods of Freidlin and Wentzell (1984) are employed to characterise the ergodic distribution. This requires calculation of the transition probabilities. Since Gaussian idiosyncrasy is employed, the transition probabilities are multivariate integrals of the normal distribution. A key result shows that these have an asympotically analytic form in the tails. Using this as noise vanishes, the main proposition of the paper follows. That is, take any 2×2 coordination game with selection of the risk-dominant equilibrium. One may add a third strategy that is (a) dominated, (b) retains the global risk-dominance of the original selected equilibrium but (c) results in the deselection of this equilibrium.

It is worth commenting on the results in the light of the Bergin and Lipman (1996) critique. The model here generates a perturbed Markov process. These authors show that perturbing an underlying non-ergodic process in an appropriate manner, any equilibrium may be selected. Essentially, if the mistake probabilities in the KMR-Young framework vary by state, then deselection of a risk-dominant equilibrium can occur. A fully general model of this kind can never produce conclusive results. The particular form of "mutations" in a model must be justified. In this paper the "mutations" arising from idiosyncrasy are reasonable. The rejection of global risk-dominance arises from a justified modelling framework, rather than arbitrary specification of state-dependent mutations.

1.4. Outline of the Paper. The argument proceeds as follows. Section 2 gives motivation to the subsequent analysis, using a simple example. The model is described in Section 3, using a simplified version of the Adaptive Play model from Myatt and Wallace (1997). Analysis begins in Section 4 with the characterisation of choice probabilities for entering players. The ergodic distribution for the Markov process of the Adaptive dynamic is examined in Section 5. A discussion of the results concludes.

2. Motivation

The arguments of this paper are best understood with reference to a simple example. Consider the following symmetric coordination game.

	1	2
1	3, 3	0, 0
2	0, 0	2, 2

(1)

The pure Nash equilibria are clearly $\{1;1\}$ and $\{2;2\}$. As a pure coordination game, the payoff-dominant equilibrium $\{1;1\}$ is also risk-dominant. To see this, a strategy frequency placing weight $2=5$ or greater on strategy 1 will yield strategy 1 as a best response. Consider now a world in which all members of a population play strategy 2. Only a minority of the population need switch to strategy 1 before best response behaviour draws the remainder of the population in. In contrast, destabilisation of an all-1 population requires a majority (3=5) of the agents to switch. In this sense, the $\{1;1\}$ equilibrium is more robust to deviations. Indeed, the models of KMR (1993) and Young (1993) specify dynamics in which the population is subject to turnover, and entrants play a best response. With a small probability of deviant choice, the ergodic behaviour of such dynamics focus almost all weight on equilibrium $\{1;1\}$ of the game (1).

The simple example of (1) may be extended with the addition of a dominated strategy:

	1	2	3
1	3, 3	0, 0	0, −100
2	0, 0	2, 2	2, −100
3	−100, 0	−100, 2	−100, −100

(2)

The third strategy always yields a penalty of -100, and so a rational player would never choose it. Notice, however, that if an opponent were to adopt strategy 3, then the best response would be strategy 2. This provides an alternative route via which a population might evolve from equilibrium $\{1;1\}$ to $\{2;2\}$. Mistaken players may accidentally choose either of strategies 2 or 3. If sufficient deviants do so, the best response process leads back to $\{2;2\}$. Essentially, in a population playing $\{1;1\}$ there are two ways for players to err. This is not the case for a population playing $\{2;2\}$ however. Accidental play of strategy 3 pushes the population back to strategy 2 once more. The mistaken play of strategy 1 is the only way to escape from $\{2;2\}$.

Inspecting the payoffs of (2), notice that a $3/5$ majority of players are still required to deviate to destabilise $\{1;1\}$. In the words of Myatt and Wallace (1997), strategy 1 has the largest basin of attraction. Moreover, requiring a deviant majority to achieve destabilisation is exactly the concept of global risk-dominance proposed by Maruta (1997).[3] In fact, both Maruta (1997) and Ellison (1996) show that global risk-dominance is sufficient for selection of an equilibrium in the dynamics of KMR (1993) and Young (1993).

It might thus be tempting to delete the dominated strategy 3 from the game (2). First, this strategy can never be part of an equilibrium profile. Second, the global risk-dominance of $\{1;1\}$ continues to ensure selection. Third, the only added effect the strategy has is to provide an extra choice that a mistaken player may take. Inspecting the payoffs of strategy 3, it seems extremely unlikely that a player would make such a mistake. Strategy 3 is clearly inferior, and no misconception of the opposing strategy frequency could persuade an agent to play it.

Such a conclusion is incorrect, however. Consider the following modification of (2):

	1	2	3
1	$3,\ 3$	$0,\ 0$	$0,\ 2\frac{1}{2}$
2	$0,\ 0$	$2,\ 2$	$2,\ -\frac{1}{2}$
3	$2\frac{1}{2},\ 0$	$-\frac{1}{2},\ 2$	$-\frac{1}{2},\ -\frac{1}{2}$

(3)

This has essentially the same structure as (2). In particular, strategy 3 continues to be strictly dominated by strategy 1. The inferiority of strategy 3 is less clear cut, however. The payoff advantage of 1 is uniformly $1/2$. Discard now the biologically driven "mistaken mutations" of KMR (1993) and Young (1993). Instead, adopt the hypothesis that individual agents are in fact idiosyncratic. Such idiosyncrasy will be reflected in the heterogeneity of payoffs for any particular individual. An individual in an all-1 population would require only an idiosyncratic preference of $1/2$ for strategy 3 in order to choose it. It is rather more likely that such an event will occur. In an all-2 population, however, there must be an idiosyncratic preference for another strategy exceeding 2 for such deviant behaviour. Global risk-dominance continues to hold. However, despite the width of $\{1;1\}$'s basin of attraction, in an intuitive sense it is shallow; it is relatively easy for

[3] Global risk-dominance is equivalent to $1/2$-dominance. The latter terminology is used by Ellison (1996).

new players to ignore the incumbent norm and play strategy 3. Once a sufficient number of heterogeneous players have entered, any population playing this game is sucked into the $\{2,2\}$ equilibrium, which remains difficult to leave. This paper accepts the conclusions of Myatt and Wallace (1997) and (1999), which argue that idiosyncrasy is a better modelling approach than mistakes. In this scenario, the addition of a dominated strategy can have dramatic effects. More generally, it is clear that global risk-dominance is not an appropriate selection criterion for larger games.

3. The Model

The model presented here is an extension of the Adaptive Play model of Myatt and Wallace (1997) with some simplifying modifications. Whereas Myatt and Wallace (1997) and (1999) focus on 2×2 symmetric games, this paper considers more general m strategy symmetric two player games. Section 3.1 outlines the trembled stage game played by agents, followed by a review of the risk-dominance concepts important in the literature in Section 3.2. The Adaptive Play dynamic is described in Section 3.3.

3.1. The Trembled Stage Game. This model retains the trembled stage game used in earlier work. Reaching beyond the 2×2 scenario, however, the base is a two player symmetric strategic form game with m actions and generic payoffs:

	1	2		m
1	a_{11}, a_{11}	a_{12}, a_{12}	$\cdots$	a_{1m}, a_{m1}
2	a_{21}, a_{12}	a_{22}, a_{22}	$\cdots$	a_{2m}, a_{m2}
	$\vdots$	$\vdots$	$\ddots$	$\vdots$
m	a_{m1}, a_{1m}	a_{m2}, a_{2m}	$\cdots$	a_{mm}, a_{mm}

Naturally, this may be represented by a $m \times m$ matrix,

Definition 1. The mean payoff matrix is defined as:

$$\Pi = [a_{ij}]_{i,j \in \{0,1,\ldots,m\}} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mm} \end{bmatrix}$$

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

The payoffs Π will be viewed as the expected payoffs for any entering player. Fixing the payoffs of a static game of complete information is doubtless a simplification. Individual agents will have idiosyncratic payoffs. This is modelled via the addition of payoff trembles. Each payoff in the normal form is subject to an independent Gaussian disturbance.

Definition 2. The payoff heterogeneity matrix is defined as:

$$\Sigma = [\sigma_{ij}]_{i,j\in\{0,1,\ldots,m\}} = \begin{bmatrix} \sigma_{11} & \sigma_{12} & \cdots & \sigma_{1m} \\ \sigma_{21} & \sigma_{22} & \cdots & \sigma_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{m1} & \sigma_{m2} & \cdots & \sigma_{mm} \end{bmatrix}$$

An entrant has trembled payoff matrix $\tilde{\Pi}$ where:

$$\tilde{\Pi} = [\tilde{\pi}_{ij}]_{i,j\in\{0,1,\ldots,m\}} = [\pi_{ij} + \sigma\varepsilon_{ij}]_{i,j\in\{0,1,\ldots,m\}}$$

where $\varepsilon_{ij} \sim N\left(0, \sigma_{ij}^2\right)$, with $E[\varepsilon_{ij}\varepsilon_{kl}] = 0$ for $ij \neq kl$ and σ is a scaling factor.

The payoff heterogeneity matrix Σ determines the relative size of the payoff trembles. In the companion work, imbalance of trembles plays a crucial rôle in the destabilisation of risk-dominant equilibria. For the results presented here, this is unnecessary. The argument is upheld with a restriction to balanced trembles, and thus only the simpler balanced tremble case is considered.

Definition 3. Trembles are balanced in the game if $\sigma_{ij} = 1$ for all i, j.

Notice that the payoff disturbances have a fully-parametric form. This, however, is a natural representation of differing payoffs across players. In particular, differences over a particular payoff might be viewed as the resulting sum of many individual idiosyncratic factors, yielding the normal distribution as a natural specification. Furthermore, this formulation allows clear closed-form results to be obtained. As argued in Myatt and Wallace (1997), Bergin and Lipman (1996) show that full generality of trembles, particularly allowing trembles to vary by state, leads to inconclusive results, and hence the approach is justified.

Bringing together the elements thus far yields the trembled stage game.

Definition 4. Define the trembled stage game G as the triple:

$$G = \langle \Pi; \Sigma; \sigma \rangle$$

The trembled stage game is thus a static Bayesian game of incomplete information.

3.2. Risk-Dominance. Concepts of risk-dominance play an important rôle in much of the evolutionary literature. An equilibrium exhibiting such a property is typically robust to a large number of deviations by other players, or equivalently a significant probability of deviating from the prescribed equilibrium strategy profiles. The initial concept of risk-dominance is due to Harsanyi and Selten (1988). The following definition applies to 2×2 symmetric coordination games, where $a_{11} > a_{22}$.

Definition 5. Strategy 1 risk-dominates Strategy 2 whenever $a_{11} - a_{21} > a_{22} - a_{21}$.

Intuitively, the best reponse to a $50:50$ strategy frequency in a 2×2 symmetric coordination game is to play the risk-dominant strategy. Destabilisation of such an equilibrium requires the deviation of more than half of the population. The standard definition of risk-dominance may be extended to m strategy symmetric two-player games via pairwise comparison of equilibria. However, the risk-dominance relation may then be cyclic, and the extension to larger games fails. More general than such pairwise risk-dominance is global risk-dominance (Maruta 1997).

Definition 6. Strategy i is globally risk-dominant if it is a best response when $x_i \geq 1/2$:[4]

$$x_i \geq 1/2 \Rightarrow \sum_{k=1}^{m} x_k a_{ik} \geq \sum_{k=1}^{m} x_k a_{jk} \quad \forall j \neq i$$

Equivalently, strategy i is $1/2$-dominant.

A globally risk-dominant strategy is extremely robust to deviations. At least half of the population must abandon a prescribed profile in order for a rational player to abandon it in turn. Moreover, this allows for any deviation by the non-conformists; in particular, they may deviate in the worst possible way. The more general notion of p-dominance indexes the robustness of equilibria by specifying the fraction p of players required for destabilisation.[5]

Notice, however, that the robustness of risk-dominant equilibria is to the number of deviations. The likelihood of such deviations is not considered. In 2×2 games, Myatt and Wallace (1997) introduce the notion of generalised risk-dominance to reflect this.

3.3. The Adaptive Play Dynamic with Solo Revisions. The dynamic considered here is a mild simplification of that described in Myatt and Wallace (1997). Take a finite

[4] x_j represents the probability placed on pure strategy j in any given mixed strategy $x = (x_k)_{k=1}^{m}$.

[5] The notion of p-dominance is employed by Morris, Rob and Shin (1995). That work is epistemological, and finds that strongly p-dominant equilibria are selected. Generally, both the evolutionary and epistemological game theory literatures have focused on equilibria with similar qualities, including risk-dominance.

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. 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 $\tilde{A}$. This player observes the strategy distribution among the incumbents, prior to the exit of the leaving player, and selects a best response to this frequency. Notice that the modification of Myatt and Wallace (1997) is that the entrant's observation is made prior to the incumbent's exit, whereas previously the dynamic specified observation of the remaining $n-1$ incumbents. This is unimportant in the results that follow, but allows for easier notation.

4. Strategy Choice and Limiting Behaviour

Section 4.1 establishes strategy choice probabilities under best response. Section 4.2 then characterises the limiting behaviour of these probabilities as idiosyncrasy vanishes.

4.1. Entrant Response. Denote by $x \in \mathbb{R}^m$ a strategy frequency vector, satisfying $x_i \geq 0$ and $\sum_{i=1}^m x_i = 1$. The following definitions will be useful.

Definition 7. The normalised mean payoff of strategy i facing frequency x is:

$$\mu_i(x) = \frac{\sum_{j=1}^m x_j a_{ij}}{\sqrt{\sum_{j=1}^m x_j^2}}$$

Definition 8. Define the normalised mean payoff advantage of i over j as:

$$\delta_{ij}(x) = \mu_i(x) - \mu_j(x)$$

Consider now an entrant facing a strategy frequency x. The following is immediate.

Lemma 1. An entrant facing strategy frequency x adopts strategy i with probability:

$$\rho_i(x) = \int_{-\infty}^{\infty} \left\{ \prod_{j \neq i} \Phi\left(z + \frac{\delta_{ij}(x)}{\sigma} \right) \right\} \phi(z)\, dz$$

where Φ and ϕ denoted the Gaussian distribution and density functions respectively.

Proof. Facing x, the payoff from strategy i is $\sum_{j=1}^m x_j \tilde{a}_{ij}$. Strategy i is chosen whenever:

$$\pi_i = \sum_{j=1}^m x_j \tilde{a}_{ij} \geq \pi_k = \sum_{j=1}^m x_j \tilde{a}_{kj} \quad \forall k \neq i$$

Recalling the assumptions of payoff idiosyncrasy:

$$\pi_i = \sum_{j=1}^m x_j a_{ij} + \sum_{j=1}^m x_j \varepsilon_{ij} \sim N\left(\sum_{j=1}^m x_j a_{ij} ; \sigma^2 \sum_{j=1}^m x_j^2 \right)$$

The assumption of balanced payoff variances simplifies the analysis. The immediate consequence is that each strategy yields a payoff with a common variance. Notice:

$$\pi_i \ge \pi_k \quad \iff \quad \frac{\pi_i}{\sqrt{\sum_{j=1}^m x_j^2}} \ge \frac{\pi_k}{\sqrt{\sum_{j=1}^m x_j^2}}$$

Denoting this normalised payoff as y_i:

$$y_i \sim N\left(\frac{\sum_{j=1}^m x_j a_{ij}}{\sqrt{\sum_{j=1}^m x_j^2}}; \sigma^2\right)$$

Strategy selection is thus a realisation of a homoskedastic multinomial probit model where option i has expectation $\mu_i(x)$. The probability of selection is thus:

$$\begin{aligned} \rho_i(\sigma) &= \Pr[y_i \ge y_j \quad \forall j \ne i] \\ &= E_{\varepsilon_i}\left[\Pr\left(\frac{\mu_i - \mu_j + \varepsilon_i}{\sigma} \ge \frac{\varepsilon_j}{\sigma} \quad \forall j \ne i \,\middle|\, \varepsilon_i\right)\right] \\ &= \int_{-\infty}^{\infty} \left\{ \prod_{j \ne i} \Phi\left(z + \frac{\delta_{ij}(x)}{\sigma}\right) \right\} \phi(z)\, dz \end{aligned}$$

which is a standard multinomial probit choice probability. ∎

The eventual aim of the analysis is to select between Nash equilibria of the game in mean payoffs. To this end, it is pertinent to consider strategy frequencies from which a best response process leads to a pure strategy Nash equilibrium.

Definition 9. Suppose that strategy profile $\{i; i\}$ forms a symmetric pure strategy Nash equilibrium in mean payoffs. Define the basin of attraction for strategy i as:

$$B_i = \left\{ x : \sum_{j=1}^m x_j a_{ij} \ge \sum_{j=1}^m x_j a_{kj} \quad \forall k \ne i \quad \text{and} \quad x_j \ge 0; \ \sum_{j=1}^m x_j = 1 \right\}$$

For a discretised frequency space, the basin is:

$$B_i = \left\{ x : \sum_{j=1}^m x_j a_{ij} \ge \sum_{j=1}^m x_j a_{kj} \quad \forall k \ne i \quad \text{and} \quad x_j \in Z^+; \ \sum_{j=1}^m x_j = n \right\}$$

If a player with mean payoffs observes a strategy frequency $x \in B_i$, then the best response is strategy i. Any further mean payoff entrant will respond with strategy i, and thus a pure best response process would lead to the Nash equilibrium $\{i; i\}$ from x.

4.2. Exponential Cost in the Multinomial Probit Model. As outlined in the previous section, the strategy choices of entering players are the realisation of a homoskedastic multinomial probit model. Unfortunately, the choice probabilities are not available in closed analytic form – numeric evaluation of multiple integrals is required. Subsequent selection analysis, however, will place great interest in the behaviour of these probabilities as payoff idiosyncrasy vanishes. Clearly, if $\mu_i > \mu_j \; \forall j \neq i$ then $\rho_i \to 1$ and $\rho_j \to 0$ as $\sigma \to 0$. Interest will focus, however, on the rate at which these latter probabilities vanish. Although analytic expressions are unavailable for fixed σ, these probabilities become parametric as idiosyncrasy vanishes. This is formalised using the following definition:

Definition 10. $f(\sigma) > 0$ has exponential cost $c > 0$ if for arbitrarily small $\varepsilon > 0$:

$$\lim_{\sigma \to 0} f(\sigma) \exp\left(\frac{c+\varepsilon}{2\sigma^2}\right) = \infty \quad \text{and} \quad \lim_{\sigma \to 0} f(\sigma) \exp\left(\frac{c-\varepsilon}{2\sigma^2}\right) = 0$$

This property is denoted $f(\sigma) = \mathcal{E}(c)$ or alternatively $c(f(\cdot)) = c$.

Thus a function has exponential cost c if it behaves as $\exp(-c/2\sigma^2)$ as σ vanishes. The exponential cost property has parallels with the standard $o(\cdot)$ and $O(\cdot)$ notation familiar from the asymptotic behaviour of functions and sequences. The main difference is that the behaviour of $f(\sigma)\exp(c/2\sigma^2)$ is undefined. Familiar properties are available, however.

Lemma 2. Exponential cost has the following properties:

$$\prod_{i=1}^{m} \mathcal{E}(c_i) = \mathcal{E}\left(\sum_{i=1}^{m} c_i\right)$$

$$\sum_{i=1}^{m} \mathcal{E}(c_i) = \mathcal{E}\left(\min_{1 \le i \le m} c_i\right)$$

$$a \times \mathcal{E}(c_i) = \mathcal{E}(c_i)$$

Further, taking ratios of functions of σ:

$$c_i > c_j \Rightarrow \lim_{\sigma \to 0} \frac{\mathcal{E}(c_i)}{\mathcal{E}(c_j)} = 0$$

Proof. Consider m functions $f_i(\sigma)$ with exponential costs $\{c_i\}$. For ε arbitrarily small, $m\varepsilon$ is arbitrarily small. Hence:

$$\exp\left(\frac{m\varepsilon + \sum_{i=1}^{m} c_i}{2\sigma^2}\right) \prod_{i=1} f_i(\sigma) = \prod_{i=1} \exp\left(\frac{\varepsilon + c_i}{2\sigma^2}\right) f_i(\sigma)$$

From this the first property of Lemma 2 follows easily. The remaining properties follow in a similar fashion. ∎

Lemma 4 will consider the exponential cost of probit probabilities. It requires the following well-known lemma (see Myatt and Wallace (1997) for a proof).

Lemma 3. The hazard $\phi(x)/(1 - \Phi(x))$ is asymptotically linear as $x \to \infty$.

This result on hazard rates may be conveniently used in a multivariate setting. The following lemma determines the exponential cost of the probit probabilities. A similar result has been proved independently by Ruud (1996).

Lemma 4. The probit ρ_i has exponential cost

$$c = \left\{ \sum_j I(\delta_{ij} \le 0) \right\} \times \mathrm{var}_{j|\delta_{ij} \le 0}(\delta_{ij})$$

where $I(\cdot)$ represents the indicator function.

Notice immediately that if strategy i is the best response in mean payoffs, then $c = 0$, since $\lim_{\sigma \to 0} \rho_i = 1$. More generally, if there are J strategies that are weakly better than i in mean payoffs, then the exponential cost is J times the variance of the mean payoff advantage of these strategies. This elegant formulation follows from the judicious choice of heterogeneity distribution. The proof follows.

Proof. Recall from Lemma 1 that:

$$\rho_i(\sigma) = \int_{-\infty}^{\infty} \left\{ \prod_{j \ne i} \Phi\left(z + \frac{\delta_{ij}}{\sigma}\right) \right\} \phi(z)\, dz \tag{4}$$

Write the product of cumulative distributions as a product of densities and hazards:

$$\prod_{j \ne i} \Phi\left(z + \frac{\delta_{ij}}{\sigma}\right) = \underbrace{\prod_{j|\delta_{ij} > 0} \Phi\left(z + \frac{\delta_{ij}}{\sigma}\right)}_{\text{cdfs} \to 1} \times \underbrace{\prod_{j \ne i|\delta_{ij} \le 0} \frac{\Phi(z + \sigma^{-1}\delta_{ij})}{\phi(z + \sigma^{-1}\delta_{ij})}}_{\text{hazards}} \times \underbrace{\prod_{j \ne i|\delta_{ij} \le 0} \phi\left(z + \frac{\delta_{ij}}{\sigma}\right)}_{\text{densities}}$$

The product of densities is combined with $\phi(z)$ to obtain:

$$\phi(z) \prod_{j \ne i|\delta_{ij} \le 0} \phi\left(z + \frac{\delta_{ij}}{\sigma}\right) = \frac{1}{(2\pi)^{J/2}} \exp\left(- \frac{z^2 + \sum_{j \ne i|\delta_{ij} \le 0} (z + \sigma^{-1}\delta_{ij})^2}{2} \right)$$

where $J = 1 + \sum_{j \ne i} I(\delta_{ij} \le 0) = \sum_j I(\delta_{ij} \le 0)$. Completing the square yields:

$$\begin{aligned} z^2 + \sum_{j \ne i|\delta_{ij} \le 0} \left(z + \frac{\delta_{ij}}{\sigma}\right)^2 &= \left(z\sqrt{J} + \frac{\sum_{j \ne i|\delta_{ij} \le 0} \delta_{ij}}{\sigma \sqrt{J}} \right)^2 + \frac{\sum_{j \ne i|\delta_{ij} \le 0} \delta_{ij}^2}{\sigma^2} - \frac{\left(\sum_{j \ne i|\delta_{ij} \le 0} \delta_{ij}\right)^2}{J\sigma^2} \\ &= \left(z\sqrt{J} + \frac{\sum_{j \ne i|\delta_{ij} \le 0} \delta_{ij}}{\sigma \sqrt{J}} \right)^2 + \frac{\kappa^2}{\sigma^2} \end{aligned}$$

where κ^2 denotes:

$$\kappa^2 = \sum_{j:\delta_{ij}\le 0} \delta_{ij}^2 - \frac{\left(\sum_{j:\delta_{ij}\le 0}\delta_{ij}\right)^2}{J} = \left\{\sum_j I(\delta_{ij}\le 0)\right\} \times \mathrm{var}_{j:\delta_{ij}\le 0}(\delta_{ij})$$

Notice that in the above summation $j:\delta_{ij}\le 0$ includes $j=i$. Adding the term $\delta_{ii}=0$ does not affect the summations, and allows the variance interpretation on the right hand side. It is convenient to economise notation as follows:

$$\prod_{j:\delta_{ij}>0}\Phi_i = \prod_{j:\delta_{ij}>0}\Phi\left(z+\frac{\delta_{ij}}{\sigma}\right) \quad\text{and}\quad \tilde\phi(z) = \phi\left(z\sqrt{J} + \frac{\sum_{j\ne i:\delta_{ij}\le 0}\delta_{ij}}{\sigma\sqrt{J}}\right)$$

The choice probability (4) is now:

$$\rho_i = \frac{\exp(-\kappa^2/2\sigma^2)}{(2\pi)^{(J-1)/2}} \int_{-\infty}^{\infty} \tilde\phi(z) \prod_{j:\delta_{ij}>0}\Phi_i \prod_{j:\delta_{ij}<0}\frac{\Phi(z+\sigma^{-1}\delta_{ij})}{\phi(z+\sigma^{-1}\delta_{ij})}dz \qquad (5)$$

Consider $c<\kappa^2$. In this case:

$$\frac{\exp(-\kappa^2/2\sigma^2)}{(2\pi)^{(j-1)/2}}\exp\left(\frac{c}{2\sigma^2}\right) = \frac{1}{(2\pi)^{(j-1)/2}}\exp\left(-\frac{(\kappa^2-c)}{2\sigma^2}\right)\to 0$$

In addition, the integrand of (5) tends to zero and hence:

$$\exp\left(\frac{c}{2\sigma^2}\right)\times\int_{-\infty}^{\infty}\left\{\prod_{i\ne j}\Phi\left(z+\frac{\delta_{ij}}{\sigma}\right)\right\}\tilde\phi(z)dz\to 0$$

Next consider $c>\kappa$. In this case:

$$\frac{\exp(-\kappa^2/2\sigma^2)}{(2\pi)^{(j-1)/2}}\exp\left(\frac{c}{2\sigma^2}\right) = \frac{1}{(2\pi)^{(j-1)/2}}\exp\left(\frac{(c-\kappa^2)}{2\sigma^2}\right)\to \infty \qquad (6)$$

This expression diverges at an exponential rate towards $+\infty$. The integral of (5) vanishes to zero, however. The rest of the proof constructs a lower bound on this integral. It will be shown that this lower bound is polynomial in σ^{-1}. Exponential terms dominate polynomials in the limit, and hence the divergent exponential part of (6) will dominate.

First bound the integral of (5) by integrating over a subset of its range. To do this, first find the best alternative to strategy i and denote its advantage over i by δ_H. Hence:

$$\delta_H = \max_{j\ne i}\delta_{ji} > 0$$

Bound the integral as follows:

$$\int_{-\infty}^{\infty}\tilde\phi(z)\prod_{j:\delta_{ij}>0}\Phi_i\prod_{j\ne i:\delta_{ij}\le 0}\frac{\Phi(z+\sigma^{-1}\delta_{ij})}{\phi(z+\sigma^{-1}\delta_{ij})}dz \ge \int_{0}^{\sigma^{-1}\delta_H}\tilde\phi(z)\prod_{j:\delta_{ij}>0}\Phi_i\prod_{j\ne i:\delta_{ij}\le 0}\frac{\Phi(z+\sigma^{-1}\delta_{ij})}{\phi(z+\sigma^{-1}\delta_{ij})}dz$$

Next bounds are sought on each term in the integrand. First, consider the product of distribution functions:

$$\prod_{j|\delta_{ij}>0} \Phi\left(z+\frac{\delta_{ij}}{\sigma}\right) \geq [\Phi(z)]^{m-J}$$

This achieves a minimum at the lower limit of integration $z = 0$, yielding:

$$\prod_{j|\delta_{ij}>0} \Phi\left(z+\frac{\delta_{ij}}{\sigma}\right) \geq \frac{1}{2^{m-J}} \quad \text{for } \frac{\delta_H}{\sigma} \geq z \geq 0$$

Next recall that the hazard ratio ϕ/Φ is decreasing in its argument. Hence:

$$\prod_{j\neq i|\delta_{ij}\leq 0} \frac{\phi(z+\sigma^{-1}\delta_{ij})}{\Phi(z+\sigma^{-1}\delta_{ij})} \leq \left[\min_{j\neq i|\delta_{ij}\leq 0} \frac{\phi(z+\sigma^{-1}\delta_{ij})}{\Phi(z+\sigma^{-1}\delta_{ij})}\right]^{J-1} = \left[\frac{\phi(z-\sigma^{-1}\delta_H)}{\Phi(z-\sigma^{-1}\delta_H)}\right]^{J-1}$$

Furthermore, on the range of integration this achieves a maximum at $z = 0$, yielding:

$$\prod_{j\neq i|\delta_{ij}\leq 0} \frac{\phi(z+\sigma^{-1}\delta_{ij})}{\Phi(z+\sigma^{-1}\delta_{ij})} \leq \left[\frac{\phi(-\sigma^{-1}\delta_H)}{\Phi(-\sigma^{-1}\delta_H)}\right]^{J-1} \quad \text{for } \frac{\delta_H}{\sigma} \geq z \geq 0$$

On taking the reciprocal, the inequality is reversed, giving a lower bound. Finally, note:

$$\begin{aligned}
\int_0^{\sigma^{-1}\delta_H} \tilde{\phi}(z)\,dz &= \int_0^{\sigma^{-1}\delta_H} \phi\left(z\sqrt{J} + \frac{\sum_{j\neq i|\delta_{ij}\leq 0}\delta_{ij}}{\sigma\sqrt{J}}\right) dz \\
&= \frac{1}{\sqrt{J}} \int_0^{\sigma^{-1}\delta_H} d\Phi\left(z\sqrt{J} + \frac{\sum_{j\neq i|\delta_{ij}\leq 0}\delta_{ij}}{\sigma\sqrt{J}}\right) \\
&= \frac{1}{\sqrt{J}}\left\{\Phi\left(\frac{\sqrt{J}}{\sigma}\left[\delta_H + \frac{1}{J}\sum_{j\neq i|\delta_{ij}\leq 0}\delta_{ij}\right]\right) - \Phi\left(\frac{1}{\sigma\sqrt{J}}\left[\sum_{j\neq i|\delta_{ij}\leq 0}\delta_{ij}\right]\right)\right\}
\end{aligned}$$

The second term vanishes:

$$\lim_{\sigma\to 0} \Phi\left(\frac{1}{\sigma\sqrt{J}}\left[\sum_{j\neq i|\delta_{ij}\leq 0}\delta_{ij}\right]\right) = 0$$

since $\delta_{ij} < 0$ for some j. Next notice that:

$$\delta_H > \frac{(J-1)\delta_H}{J} > \frac{\sum_{j\neq i|\delta_{ij}\leq 0}\delta_{ji}}{J} = -\frac{\sum_{j\neq i|\delta_{ij}\leq 0}\delta_{ij}}{J}$$

and thus:

$$\lim_{\sigma\to 0} \Phi\left(\frac{\sqrt{J}}{\sigma}\left[\delta_H + \frac{1}{J}\sum_{j\neq i|\delta_{ij}\leq 0}\delta_{ij}\right]\right) = 1$$

Having obtained a bound for $\int_0^{\sigma^{-1}\delta_H} \phi(z)\,dz$, the bounding components are assembled:

$$
\begin{aligned}
\int_{-\infty}^{\infty} \phi(z) \prod_{j:\delta_{ij}>0} \Phi_{-} \prod_{j\neq i:\delta_{ij}\le 0} \frac{\Phi(z+\sigma^{-1}\delta_{ij})}{\phi(z+\sigma^{-1}\delta_{ij})}\,dz
&\ge \int_0^{\sigma^{-1}\delta_H} \phi(z) \prod_{j:\delta_{ij}>0} \Phi_{-} \prod_{j\neq i:\delta_{ij}\le 0} \frac{\Phi(z+\sigma^{-1}\delta_{ij})}{\phi(z+\sigma^{-1}\delta_{ij})}\,dz \\
&\ge \frac{1}{2^{m-j}} \int_0^{\sigma^{-1}\delta_H} \phi(z) \prod_{j\neq i:\delta_{ij}\le 0} \frac{\Phi(z+\sigma^{-1}\delta_{ij})}{\phi(z+\sigma^{-1}\delta_{ij})}\,dz \\
&\ge \frac{1}{2^{m-j}} \left[\frac{\Phi(-\sigma^{-1}\delta_H)}{\phi(-\sigma^{-1}\delta_H)}\right]^{J-1} \int_0^{\sigma^{-1}\delta_H} \phi(z)\,dz \\
&\to \frac{\sigma^{J-1}}{2^{m-J}\sqrt{J}\,\delta_H^{J-1}}
\end{aligned}
$$

where the last step employs the asymptotic linearity of hazards, taken from Lemma 3. The expression is polynomial in σ^{-1}. This is dominated asymptotically by the exponential term, so that

$$
\int_{-\infty}^{\infty} \left\{ \prod_{i\neq j} \Phi\left(z + \frac{\delta_{ij}}{\sigma}\right) \right\} \phi(z)\,dz \times \exp\left(\frac{c}{2\sigma^2}\right) \to \infty
$$

which completes the proof. ∎

This proposition gives a convenient form for "mutations" generated from a idiosyncrasy-driven adaptive learning model, and is the central contribution of this paper. Further commentary on its applicability is given in Section 6.

5. Ergodic Analysis of Adaptive Play

This section considers the Adaptive Play dynamic described in Section 3.3. As in Myatt and Wallace (1997), the now-standard graph theoretic methods of Freidlin and Wentzell are employed to find the ergodic distribution. First the process is formalised and the Markovian transition probabilities calculated. A review of the rooted tree method follows. The effect of dominated strategies on equilibrium selection is then described, and formalised as Proposition 1.

5.1. The State Space and Markov Transitions for Adaptive Play. The Adaptive Play dynamic yields a Markov chain. Behaviour of entrants is conditional only on the strategy frequency among incumbents. The current state may be characterised by the number of agents playing each strategy i. This is formalised as follows.

Definition 11. Define the Markov state space S as:

$$S = \left\{ s \in \mathbb{Z}_+^m \quad \text{such that} \quad \sum_{i=1}^{m} s_i = n \right\}$$

where $\mathbb{Z}$ represents the non-negative integers.

The state space is thus the discretised unit simplex. It is clear from the description of the dynamic in Section 3.3 that the population may only move one step at a time. It is useful to introduce the following notation. Denote by e_i a $m \times 1$ vector, with zero elements except for the ith element, which is set to 1. Characterisation of the Markov transition probabilities follows.

Lemma 5. The Markov transition probabilities satisfy:

$$p_{ss'} = \begin{cases} \frac{i}{n}\rho_j\left(\frac{s}{n}\right) & s' = s - e_i + e_j \\ \sum_{i=1}^{m} \frac{i}{n}\rho_i\left(\frac{s}{n}\right) & s' = s \\ 0 & s' \neq s - e_i + e_j \quad \forall i, j \end{cases}$$

Proof. Start in state s. The population updates one player at a time, and hence must lose a strategy and gain a (possibly identical) strategy. It loses strategy i with probability i/n, and the entrant adopts strategy $j \neq i$ with probability $\rho_j\left(\frac{s}{n}\right)$. Alternatively, the process may not shift, requiring any strategy to be lost and replaced by the same, yielding the summation. ∎

From Lemma 5, a one-step movement occurs with probability $(i/n)\rho_j(s/n)$. The term i/n is a multiplicative factor independent of the payoff idiosyncrasy σ. Limiting behaviour as $\sigma \to 0$ is determined entirely by the multinomial choice probability $\rho_j(s/n)$.

Corollary 1. For $s \neq s' = s - e_j + e_i$ the transition probability has exponential cost

$$p_{ss'} = \Theta\left(\sum_j I(\delta_{ij} \leq 0) \times \mathrm{var}_{j:\delta_{ij} \leq 0}(\delta_{ji}) \right)$$

Proof. The multiplicative factor j/n has no affect on the exponential cost and the result follows from Lemma 4. ∎

In the limit, as $\sigma \to 0$, the key determinant of ergodic behaviour will be the exponential cost of the transitions.

5.2. The Ergodic Distribution and Rooted Trees. For 2×2 games, the state space reduces to the integers $\{0, 1, \ldots, n\}$. The transition probabilities characterised above yield a convenient tridiagonal form for the Markov matrix. For $m > 2$, however, the analysis is more complex. Following previous work, including KMR (1993) and Young (1993), the graph-theoretic approach of Freidlin and Wentzell (1984) is employed.

The Freidlin and Wentzell (1984) approach constructs a directed graph on the state space S with edge weights corresponding to Markov transition probabilities. The directed edge set $E \subseteq S \times S$ has weights $p : E \mapsto \mathbb{R}^+$, where the first and second coordinates represent source and target nodes respectively. A tree rooted at s is a set of edges $h \subseteq S \times S$ such that each node $i \neq s$ has a unique successor. All sequences of edges lead to s, which has no successor. The collection of trees rooted at s is H_s. The weight of such a tree h is:

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

Sum over all trees rooted at s to obtain:

$$q_s = \sum_{h \in H_s} w_h \tag{7}$$

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 s. The following is due to Freidlin and Wentzell (1984, Chapter 6, Lemma 3.1).

Lemma 6. The invariant distribution μ satisfies:

$$\mu_s = \frac{q_s}{\sum_{s' \in S} q_{s'}} = \frac{\sum_{h \in H_s} \prod_{(i,j) \in h} p_{ij}}{\sum_{s' \in S} \sum_{h \in H_{s'}} \prod_{(i,j) \in h} p_{ij}}$$

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 $q_s / q_{s'}$. Notice that the weight of a tree is the product of transition probabilities. The exponential cost of this product may be calculated by summing the component costs.

Lemma 7. The exponential cost of a tree weight satisfies:

$$w_h = e\left(\sum_{(s,s') \in h} c(s, s')\right)$$

Proof. The proof follows directly from Lemma 2. ∎

Lemma 8. Generically, only the least cost rooted trees matter:

$$\lim_{\sigma \to 0} \frac{q_s}{q_{s'}} = 0 \iff \min_{h \in H_s} \left\{ \sum_{(i,j) \in h} c(i,j) \right\} > \min_{h \in H_{s'}} \left\{ \sum_{(z,z') \in h} c(i,j) \right\}$$

Proof. From Lemma 6 it is clear that:

$$\frac{q_s}{q_{s'}} = \frac{\sum_{h \in H_s} \prod_{(i,j) \in h} p_{ij}}{\sum_{h \in H_{s'}} \prod_{(i,j) \in h} p_{ij}}$$

Using Lemma 2:

$$\lim_{\sigma \to 0} \frac{q_s}{q_{s'}} = 0 \iff c\left(\sum_{h \in H_s} \prod_{(i,j) \in h} p_{ij} \right) > c\left(\sum_{h \in H_{s'}} \prod_{(i,j) \in h} p_{ij} \right)$$

Now, again using Lemma 2, take the first term:

$$\begin{aligned} \sum_{h \in H_s} \prod_{(i,j) \in h} p_{ij} &= \sum_{h \in H_s} \prod_{(i,j) \in h} \Theta(c(i,j)) \\ &= \sum_{h \in H_s} \Theta\left(\sum_{(i,j) \in h} c(i,j) \right) \\ &= \Theta\left(\min_{h \in H_s} \sum_{(i,j) \in h} c(i,j) \right) \end{aligned}$$

From this the result clearly follows. ∎

5.3. Strategy Selection for Vanishing Idiosyncrasy. A 'least cost tree' rooted at a stable population profile gives paths from all other nodes to the target node. In particular, it provides an escape path. With such paths in mind, consider the 2×2 game (1) of Section 2. Here it is represented by its mean payoff matrix:

$$\Pi = \begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix}$$

Denote by node 1 the state where all n players adopt strategy 1, and similarly for node 2. The observations of Section 2 are reviewed here. Given that all players adopt strategy 1, a mean payoff entrant will follow, since $3 > 0$. For an entrant to be contrarian, the player's idiosyncratic preference for strategy 2 over 1 must exceed 3. Note also, that it is a best response in mean payoffs to choose strategy 1 when a fraction $2/5$ or more of the population adopt strategy 1. To dislodge node 1, it is thus necessary to observe approximately $\frac{3}{5}n$ contrarians. Suppose instead that the process lies at node 2. An entrant need have an idiosyncratic preference of only 2 to play against the current standard. Moreover, only $\frac{2}{5}n$ contrarians are required to dislodge this strategy. It is thus easy to see why the risk-dominant profile $\{1;1\}$ is selected by the adaptive play dynamic. Now consider instead the extension to the 3 strategy game (3):

$$\Pi' = \begin{bmatrix} 3 & 0 & 0 \\ 0 & 2 & 2 \\ 2\frac{1}{2} & -\frac{1}{2} & -\frac{1}{2} \end{bmatrix}$$

The third strategy is strictly dominated, and hence no mean payoff entrant would choose it. Notice that to escape from node 1 still requires $\frac{3}{5}n$ contrarians to play either strategy 2 or strategy 3. An entrant, however, requires an idiosyncratic preference for strategy 3 over 1 of only $\frac{1}{2}$ in order to move away from node 1. Movement away from node 2, however, still requires an idiosyncratic preference of at least 2. The key insight here is similar to that of Myatta and Wallace (1997). Intuitively, whereas the basin of attraction for node 1 is wide, it also shallow. A large number of contrarians are required to escape, but these occur with relatively high probability. This intuition depends in no important way on vanishing noise. Moreover, the payoffs of the added strategy 3 need not be particulary close to those of strategy 1. The intuition is formalised in Proposition 1, which forms the main result for adaptive play. The proof is obtained by taking a general version of the examples presented here and ensuring the payoffs of the dominated strategy are sufficiently close to those of strategy 1.

Proposition 1. Take a 2×2 symmetric coordination game where strategy 1 is risk-dominant and add a third strategy. The additional strategy may be constructed so that:

1. The strategy is strictly dominated in mean payoffs.
2. Strategy 1 remains globally risk-dominant.
3. Strategy 2 is selected for vanishing heterogeneity.

Proof. Begin with the 2×2 coordination game:

$$\Pi = \begin{bmatrix} a_{11} & 0 \\ 0 & a_{22} \end{bmatrix}$$

where $a_{11} > a_{22}$ and the off-diagonal payoffs are normalised to zero without loss of generality. Since this is a pure coordination game, and $a_{11} > a_{22}$, it immediately follows that the profile $\{1;1\}$ is the risk-dominant Nash equilibrium. For small $\epsilon > 0$, extend this to Π':

$$\Pi' = \begin{bmatrix} a_{11} & 0 & 0 \\ 0 & a_{22} & a_{22} \\ a_{11} - \epsilon & -\epsilon & -\epsilon \end{bmatrix}$$

yielding a three strategy game. Notice immediately that strategy 1 strictly dominates the added strategy 3. Global risk dominance must be checked. In particular, the payoff from strategy 1 must exceed that from strategy 2 for any opposing strategy frequency x satisfying $x_1 \geq 1/2$. For such a frequency, the payoff from strategy 1 is $x_1 a_{11}$, and from strategy 2 is $x_2 a_{22} + x_3 a_{22} = (1 - x_1) a_{22} \leq x_1 a_{11}$ and hence global risk dominance is satisfied. Notice that strategy 3 is a near identical alternative to strategy 1; however, given a population with a large fraction playing strategy 3, strategy 2 is a better response than strategy 1.

Basins of attraction are the next focus. Since strategy 3 is dominated, the best response to any frequency is either strategy 1 or 2, each of which form Nash equilibrium. All states are thus in either B_1 or B_2. Consider frequency x. For $x \in B_1$:

$$x_1 a_{11} > (1 - x_1) a_{22} \iff x_1 > \frac{a_{22}}{a_{11} + a_{22}} < \frac{1}{2}$$

It is clear that $x \in B_1$ if and only if $x_1 > a_{22}/(a_{11} + a_{22}) = x_1^*$. The tied case is omitted without loss of generality.

Next construct rooted trees. Denote by nodes 1 and 2 the states where all n agents play strategies 1 and 2 respectively. Begin the construction of a tree rooted at node 2. To do this, construct a sequence of edges beginning at node 1. At each state, map an edge to the state involving one less strategy 1 player and one more strategy 3 player. Consider a node s on this path. The transition probability is:

$$\frac{s_1}{n} \rho_3 \left(\frac{s}{n} \right) \tag{8}$$

Continue this path until a state inside B_2 is reached. For all other nodes, map to a node with one more player 1 if $s/n \in B_1$, and one more player 2 if $s/n \in B_2$. These latter edges all have zero exponential cost. The only costly edges are the ones leading away from node 1. It is easy to check that this mapping does indeed form a tree rooted at node 2.

Next, consider the transition probability (8) and denote $x = s/n$. This transition involves one more strategy 3 player. Using Lemma 4, this has exponential cost:

$$\sum_j I(\delta_{3j} \leq 0) \times \mathrm{var}_{j|\delta_{3j} \leq 0}(\delta_{j3})$$

Recalling earlier definitions:

$$\delta_{3j}(x) = \mu_3(x) - \mu_j(x) = \frac{\sum x_k (a_{3k} - a_{jk})}{\sum_{k=1}^{3} x_k^2}$$

The mean payoff advantages are only required for $j : \delta_{3j} \leq 0$. On the discretised simplex with sufficiently small ε, for all source nodes it is easily established that $\delta_{32} > 0$. Thus

$\sum_j I(\delta_{3j} \leq 0) = 2$. Naturally $\delta_{33} = 0$, and hence the exponential cost satisfies:

$$\sum_j I(\delta_{3j} \leq 0) \times \mathrm{var}_{j|\delta_{3j} \leq 0}(\delta_{j3}) = 2 \times \left\{ \frac{\delta_{13}^2}{2} - \left(\frac{\delta_{13}}{2} \right)^2 \right\} = \frac{\delta_{13}^2}{2}$$

Notice now that $a_{3k} - a_{1k} = -\varepsilon$. Moreover, $\sum_{k=1}^{3} x_k^2 \geq 1/3$ so that

$$\sum_j I(\delta_{3j} \leq 0) \times \mathrm{var}_{j|\delta_{3j} \leq 0}(\delta_{j3}) \leq \frac{3\varepsilon^2}{2}$$

Each positive cost edge thus has a cost not exceeding $3\varepsilon^2/2$. A sufficient number of transitions is required to escape B_1. A fraction $a_{11}/(a_{11} + a_{22})$ of the population is sufficient. Hence the total cost of the tree rooted at 2 does not exceed:

$$\left[\frac{n a_{11}}{a_{11} + a_{22}} \right] \times \frac{3\varepsilon^2}{2} \qquad (9)$$

This provides an upper bound on the minimum cost tree. Construct now a tree rooted at node 1. Such a tree requires a transition away from node 2, at which $x_2 = 1$. This transition will involve strictly positive exponential cost. For instance, suppose this transition involves an added strategy 1. Then $\sum_j I(\delta_{1j} \leq 0) = 2$. Moreover:

$$\delta_{12} = a_{12} - a_{22} = -a_{22}$$

It is simple to show that the cost of this edge will be $a_{22}^2/2$. The cost is higher for a transition away involving an added strategy 3. Thus any tree rooted at node 2 must have a weight of at least $a_{22}^2/2$, providing a lower bound on the cost of the minimum rooted tree. But equation (9) is arbitrarily small for small ε. It follows that node 2 is selected. ∎

6. Discussion

6.1. Proximity of the Dominated Strategy. The formal proof of Proposition 1 requires the dominated strategy 3 to be sufficiently close to strategy 2. How close? Recall that the total cost of the tree rooted at node 2 did not exceed:

$$\underbrace{\left[\frac{n a_{11}}{a_{11} + a_{22}} \right]}_{\text{no. of steps}} \times \underbrace{\frac{3\varepsilon^2}{2}}_{\text{step size}}$$

Next move to a tree rooted at node 1. The first step away from node 2 involved an exponential cost of at least $a_{22}/2$. Of course, further steps will involve less exponential cost – the basin depth is variable in the adaptive dynamic. Suppose for simplicity that the basin depth were constant. This is the case in the Sophisticated Play dynamic of Myatt and Wallace (1999), and provides a reasonable approximation. How many steps are required to reach node 1? The approximate number is $n a_{22}/(a_{11} + a_{22})$. The required (approximate) inequality for the selection of strategy 2 is then:

$$3\varepsilon^2 a_{11} < a_{22}^3 \quad , \quad \varepsilon < \sqrt{\frac{a_{22}^3}{3a_{11}}}$$

For the case of $a_{11} = 3$ and $a_{22} = 2$, this translates to $\varepsilon < \sqrt{8/9} \approx 0.94$. Such a case is illustrated in the following game:

$$\Gamma = \begin{bmatrix} 3 & 0 & 0 \\ 0 & 2 & 2 \\ 2\frac{1}{4} & -\frac{3}{4} & -\frac{3}{4} \end{bmatrix}$$

It follows that the dominated strategy need not be particularly close.

6.2. **Future Applicability.** This paper has demonstrated that dominated strategies can play a key rôle in equilibrium selection. In particular, the replacement of mutations with idiosyncrasy calls into the question the robustness of global risk dominance as a selection criterion. Why is this? The difficulty of moving from one Nash equilibrium to another is determined by both the number of steps, and the probability of taking these steps. This is clear in the analysis of Myatt and Wallace (1997) and (1999), where basin depth and width both contribute to selection. In the 2×2 case with balanced trembles, the width and depth are both determined by the same term. It follows that risk-dominance retains its selection power. This is not the case here. Consider a sequence of transitions from strategy 1 to strategy 2. The required number of transitions is determined by the relative payoffs of strategy 1 and strategy 2. The probability of each transition, however, is instead determined by the relative payoffs of strategy 1 and strategy 3. Depth and width of escape paths are determined by two different criteria. This is not present in the mutation-driven models — the transition probabilities are uniform.

Adopting the idiosyncratic methodology suggested in this paper may warrant further investigation of larger games. It is here that the contribution is most valuable. Lemma 4 gives a convenient closed form for the exponential cost of probit transition probabilities, hence allowing easy use of the Freidlin and Wentzell (1984) rooted tree approach. Companion research applies the techniques of this paper to the Sophisticated Play paradigm in larger games.

References

Bergin, J., and B. Lipman (1996): "Evolution with State-Dependent Mutations," Econometrica, 64(4), 943–956.

Ellison, G. (1996): "Basins of Attraction, Long Run Equilibria and the Speed of Step-by-Step Evolution," mimeo, MIT.

Freidlin, M. I., and A. D. Wentzell (1984): Random Perturbations of Dynamical Systems. Springer-Verlag, Berlin and New York.

Harsanyi, J. C., and R. Selten (1988): A General Theory of Equilibrium Selection in Games. MIT Press, Cambridge.

Hofbauer, J., and J. W. Weibull (1996): "Evolutionary Selection Against Dominated Strategies," Journal of Economic Theory, 71, 558–575.

Kandori, M., G. J. Mailath, and R. Rob (1993): "Learning, Mutation and Long-Run Equilibria in Games," Econometrica, 61(1), 29–56.

Maruta, T. (1997): "On the Relationship Between Risk-Dominance and Stochastic Stability," Games and Economic Behavior, 19(2), 221–234.

Morris, S., R. Rob, and H. S. Shin (1995): "Dominance and Belief Potential," Econometrica, 63(1), 145–157.

Myatt, D. P., and C. Wallace (1997): "Adaptive Dynamics with Idiosyncratic Agents," Nuffield College Economics Discussion Paper W31.

——— (1999): "Sophisticated Play by Idiosyncratic Agents," Nuffield College Economics Discussion Paper 1999-W1.

Palomino, F. (1995): "On the Survival of Strictly Dominated Strategies," IAE Working Paper 308.95, University of Tilburg.

Rubinstein, A. (1998): Modeling Bounded Rationality. MIT Press, Cambridge, MA.

Ruud, P. A. (1996): "Approximation and Simulation of the Multinomial Probit Model: An Analysis of Covariance Matrix Estimation," mimeo UC Berkeley.

Young, H. P. (1993): "The Evolution of Conventions," Econometrica, 61(1), 57–84.