

The Evolution of Coordination under Inertia*

Thomas Norman
Nuffield College, Oxford

January 31, 2003

Abstract

This paper models the phenomenon of inertia driven by individual strategy switching costs in a stochastic evolutionary context. Kandori, Mailath, and Rob's (1993) model of a finite population of agents repeatedly playing a 2×2 symmetric coordination game is extended to allow for such inertia. Taking noise to the limit, a number of new short- to medium-run equilibria emerge, centred around the mixed-strategy equilibrium. Thus, unusually, an evolutionary model is seen to provide some justification for the controversial concept of mixed-strategy equilibrium. However, Kandori, Mailath, and Rob's long-run selection of the risk-dominant equilibrium continues to hold, both under fixed-rate mutations and under state-dependent mutations driven by stochastic switching costs. The key to this is the satisfaction of Blume's (1999) "skew-symmetry" of the noise process, which is shown to be crucial even under simultaneous strategy revisions. In fact, the presence of the new short-run equilibria can under certain conditions serve to *reduce* the expected waiting time before the risk-dominant equilibrium is reached - an instance of Ellison's (2000) idea that evolution is more rapid when it can proceed via a series of small "steps" between extremes. This suggests inertia to be a surprisingly efficient phenomenon, and also serves to moderate the force of the Ellison (1993) critique of excessively long transition times in models with vanishing noise.

*Discussions with my supervisors, David Myatt and Chris Wallace, as well as with Margaret Meyer, Kevin Roberts and seminar participants at the Department of Economics, Oxford University, are gratefully acknowledged. The usual disclaimer applies.

1 Introduction

“Change is not made without inconvenience, even from worse to better.”

Samuel Johnson, *A Dictionary of the English Language* (1755)

In the last ten years, developments in evolutionary game theory have promised to deliver mainstream game theory from its perpetual demon: moving from deterministic to stochastic evolutionary models solves the equilibrium selection problem for repeated games in the long run. The introduction of persistent randomness (or “noise”) into an evolutionary model allows the use of the probabilistic toolkit of Markov chains, the most useful product of which is a long-run ergodic distribution over the model’s outcomes which selects between them as the noise goes to zero. This provides a high return on an evolutionary approach which both complements and augments perfectly rational game theory. Evolutionary models can admit a variety of boundedly rational behaviour and generally end up converging to rational outcomes. If they can both provide realistic behavioural foundations for game theory, and sharpen its analytical predictions, then they will have made a significant contribution indeed.

However, in the majority of existing evolutionary models, deterministic and stochastic alike, the forms of bounded rationality admitted do not extend to the possibility of player inertia. Individual players are assumed to be able costlessly to change their strategies, and hence they do so whenever there is even the smallest gain to be made. Even in those models that acknowledge and incorporate the presence of inertia, its causes are generally left exogenous. This paper (and the companion paper, Norman (2003a)) studies the effects of generalising stochastic evolutionary models to allow for the existence of “switching costs” to changing strategies from one period to the next, and thus the possibility of (endogenously determined) inertia. Such costs are intuitively realistic. As Lipman and Wang (2000) point out, changing actions involves real costs in many economic contexts, such as a firm’s investment decisions (“set up” or “shut down” costs) or price-setting games (“menu” costs). But even in the absence of such tangible costs, if playing a given strategy is complex, then switching strategies may be “hard”, imposing learning and implementation costs on the individual player. It is thus an important and natural “robustness check” of existing models to study how they change in the presence of these “switching costs”.

The motivation is provided by coordination games, with the existence of inertia representing a possible explanation for the persistence of suboptimal coordination in real-life situations such as David’s (1985) choice between QWERTY and Dvorak keyboard layouts, and the contemporary issue of Microsoft versus Linux operating systems. Thus, a fixed switching cost is introduced into the popular Kandori, Mailath, and Rob (1993) model of a finite population of agents repeatedly playing a 2×2 symmetric coordination game. Analysis of the limit of the ergodic distribution as the noise in the system goes to zero reveals that Kandori, Mailath and Rob’s long-run selection of the risk-dominant

equilibrium continues to hold in the presence of switching cost-driven inertia, until the switching costs become question-beggingly high. However, the model acquires a number of new short- to medium-run equilibria centred around (and including) the mixed-strategy equilibrium. The relative importance of these two findings is determined by the transition times between equilibria. Ellison (1993) and others have drawn attention to the fact that transition times become prohibitively large when noise is taken to the limit. This suggests that a significant amount of time could be spent in and around the mixed-strategy equilibrium in the presence of inertia. This finding is particularly interesting in view of the fact that mixed-strategy equilibria have been almost universally ruled unstable in the stochastic adjustment dynamics literature to date - a regularity which has served further to inter this long since controversial concept. Its emergence here as a short- to medium-run equilibrium of the Kandori, Mailath and Rob (henceforth KMR) model under inertia would appear to offer it some substantial evolutionary justification, particularly if “in the long run we are all dead”.

It need not, however, be the case that the long run is long enough to invoke Keynes’ famous dictum. Indeed, whilst the presence of player inertia and the resulting new mixed short- to medium-run equilibria might at first appear likely to *slow down* convergence to the risk-dominant equilibrium, an analysis of the transition times in fact yields the opposite result: the presence of the new mixed short- to medium-run equilibria serves to *reduce* the expected waiting time before the risk-dominant equilibrium is reached. The reason for this result is that, as Ellison (2000) has noted, evolution is more rapid when it may proceed via a series of small “steps” between intermediate steady states, such as those provided by the inertia-born mixed equilibria. The counterintuitive conclusion is thus established that the presence of inertia serves to *speed up* evolution, thus showing inertia to play a surprisingly efficient evolutionary role and also moderating the Ellison (1993) “very long run” critique of models with vanishing noise. Moreover, the force of the Ellison (1993) critique is further reduced by the finding that the presence of switching costs serves to attenuate the positive effect of the population size on transition times. This is again due to the presence of the new mixed absorbing states, and specifically the fact that they become more numerous - thus providing more evolution-facilitating “steps” - as the size of the population increases.

To this point, the analysis is kept as simple as possible by assuming the “switching cost” of strategy change to be fixed over time and across players. However, it is clear that, in reality, different players (with different priorities, abilities and constraints) will vary in the size of their switching costs, and any given player’s switching cost will fluctuate in size over time (as his priorities, abilities and constraints change). Thus, in section 4, the model is altered to make the switching cost itself stochastic across players and time. The introduction of this stochastic switching cost obviates the need to introduce the state-independent mutations of the KMR-style model, thus also serving to address the criticisms of Bergin and Lipman (1996) by building a model with endogenously generated “state-dependent mutations”. The earlier long-run selection results are confirmed for this more realistic model, but the findings for transition times

are qualified: whilst the “step-by-step” effect is still present, the switching costs now make the extreme states even harder to escape. Nonetheless, it is shown that there exists a threshold level of the mean switching cost, below which limiting transition times are reduced by inertia, given a large enough population. This threshold mean switching cost is increasing in the size of the population, and is also crucially affected by the assumed switching cost distribution.

2 Relevant Literature

The stochastic adjustment dynamics literature was born in the papers of Foster and Young (1990), Kandori, Mailath, and Rob (1993), and Young (1993). Foster and Young (1990), drawing on the Markovian techniques of Freidlin and Wentzell (1984), introduced the concept of stochastic stability in continuous dynamical systems into evolutionary biology, and provided a method for the analytical computation of the stochastically stable set. Kandori, Mailath, and Rob (1993), and Young (1993) then brought the somewhat simpler discrete analysis into the realm of economics, with immediate dramatic effect. The multiplicity of equilibria in 2×2 coordination games was overturned by KMR (1993) in an adaptive population setting with random mutations (“noise”); the unique long-run (stochastically stable) equilibrium as the noise vanished was found to be the risk-dominant one. Young (1993), meanwhile, also found risk-dominance to determine long-run selection in a quite different model of a 2×2 coordination game in which each player would choose an optimal strategy based on a sample of information about what the other player had done in the past. These papers provided the basis of a new and fruitful direction for evolutionary game theory, appearing to offer a solution to the persistent problem of multiple equilibria and path dependence: the long-run equilibrium of a stochastic model was unique for generic games with strict Nash equilibria.

Nonetheless, the new stochastic evolutionary game theory was not without its critics. The first main focus of criticism was the time required for transitions between the equilibria of an evolutionary model. Ellison (1993) and others have drawn attention to the fact that transition times become prohibitively large when noise is taken to the limit; hence, the long-run selection results of models such as KMR might only apply in the “very long run” (or “ultra-long run”), a problem which grows with the size of the population. Thus it seemed that, whilst long-run selection results for vanishing noise might be suited to small populations, initial conditions were likely to be a more informative object of analysis for small noise and large populations. Theorists have attempted to tackle this “very long run” problem in a number of ways. Ellison (1993), and Robson and Vega-Redondo (1996) have reduced transition times by altering the matching process which describes how players meet. Binmore and Samuelson (1997), meanwhile, study a model of “noisy learning” with *non*-vanishing noise, which also serves to speed up evolution. More recently, Ellison (2000) has drawn attention to a third effect which can reduce waiting times: evolution tends to be more rapid when it may proceed via a series of small “steps” between

intermediate steady states (rather than requiring sudden large changes). The effects of such “step-by-step evolution” are readily captured by Ellison’s “radius-modified coradius” approach to stochastic adjustment models.

The second main criticism of the early models of stochastic adjustment dynamics concerned the mutation mechanism employed. Bergin and Lipman (1996) criticise the arbitrariness of mutations occurring at a rate independent of the current state of the system. Such “state-independent mutations”, embodied in the fixed mutation rate ε of KMR and others, imply that players make mistakes (or experiment, etc.) with the same probability irrespective of the current strategy frequencies, and thus of the expected payoffs at stake. Bergin and Lipman demonstrate that, given any model of the effect of mutations, any invariant distribution of the “mutationless” process is close to an invariant distribution of the process with appropriately chosen small mutation rates. This implies that any strict Nash equilibrium of a strategic form game is selected under some suitably chosen mutation model. Bergin and Lipman’s paper highlights the importance of developing models or other criteria to determine “reasonable” classes of “state-dependent mutations”. Myatt and Wallace (1998) present a candidate for just such a “reasonable” mutation process; their model of state-dependent mutations driven by payoff heterogeneity (rather than mistakes, or experimentation) remains broadly supportive of risk-dominance. Risk-dominance is further strengthened by Lee, Szeidl, and Valentinyi’s (2001) result that, for local interaction games, the risk-dominant equilibrium is uniquely selected given a sufficiently large population. More generally, Blume (1999) alleviates the indeterminacy of the Bergin and Lipman critique with his finding that the known stochastic stability results are preserved for the (large) class of noise processes satisfying a certain symmetry condition.

The development of the stochastic evolutionary field has clearly been rapid then, but numerous important avenues have yet to be explored. One such avenue is player inertia driven by individual strategy switching costs, the subject of this paper. The idea of inertia driven by switching costs in repeated game contexts is of course not a new one. Within the traditional perfectly rational paradigm, Klemperer’s (1987a, 1987b) consumer switching costs¹ and Radner’s (1980) ε -Nash equilibrium are two obvious examples. In particular, Farrell and Klemperer’s (2001) recent survey of competition with switching costs and network effects considers settings of a similar nature to the model presented here. The closest parallel to the present paper in the perfectly rational literature, however, is Lipman and Wang (2000), who study the effects of switching costs in general repeated game contexts. They add small costs of changing actions and frequent repetition to finitely repeated games, and find that doing so makes credible certain commitments which then serve to overturn all the standard results for finitely repeated games.

By contrast with the perfectly rational literature, however, it would seem that inertia has for the most part been left unmodelled in evolutionary contexts.

¹On consumer switching costs see, for example, Klemperer (1995), Beggs and Klemperer (1992), Farrell (1987), Farrell and Shapiro (1988, 1989), and Padilla (1995).

In reinforcement models, for example, the probability of taking an action in the present increases with the payoff that resulted from taking that action in the past.² Such models admit an intuitive role for inertia, but this inertia is assumed exogenous, its root causes left unmodelled. Meanwhile, inertia plays an explicit role in the KMR (1993) model, with their weakly monotonic selection dynamic capturing the idea that only some (as opposed to all) players need be adjusting their behaviour in any given period. But this ignores the important possibility that *no* player adjusts his behaviour in a given period, and moreover, there is again no endogenous determination of the inertia.

There are admittedly some evolutionary papers that model costly play of some sort. One example is Sethi (1998), whose model of “strategy-specific barriers to learning” in the replicator dynamics explores the consequences of strategies varying in the ease with which they can be learned. Another example is van Damme and Weibull’s (1998) model of “mutations driven by control costs”, which has mutation rates determined by individual mistake probabilities, which players can control at some cost. State-dependent mutations are thus based on an economically justified model here, in the manner suggested by Bergin and Lipman (1996).³ But whilst both of these papers model important ways in which strategy adoption might be costly, neither captures the idea that strategy *change* is costly compared to the (cost-free) status quo. The *switching costs* postulated in this paper, by contrast, draw attention to the learning and implementation costs incurred when a player switches to a new strategy - a plausible source of individual-level inertia.

3 Fixed-Rate Mutations

3.1 The Model

Following Kandori, Mailath, and Rob (1993), the model focuses on a finite number N (N an even number) of players who are repeatedly matched to play the symmetric 2×2 stage game defined by the general payoff matrix

$$\Lambda = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \quad (1)$$

and who adjust their behaviour over time. Actions are taken in discrete time, $t = 1, 2, \dots$, each player choosing his pure strategy for a period t at the beginning of that period. The number z^t of players playing strategy 1 at time t defines the state of the dynamical system, the state space being $Z = \{0, 1, \dots, N\}$. The

²See Bush and Mosteller (1955), Suppes and Atkinson (1960), Arthur (1993), Roth and Erev (1995), Börgers and Sarin (1995, 1997).

³This model again selects the risk-dominant equilibrium of coordination games.

average payoff of a player with strategy i , $\pi_i(z^t)$, is assumed to be⁴

$$\pi_1(z) = \frac{(z-1)}{(N-1)}a_{11} + \frac{(N-z)}{(N-1)}a_{12}, \quad (2)$$

$$\pi_2(z) = \frac{z}{(N-1)}a_{21} + \frac{(N-z-1)}{(N-1)}a_{22} \quad (3)$$

These become expected payoffs under the assumption that players are myopic.

In the KMR analysis, the selection dynamics studied are all those under which better strategies are better represented in the population in the next period (the so-called ‘‘Darwinian’’ property), this class including the best-response dynamic and the replicator dynamics. The quite general flavour of this analysis is retained, but the dynamics are modified to include a fixed switching cost c , incurred whenever a player switches strategies.

Definition 1 A 1-incumbent (*resp.* 2-incumbent) in period t is a player who was a 1-strategist (*resp.* 2-strategist) in period $t-1$.

Definition 2 The cost-adjusted payoff matrices are

$$\Lambda_1 = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} - c & a_{22} - c \end{bmatrix} \quad (4)$$

and

$$\Lambda_2 = \begin{bmatrix} a_{11} - c & a_{12} - c \\ a_{21} & a_{22} \end{bmatrix} \quad (5)$$

for 1-incumbents and 2-incumbents, respectively.

Given the presence of this switching cost, KMR’s underlying *deterministic dynamic*,

$$z^{t+1} = b(z^t) \quad (6)$$

now has the following ‘‘modified Darwinian’’ property:^{5,6}

$$(b(z) - z) \text{ is } \left\{ \begin{array}{ll} \text{strictly negative} & \text{iff } \pi_1(z) < \pi_2(z) - c \\ \text{nonnegative} & \text{iff } \pi_1(z) \geq \pi_2(z) - c \\ \text{nonpositive} & \text{iff } \pi_1(z) - c \leq \pi_2(z) \\ \text{strictly positive} & \text{iff } \pi_1(z) - c > \pi_2(z) \end{array} \right\} \text{ for } z \neq 0, N. \quad (7)$$

Taking equation (7) line by line, it simply says that the number z of 1-strategists will decrease if and only if 1-incumbents are changing strategies; that z cannot

⁴See KMR (1993), p. 37, for models of matching which generate these payoffs. Note that a player will ignore his own play in forming his beliefs from the current population strategy frequency.

⁵The strict and weak inequalities are assigned according to the assumption that indifference leads to inertia at the individual level. This assumption is unimportant for the results.

⁶For the extreme states, 0 and N , it is assumed that $b(0) > 0$ iff $\pi_1(0) - c > \pi_2(0)$, and $b(0) = 0$ iff $\pi_1(0) - c \leq \pi_2(0)$. Similarly, $b(N) < N$ iff $\pi_1(N) < \pi_2(N) - c$, and $b(N) = N$ iff $\pi_1(N) \geq \pi_2(N) - c$.

decrease if 1-incumbents are *not* changing strategies; that z cannot increase if 2-incumbents are not changing strategies; and that z will increase if and only if 2-incumbents are changing strategies. Putting this information together, it follows that⁷

$$(b(z) - z) \text{ is } \left\{ \begin{array}{ll} \text{strictly negative} & \text{iff } \pi_1(z) - \pi_2(z) < -c \\ 0 & \text{iff } -c \leq \pi_1(z) - \pi_2(z) \leq c \\ \text{strictly positive} & \text{iff } \pi_1(z) - \pi_2(z) > c \end{array} \right\}$$

for $z \neq 0, N$. (8)

The model has thus been generalised so that KMR's analysis is now the special case where $c = 0$.

The model is now made stochastic by introducing some noise ("mutations") into the system. Following KMR, it is assumed that each player's strategy "flips" with probability ε in each period (i.i.d. across players and over time).⁸ This yields the nonlinear stochastic difference equation

$$z^{t+1} = b(z^t) + x^t - y^t, \quad (9)$$

where x^t and y^t have the binomial distributions:

$$x^t \sim \text{Bin}(N - b(z^t), \varepsilon) \quad \text{and} \quad y^t \sim \text{Bin}(b(z^t), \varepsilon)$$

The dynamical system in equation (9) defines a Markov chain on the finite state space $Z = \{0, 1, \dots, N\}$. $P^\varepsilon = [p_{ij}]$ is the Markov matrix, with the transition probabilities given by

$$p_{ij} = \Pr(z^{t+1} = j \mid z^t = i) \quad (10)$$

All elements in the matrix P^ε are strictly positive under the above assumptions.

The long-run behaviour of the Markov chain in equation (9) is captured by the *stationarity equations* $\mu^\varepsilon P^\varepsilon = \mu^\varepsilon$, the solution of which is the distribution μ^ε (over states) that is stationary under P^ε . However, it is not immediately clear whether this solution is unique. If it is not, long-run behaviour will be sensitive to initial conditions, so that the path dependence of deterministic models remains. However, for an *ergodic* Markov process, the stationarity equations *will* have a unique solution and the long-run behaviour embodied in μ^ε will be independent of initial conditions. The unique invariant distribution μ^ε may then be used to explore the equilibrium selection issue by providing information on whether some outcomes are much more likely than others. Indeed, μ^ε can be interpreted as the proportion of time that the society spends in each state in the long run.

Lemma 1 *An irreducible and aperiodic Markov chain is ergodic.*

⁷For the extreme states, similarly, $b(0) > 0$ iff $\pi_1(0) - \pi_2(0) > c$, and $b(0) = 0$ iff $\pi_1(0) - \pi_2(0) \leq c$; $b(N) < N$ iff $\pi_2(N) - \pi_1(N) > c$, and $b(N) = N$ iff $\pi_2(N) - \pi_1(N) \leq c$.

⁸The usual stories of "experimentation" or of players dying (with probability 2ε) and being replaced by ignorant newcomers may be appealed to. For example, see KMR (1993), pp. 38-9.

This is a standard result - see, for instance, Theorem 11.2 of Stokey, Lucas, and Prescott (1989).

Proposition 1 *The adaptive response dynamic defined by the transition probabilities p_{ij} in equation (10) is an irreducible, aperiodic Markov process on the finite state space Z . Consequently, it has a unique invariant (ergodic) distribution μ^ε .*

Proof. Since all elements in the Markov matrix P^ε are strictly positive, i.e. any state is accessible from any other state in a single period, it follows that the process is irreducible. Moreover, since in every state there is a positive probability of the system remaining in that state in the next period, the process is aperiodic. Lemma 1 then implies that the process has a unique invariant distribution. ■

The stage game Each time the players are matched to play the game Λ in (1), they will play a symmetric coordination game. The class of symmetric coordination games is the set of all games satisfying the conditions $a_{11} > a_{21}$ and $a_{22} > a_{12}$. Such a game evidently has two symmetric strict⁹ pure-strategy Nash equilibria, (1, 1) and (2, 2), with associated security (or miscoordination) payoffs a_{12} and a_{21} . Thus, the players would like to coordinate on one or other of the strategies. If $a_{11} \neq a_{22}$, then one of the pure-strategy Nash equilibria Pareto dominates the other, and the game is described as being one of *common interest*. There is also, however, a symmetric mixed-strategy Nash equilibrium entailing mixing probabilities of $[\rho, 1 - \rho]$, where

$$\rho = \frac{a_{22} - a_{12}}{(a_{22} - a_{12}) + (a_{11} - a_{21})}$$

There is thus a set of three Nash equilibria, $\Theta^{NE} = \{(1, 1), (2, 2), (\rho, \rho)\}$, with no apparent way of selecting between them - the classic game-theoretic problem of multiple equilibria.¹⁰ If players could coordinate, they would presumably coordinate on the Pareto-dominant equilibrium, but in the absence of explicit coordination it is not obvious that this will occur.

The most prominent solution to this equilibrium selection problem in 2×2 coordination games is Harsanyi and Selten's (1988) notion of "risk dominance", which captures the tradeoff between high payoffs and high risk faced by individual players in a noncooperative setting.¹¹ The risk-dominant equilibrium for a 2×2 game is the one that minimises the product of the players' losses associated with unilateral deviations. In terms of the stage game Λ , equilibrium (1, 1) risk-dominates (2, 2) if and only if $a_{11} - a_{21} > a_{22} - a_{12}$, corresponding exactly to $\rho < \frac{1}{2}$.

⁹A Nash equilibrium is *strict* if each player's strategy therein is his *unique* best response to those of his opponents.

¹⁰It is worth noting that the principal refinement of *deterministic* evolutionary game theory, namely *evolutionary stability*, is of some, limited use here: it is straightforward to check that the evolutionarily stable strategies (ESS) of a coordination game are the two pure-strategy equilibria.

¹¹This tradeoff was pointed out by Aumann (1987).

Example 1 *Risk dominance accounts well for persistent suboptimal coordination when the “optimal” strategy is “riskier” than the other, in the sense of being associated with a greater payoff loss if the other player miscoordinates. To borrow the example of Myatt and Wallace (1998), consider two players who play the following game in deciding whether to adopt the IBM or Apple Macintosh (Mac) standards in their personal computer work:*

	<i>IBM</i>	<i>Mac</i>
<i>IBM</i>	5	2
	5	4
<i>Mac</i>	4	6
	2	6

The players benefit from compatibility, but given compatibility, Mac adoption results in higher productivity than IBM adoption - (Mac, Mac) is the Pareto-dominant equilibrium. However, the loss from incompatibility is more severe for Mac users than for IBM users due, say, to wider outside support for the IBM standard - (IBM, IBM) is the risk-dominant equilibrium ($\rho = \frac{2}{5} < \frac{1}{2}$). The divergence of the Pareto-dominant and risk-dominant equilibria characterises the game as a Rousseau (1913) “Stag Hunt”, and if the risk-dominant equilibrium is played, it is clearly suboptimal.

However, when the “optimal” strategy is not “riskier” in the given sense, persistent suboptimal coordination *cannot* be explained by risk dominance. For example, the class of *pure* coordination games, characterised by zero security payoffs $a_{12} = a_{21} = 0$, sees the condition for risk-dominance of the (1, 1) equilibrium reduced to $a_{11} > a_{22}$, so that the risk-dominant and Pareto-dominant equilibria coincide. In view of this inability of risk dominance always to explain persistent suboptimal coordination, it is natural to consider what stochastic evolutionary models tell us about selection under realistic embellishments and generalisations, such as switching cost-driven inertia.

3.2 Analysis

3.2.1 Long-run equilibrium with vanishing noise

By rearranging the average payoffs in equations (2) and (3) to express them explicitly as linear functions of z ,

$$\pi_1(z) = \frac{(a_{11} - a_{12})}{(N - 1)}z + \frac{Na_{12} - a_{11}}{(N - 1)} \quad (11)$$

$$\pi_2(z) = \frac{(a_{21} - a_{22})}{(N - 1)}z + a_{22} \quad (12)$$




Figure 1: KMR (1993) model: coordination games

the two functions can easily be graphed, as in figure 1. Note that the $\pi_1(z)$ line will be steeper than the $\pi_2(z)$ line if and only if

$$\frac{(a_{11} - a_{12})}{(N - 1)} > \frac{(a_{21} - a_{22})}{(N - 1)}$$

$$\Leftrightarrow a_{11} - a_{21} > a_{12} - a_{22}$$

and that this holds necessarily for coordination games given that $a_{11} > a_{21}$ and $a_{22} > a_{12}$.¹² Moreover, these same two defining conditions of coordination games imply that $\pi_1(z) < \pi_2(z)$ for $z = 0$ and that $\pi_1(z) > \pi_2(z)$ for $z = N$ (as is clear from equations (2) and (3)), so that there will certainly exist a value $z^* \in Z$ where the two average payoff lines cross (though this z^* need not, of course, be an integer). This is KMR's critical level of population z^* for which

$$\text{sign}(\pi_1(z) - \pi_2(z)) = \text{sign}(z - z^*)$$

As KMR point out, z^* essentially corresponds to the mixed-strategy equilibrium of the coordination stage game, which puts probability $\rho = \frac{(a_{22} - a_{12})}{a_{11} - a_{21} + a_{22} - a_{12}}$ on

¹²There is, however, no implication that $\pi_1(z)$ need have a positive slope and $\pi_2(z)$ a negative slope.

strategy 1.¹³ It is straightforward to solve for the value of z^* :

$$\begin{aligned}
0 &= \pi_1(z^*) - \pi_2(z^*) \\
\stackrel{(2),(3)}{\implies} 0 &= \left(\frac{(z^* - 1)}{(N - 1)} a_{11} + \frac{(N - z^*)}{(N - 1)} a_{12} \right) - \\
&\quad \left(\frac{z^*}{(N - 1)} a_{21} + \frac{(N - z^* - 1)}{(N - 1)} a_{22} \right) \\
\Rightarrow z^* &= \frac{N(a_{22} - a_{12}) + a_{11} - a_{22}}{\theta}, \tag{13}
\end{aligned}$$

where $\theta = (a_{11} - a_{21} + a_{22} - a_{12})$ is the sum of the normalised stage game payoffs, $a_1 = a_{11} - a_{21}$ and $a_2 = a_{22} - a_{12}$.

In KMR's model, there are just two absorbing states, 0 and N , corresponding to the two pure-strategy Nash equilibria of the stage game Λ in equation (1).¹⁴ These states have basins of attraction given by $\{z < z^*\}$ and $\{z > z^*\}$ respectively, and the state with the larger basin of attraction (N if $z^* < \frac{N}{2}$; 0 if $z^* > \frac{N}{2}$) is the unique *stochastically stable outcome*.¹⁵ This outcome corresponds to the risk-dominant equilibrium of the underlying stage game Λ .¹⁶ The selection criterion of *stochastic stability* identifies which outcome(s) receive positive weight in the ergodic distribution as the amount of noise tends to zero: a state z is *stochastically stable* (Young 1993) if

$$\lim_{\varepsilon \rightarrow 0} \mu^\varepsilon(z) > 0$$

Intuitively, stochastically stable states are those that are most likely to be observed over the long run when noise is small, and they are thus sometimes referred to as the *long-run equilibria* of a system. Thus, the long-run equilibrium of the KMR coordination model is the risk-dominant equilibrium.

It will be seen below that the addition of switching costs to the KMR model does not alter the long-run selection of the risk-dominant equilibrium, but it does create new short- to medium-run equilibria in the form of a set of new absorbing states.

Proposition 2 *In the presence of a switching cost $c > 0$, there exists a multiplicity of mixed absorbing states¹⁷ of the unperturbed process P^0 in addition to the set of pure absorbing states $\{0, N\}$ present when $c = 0$.*

¹³In this population context, the *exact* analogue of the stage game mixed-strategy equilibrium is the state ρN (which again need not be an integer) where a fraction ρ of all players are playing strategy 1. Note that z^* will not be exactly equal to ρN because 1-incumbents and 2-incumbents face slightly different strategy distributions due to the finiteness of the population. The difference between z^* and ρN does, however, vanish as the population size becomes large.

¹⁴Note that z^* would also be an (extremely unstable) absorbing state if it happened to be an integer.

¹⁵See KMR's (1993) Theorem 3, p. 44.

¹⁶See KMR's (1993) Corollary 1, p. 46.

¹⁷A *mixed* state is one where both strategies are being played by some strictly positive number of players (as opposed to the two *pure* states, 0 and N). This has also been termed a *polymorphic profile* (by contrast with a *monomorphic profile*) in the literature (e.g. Robson and Vega-Redondo (1996)).

Proof. The modified Darwinian property of the selection dynamics in equation (8) says that there is a range of the average payoff difference between the two strategies, $(\pi_1(z) - \pi_2(z)) \in [-c, c]$, for which $(b(z) - z) = 0$. Since the average payoff lines in figure 1 will certainly cross at some $z^* \in Z$ (as was shown above), it follows that there will be a range of z , $[z_L, z_H] \in Z$, for which $(b(z) - z) = 0$. Now, it is clear from the deterministic dynamic in equation (6) that any state for which $(b(z) - z) = 0$ is an absorbing state of the unperturbed process P^0 (since then $z^{t+1} = z^t$ for $\varepsilon = 0$). Thus, there is a range of $z \in \{Z \setminus \{0, N\}\}$, each (integer) element of which is a (mixed) absorbing state. ■

In addition to the regular KMR absorbing states, $E_\emptyset = \{0\}$ and $E_N = \{N\}$, there is thus now also the new set of mixed absorbing states, $E_M = \{z \in Z \mid (\pi_1(z) - \pi_2(z)) \in [-c, c]\}$, providing the system with additional rest points. This means that the model now has several new *mixed* equilibria in the short to medium run, any of which is likely to persist for a significant period of time once reached. Thus, if the initial conditions dictate that the system is in one of the mixed absorbing states at the outset, then it is likely to remain there for a nontrivial length of time, and to remain within the *region* of mixed absorbing states E_M for considerably longer.¹⁸

The new mixed absorbing states also alter the KMR analysis of long-run equilibrium. Both here and in the KMR model, this analysis faces the problems associated with a finite population, for which reason KMR make use of the integers around z^* :

$$\begin{aligned}\alpha &= \min \{z \in Z \mid \pi_1(z) > \pi_2(z)\} = \lceil z^* \rceil, \text{ and} \\ \beta &= \max \{z \in Z \mid \pi_1(z) < \pi_2(z)\} = \lfloor z^* \rfloor\end{aligned}$$

These integers are illustrated in figure 2. In order to consider the new set of mixed absorbing states, E_M , use must also be made of the integers around z_L and z_H , the limits of the range of z for which $(b(z) - z) = 0$.

Definition 3 *Given equation (8), the lower limit z_L of the range of z for which $(b(z) - z) = 0$ is defined as*

$$\begin{aligned}-c &= \pi_1(z_L) - \pi_2(z_L) \\ \xrightarrow{(2),(3)} z_L &= \frac{N(a_{22} - a_{12} - c) + a_{11} - a_{22} + c}{\theta},\end{aligned}\tag{14}$$

whilst the upper limit z_H is defined as

$$\begin{aligned}c &= \pi_1(z_H) - \pi_2(z_H) \\ \xrightarrow{(2),(3)} z_H &= \frac{N(a_{22} - a_{12} + c) + a_{11} - a_{22} - c}{\theta}\end{aligned}\tag{15}$$

¹⁸How long will of course depend upon the precise parameters of the game.




Figure 2: KMR coordination games with switching costs

Definition 4 *The integers around z_L are defined as*

$$\alpha_L = \min \{z \in Z \mid \pi_1(z) - \pi_2(z) > -c\} = \lceil z_L \rceil, \text{ and}$$

$$\beta_L = \max \{z \in Z \mid \pi_1(z) - \pi_2(z) < -c\} = \lfloor z_L \rfloor$$

whilst those around z_H are defined as

$$\alpha_H = \min \{z \in Z \mid \pi_1(z) - \pi_2(z) > c\} = \lceil z_H \rceil, \text{ and}$$

$$\beta_H = \max \{z \in Z \mid \pi_1(z) - \pi_2(z) < c\} = \lfloor z_H \rfloor$$

These integers are again illustrated in figure 2, and they immediately allow the following characterisation of the set of mixed absorbing states:

$$E_M = \{\alpha_L, \alpha_L + 1, \dots, \beta, \alpha, \dots, \alpha_H, \beta_H\}$$

$$= \{\alpha_L + j\}, \quad j = 0, 1, \dots, (\beta_H - \alpha_L) \quad (16)$$

The number m of these mixed absorbing states is clearly given by

$$m = \beta_H - \alpha_L + 1$$

$$= \beta_H - \beta_L = \alpha_H - \alpha_L \quad (17)$$

Defining $\zeta_L = \alpha_L - z_L$ and $\zeta_H = \alpha_H - z_H$, m can then be expressed as a function of the parameters of the model:

$$m = (z_H + \zeta_H) - (z_L + \zeta_L)$$

$$= \frac{2(N-1)c}{\theta} + (\zeta_H - \zeta_L) \quad (18)$$

The number of mixed absorbing states m is thus increasing in N and c , but decreasing in the sum of the normalised coordination payoffs θ . The total number n of absorbing states (including the extreme states 0 and N) is of course $n = m + 2$.

In order to investigate the effect of the new mixed absorbing states on the long-run equilibrium of the model, use is made of Young's (1993) method¹⁹ for computing the stochastically stable states of a regular²⁰ perturbed Markov process. This method is based on the notion of *rooted trees* constructed on the set of the recurrent class(es) E_1, E_2, \dots, E_K of the unperturbed process P^0 . For a given pair of (distinct) recurrent classes E_i and E_j , an ij -path is a sequence of states that begins in E_i and ends in E_j . The *resistance* of this path is then the number of mutations required to transit from recurrent class E_i to recurrent class E_j along this path, and the minimum resistance over all possible ij -paths is denoted r_{ij} .²¹ Construct a complete directed graph with K vertices (one for each recurrent class), and weight each directed edge $i \rightarrow j$ with the appropriate minimum resistance r_{ij} . A tree rooted at a particular recurrent class E_j 's vertex j (a j -tree) is then a set of $(K - 1)$ directed edges such that there is a unique directed path in the tree to j from every vertex (i.e. recurrent class) other than j . The *resistance of a rooted tree* T is the sum of the resistances r_{ij} on T 's $(K - 1)$ edges, and the minimum resistance over all trees rooted at j is called the *stochastic potential* γ_j of the recurrent class E_j .

Lemma 2 (Young (1993)) *If P^ε is a regular perturbed Markov process, and μ^ε is its unique stationary distribution for each $\varepsilon > 0$, then $\lim_{\varepsilon \rightarrow 0} \mu^\varepsilon(z) = \mu^0$ exists, and the limiting distribution μ^0 is a stationary distribution of the unperturbed process P^0 . The stochastically stable states are precisely those states contained in the recurrent class(es) of P^0 having minimum stochastic potential.*

The intuition for this result is that, for a small and positive noise level ε , the process is most likely to follow paths leading towards the recurrent classes having minimum potential.

Lemma 3 *The most efficient (i.e., lowest resistance) j -tree for any given recurrent class j of the KMR coordination game model under switching costs involves direct jumps out of the basins of attraction of the extreme state/s, but one-step*

¹⁹See Young (1998), section 3.4, for more detail and illustrative examples.

²⁰Let P^ε be a Markov process on Z for each ε in some interval $[0, \varepsilon^*]$. P^ε is described as a *regular* perturbed Markov process if P^ε is irreducible for every $\varepsilon \in (0, \varepsilon^*]$, and for every $i, j \in Z$, p_{ij}^ε approaches p_{ij}^0 at an exponential rate, that is,

$$\lim_{\varepsilon \rightarrow 0} p_{ij}^\varepsilon = p_{ij}^0,$$

and

$$\text{if } p_{ij}^\varepsilon > 0 \text{ for some } \varepsilon > 0, \text{ then } 0 < \lim_{\varepsilon \rightarrow 0} \frac{p_{ij}^\varepsilon}{\varepsilon^{r(i,j)}} < \infty \text{ for some } r(i,j) \geq 0.$$

The real number $r(i, j)$ is called the *resistance* of the transition $i \rightarrow j$. Note that transitions that can occur under P^0 have zero resistance (i.e. $p_{ij}^0 > 0$ if and only if $r(i, j) = 0$).

²¹Note that other evolutionary models (such as KMR) employ the analogous notion of the *cost* of a transition path, rather than the *resistance*.

jumps in the direction of j 's basin of attraction within the region of mixed absorbing states.

Proof. Under vanishing noise, the number of mutations should be minimised in order to identify the most efficient j -tree. Given the retarding effect of the deterministic dynamic $b(z)$, this clearly entails direct jumps²² out of the basins of attraction of the extreme states 0 and N .²³ But this retarding effect is not at work within the region of mixed absorbing states E_M , where the deterministic dynamic dictates inertia, and thus a sequence of gradual steps more probable than individual direct jumps can be utilised. In particular, note that there must be a transition path (or “branch”) rooted at each recurrent class in a j -tree. Minimising mutations once more dictates that the branches rooted at mixed absorbing states should be “one-step” transitions into an adjacent state.

Now, the most efficient way for the j -tree to enter j 's basin of attraction is via the mixed absorbing states, since their branches must exist anyway, and their resistance is unaffected by their direction. This “saves” a number of direct jumps out of the extreme state/s. Hence the one-step transitions rooted in the E_M region should be directed towards j 's basin of attraction. ■

Lemma 3 implies most efficient j -trees of the form illustrated in figure 3.²⁴ In particular, consider the E_M region of the most efficient 8-tree in figure 3(b). Clearly the path of least resistance from state 5 to state 8 is that where a single mutation takes the system into state 6, and thus the basin of attraction of state 8. Conditional on this path existing, the path of least resistance from state 4 to state 8, meanwhile, involves a single mutation into state 5, from where the path to state 8 already exists. Given the stability of state 5, there is nothing to be gained (in terms of efficiency) by jumping directly from state 4 to state 6, and indeed it is counter-productive (less probable) to do so. Similar reasoning applies to state 3 conditional on the existence of the $4 \rightarrow 8$ path.

Proposition 3 *Provided that $z^* \neq \frac{N}{2}$ and $N \geq 2$, the risk-dominant equilibrium remains the unique stochastically stable state (and is thus selected with probability one by the limit distribution) in the presence of a fixed switching cost c , until c exceeds the maximum payoff gain at stake in the game.*

Proof. The resistances of Lemma 3's most efficient j -trees (i.e., the stochastic potentials γ_j of the system's recurrent classes (see page 15 above)) are straightforward to calculate by “counting mutations”. Thus, consulting figure

²²By “direct jumps” is meant just enough simultaneous mutations to move between the two states concerned in one period.

²³This is precisely the argument in the KMR model without switching costs - see KMR (1993), pp. 44-46.

²⁴In figure 3, each square corresponds to a state, the shaded squares being absorbing states of the unperturbed process. Solid arrows constitute branches of the relevant j -tree, whilst dashed arrows represent the deterministic dynamic (paths of zero resistance).

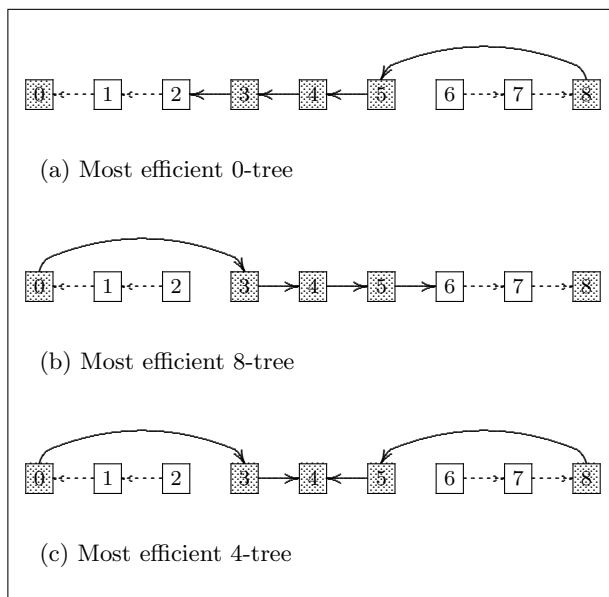


Figure 3: Most efficient j -trees, $N = 8$

2, for the two pure absorbing states:

$$\begin{aligned}\gamma_\emptyset &= m + (N - \beta_H) \\ &= N - \beta_L\end{aligned}\tag{19}$$

and

$$\begin{aligned}\gamma_N &= \alpha_L + m \\ &= \alpha_H\end{aligned}\tag{20}$$

The stochastic potential of any given mixed absorbing state²⁵ $z_{M,k} \in E_M$, meanwhile, is given by

$$\begin{aligned}\gamma_k &= \alpha_L + (N - \beta_H) + (m - 1) \\ &= N\end{aligned}\tag{21}$$

Clearly $\gamma_k > \gamma_\emptyset, \gamma_N$, so that the mixed absorbing states $z_{M,k} \in E_M$ cannot be stochastically stable, by Lemma 2 (page 15). Comparing the two pure absorbing states, meanwhile, it is clear that N has the lower (and thus minimum) stochastic potential if and only if

$$\begin{aligned}\gamma_N &< \gamma_\emptyset \\ \stackrel{(19),(20)}{\iff} \alpha_H &< N - \beta_L\end{aligned}$$

²⁵ $z_{M,k} = \{\alpha_L + k\}$, where k can equal $0, 1, \dots, (\beta_H - \alpha_L)$.

Ignoring the ζ 's, this condition essentially corresponds to

$$\begin{aligned}
& z_H < N - z_L \\
\stackrel{(14),(15)}{\implies} & \frac{N(a_{22} - a_{12} + c) + a_{11} - a_{22} - c}{\theta} < \\
& N - \frac{N(a_{22} - a_{12} - c) + a_{11} - a_{22} + c}{\theta} \\
& \frac{N(a_{22} - a_{12}) + a_{11} - a_{22}}{\theta} < \frac{N}{2} \\
& \stackrel{(13)}{\implies} z^* < \frac{N}{2} \tag{22}
\end{aligned}$$

But this is just the condition for N being the risk-dominant equilibrium. Thus, N has minimum stochastic potential - and hence, by Lemma 2, is (uniquely) stochastically stable - under precisely the same condition as in KMR (1993, Theorem 3, p. 44), i.e. that 1 be the risk-dominant strategy.

However, this analysis assumes that the set of mixed absorbing states E_M is wholly contained within the state space, i.e. $E_M \in (0, N)$. For large enough values of c , β_L will fall below zero and α_H will rise above N , so that E_θ and E_N are absorbed within E_M . Assume, without loss of generality, that 1 is the risk-dominant strategy in the stage game Λ (equation (1)), so that N is the population analogue of the risk-dominant equilibrium (and $z^* < \frac{N}{2}$). In this case, there will be a range of c ($a_{22} - a_{12} < c < a_{11} - a_{21}$) for which E_θ has been absorbed into E_M but E_N has not; over this range, $\gamma_\theta = N$, whilst γ_N remains α_H , and the risk-dominant equilibrium N thus remains the unique stochastically stable outcome. Since the only restriction on a_{12} is that it must be less than a_{22} , this possibility should not be ignored, although it is of secondary interest here given that the focus is on small switching costs. Of even less interest is the case where c is high enough ($c > a_{11} - a_{21}$) to outweigh the maximum possible payoff gain from a strategy switch, from which point $\alpha_H > N$, and E_N is thus also absorbed into E_M . In this case, γ_N too becomes N , and all states $z \in Z$ become stochastically stable. The risk-dominant equilibrium is thus overturned as the unique long-run equilibrium at this point, but only by begging the question with the size of the switching costs - absolute inertia prevails.²⁶ ■

Thus, the long-run equilibrium of the KMR model under vanishing noise is unchanged in the presence of switching costs, and it would seem that risk-dominance continues to reign supreme over time. Having said this, however, it is easy to imagine a variety of possible modifications to the model which might overthrow the long-run selection of the risk-dominant equilibrium. For example, with sufficiently asymmetric switching costs (so that switching from strategy 2 to strategy 1 were harder than the reverse), the risk-dominated equilibrium would

²⁶Note that if the players' myopia were demonstrated to be rational, rather than simply assumed, then these larger values of c would be of more interest, and the results would provide an explanation for persistent suboptimal coordination in the long run. For investigations of the justification of the myopia assumption, see Ellison (1997) and Blume (1995).

clearly be selected. Indeed, one plausible possibility that might be considered is a one-way switching cost, so that it is costly to switch from strategy 2 to strategy 1, but not to switch back. This might be a feature of a model of technological standards, where there is a cost of technological advance but no cost of regress. This could clearly lead to a situation of “lock-in” at the inferior standard in the long run.²⁷

KMR’s long-run selection results have thus become vulnerable in a way that was impossible in their model. This is the case because the model with switching costs modifies the KMR deterministic dynamic in a way not allowed for by their assumptions. Despite the generality of KMR’s selection dynamics, then, a limit to their applicability has been found (in the same way as pure rationality lies outside the opposite limit). Admittedly, though, the sort of modification required to overthrow the risk-dominant equilibrium in the long run might be considered less than appealing. Asymmetric switching costs, for instance, would appear to rob the concept of risk-dominance of its essence, so that its predictive content should be expected to follow in its wake.

3.2.2 Transition times

It would thus seem that the effects of inertia are unlikely to be significant for long-run selection. However, unsurprisingly, inertia *is* likely to be significant in determining *how long* is the long run. Moreover - and somewhat more surprisingly - the presence of inertia serves, not to increase, but to *decrease* transition times, due to an unexpected side-effect of the new inertia-born mixed absorbing states E_M .

In order to demonstrate this, Ellison’s (2000) “radius-modified coradius” approach (mentioned in section 2 above) is employed. Ellison defines the *radius* $R(j)$ of the basin of attraction of a recurrent class j as the minimum number of mutations necessary to escape that basin of attraction; this serves as a measure of j ’s persistence once reached. The *coradius* $CR(j)$ of the basin of attraction of j , meanwhile, is defined as the maximum over all other states of the minimum number of mutations necessary to reach j ; this serves as a measure of j ’s tendency to be reentered once left. A better measure of this tendency, however, is provided by Ellison’s *modified coradius* $CR^*(j)$ of j ’s basin of attraction, which captures the extent to which large evolutionary changes will occur more quickly when they can be achieved by passing through a number of intermediate steady states. This modified coradius is found by subtracting from the coradius, $CR(j)$, the radius of each of the intermediate recurrent classes $i \in I$ through which the path (from the coradius’s root to j) passes. In terms of Young’s (1993) resistance notation (see page 15 above),

$$CR^*(j) = \max_{z \notin j} r_{zj}$$

²⁷Note, however, that such a conclusion would be somewhat dubious here, given that the long run involves time spent in each equilibrium, and presumably such a switching cost should be incurred only once.

and

$$CR^*(j) = CR(j) - \sum_{i \in I} R(i)$$

Lemma 4 (Ellison (2000)) *If $R(j) > CR^*(j)$, then: (a) the (long-run) stochastically stable set of the model is contained in j ; and (b) for any $z \notin j$ the expected waiting time $W(z, j, \varepsilon)$ until a state belonging to j is first reached, given that play in the ε -perturbed model begins in state z , is $W(z, j, \varepsilon) = O(\varepsilon^{-CR^*(j)})$ as $\varepsilon \rightarrow 0$.²⁸*

Proposition 4 *The presence of switching costs serves to reduce the expected waiting time before the long-run equilibrium is reached. Moreover, this expected waiting time is decreasing in the size of the switching cost c .*

Proof. In the KMR model with switching costs, the radius of the basin of attraction of the risk-dominant equilibrium N is clearly²⁹

$$R(N) = N - \beta_H \tag{23}$$

whilst the coradius is

$$CR(N) = \alpha_H$$

The set I of intermediate recurrent classes through which the path from (the coradius's root at) 0 to N passes is simply the set of mixed absorbing states E_M , each of which has a radius of 1. The modified coradius is thus

$$\begin{aligned} CR^*(N) &= \alpha_H - \sum_{k=0}^{\beta_H - \alpha_L} R(M_k) \\ &= \alpha_H - m \\ &= \alpha_L \end{aligned} \tag{24}$$

It is thus clear from equations (23) and (24) that $R(N) > CR^*(N)$ if and only if

$$N - \beta_H > \alpha_L$$

Ignoring the ζ 's and rearranging, this condition essentially corresponds to

$$N - z_L > z_H$$

This is precisely the condition for long-run selection of the risk-dominant equilibrium N in equation (22). Since the first part of Lemma 4 dictates that N is the (long-run) stochastically stable outcome if $R(N) > CR^*(N)$, this confirms Proposition 3.

²⁸ $W(z, j, \varepsilon)$ is of order $\frac{1}{\varepsilon^\delta}$, denoted $O(\varepsilon^{-\delta})$, as $\varepsilon \rightarrow 0$ if and only if $\text{plim } \varepsilon^\delta W(z, j, \varepsilon)$ is a finite nonzero constant. It is of order less than $\frac{1}{\varepsilon^\delta}$, denoted $o(\varepsilon^{-\delta})$, meanwhile, if and only if $\text{plim } \varepsilon^\delta W(z, j, \varepsilon)$ equals zero.

²⁹Refer to figure 2 (page 14).

Now, assuming (without loss of generality) N to be the risk-dominant equilibrium once again, it follows that $R(N) > CR^*(N)$. The second part of Lemma 4 thus implies that, for any $z \neq N$ the expected waiting time until N is first reached given that play in the ε -perturbed model begins in state z is $W(z, N, \varepsilon) = O(\varepsilon^{-\alpha_L})$ as $\varepsilon \rightarrow 0$. It should be clear from figure 1 (page 11) that, in the KMR model *without* switching costs, $R(N) = N - \beta$ and $CR^*(N) = CR(N) = \alpha$. The introduction of switching costs thus changes the expected waiting time from $O(\varepsilon^{-\alpha})$ to $O(\varepsilon^{-\alpha_L})$, which is clearly a reduction given that $\alpha_L < \alpha$. Moreover, since $\alpha_L \approx z_L$, and

$$\frac{\partial z_L}{\partial c} = - \left(\frac{N-1}{\theta} \right) < 0$$

(by equation (14)) it follows that the expected waiting time $W(z, N, \varepsilon)$ before the risk-dominant equilibrium is reached is a decreasing function of the switching cost c .³⁰ ■

Thus, the surprising conclusion is reached that the presence of individual-level inertia serves to *improve* the overall efficiency of the evolutionary process by speeding the transition time to long-run equilibrium.³¹ Whilst initially counter-intuitive, this result is elucidated by Ellison's (2000) account of the role of intermediate stable "steps" in speeding evolution between extremes, discussed in section 2 above. In this case, the new mixed absorbing states E_M provide resting points for the population on its way from 0 to N , thus increasing the likelihood of this path and reducing the expected time taken before it is observed. This also, of course, applies in reverse: Ellison's (2000, pp. 31-32) result that the expected waiting time for leaving any given recurrent class j 's basin of attraction is $O(\varepsilon^{-R(j)})$ implies that less undisturbed time is on average spent in the state N (and the state 0) in the presence of switching costs (since both of their radiuses fall under switching costs). All of this serves to reduce the force of the Ellison (1993) "very long run" critique that transition times become prohibitively large as noise tends to zero.

The Ellison (1993) critique is also questioned from another direction by the introduction of switching costs, specifically in its contention that the "very long run" problem becomes more acute the larger the size of the population N . This is typically true in existing models, given that convergence times tend to be rapidly increasing in population size, but this problem is diluted somewhat in the presence of switching costs.

³⁰Note again, however, the complications when E_M is not wholly contained within the state space Z . Once c has risen high enough that β_L has fallen below 0 ($c > a_{22} - a_{12}$), the modified coradius (and thus the exponent of the order of the expected waiting time) of N will remain at 0 until c has risen to such a point ($c > a_{11} - a_{21}$) that α_H exceeds N . At this point it is no longer the case that $R(N) > CR^*(N)$ and, as was seen above, N ceases to be the unique stochastically stable outcome. Again, though, these cases (and particularly the latter case) are of limited interest given the size of the switching costs required.

³¹Note that all references to reduced expected waiting (or transition) times in fact refer to a reduction in the *order* of expected waiting times. This is common practise in the literature on stochastic adjustment dynamics, and indeed the order is all that matters as noise vanishes.

Proposition 5 *The presence of switching costs serves to attenuate the positive effect of population size on transition times.*

Proof. Continuing along the lines of the proof of Proposition 4, it is clear from equation (14) that

$$\frac{\partial z_L}{\partial N} = \frac{a_{22} - a_{12} - c}{\theta} \quad (25)$$

Since $a_{22} > a_{12}$ by definition for coordination games, the positive effect of N on z_L (and thus on the convergence time to N given that $CR^*(N) = \alpha_L \approx z_L$) in the KMR case of zero switching costs is manifest. A positive value for the switching cost, however, clearly attenuates this positive effect, to a degree increasing in the size of the switching cost c .³² ■

The intuition for this result is again based on the accelerated “step-by-step” evolution allowed by the presence of the new mixed absorbing states. As the size of the population N grows, the number of mixed absorbing states m also grows, providing more intermediate stable “steps” for the population to evolve through. It is still the case that a larger N necessitates more mutations in order for the population to leave the extreme absorbing states’ basins of attraction, but these required numbers of mutations increase more slowly with N in the presence of the mixed absorbing states E_M . Moreover, the more mixed absorbing states there are (i.e., the higher the switching cost c), the stronger this attenuating effect will be. Thus, under switching costs, the Ellison (1993) contention that transition times under vanishing noise become unreasonably long as N becomes large is questioned to a degree increasing in the size of the switching costs.

4 State-Dependent Mutations

The principal doubts over the conclusions reached in the previous section are raised by two of the main criticisms of the model employed, namely the implausibility of a *fixed* and *uniform* switching cost, and the problems associated with *state-independent mutations*. Whilst it is highly realistic that individual players will face switching costs of changing their strategies, it is far less realistic that there will be a single time-invariant switching cost that is the same across all players, as was the case with c in the above KMR-style model. Moreover, the model is clearly subject to the Bergin and Lipman (1996) critique of models with state-independent mutations, the mutation rate ε remaining constant across states. Since mutations may be chosen such that any invariant distribution of the unperturbed process is the limiting ergodic distribution of a perturbed process, Bergin and Lipman argue that state-independent mutations are arbitrary, and that an economically justified model of *state-dependent* mutations should instead be employed. It is these criticisms that motivate the next model.

³²Although once $c > a_{22} - a_{12}$, β_L falls below 0 and $CR^*(N)$ is fixed at 0 until α_H rises above N ($c > a_{11} - a_{21}$), as discussed in note 30.

The obvious response to the first criticism is to drop the assumption of a fixed switching cost, and instead allow c to be stochastic, determined as the realisation of a random variable C . Each individual player could then take a draw from C each period to determine his switching cost for that period, yielding the realistic feature of *idiosyncratic, time-varying* switching costs. This sort of *stochastic switching costs* model would explicitly incorporate player heterogeneity, with players varying in their switching costs according to abilities, situations, priorities and so on.³³ Happily, this step also has the potential to address the second criticism by endogenising the mutation rate in the manner suggested by Bergin and Lipman. The stochastic switching costs can provide the “error” necessary to yield an ergodic Markov chain, with this “error” now being interpreted as players behaving differently to what one would expect from the *cost-less* payoff matrix Λ in equation (1).

However, if - as the intuition for switching costs would at first suggest - the support of the random variable C were to be restricted to the positive real line \mathbb{R}_+ , then the resulting Markov process would not be irreducible; “mutations” against the flow of the deterministic dynamic would have zero probability since, with a positive switching cost, no player will ever switch away from the strategy which currently has the higher expected payoff.³⁴ A reducible Markov chain remains subject to path-dependence, and does not yield an ergodic long-run distribution. To attain this, the state-independent mutation rate ε of the earlier chapters could be re-introduced, but this seems ad hoc and defeatist. Instead, the support of the switching cost random variable C could be extended to the whole real line \mathbb{R} , allowing the possibility of switching *benefits*. The existence of switching benefits in some players, grounded for example in an urge for creativity or nonconformity, is a realistic feature to incorporate within the model, and can be kept relatively improbable by assuming that C is distributed with positive mean. This will suffice to deliver an irreducible Markov chain - all states now being accessible from all others - and thus the desired ergodic long-run distribution over states.

However, the move to stochastic switching costs is shown to have no effect on the long-run selection results of section 3, thus confirming the finding that (small) switching costs do not alter the long-run equilibrium of evolutionary coordination games. Transition times, on the other hand, *are* affected by the move to the stochastic switching costs model. In particular, the counter-intuitive result that inertia speeds evolution is qualified in the model with state-dependent mutations; the “step-by-step” effect of section 3 is still at work, but a counter-vailing force emerges that slows evolution. Nonetheless, the net effect is still to speed up evolution, provided that the mean switching cost is not “too high”.

³³The model would thus have parallels with the Myatt and Wallace (1998) model of adaptive play by idiosyncratic agents.

³⁴This observation illustrates well the Bergin and Lipman (1996) criticism of the arbitrariness of state-independent mutations: there might be no reason to expect certain sorts of mutations under certain circumstances, and to ignore this is to assume away path-dependence when it may be an essential feature of the real-life process. The model here seeks to avoid this problem by providing an *economically justified* model of mutations which generates irreducibility.

4.1 The Model

The basic KMR structure of a finite population of N players repeatedly playing the 2×2 game Λ (equation (1)) is retained, with associated cost-adjusted payoff matrices Λ_1 and Λ_2 (equations (4) and (5)). The first major change to the model is that the switching cost c in any given period is now determined for each player individually as an independent and identically distributed draw from the switching cost random variable C , which is assumed to have a cumulative distribution function (cdf) F with a positive mean $\bar{c} > 0$ and variance σ^2 . This is a natural, general representation of differing switching costs across players, the positive mean focusing attention on switching *costs* but the infinite support delivering a small probability of switching *benefits*.

The expected payoffs are as they were in section 3 (see equations (2) and (3), page 7). Define $\varpi_1(z) = \pi_1(z) - \pi_2(z)$, and $\varpi_2(z) = \pi_2(z) - \pi_1(z)$, to be the expected payoff gains at stake from strategy switches in state z . Clearly a 1-incumbent will switch to being a 2-strategist if and only if

$$\varpi_2(z) > c$$

whilst a 2-incumbent will switch to being a 1-strategist if and only if

$$\varpi_1(z) > c$$

Defining s_l^t to be the strategy of player l in period t , the switching probabilities conditional on player l having been selected for review are then immediate:

$$\begin{aligned} \Pr(s_l^{t+1} = 2 \mid s_l^t = 1) &= F(\varpi_2(z) - c) \\ \Pr(s_l^{t+1} = 1 \mid s_l^t = 2) &= F(c - \varpi_1(z)) \end{aligned}$$

The probability of a selected 2-incumbent switching to strategy 1 in a given state is illustrated in figure 4 ($f(\cdot)$ representing the switching cost random variable C 's probability density function (pdf)).³⁵

The second major change to the KMR-style model of section 3 is in the selection dynamic employed. In this more complicated setting, it is desirable for the sake of mathematical convenience to focus attention on the best-response dynamic $B(z)$ - where players play their *best response* to current strategy frequencies - rather than KMR's more general "Darwinian" dynamic $b(z)$. This simplifies the analysis and also serves to isolate the effect of inertia on the system's review rate, without the distraction of having other forms of bounded rationality at work. The effect of the selection dynamic can still be analysed in this setting by comparing the extreme cases of the simultaneous-revisions dynamic of KMR (1993) and others, and the "one-step" single-revisions dynamic favoured by Binmore and Samuelson (1997), Blume (1999), and Myatt and Wallace (1998). Under the "one-step" dynamic, a *single* randomly selected member of the population has the opportunity to revise his strategy at the end

³⁵Whilst the pdf depicted in figure 4 is a Normal density, $f(\cdot)$ is a general pdf in the model, and thus can take any form.




Figure 4: The switching cost C 's pdf

of each period. Given his realised value of c , this player observes the strategy distribution among the incumbent population and selects a best response to this frequency.³⁶ In the simultaneous-revisions model by contrast, *each* player has this opportunity to revise his strategy each period. Evolution can thus proceed far more rapidly under simultaneous revisions, *ceteris paribus*.³⁷

4.2 Analysis

4.2.1 Single revisions

In this subsection, use will be made of the “one-step” best-response dynamic $B^1(z)$, where just one player at a time has the opportunity to revise his strategy. This gives a “birth-death” process, which is straightforward to analyse. Given that the reviewing player l is randomly selected, the probability that he is a 1-incumbent is simply $\frac{z}{N}$, implying the following transition probabilities between states.

Lemma 5 *The transition probabilities p_{ij} of the Markov matrix P satisfy:*

$$p_{ij} = \left\{ \begin{array}{ll} \left(\frac{i}{N} \right) F(\varpi_2(i-1)) & j = i-1 \\ \left(\frac{i}{N} \right) (1 - F(\varpi_2(i-1))) + \left(\frac{N-i}{N} \right) (1 - F(\varpi_1(i))) & j = i \\ \left(\frac{N-i}{N} \right) F(\varpi_1(i)) & j = i+1 \end{array} \right\}$$

and are zero elsewhere.

³⁶An equivalent scenario is the familiar story of player exit and entry, whereby a randomly selected player leaves at the end of each period, and is replaced by another player with a new draw from C .

³⁷Both the single- and simultaneous-revisions dynamics require the implicit assumption that the switching cost of an updating individual will have changed since the last revision, since he must take a fresh draw from C . This is a reasonable assumption when the noise is small, and thus unproblematic for vanishing heterogeneity results obtained as σ tends to zero. More generally, the procedure can be justified by noting that players are more likely to revise their strategy whenever their switching costs change.

Proof. Given the single-revisions framework, the process cannot move from state i to $j < i - 1$ or $j > i + 1$. A move to state $i + 1$ requires that a 2-incumbent be selected for review, and that having been selected his best response to the current strategy frequency be to switch to strategy 1. The former occurs with probability $(N - i)/N$; the latter with probability $F(\varpi_1(i))$. Similar arguments apply for the cases $j = i$ and $j = i - 1$. ■

Proposition 6 *The adaptive response dynamic defined by the transition probabilities p_{ij} in Lemma 5 is an irreducible, aperiodic Markov process on the finite state space Z . Consequently, it has a unique invariant distribution.*

Proof. Since the normal distribution has full support, either strategy may be chosen by any reviewing player. The process can thus move in either direction from any state i (except the extreme states), as formalised in Lemma 5, so that every state is accessible from all others in finite time - i.e., the process is irreducible. Moreover, since in every state there is a positive probability of the system remaining in that state in the next period, the process is aperiodic. The process thus has a unique invariant distribution by Lemma 1. ■

Long-run equilibrium with vanishing heterogeneity In the previous section, long-run equilibrium results were obtained by analysing the limit of the ergodic distribution as the probability of mutation ε tended to zero. The analogue in this model, as in the model of Myatt and Wallace (1998), is vanishing player heterogeneity over switching costs - i.e., taking $\sigma^2 \rightarrow 0$.

Given the departure from a uniform mutation rate ε , the simple Young stochastic potential technique of Lemma 2 cannot be applied here. Fortunately, Young's method fits into a wider *graph-theoretic* approach to the analysis of the long-run behaviour of perturbed Markov chains. This graph-theoretic approach is in turn derived from general Markovian theory.³⁸ If ε were known precisely, it would (in theory) be possible to compute the actual distribution μ^ε by simply solving the stationarity equations $\mu^\varepsilon P^\varepsilon = \mu^\varepsilon$. However, in most applications of interest in economics, the size of the state space would make this a very cumbersome task. This fact led to the importing of the simplified graph-theoretic techniques of Freidlin and Wentzell (1984) into the economics discipline by Foster and Young (1990), KMR (1993), and Young (1993), and the resulting birth of stochastic adjustment dynamics.

Like the computation of Young's stochastic potential function, the graph-theoretic computation of μ^ε is based on the notion of rooted trees, but this time constructed on the whole state space Z rather than merely on the set of recurrent classes. Specifically, consider a directed graph whose vertex set is the state space Z . The edges of this graph form a *z-tree* (for some particular $z \in Z$) if it consists of $|Z| - 1$ edges and from every vertex $i \neq z$ there is a unique directed path from i to z . A *z-tree's* edges are weighted with the appropriate Markov transition probabilities p_{ij} . Representing any given directed edge $i \rightarrow j$ by the ordered pair of vertices (i, j) , a *z-tree* T can then be represented as a

³⁸A standard reference for Markovian theory is Karlin and Taylor (1975).




Figure 5: Unique positively weighted 3-tree, $N = 8$

subset of ordered pairs. Let \mathcal{T}_z be the family of all z -trees for a given z . Define the *likelihood* of z -tree $T \in \mathcal{T}_z$ to be

$$p(T) = \prod_{(i,j) \in T} p_{ij}$$

Lemma 6 (Freidlin and Wentzell (1984)) *Let P be an irreducible Markov process on a finite state space Z .³⁹ Its stationary distribution μ has the property that the probability $\mu(z)$ of each state z is proportional to the sum of the likelihoods of its z -trees, that is,*

$$\mu(z) = \frac{v(z)}{\sum_{i \in Z} v(i)}, \text{ where } v(z) = \sum_{T \in \mathcal{T}_z} p(T) \quad (26)$$

This result allows computation of an *exact* estimate of a system’s ergodic distribution μ^ε for each $\varepsilon > 0$;⁴⁰ it is not a limiting result as ε tends to zero.

The additional analytical power of the Freidlin-Wentzell method in providing an immediate closed form for the invariant distribution μ comes at the price of the greater complexity inherent in constructing trees on the whole state space (rather than on the set of recurrent classes, as in Young’s method⁴¹). The number of z -trees to be considered soon becomes prohibitively large as the state space grows. This is where the use of the single-revisions dynamic becomes useful, since it means that $v(z)$ in equation (26) takes a very simple form. With only one revision at a time, there is only one possible positively weighted z -tree for any given state z - that z -tree involving successive one-step jumps from every state in the direction of z , as illustrated in figure 5.⁴² Any other one-step z -tree violates the requirement of a unique path to z from every other state. It follows that $v(z)$ is given by

$$v(z) = \prod_{0 \leq i < z} p_{i(i+1)} \prod_{z < i \leq N} p_{i(i-1)}$$

³⁹Note that Lemma 2’s condition that P be a *regular* perturbed process is no longer required.

⁴⁰See Young (1998), section 3.4, for illustrative examples.

⁴¹Young’s technique in fact follows from that of Freidlin and Wentzell, taking advantage of uniform mutation rates, and of the zero resistance of paths along the deterministic dynamic, in order to achieve greater analytical simplicity by simply “counting mutations” between recurrent classes.

⁴²For a proof, see Myatt and Wallace’s (1998) Lemma 3.

which implies, in combination with the transition probabilities in Lemma 5, that

$$v(z) = \frac{1}{N^N} \prod_{0 \leq i < z} (N - i) F(\varpi_1(i)) \prod_{z < i \leq N} i F(\varpi_2(i - 1)) \quad (27)$$

The unperturbed process P^0 in this model is that where there is no heterogeneity in switching costs, $\sigma^2 = 0$. In this case, the switching cost pdf (figure 4, for example) collapses to a point mass on the mean switching cost \bar{c} , and the model becomes that of section 3, with $c = \bar{c}$ (see figure 2, page 14). The unperturbed process thus has a set of recurrent classes consisting of the two extreme absorbing states, 0 and N , as well as the mixed absorbing states E_M of Proposition 2 (page 12). As usual, the ergodic distribution will focus all weight on these states as perturbations go to zero, but to select between them it is necessary to consider their *relative weight* in the ergodic distribution. Note that Lemma 6 implies that the relative weight of any two states z and z' in the ergodic distribution μ may be assessed by considering the ratio

$$\frac{\mu(z)}{\mu(z')} = \frac{v(z)}{v(z')} \quad (28)$$

Following Myatt and Wallace, a state z will be said to *dominate* another state z' for vanishing heterogeneity whenever $\lim_{\sigma^2 \rightarrow 0} \frac{v(z)}{v(z')} = \infty$. If a state dominates all others in this sense, it is clearly the unique stochastically stable state of the system.⁴³ Propositions 7 and 8 below establish that this state is the risk-dominant equilibrium for a wide class of switching cost distributions.

Definition 5 *A distribution F with mean \bar{y} has a likelihood ratio unbounded in the tails if*

$$\begin{aligned} \lim_{\sigma^2 \rightarrow 0} \frac{F(y)}{F(y - \epsilon)} &\rightarrow \infty, & \forall y \leq \bar{y}, \forall \epsilon > 0 \\ \text{and} \quad \lim_{\sigma^2 \rightarrow 0} \frac{1 - F(y)}{1 - F(y + \epsilon)} &\rightarrow \infty, & \forall y > \bar{y}, \forall \epsilon > 0 \end{aligned}$$

Proposition 7 *Assume that the switching cost distribution F has a likelihood ratio unbounded in the tails. Then $z \in \{0, N\}$ dominates any mixed state $z \in (0, N)$ for vanishing heterogeneity.*

Proof. For $z > z^*$, compare the limiting weights of states z and N in the ergodic distribution using equations (27) and (28):

$$\begin{aligned} \lim_{\sigma^2 \rightarrow 0} \frac{v(N)}{v(z)} &= \lim_{\sigma^2 \rightarrow 0} \frac{\prod_{z \leq i < N} (N - i) F(\varpi_1(i))}{\prod_{z < i \leq N} i F(\varpi_2(i - 1))} \\ &= \lim_{\sigma^2 \rightarrow 0} \prod_{z \leq i < N} \frac{(N - i) F(\varpi_1(i))}{(i + 1) F(\varpi_2(i))} \end{aligned} \quad (29)$$

⁴³Refer to the definition of stochastic stability on page 12.

where the shared branches of the z - and N -trees for $0 \leq i < z$ lead to the corresponding terms being cancelled, and the second step merely re-indexes the denominator.

Now, as $\sigma^2 \rightarrow 0$ (i.e. the game approaches that of section 3)⁴⁴, $F(\varpi_1(i))$ (the probability of a selected 2-incumbent switching to strategy 1) tends to 1 for $i > z_H$, but to 0 for $i \leq z_H$.⁴⁵ Similarly, $F(\varpi_2(i))$ (the probability of a selected 1-incumbent switching to strategy 2) tends to 1 in the limit for $i < z_L$, but to 0 for $i \geq z_L$. It is thus clear that, for $z \in (z_H, N)$, the numerator of equation (29) tends to $(N - z)!$ as $\sigma^2 \rightarrow 0$, whilst the denominator tends to 0; hence, $\lim_{\sigma^2 \rightarrow 0} (v(N)/v(z)) = \infty$, and state N dominates all states $z \in (z_H, N)$.

Meanwhile, for $z \in (z^*, z_H]$, the additional terms in the range $z \leq i \leq z_H$ must also be considered. All of these terms tend to zero in the limit in both the numerator and the denominator, so that the numerator and the denominator themselves both tend to zero. However, consider the likelihood ratio term $(F(\varpi_1(i))/F(\varpi_2(i)))$ for a given $i \in (z^*, z_H]$. Since $\varpi_1(i) > \varpi_2(i)$ for $i > z^*$, a sufficient condition for this ratio to tend to infinity as $\sigma^2 \rightarrow 0$ is that the switching cost distribution F have a likelihood ratio unbounded in the tails, as defined in Definition 5. For such switching cost distributions then, all of the likelihood ratio terms $(F(\varpi_1(i))/F(\varpi_2(i)))$ for $i \in (z^*, z_H]$ will tend to infinity as heterogeneity vanishes. Meanwhile, the likelihood ratio terms for $i \in (z_H, N)$ clearly tend to infinity in the limit, and thus, so does the whole of equation (29). Hence, $\lim_{\sigma^2 \rightarrow 0} (v(N)/v(z)) = \infty$ again, and state N dominates all states $z \in (z^*, z_H]$.

For $z \leq z^*$,⁴⁶ a similar comparison of the limiting weights of states z and 0 in the ergodic distribution delivers 0's dominance over this range. ■

Proposition 8 *For switching cost distributions with a likelihood ratio unbounded in the tails, the risk-dominant equilibrium remains the unique stochastically stable state.*

Proof. Compare the limiting weights of states 0 and N in the ergodic distribution using equations (27) and (28):

$$\begin{aligned} \lim_{\sigma^2 \rightarrow 0} \frac{v(N)}{v(0)} &= \lim_{\sigma^2 \rightarrow 0} \frac{\prod_{0 \leq i < N} (N - i) F(\varpi_1(i))}{\prod_{0 < i \leq N} i F(\varpi_2(i - 1))} \\ &= \lim_{\sigma^2 \rightarrow 0} \prod_{0 \leq i < N} \frac{(N - i) F(\varpi_1(i))}{(i + 1) F(\varpi_2(i))} \\ &= \lim_{\sigma^2 \rightarrow 0} \prod_{0 \leq i < N} \frac{F(\varpi_1(i))}{F(\varpi_2(i))} \end{aligned} \tag{30}$$

Now, both the numerator and the denominator of equation (30) contain terms tending to zero as $\sigma^2 \rightarrow 0$, so that the numerator and denominator themselves

⁴⁴Refer in particular to figure 2 (page 14).

⁴⁵ z_L and z_H are as defined in equations (14) and (15) (page 13), with $c = \bar{c}$.

⁴⁶If it occurs, the term $i = z^*$ of course tends to zero at the same rate in both numerator and denominator, and thus has no effect on the overall limiting relative weight.

both converge to zero in the limit. Hence, the individual likelihood ratio terms must once again be considered.

Assume without loss of generality that N is the risk-dominant equilibrium, and re-express equation (30) as

$$\lim_{\sigma^2 \rightarrow 0} \frac{v(N)}{v(0)} = \lim_{\sigma^2 \rightarrow 0} \frac{\prod_{0 \leq i \leq z^*} F(\varpi_1(i)) \prod_{z^* < i \leq \hat{z}} F(\varpi_1(i))}{\prod_{z^* < i \leq \hat{z}} F(\varpi_2(i)) \prod_{0 \leq i \leq z^*} F(\varpi_2(i))} \frac{\prod_{\hat{z} < i < N} F(\varpi_1(i))}{\prod_{\hat{z} < i < N} F(\varpi_2(i))} \quad (31)$$

where $\hat{z} = 2z^*$ is defined by $\varpi_1(\hat{z}) = \varpi_2(0)$ (and/or $\varpi_2(\hat{z}) = \varpi_1(0)$), as illustrated in figure 6. Now observe that $\varpi_1(i) = \varpi_2(2z^* - i)$ for all i , so that the terms in the numerator and denominator of the first two ratios on the RHS of equation (31) will be almost identical, and will thus approximately cancel each other out. This can be seen intuitively, in figure 6, to be a consequence of the linearity of the expected payoffs; each $\varpi_1(i)$ term in the range $z^* < i \leq \hat{z}$ has an almost identical $\varpi_2(i)$ term in the range $0 \leq i < z^*$ (reflected in the line $z = z^*$), and vice versa. The only imprecision arises from integer problems, which wash out as N becomes large.

Thus, equation (31) becomes

$$\lim_{\sigma^2 \rightarrow 0} \frac{v(N)}{v(0)} \approx \lim_{\sigma^2 \rightarrow 0} \prod_{\hat{z} < i < N} \frac{F(\varpi_1(i))}{F(\varpi_2(i))}$$

As long as $\beta_L > 0$ (i.e. $\bar{c} < a_{22} - a_{12}$), all the terms in the numerator tend to 1 as $\sigma^2 \rightarrow 0$ and all those in the denominator tend to 0, so that $\lim_{\sigma^2 \rightarrow 0} (v(N)/v(z)) = \infty$, and state N dominates state 0.⁴⁷ By Proposition 7, state N thus dominates all others, and is the unique stochastically stable state.

Long-run selection is thus determined by whether there are more terms in the region $(z^*, N]$ or the region $[0, z^*)$, i.e. which is the greater of $(N - z^*)$ and z^* . Thus, once again, the long-run equilibrium is the *risk-dominant equilibrium*. Intuitively, in figure 6, the risk-dominant equilibrium N gets the terms in the shaded region that are not cancelled out by corresponding terms on the other side of z^* ; all other terms cancel each other out, so that this shaded region is decisive. ■

Thus, for any switching cost distribution satisfying the property of a likelihood ratio unbounded in the tails, the risk-dominant equilibrium remains the unique long-run equilibrium as heterogeneity vanishes. And a wide variety of reasonable distributions *do* satisfy this property - for example, the Normal, Student's t , exponential, logistic and gamma distributions. Moreover, the likelihood

⁴⁷Indeed, even if $\beta_L < 0$, so that $\hat{z} \leq z_H$, it is still the case that $\varpi_1(i) > \varpi_2(i)$ for all $i \in (\hat{z}, N)$, so that $\lim_{\sigma^2 \rightarrow 0} (v(N)/v(z)) = \infty$ still holds by the arguments in the proof of Proposition 7 (provided F has a likelihood ratio unbounded in the tails).




Figure 6: Expected payoff differences

ratio unbounded in the tails property is a sufficient, but not a necessary, condition for the result, so that an even wider class of distributions is admissible.⁴⁸

In fact, the above result is a particular instance of Blume’s (1999) result that, with a random matching birth-death process, the risk-dominant equilibrium is the unique stochastically stable state for 2×2 coordination games under *skew-symmetric* noise processes. Roughly speaking, a skew-symmetric noise process is one under which only payoff differences ($\varpi_1(z)$ and $\varpi_2(z)$ here), and not the “names” of strategies matter to choice. In other words, the probability of a player choosing strategy 2 when strategy 1 has an expected payoff advantage of δ and he is currently playing 1 is the same as that of him choosing 2 when he is currently playing 1 and 2 has an expected payoff advantage of δ .

⁴⁸Note, however, that the uniform mutation rate ε of section 3 does *not* satisfy this property.

4.2.2 Simultaneous revisions

The selection dynamic employed in this subsection is the simultaneous-revisions best-response dynamic $B^N(z)$. Under this dynamic, all players have the opportunity to revise their strategies each period; each player takes an independent and identically distributed draw from the switching cost random variable C , and then best-responds with respect to the current strategy frequency and his c draw in deciding whether or not to switch strategies.

Lemma 7 *Under the simultaneous-revisions dynamic, the new transition probabilities p_{ij} which constitute the perturbed Markov matrix P^{σ^2} are*

$$p_{ij} = \sum_{k=\max\{j-i,0\}}^{\min\{j,N-i\}} \binom{i}{i+k-j} \binom{N-i}{k} \times F(\varpi_2(i-1))^{i+k-j} (1-F(\varpi_2(i-1)))^{j-k} \times F(\varpi_1(i))^k (1-F(\varpi_1(i)))^{N-i-k}$$

Proof. There are $(\min\{j, N-i\} - \max\{j-i, 0\})$ possible combinations of strategy switches in a period that will lead the system from state i to state j , each of which has a number of permutations (essentially relabelling players) given by the product of the two binomial terms. The product of F terms is then the probability of each permutation. ■

Proposition 9 *The adaptive response dynamic defined by the transition probabilities p_{ij} in Lemma 7 is an irreducible, aperiodic Markov process on the finite state space Z . Consequently, it has a unique invariant distribution.*

Proof. Since F has full support, either strategy may be chosen by any reviewing player. The process can thus move in either direction from any state i (except the extreme states). Moreover, given that all players have the opportunity to revise their strategies in each period, the process can move anything from 0 to N states in one period. Thus, every state is accessible from all others within one period (i.e. every entry p_{ij} in the perturbed Markov matrix P^{σ^2} of Lemma 7 is strictly positive), so that the process is irreducible. Moreover, since in every state there is a positive probability of the system remaining in that state in the next period, the process is aperiodic. By Lemma 1, the process thus has a unique invariant distribution. ■

Long-run equilibrium with vanishing heterogeneity As under single revisions, the unperturbed Markov process P^0 under simultaneous revisions is that where there is no heterogeneity in switching costs, $\sigma^2 = 0$. As this limit is approached, the model again becomes that of section 3, with recurrent classes 0, N and E_M . As before, limiting relative weights must be considered in order to select between these candidate equilibria.

However, this task is considerably less straightforward than under the single-revisions dynamic. The transition probabilities of Lemma 7 provide a stark illustration of the potential complexity of the Freidlin-Wentzell approach to long-run equilibrium analysis. Every branch of every possible z -tree is weighted by a transition probability p_{ij} , and under simultaneous revisions the number of possible z -trees per state soon becomes prohibitively large as the state space grows. Fortunately, some simplifications are available.

Lemma 8 *The probability p_{ij} of the transition between states i and j is of the order of*

$$\begin{aligned} \varrho_{ij} = & \max_{k \in [\max\{j-i, 0\}, \min\{j, N-i\}]} \binom{i}{i+k-j} \binom{N-i}{k} \\ & \times F(\varpi_2(i-1))^{i+k-j} (1 - F(\varpi_2(i-1)))^{j-k} \\ & \times F(\varpi_1(i))^k (1 - F(\varpi_1(i)))^{N-i-k} \end{aligned} \quad (32)$$

as $\sigma^2 \rightarrow 0$. Moreover, for switching cost distributions satisfying the condition $F(-x) < (1 - F(x))$, and with a likelihood ratio unbounded in the tails, this maximum is achieved at $k = \max\{j - i, 0\}$.

Proof. The first part of the result is immediate from Lemma 7, given that the order of a summation of terms is determined by the highest order term. To see the second part, consider the effect of increasing k by 1 in equation (32): $(1 - F(\varpi_2(i-1)))(1 - F(\varpi_1(i)))$ is removed from the expression, and replaced by the strictly lower $F(\varpi_2(i-1))F(\varpi_1(i))$ (under the assumption that $F(-x) < (1 - F(x))$). A likelihood ratio unbounded in the tails is then sufficient to guarantee that the maximand in equation (32) is decreasing in k as $\sigma^2 \rightarrow 0$. ■

Intuitively, the minimum value of k is selected in Lemma 8 because this minimizes the number of strategy switches used to effect a given transition. Two strategy switches in opposite directions merely cancel each other out, and two players remaining inert is more probable than two players switching in opposite directions; hence, anything above the minimum number of strategy switches serves to reduce a transition path's probability. This is a natural feature of a model of evolution under inertia, and indeed a wide class of distributions satisfy the conditions of the Lemma given a positive mean. For example, symmetry about \bar{c} is sufficient (but not necessary) to ensure that $F(-x) < (1 - F(x))$ for $\bar{c} > 0$.

Lemma 9 *Under the adaptive response dynamic defined by the transition probabilities p_{ij} in Lemma 7, state z dominates another state z' for vanishing heterogeneity whenever*

$$\lim_{\sigma^2 \rightarrow 0} \frac{\varrho(z)}{\varrho(z')} = \infty,$$

where

$$\varrho(z) = \prod_{(i,j) \in T_z^{\max}} \varrho_{ij}$$

and $T_z^{max} = \arg \max_{T \in \mathcal{T}_z} p(T)$ is the most probable z -tree for a given state z .

Proof. Recall that a state z dominates another state z' for vanishing heterogeneity whenever $\lim_{\sigma^2 \rightarrow 0} (v(z)/v(z')) = \infty$. Recall also from Lemma 6 that $v(z)$ is given by the sum of the likelihoods of all possible z -trees for a given state z . Thus,

$$\lim_{\sigma^2 \rightarrow 0} \frac{v(z)}{v(z')} = \lim_{\sigma^2 \rightarrow 0} \frac{\sum_{T \in \mathcal{T}_z} p(T)}{\sum_{T \in \mathcal{T}_{z'}} p(T)},$$

where $p(T) = \prod_{(i,j) \in T} p_{ij}$ is the likelihood of the tree T , which belongs to a family of z -trees \mathcal{T}_z (for a given state z).

Now, $\lim_{\sigma^2 \rightarrow 0} (v(z)/v(z')) = \infty$ if and only if $v(z)$ is of higher order than $v(z')$ (i.e. $v(z') = o(v(z))$).⁴⁹ Hence

$$\lim_{\sigma^2 \rightarrow 0} \frac{v(z)}{v(z')} = \infty \quad \Leftrightarrow \quad \lim_{\sigma^2 \rightarrow 0} \frac{\max_{T \in \mathcal{T}_z} p(T)}{\max_{T \in \mathcal{T}_{z'}} p(T)} = \infty$$

Since this will in turn be true if and only if

$$\max_{T \in \mathcal{T}_{z'}} p(T) = o(\max_{T \in \mathcal{T}_z} p(T)),$$

dominance is seen to be determined by a comparison of the order of the likelihood of each state's highest-order z -tree.⁵⁰ The likelihoods of the most probable z -trees are themselves products of the transition probabilities p_{ij} of Lemma 7, each of which is of the order of ϱ_{ij} as $\sigma^2 \rightarrow 0$ by Lemma 8. ■

In order to employ Lemma 9 precisely, it is first necessary to identify the most probable z -tree T_z^{max} for each state z . This task is far from straightforward, and is addressed in Norman (2003b). There it emerges that the most probable way of escaping a given basin of attraction depends on the assumed “noise model” - in this case the density of switching costs - but in general is unlikely to involve either one-step transitions or direct jumps. As a result, fully operationalizing Lemma 9 is a complex task, feasible only for particular noise models. However, for the purposes required here, all that matters once again is the skew-symmetry of the noise process.

Proposition 10 *Assume that the switching cost distribution F has a likelihood ratio unbounded in the tails. Then $z \in \{0, N\}$ dominates any mixed state $z \in (0, N)$ for vanishing heterogeneity.*

Proof. For $z \in (z_H, N)$, consider $\varrho(N)/\varrho(z)$, and call this the *order ratio* for convenience. Once shared terms have been cancelled, this ratio is simply

⁴⁹In general, for two functions $f(x)$ and $g(x)$, if $(f(x)/g(x)) \rightarrow 0$ as $x \rightarrow \infty$, then f is of smaller order than g , denoted $f(x) = o(g(x))$. If, on the other hand, $\lim_{x \rightarrow \infty} (f(x)/g(x)) \leq \text{constant}$, then f is of the same order as g , denoted $f(x) = O(g(x))$.

⁵⁰This is a key result in graph-theoretic Markovian theory which, for example, underlies Young's stochastic potential technique.

$\varrho_{zN}/\Upsilon_{Nz}$, where Υ_{Nz} is the (unknown) product of ϱ_{ij} terms leading from state N to state z . The numerator of this order ratio tends to 1 as $\sigma^2 \rightarrow 0$ whilst the denominator tends to 0. Hence $\lim_{\sigma^2 \rightarrow 0} \varrho(N)/\varrho(z) = \infty$, and state N dominates all states $z \in (z_H, N)$ by Lemma 9.

For $z \in (z^*, z_H]$, meanwhile, the order ratio $\varrho(N)/\varrho(z)$ becomes somewhat more complicated:

$$\frac{\varrho(N)}{\varrho(z)} = \frac{\Upsilon_{zN}}{\Upsilon_{Nz}}$$

Since z now lies within the region of mixed absorbing states E_M , both numerator and denominator of the order ratio will now tend to 0 as $\sigma^2 \rightarrow 0$. However, since the transition from z to N is in the more probable direction, if F has a likelihood ratio unbounded in the tails, Υ_{zN} will dominate Υ_{Nz} . Thus, $\lim_{\sigma^2 \rightarrow 0} \varrho(N)/\varrho(z) = \infty$, and by Lemma 9 state N dominates all states $z \in (z^*, z_H]$.

For $z \leq z^*$, similar arguments hold for $\varrho(0)/\varrho(z)$. ■

Definition 6 *A given z -tree T_z and a given z' -tree $T_{z'}$ are said to be k -symmetric if all paths rooted in the region $[z^* - k, z^*)$ of each tree have symmetric counterparts rooted in the region $(z^*, z^* + k]$ of the other tree.*

Absent integer problems, all order ratio terms in a k -symmetric region cancel for skew-symmetric noise processes.

Lemma 10 *The most probable path from 0 to N involves one-step transitions from α_L to α_H . Similarly, the most probable path from N to 0 involves one-step transitions from β_H to β_L .*

Proof. This follows directly from Norman's (2003b) Proposition 3. ■

It follows from Lemma 10 that the most probable N - and 0-trees will be k -symmetric, where $k \in [z^* - z_L, z^*]$. Clearly if $k = z^*$, then all order ratio terms in the region $[0, \hat{z}]$ cancel (assuming N risk-dominant), leaving a situation similar to that of Proposition 8. Under simultaneous revisions, however, it is possible that $k < \min\{z^*, N - z^*\}$, and more care is required.

Proposition 11 *For switching cost distributions with a likelihood ratio unbounded in the tails, the risk-dominant equilibrium remains the unique stochastically stable state.*

Proof. The only candidates for long-run equilibrium are states 0 and N by Proposition 10. Assume without loss of generality that N is the risk-dominant equilibrium, and consider the most probable 0-tree T_0^{\max} . Now construct its k -symmetric N -tree \tilde{T}_N with k as high as possible, a direct jump from 0 to $z^* - k$ (if $z^* - k \neq 0$), and transitions with a constant limiting probability (i.e., branches of the underlying deterministic dynamic) elsewhere. Assuming

the switching cost distribution has a likelihood ratio unbounded in the tails, \tilde{T}_N clearly dominates T_0^{\max} for vanishing heterogeneity, since

$$\lim_{\sigma^2 \rightarrow 0} \frac{\varrho_{0(z^*-k)}}{\Upsilon_{N(z^*+k)}} = \infty$$

And, by definition, the most probable N -tree T_N^{\max} dominates \tilde{T}_N for vanishing heterogeneity. Hence T_N^{\max} dominates T_0^{\max} in the limit. ■

So, the power of Blume’s (1999) skew-symmetry is seen to extend beyond his single-revisions dynamic to the case of simultaneous revisions, confirming the selection of the risk-dominant equilibrium. Hence, when mutations are driven by stochastic switching costs, the anxieties of indeterminacy raised by Bergin and Lipman (1996) turn out to have limited bite.

4.2.3 Transition times

Given the move to a state-dependent mutations setting, Ellison’s (2000) “radius-modified coradius” techniques cannot be directly applied. However, Norman (2003b) extends Ellison’s key theorem to a general state-dependent mutations setting, allowing the verification of the above stochastic stability results, and the calculation of the order of transition times.

Under single revisions, switching-cost-driven inertia straightforwardly increases expected waiting times by reducing each individual transition probability $F(\cdot)$. However, the “step-by-step” evolution of subsection 3.2.2 has no effect here, since the single-revisions dynamic only allows one possible evolutionary path between equilibria - the one-step-at-a-time path (see figure 5, page 27). Only under simultaneous revisions, then, can the question of transition times under inertia be resolved.

In order to employ Norman’s (2003b) theorem, it is first necessary to identify the least probable (over all recurrent classes) of the most probable paths from each recurrent class to the risk-dominant equilibrium N . The most probable path from 0 to N is clearly less probable than that from any of the mixed absorbing states E_M to N , but its *nature* is more difficult to ascertain (Norman 2003b). However, Lemma 10 (derived from Norman’s (2003b) Proposition 3) does show that the path should contain one-step transitions towards N within the region of mixed absorbing states. This effectively says that, if Ellison’s (2000) “intermediate steps” are in place (i.e., given a positive value of \bar{c}), the most probable path from 0 to N exploits each step in the “step-by-step” fashion. However, it is not yet clear whether it is desirable for these steps to be in place (i.e., whether positive \bar{c} speeds evolution).

Lemma 11 *The expected waiting time before N is reached, starting from any $z \neq N$ in the σ^2 -perturbed model with simultaneous revisions, is*

$$W(z, N, \sigma^2) = O\left(\left[\Upsilon_{0\alpha_L} \prod_{i \in [\alpha_L, z^*]} \frac{\varrho_{i(i+1)}}{\varrho_{i(i-1)}}\right]^{-1}\right) \quad (33)$$

as $\sigma^2 \rightarrow 0$.

Proof. The appropriate path to consider involves some (unknown) sequence of jumps leading from 0 to α_L (of the order of $\Upsilon_{0\alpha_L}$ by definition), followed by one-step transitions through the set of mixed absorbing states E_M by Lemma 10. Meanwhile, the most probable path of escape from the basin of attraction of any given mixed absorbing state is a one-step transition in the more probable direction (towards 0 for $z \in [\alpha_L, z^*]$; towards N for $z \in (z^*, \beta_H]$). Part (b) of Norman’s (2003b) Theorem 1 then delivers the desired result, after cancelling shared terms in the range $(z^*, \beta_H]$. ■

Proposition 12 *Under simultaneous revisions, the expected waiting time until the risk-dominant equilibrium is reached from any other state is decreasing in the mean switching cost over some range $\bar{c} \in [0, \hat{c})$, given sufficiently large N (and as $\sigma^2 \rightarrow 0$). Moreover, \hat{c} is increasing in N .*

Proof. When $\bar{c} = 0$, there are no mixed absorbing states, and the expected waiting time in equation (33) is merely $O(\Upsilon_{0\alpha_L})$. Now consider raising \bar{c} just enough such that there is only one mixed absorbing state below the mixed-strategy equilibrium z^* ; in other words, $\alpha_L = \lfloor z^* \rfloor$. This gives one term in the product in equation (33), and this term will tend to 1 as N gets large. At the same time, one improbable transition is removed from the $\Upsilon_{0\alpha_L}$ term by the shortening of the distance from 0 to α_L . Each of the remaining transitions in $\Upsilon_{0\alpha_L}$ also becomes less probable, but vanishingly so as N becomes large (and so, the required \bar{c} becomes small). It follows that the small rise in \bar{c} above zero lowers the order of the expected waiting time in equation (33) for sufficiently large N . If \bar{c} is raised a little more, such that there are now two mixed absorbing states below the mixed-strategy equilibrium - i.e., $\alpha_L = (\lfloor z^* \rfloor - 1)$ - then another improbable transition is removed from $\Upsilon_{0\alpha_L}$, the remaining transitions become a little less probable, and another product term is created in equation (33). Similar arguments then yield a lower expected waiting time again for sufficiently large N . And so on, with each successive new mixed absorbing state as \bar{c} rises. However, holding N constant, increasing \bar{c} yields successively decreasing new product terms, and successively lower transition probabilities, so that there will come a point, \hat{c} , where increasing \bar{c} further may increase the expected waiting time. ■

Thus, in the more realistic state-dependent mutations setting, it is confirmed that transition times fall under switching costs, at least initially. The effects at work are, however, more complicated than in section 3. Under stochastic switching costs, if the mean switching cost \bar{c} rises, the number of mixed absorbing states increases, which has two conflicting effects. First, it shortens the escape from 0’s basin of attraction (i.e., reducing α_L) in favour of the less improbable one-step transitions through the mixed absorbing states. This is the analog of the “step-by-step” effect at work in Section 3, and it serves to reduce the expected waiting time *ceteris paribus*. However, a countervailing

effect is also at work, in that a higher mean switching cost makes every transition less likely - the only effect at work under single revisions. Proposition 12 demonstrates that, given a large enough population, the “step-by-step” effect dominates for “small” \bar{c} and vanishing noise, so that inertia speeds up evolution. However, there exists a threshold level of the mean switching cost \hat{c} , above which the expected waiting time until long-run equilibrium is reached is no longer necessarily decreasing in \bar{c} . How high or low is \hat{c} will vary with the assumed noise model and the population size.

5 Conclusion

This paper is motivated by the belief that player inertia is an important phenomenon in repeated game contexts, and that it is driven in large part by the presence of switching costs to changes in behaviour. Such switching costs are introduced within the natural environment of a stochastically adaptive population repeatedly playing a coordination game. There are three main findings to note from this step. Firstly, the risk-dominant equilibrium remains the unique long-run equilibrium under all but the highest levels of inertia. Secondly, the concept of mixed-strategy equilibrium is rejuvenated by the introduction of inertia, which transforms it from an unstable knife-edge to the focus of a new group of short- to medium-run equilibria. Finally, evolution is, under certain conditions, speeded by the presence of switching costs.

The first finding further strengthens risk-dominance as the solution concept of choice in 2×2 evolutionary coordination games. Indeed, its long-run selection holds under quite general conditions even with state-dependent mutations driven by stochastic switching costs, thanks to Blume’s (1999) skew-symmetry property of the noise process being satisfied. Thus, Bergin and Lipman’s (1996) fears of the indeterminacy of stochastic adjustment models are revealed to have limited bite in this case. In fact, the extension of Blume’s skew-symmetry to the case of the simultaneous-revisions dynamic of subsection 4.2.2 suggests his results to hold very generally. Moreover, it would seem that persistent suboptimal coordination is unlikely to be founded in the existence of individual strategy switching costs, unless they outweigh the maximum possible single-period payoff gain. Given the assumption of myopic players, this would be begging the question somewhat, unless myopia could be shown to be rational in such a context.

The second finding is more novel. The introduction of switching costs means that where the payoff gain at stake from a strategy change is small (i.e., close to the mixed-strategy equilibrium), players can no longer “be bothered” to switch strategies. This effect allows the mixed-strategy equilibrium and nearby states to assume a role as short- to medium-run equilibria. Contrary to most evolutionary models then, in the presence of switching costs the system will rest for significant periods of time in the states around the mixed-strategy equilibrium. However, it is clear from the proof of Proposition 3 that none of the mixed absorbing states, including the mixed-strategy equilibrium, are candidates for

selection in the long run; indeed, more time is spent in the risk-*dominated* equilibrium in the long run than in any given mixed absorbing state. This is due to the advantage accorded the extreme states by the deterministic dynamic in coordination games. The natural question is then raised as to what the effect of switching costs would be on a game in which the deterministic dynamic acted *in favour of* the mixed-strategy equilibrium. This question is addressed in Norman (2003a).

The most interesting consequence of inertia to emerge in this paper, however, is the third finding - the possible increased speed of evolution thanks to the “step-by-step” (Ellison 2000) effect of the new mixed absorbing states. This finding is straightforward under uniform mutations, but the stochastic switching costs model of section 4 reveals the importance of the mean switching cost, the population size and the noise model to the result. In general though, switching costs with a “small” mean will speed evolution. This finding serves to qualify the Ellison (1993) critique of unreasonably long transition times under vanishing noise, offering an alternative escape to Ellison’s own local interaction or the retreat to a model with non-vanishing noise. Given that models with vanishing noise are far easier to work with than those with non-vanishing noise, and that switching costs are likely to be a widespread phenomenon, this is potentially significant.

On a more practical level, the altered speed of evolution has interesting implications for the applications of evolutionary coordination games. For example, consider the evolution of political institutions. There is clearly a great deal of inertia in such institutions, driven by exactly the sorts of costs of institutional change modelled in this paper. Such inertia is generally seen by political scientists as an impediment to institutional change and an obstacle to the sort of efficient selection generally predicted by the game-theoretic literature. The modelling of inertia in this paper, however, suggests a more complex picture. The stickiness of particular institutional arrangements can serve to speed up evolution by providing Ellison’s (2000) facilitating intermediate links in the evolutionary chain. To present the intuition in this political context, the transformation from a monarchy to a democracy may proceed much more quickly if a King can coexist with a legislature for some period of time without the system reverting back to pure monarchy; a strong French Constitutional Council might emerge far sooner if there can be a stable interim where it is powerless; and a PR electoral system could replace SMSP far sooner if a compromise system can exist in the interim. However, inertia will also make the abandonment of a particular set of institutional arrangements less likely, so that the overall effect on the speed of change will depend on the particulars of the institutional environment.

In conclusion, this paper has sought to fill a gap in the rapidly evolving literature on the emergence of conventions and institutions in stochastic adaptive contexts. Player inertia has long been recognised as an important phenomenon, but its intuitive cause in costly individual strategy changes has been left unmodelled. The consequences of introducing such costs have been shown to be nontrivial, but not necessarily in the inefficient manner suggested by popular

opinion; rather, the presence of such inertia can, under certain conditions, speed up evolution to the risk-dominant equilibrium by allowing the system to rest in the stabilised mixed-strategy equilibrium along the way.

References

- ARTHUR, W. B. (1993): "On Designing Economic Agents that Behave Like Human Agents," *Journal of Evolutionary Economics*, 3, 1–22.
- AUMANN, R. J. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1–18.
- BEGGS, A., AND P. KLEMPERER (1992): "Multi-Period Competition with Switching Costs," *Econometrica*, 60, 651–666.
- BERGIN, J., AND B. LIPMAN (1996): "Evolution with State-Dependent Mutations," *Econometrica*, 64, 943–956.
- BINMORE, K., AND L. SAMUELSON (1997): "Muddling Through: Noisy Equilibrium Selection," *Journal of Economic Theory*, 74, 235–265.
- BLUME, L. (1995): "Evolutionary Equilibrium with Forward-Looking Players," Cornell University Working Paper.
- (1999): "How Noise Matters," Cornell University Working Paper.
- BÖRGERS, T., AND R. SARIN (1995): "Naive Reinforcement Learning with Endogenous Aspirations," Discussion Paper, University College, London.
- (1997): "Learning through Reinforcement and Replicator Dynamics," *Journal of Economic Theory*, 77, 1–14.
- BUSH, R., AND F. MOSTELLER (1955): *Stochastic Models for Learning*. Wiley, New York.
- DAVID, P. A. (1985): "Clio and the Economics of QWERTY," *AEA Papers and Proceedings*, 75, 332–337.
- ELLISON, G. (1993): "Learning, Local Interaction and Coordination," *Econometrica*, 61, 1047–1071.
- (1997): "Learning from Personal Experience: One Rational Guy and the Justification of Myopia," *Games and Economic Behavior*, 19, 180–210.
- (2000): "Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution," *Review of Economic Studies*, 67, 17–45.
- FARRELL, J. (1987): "Competition with Lock-In," University of California, Berkeley: Mimeo.

- FARRELL, J., AND P. KLEMPERER (2001): “Coordination and Lock-In: Competition with Switching Costs and Network Effects,” Working Paper.
- FARRELL, J., AND C. SHAPIRO (1988): “Dynamic Competition with Switching Costs,” *RAND Journal of Economics*, 19, 123–137.
- (1989): “Optimal Contracts with Lock-in,” *American Economic Review*, 79, 51–68.
- FOSTER, D. P., AND H. P. YOUNG (1990): “Stochastic Evolutionary Game Dynamics,” *Theoretical Population Biology*, 38, 219–232.
- FREIDLIN, M. I., AND A. D. WENTZELL (1984): *Random Perturbations of Dynamical Systems*. Springer-Verlag, Berlin.
- HARSANYI, J., AND R. SELTEN (1988): *A General Theory of Equilibrium Selection in Games*. The MIT Press, Cambridge, Massachusetts.
- KANDORI, M., G. J. MAILATH, AND R. ROB (1993): “Learning, Mutation and Long-Run Equilibria in Games,” *Econometrica*, 61, 29–56.
- KARLIN, S., AND H. M. TAYLOR (1975): *A First Course in Stochastic Processes*. Academic Press, San Diego, second edn.
- KLEMPERER, P. (1987a): “The Competitiveness of Markets with Switching Costs,” *RAND Journal of Economics*, 18, 138–151.
- (1987b): “Markets with Consumer Switching Costs,” *Quarterly Journal of Economics*, 102, 375–394.
- (1995): “Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade,” *Review of Economic Studies*, 62, 515–539.
- LEE, I. H., A. SZEIDL, AND A. VALENTINYI (2001): “Contagion and State-Dependent Mutations,” Working Paper.
- LIPMAN, B. L., AND R. WANG (2000): “Switching Costs in Frequently Repeated Games,” *Journal of Economic Theory*, 93, 149–190.
- MYATT, D. P., AND C. WALLACE (1998): “Adaptive Play by Idiosyncratic Agents,” Discussion Paper, Nuffield College, Oxford.
- NORMAN, T. (2003a): “The Evolution of Conflict under Inertia,” Working Paper.
- (2003b): “Step-By-Step Evolution with State-Dependent Mutations,” Working Paper.
- PADILLA, A. J. (1995): “Revisiting Dynamic Duopoly with Consumer Switching Costs,” *Journal of Economic Theory*, 67, 520–530.

- RADNER, R. (1980): "Collusive Behavior in Oligopolies with Long but Finite Lives," *Journal of Economic Theory*, 22, 136–156.
- ROBSON, A., AND F. VEGA-REDONDO (1996): "Efficient Equilibrium Selection in Evolutionary Games with Random Matching," *Journal of Economic Theory*, 70, 65–92.
- ROTH, A. E., AND I. EREV (1995): "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, 164–212.
- ROUSSEAU, J.-J. (1913): "The Inequality of Man," in *Rousseau's Social Contract and Discourses*, ed. by G. Cole, pp. 157–246. J. M. Dent, London, first published 1755.
- SETHI, R. (1998): "Strategy-Specific Barriers to Learning and Nonmonotonic Selection Dynamics," *Games and Economic Behavior*, 23, 284–304.
- STOKEY, N. L., R. E. LUCAS, AND E. C. PRESCOTT (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, Massachusetts.
- SUPPES, P., AND R. ATKINSON (1960): *Markov Learning Models for Multiperson Interactions*. Stanford University Press, Stanford.
- VAN DAMME, E., AND J. W. WEIBULL (1998): "Evolution with Mutations Driven by Control Costs," Tilburg University: Mimeo.
- YOUNG, H. P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57–84.
- (1998): *Individual Strategy and Social Structure*. Princeton University Press, Princeton, New Jersey.