



DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

RAPID EVOLUTION UNDER INERTIA

Thomas W.L. Norman

Number 299

December 2006

Manor Road Building, Oxford OX1 3UQ

Rapid Evolution under Inertia*

Thomas W. L. Norman

All Souls College, Oxford OX1 4AL, UK

First version: February 11, 2004

This version: December 21, 2006

Abstract

This paper demonstrates that inertia driven by switching costs leads to more rapid evolution in a class of games that includes $m \times m$ pure coordination games. Under the best-response dynamic and a fixed rate of mutation, the expected waiting time to reach long-run equilibrium is of lower order in the presence of switching costs, due to the creation of new absorbing “inertia” states that allow Ellison’s (*Review of Economic Studies* **67**, 2000, 17–45) “step-by-step” evolution to occur. *Journal of Economic Literature* Classification Number: C72, C73.

Key Words: evolution; mutations; long-run equilibrium; waiting times; inertia; switching costs.

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

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

1 Introduction

Stochastic evolutionary models are appealing on at least two levels: as a sub-rational justification for game-theoretic equilibria; and as a mechanism for selection amongst these equilibria. However, on neither of these levels are such models wholly compelling. The assumed “myopia” of evolutionary players in treating their environment as stationary is an unattractive feature for many, whilst the length of time required to reach the selected “long-run” equilibrium is often argued to be too long. This paper endogenously models a possible justification for myopia—inertia—and finds that the effect can be to reduce the order of the time taken to reach the “long run.”

The literature on stochastic evolution was born in the papers of Foster and Young (1990), Kandori, Mailath, and Rob (1993), and Young (1993). By introducing an explicit noise mechanism into evolutionary models, they provide a path-independent way to measure the likelihood of a game’s equilibria that selects between them as the noise vanishes: the long-run (“stochastically

*I thank David Myatt and Chris Wallace for numerous discussions, as well as Margaret Meyer, Kevin Roberts, Peyton Young, and seminar participants at Oxford University, University College London, the CEPR European Summer Symposium in Economic Theory 2003, Gerzensee, and the Second World Congress of the Game Theory Society 2004, Marseille. I am also grateful for the helpful comments of the associate editor and referee. The usual disclaimer applies. Email thomas.norman@all-souls.oxford.ac.uk.

stable”) equilibrium of an evolutionary model is unique for generic games with strict Nash equilibria. However, Ellison (1993) and others draw attention to the fact that transition times between equilibria become prohibitively large when noise vanishes—a problem that grows with the size of the population. Theorists have attempted to tackle this “very long run” problem in a number of ways: Ellison (1993) and Robson and Vega-Redondo (1996) reduce transition times by altering the matching process; Young (1993) has players sample from a limited memory of past play; and Binmore and Samuelson (1997) study a model of “noisy learning” with *non*-vanishing noise. More recently, Ellison (2000) has drawn attention to another effect which can reduce waiting times: “evolution in a given direction tends to be more rapid (and hence more powerful) when it may proceed via a sequence of smaller steps (rather than requiring sudden large changes).”¹ The effects of such “step-by-step” evolution are captured by Ellison’s “radius–modified coradius” approach to stochastic evolutionary models.

An even more fundamental criticism of such models, however, is that—in order for the dynamics to satisfy the Markov property essential to the techniques employed—they arguably place unreasonable restrictions on individual behavior. In particular, in assuming that beliefs about the next period’s play are based solely on the current state, players are required to treat their environment as stationary, with all other players leaving their strategies unchanged. This means that they must ignore the fact that other players will be conducting a similar reasoning process to themselves—a significant departure from the empathetic reasoning of perfect rationality.² Such “myopia” can, however, be viewed as reasonable on a number of possible grounds, principal among which is the presence of substantial inertia. For Kandori, Mailath, and Rob (1993) (henceforth KMR), for example, under inertia the small fraction of agents who are changing their strategies in any period are justified in acting myopically: “they know that only a small segment of the population changes its behavior at any given point in time and, hence, strategies that proved to be effective today are likely to remain effective for some time in the future.”³

Inertia is a prominent feature of a number of evolutionary models. Reinforcement models (e.g. Börgers and Sarin 1997) incorporate exogenous inertia in the sense that some weight is given to past strategy mixtures when choosing current mixtures. KMR’s (1993) weakly monotonic “Darwinian” selection dynamic, meanwhile, captures the idea that only some (as opposed to all) players need be optimally adjusting their behavior in any given period. Samuelson’s (1994) “best-responses dynamics with inertia” makes a similar requirement, highlighting “situations where calculating optimal actions is costly, difficult, or time consuming, so that agents will usually simply repeat the actions they have taken in the past.”⁴ Indeed, “inertia” has generally been interpreted in the evolutionary literature as a failure to best-respond on the part of some players. However, its causes have hitherto been left exogenous.

Whilst one can imagine a number of possible ways to explain and endogenize this sort of inertia,

¹Ellison (2000), p. 18.

²For an investigation of the justification of this “myopia” assumption, see Ellison (1997).

³KMR (1993), p. 31.

⁴Samuelson (1994), p. 39.

a natural and simple candidate is the presence of “switching costs” to changing strategies from one period to the next. Such costs are intuitively realistic; as Lipman and Wang (2000) point out, changing strategies 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).⁵ 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. The omission of such costs from an evolutionary model with inertia is not innocuous; if players are assumed to be able to change their strategies costlessly, at least some players do so even when there is only the smallest gain to be made. The possibility of system-wide inertia is thus relegated to a knife-edge at best.

Some evolutionary papers have modelled certain forms of costly play: Sethi’s (1998) 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, whilst van Damme and Weibull’s (2002) model of “mutations driven by control costs” has mutation rates determined by individual mistake probabilities that players can control at some cost. 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 implications of such costly strategy change are explored in the present paper in the familiar context of a population of agents repeatedly matched at random to play a symmetric two-player game, with myopic best responses to current population strategy frequencies shaping the evolution of play in discrete time. It is found that, in the presence of switching costs, it becomes possible that no player has an incentive to change strategies in some states, so that there is system-wide inertia and a new set of short-run equilibria. These new absorbing “inertia states” provide Ellison’s (2000) intermediate steady states through which step-by-step evolution can occur. Thus, the surprising conclusion is reached that, in a class of games that includes pure coordination games, the presence of individual-level inertia serves to weakly lower the order of the transition time to long-run equilibrium under the fixed mutation rate of KMR (1993), Young (1993) and others. Moreover, the reduction is strict for switching costs exceeding a certain “small” threshold.

Following the next section’s recap of Ellison (2000) and Section 3’s delineation of the model, Section 4 presents an example of this effect at work in a setting similar to the KMR coordination model. Section 5 then presents a more general class of games, characterizes the new inertia states present under switching costs, and gives the main result on waiting times. Finally, Section 6 discusses the effect of making the mutation probability state-dependent and sensitive to switching costs, and it is shown that faster evolution will still result if the effect of switching costs on the mutation rate is not too sharply negative.

⁵See Klemperer (1995) for a survey of the developed industrial organization literature on switching costs, and Lipman and Wang (2000) for a study of the effects of such costs in general repeated-game contexts.

2 Preliminaries

Ellison (2000) provides a very general framework for stochastic evolutionary models, which we now summarize.

Definition 1 A model of evolution with noise is a triple $(Z, P, P(\varepsilon))$ consisting of:

1. a finite set Z referred to as the state space of the model;
2. a Markov transition matrix $P = (P_{ij})$ on Z , consisting of the $|Z| \times |Z|$ transition probabilities

$$P_{ij} := \text{Prob}\{z_{t+1} = j \mid z_t = i\},$$

where $\{z_t : t \in \mathbb{N}\}$ are the realizations of a discrete-time stochastic process taking values in Z ; and

3. a family of Markov transition matrices $P(\varepsilon)$ on Z indexed by a parameter $\varepsilon \in [0, \bar{\varepsilon})$ such that:
 - (a) $P(\varepsilon)$ is ergodic for each $\varepsilon > 0$;
 - (b) $P(\varepsilon)$ is continuous in ε with $P(0) = P$;
 - (c) there exists a (possibly asymmetric) cost function $c : Z \times Z \rightarrow \mathbb{R}_+ \cup \{\infty\}$ such that for all pairs of states $z, z' \in Z$, $\lim_{\varepsilon \rightarrow 0} P_{zz'}(\varepsilon)/\varepsilon^{c(z, z')}$ exists and is strictly positive if $c(z, z') < \infty$ (with $P_{zz'}(\varepsilon) = 0$ for sufficiently small ε if $c(z, z') = \infty$).

The “cost” function of 3c here measures any given transition’s order of probability as ε vanishes, and should not be confused with the switching costs to be introduced below. In the common population setting—to which we specialize in Section 3—a large finite population of players are randomly matched to play some game; the state space Z then describes their play, which in turn responds to the state according to given behavioral assumptions on the players that determine P , which becomes $P(\varepsilon)$ in the presence of some random noise. The framework of Definition 1, however, captures a more general class of stochastic evolutionary settings (including Young 1993, for instance).

For any fixed $\varepsilon > 0$, the Markov process $P(\varepsilon)$ has an ergodic distribution μ^ε . Moreover, under the conditions of Definition 1, $\lim_{\varepsilon \rightarrow 0} \mu^\varepsilon(z)$ exists and is a stationary distribution of P . The selection criterion of “stochastic stability” identifies which outcome(s) receive positive weight in this limit distribution; a state z is *stochastically stable* 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. These long-run equilibria must belong to the model’s recurrent classes (Young 1993).

Recall that $\Omega \subset Z$ is a *recurrent class* of (Z, P) —or equivalently, a *limit set* L of $(Z, P, P(\varepsilon))$ —iff, $\forall \omega \in \Omega$, $\sum_{x \in \Omega} P_{\omega x} = 1$, and for all $\omega, \omega' \in \Omega$ there exists $k > 0$ and a sequence of states $(z_1, z_2, \dots, z_{k-1})$ such that $P_{\omega z_1} P_{z_1 z_2} \dots P_{z_{k-1} \omega'} > 0$. A singleton limit set is termed an *absorbing state*. The set of all limit sets of $(Z, P, P(\varepsilon))$ is denoted \mathcal{L} . $W(x, \Omega, \varepsilon)$ denotes the *expected wait* until a state belonging to Ω is first reached, given that play in the ε -perturbed model begins in state x .

A cornerstone of Ellison’s (2000) approach is that, if Ω is the long-run equilibrium of a model of evolution with noise, examining $\max_{x \in Z} W(x, \Omega, \varepsilon)$ when ε is small can address the issue of how quickly the system converges to this equilibrium. Of course, $W(\cdot)$ will in general tend to infinity as ε goes to zero, but the speed of convergence can still be judged according to how quickly the waiting times increase as ε vanishes. Ellison’s “Radius–Modified Coradius” Theorem (Lemma 1 below) exploits the connection between stochastic stability and transition times to provide a new technique for finding the long-run equilibrium of an evolutionary model, and also bounds the speed with which evolutionary change occurs. The intuition behind the theorem is that “if a social convention tends to *persist* for a long time after it is established and is sufficiently *attractive* in the sense of being likely to emerge relatively soon after play begins in any other state, then in the long run that convention will prevail most of the time.”⁶ Two new measures are employed to exploit this intuition: the “radius” and the “(modified) coradius.”

Some technical apparatus is required in the construction of these measures. A *path* from a set X to a set Y is defined to be a finite sequence of distinct states (z_1, z_2, \dots, z_T) with $z_1 \in X$, $z_t \notin Y$ for $2 \leq t \leq T-1$, and $z_T \in Y$. $\rho(X, Y)$ is the set of all paths from X to Y , and $c(z_1, z_2, \dots, z_T) = \sum_{t=1}^{T-1} c(z_t, z_{t+1})$ is the cost of a given path. Let Ω be a union of one or more of the limit sets of $(Z, P, P(\varepsilon))$. The *basin of attraction* of Ω is

$$D(\Omega) := \{z \in Z \mid \text{Prob} \{ \exists T \text{ s.t. } z_t \in \Omega, \forall t > T \mid z_0 = z \} = 1 \}.$$

This is the set of states from which the unperturbed process converges to Ω with probability one. The *radius* of the basin of attraction of Ω , $R(\Omega)$, is the minimum cost of any path from Ω out of $D(\Omega)$,

$$R(\Omega) = \min_{(z_1, \dots, z_T) \in \rho(\Omega, Z \setminus D(\Omega))} c(z_1, z_2, \dots, z_T).$$

The *coradius* of the basin of attraction of Ω , $CR(\Omega)$, is the maximum over all other states of the minimum cost of any path to Ω ,

$$CR(\Omega) = \max_{x \notin \Omega} \min_{(z_1, \dots, z_T) \in \rho(x, \Omega)} c(z_1, z_2, \dots, z_T).$$

Given a path (z_1, z_2, \dots, z_T) from x to Ω , let (L_1, L_2, \dots, L_r) be the sequence of limit sets through which the path passes consecutively (with no successive appearances by any particular limit set). Note that x may belong to L_1 , and that $L_i \not\subseteq \Omega$ for $i < r$ but $L_r \subseteq \Omega$. A *modified cost*

⁶Ellison (2000), pp. 17–18, emphasis added.

function, c^* , is obtained by subtracting from the cost of the path the radius of the *intermediate* limit sets L_2, \dots, L_{r-1} (excluding L_1 and L_r) through which the path passes,

$$c^*(z_1, z_2, \dots, z_T) = c(z_1, z_2, \dots, z_T) - \sum_{i=2}^{r-1} R(L_i).$$

The *modified coradius* of the basin of attraction of Ω is then

$$\begin{aligned} CR^*(\Omega) &= \max_{x \notin \Omega} \min_{(z_1, \dots, z_T) \in \rho(x, \Omega)} c^*(z_1, z_2, \dots, z_T) \\ &= \max_{x \notin \Omega} \min_{(z_1, \dots, z_T) \in \rho(x, \Omega)} c(z_1, z_2, \dots, z_T) - \sum_{i=2}^{r-1} R(L_i). \end{aligned} \quad (1)$$

The “coradius” of $D(\Omega)$ is thus shortened by incorporating the effect of “step-by-step” evolution through intermediate steady states, with the resulting “modified coradius” providing a bound on the *attractiveness* of Ω . The radius of $D(\Omega)$, meanwhile, provides a bound on the *persistence* of the set Ω .

These bounds are exploited to give Ellison’s main result:

Lemma 1 (Ellison 2000) *Let $(Z, P, P(\varepsilon))$ be a model of evolution with noise, and suppose that for some set Ω which is a union of limit sets $R(\Omega) > CR^*(\Omega)$. Then:*

1. *the long-run stochastically stable set of the model is contained in Ω ; and*
2. *for any $y \notin \Omega$, $W(y, \Omega, \varepsilon) = O(\varepsilon^{-CR^*(\Omega)})$ as $\varepsilon \rightarrow 0$.*

As in Ellison (2000), the notation “ $f(x) = O(g(x))$ as $x \rightarrow 0$ ”—to be read “ $f(x)$ is of the order of $g(x)$ ”—is shorthand for “there exists $C, \bar{x} > 0$ such that $|f(x)| < C|g(x)|$ for all $x \in (0, \bar{x})$.”

Part 1 of Lemma 1 provides a new way of finding a model’s long-run equilibrium, highlighting in the process the essential, simple requirement that such a state(s) be easy to enter and difficult to escape. Part 2 of the lemma, meanwhile, measures the maximum expected waiting time before the stochastically stable set is reached in the limit. This gives important information on the validity of the long-run equilibrium analysis: “If this maximum wait is small, convergence is fast and Ω can be regarded as a good prediction for what we might expect to see in the medium run. If the maximum wait is large, one should be cautious in drawing conclusions from an analysis of long run stochastic stability.”⁷ By shortening the long-run equilibrium’s coradius, “step-by-step” evolution can serve to reduce this maximum wait, and hence to enhance the appeal of stochastic stability.

That intermediate limit sets can speed evolution is quite intuitive; rather than having to make a sudden, dramatic change, the process can exploit a more gradual sequence of smaller, less improbable steps, resting in the intermediate limit sets along the way. Ellison’s biological analogy is evocative:

⁷Ellison (2000), p. 22.

“Think of how a mouse might evolve into a bat. If the process of growing a wing required ten distinct independent genetic mutations and a creature with anything less than a full wing was not viable, we would have to wait a very, very long time until one mouse happened to have all ten mutations simultaneously. If instead a creature with only one mutation was able to survive equally well (or had an advantage, say, because a flap of skin on its arms helped it keep cool), and a second mutation at any subsequent date produced another viable species, and so on, then evolution might take place in a reasonable period of time.”⁸

3 The Model

Consider now the application of this framework to a single population of N agents, who are repeatedly matched at random in periods $t = 1, 2, \dots$ to play a 2-player symmetric normal-form game $\Gamma \in \mathcal{G}$, with finite strategy set S of cardinality m and payoff function $\pi : S^2 \rightarrow \mathbb{R}$. The agents cannot play mixed strategies. The state space is $Z = \{(z(s_1), \dots, z(s_m)) \in \mathbb{Z}_+^m : z(s_1) + \dots + z(s_m) = N\}$, the state $z_t \in Z$ recording the number of agents $z_t(s)$ playing each pure strategy $s \in S$ in period t .

Assume that the agents are myopic, in the sense that they believe that the current state z will persist next period, and they are not concerned with subsequent periods. Note the implicit assumption that the population is large enough for each agent to ignore his own impact on the state next period.⁹ Then all agents have expected payoffs $\pi(s, \sigma)$, $s \in S$, equal to those of a player facing an opponent with the mixed strategy $\sigma = z/N$, where $\pi(\cdot)$ is extended to mixed strategies in the usual manner. Let $\Sigma(S)$ denote the set of all mixed strategies, $B(\sigma) \subseteq S$ the set of pure-strategy best responses to the mixed strategy $\sigma \in \Sigma(S)$, and $C(\sigma) := \{s \in S \mid \sigma(s) > 0\}$ the support of σ .

Definition 2 *An (unperturbed) model of evolution (Z, P) has the best-response dynamic if, for all $z, z' \in Z$,*

$$[P_{zz'} > 0] \Leftrightarrow [s \notin B(z/N) \Rightarrow z'(s) = 0].$$

In this case, let $BR(z) := \{z' \in Z \mid P_{zz'} > 0\}$ be the possibility set of the dynamic P in state z .

Under the best-response dynamic then, a transition has positive probability if and only if it leads to a state where every agent is playing a best response to the original state.

Into this setting we introduce a switching cost $\delta \in \mathbb{R}_+$ of individual strategy change. This has the effect that, in states where there previously existed profitable strategy switches, *inertia*—retaining one’s current strategy rather than switching to $s \in B(z/N)$ —may now be optimal in the presence of switching costs. To make this precise, we require a concept of an (unperturbed)

⁸*Ibid.*, p. 19.

⁹This is to simplify the exposition; the results to follow carry over to the small-population setting with minimal modification.

model of evolution under inertia (Z, P_δ) that is equivalent to our unperturbed best-response model (Z, P) with each player incurring a switching cost δ if he changes his strategy from one period to the next.

Definition 3 *If (Z, P) has the best-response dynamic, then the corresponding model of evolution under inertia (Z, P_δ) , $\delta \in \mathbb{R}_+$, has:*

$$[P_{\delta zz'} > 0] \Leftrightarrow \left[\forall s \in S : s \notin B(z/N) \Rightarrow \left[\left[\pi(B(z/N), z/N) - \pi(s, z/N) > \delta \text{ and } z'(s) = 0 \right] \right. \right. \\ \left. \left. \text{or } \left[\pi(B(z/N), z/N) - \pi(s, z/N) \leq \delta \text{ and } z'(s) = z(s) \right] \right] \right].$$

In this case, let $BR_\delta(z) := \{z' \in Z \mid P_{\delta zz'} > 0\}$ be the possibility set of the dynamic P_δ in state z .

In such a model of evolution under inertia then, a transition has positive probability if and only if it involves every agent:

- switching to a best response $s \in B(z/N)$ if it yields him a payoff gain exceeding δ ; and
- leaving his strategy unchanged otherwise.

Note that $P_0 = P$.

The most popular way to introduce noise into an evolutionary model is to follow KMR (1993), Young (1993) and others in specifying a fixed rate of “mutation.” Under the best-response dynamic, this operates as follows.

Definition 4 *If (Z, P) has the best-response dynamic, then $(Z, P, P(\varepsilon))$ has KMRY noise if, for all $z \in Z$, $P(\varepsilon)$ has each player independently playing a best response $s \in B(z/N)$ with probability $1 - \varepsilon$, but choosing any strategy $s \in S$ at random with the complementary probability ε .*

Similarly, if (Z, P_δ) is a model of evolution under inertia, then the perturbed model with KMRY noise $(Z, P_\delta, P_\delta(\varepsilon))$ has each player optimally best-responding or remaining inert with probability $1 - \varepsilon$, and choosing at random otherwise. This perturbed model has associated cost function $c_\delta(\cdot)$, expected waiting time $W_\delta(\cdot)$, and so on. Ellison’s radius and modified coradius are particularly straightforward to calculate under KMRY noise, since the probabilistic cost $c(z_1, z_2, \dots, z_T)$ of a path is then given by simply counting the number of “mutations” (ε -probability events) required to effect the constituent transitions.

4 Example: 2×2 Coordination Games

We now present a simple example of switching costs speeding evolution, before proceeding to the general result in the next section. Consider the KMR-like case where Γ is a 2×2 coordination

game with strategy (action) set $S = \{s_1, s_2\}$ and payoff matrix

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix},$$

where $a_{11} > a_{21}$ and $a_{22} > a_{12}$. Such a game of course has multiple equilibria: two symmetric, strict pure-strategy Nash equilibria, (s_1, s_1) and (s_2, s_2) , and a symmetric mixed-strategy equilibrium where each player plays strategy s_1 with probability

$$\sigma^* = \frac{a_{22} - a_{12}}{(a_{22} - a_{12}) + (a_{11} - a_{21})}.$$

The most prominent solution to the resulting equilibrium selection problem is Harsanyi and Selten's (1988) "risk dominance." The risk-dominant equilibrium for a 2×2 game is the one that minimizes the product of the players' losses associated with unilateral deviations. Equivalently, the risk-dominant strategy is the optimal strategy for a player who believes that his opponent will play each strategy with probability $1/2$. In terms of the stage game Γ , equilibrium (s_1, s_1) risk-dominates (s_2, s_2) if and only if $a_{11} - a_{21} > a_{22} - a_{12}$, corresponding exactly to $\sigma^* < \frac{1}{2}$.

Since the system is essentially one-dimensional, let the number $z(s_1)$ of players playing strategy s_1 define the state z of the dynamical system, the simplified state space thus being $Z = \{0, 1, \dots, N\}$. The possibility set of the best-response dynamic is then

$$BR(z) = \begin{cases} 0 & \text{if } \pi(s_1, z/N) < \pi(s_2, z/N) \\ Z & \text{if } \pi(s_1, z/N) = \pi(s_2, z/N) \\ N & \text{if } \pi(s_1, z/N) > \pi(s_2, z/N) \end{cases},$$

giving the unperturbed dynamics depicted in Figure 1.¹⁰ Given that $a_{11} > a_{21}$ and $a_{22} > a_{12}$ for coordination games, the $\pi(s_1, z/N)$ line must be steeper than the $\pi(s_2, z/N)$ line, and there must exist a value $z^* \in (0, N)$ where the two lines cross,

$$z^* = \frac{N(a_{22} - a_{12})}{a_{11} - a_{21} + a_{22} - a_{12}} = \sigma^* N,$$

corresponding to the mixed-strategy equilibrium of the coordination stage game.

There are just two recurrent classes for this model—the absorbing states 0 and N corresponding to the two pure-strategy Nash equilibria of the coordination stage game. These absorbing states have basins of attraction $D(0) = \{z \mid z < z^*\}$ and $D(N) = \{z \mid z > z^*\}$, and the state with the larger basin of attraction (N if $z^* < N/2$; 0 if $z^* > N/2$) is the unique stochastically stable outcome under KMR noise (KMR's Theorem 3).¹¹ This outcome corresponds to the risk-dominant

¹⁰KMR actually consider a more general class of "Darwinian" dynamics, under which a best response in one period must be better represented in the next period. They also have $BR(z) = z$ if $\pi(s_1, z/N) = \pi(s_2, z/N)$, which is inconsequential for our purposes.

¹¹If $z^* = N/2$, then both 0 and N are stochastically stable and receive equal weight in the limit distribution (KMR's Theorem 4).




Figure 1: KMR-style model

equilibrium of the underlying stage game (KMR's Corollary 1), and is henceforth assumed to be N without loss of generality. Selection of the risk-dominant equilibrium can be confirmed using Ellison's (2000) radius-coradius approach. Clearly $R(N) = N - \beta$ and $CR^*(N) = CR(N) = \alpha$, where α and β are the integers above and below z^* :

$$\begin{aligned}\alpha &:= \min \{z \in Z \mid \pi(s_1, z/N) > \pi(s_2, z/N)\} \\ \beta &:= \max \{z \in Z \mid \pi(s_1, z/N) \leq \pi(s_2, z/N)\}.\end{aligned}$$

Since $N - \beta > \alpha$ if N is risk dominant (ignoring integer problems given the assumed large population), part 1 of Lemma 1 implies that N is the unique long-run equilibrium.

Now, in the presence of a fixed switching cost $\delta \in \mathbb{R}_+$, the possibility set of the dynamic becomes

$$BR_\delta(z) = \begin{cases} 0 & \text{if } \pi(s_2, z/N) - \pi(s_1, z/N) > \delta \\ z & \text{if } |\pi(s_1, z/N) - \pi(s_2, z/N)| \leq \delta \\ N & \text{if } \pi(s_1, z/N) - \pi(s_2, z/N) > \delta \end{cases}.$$

Since the expected-payoff lines in Figure 1 certainly cross at z^* , it follows that there is an interval $[z_L, z_H] \subseteq [0, N]$ in a neighborhood of z^* for which $BR_\delta(z) = z$. If δ exceeds a certain threshold




Figure 2: KMR-style model with switching cost δ

δ_I , then $[z_L, z_H] \cap Z \neq \emptyset$ and there exists a nonempty set

$$\mathcal{L}_I = \{z \in Z \mid \pi(s_1, z/N) - \pi(s_2, z/N) \in [-\delta, \delta]\}$$

of new absorbing states of the unperturbed process P_δ , in addition to the set of extreme absorbing states $\{0, N\}$ present when $\delta = 0$.¹² We will refer to these new absorbing states—where no strategy switch generates an expected payoff gain in excess of the switching cost δ —as “inertia states.” The required δ_I is

$$\delta_I = \min \{ \min \{ \delta \in \mathbb{R}_+ \mid z_L \leq \beta \}, \min \{ \delta \in \mathbb{R}_+ \mid z_H \geq \alpha \} \}.$$

Since the integers α and β converge to z^* as $N \rightarrow \infty$, this threshold δ_I tends to 0, and is thus to be considered “small.” The resulting dynamics of the KMR-style model under inertia are illustrated in Figure 2.

The limits of the range of z for which $BR_\delta(z) = z$,

$$z_L = \frac{N(a_{22} - a_{12} - \delta)}{a_{11} - a_{21} + a_{22} - a_{12}} \quad \text{and} \quad z_H = \frac{N(a_{22} - a_{12} + \delta)}{a_{11} - a_{21} + a_{22} - a_{12}},$$

¹²For sufficiently large δ , the extreme absorbing states will also be subsumed within \mathcal{L}_I .

along with their surrounding integers,

$$\begin{aligned}\alpha_L &:= \min \{z \in Z \mid \pi(s_1, z/N) - \pi(s_2, z/N) \geq -\delta\} \\ \beta_L &:= \max \{z \in Z \mid \pi(s_1, z/N) - \pi(s_2, z/N) < -\delta\} \\ \alpha_H &:= \min \{z \in Z \mid \pi(s_1, z/N) - \pi(s_2, z/N) > \delta\} \\ \beta_H &:= \max \{z \in Z \mid \pi(s_1, z/N) - \pi(s_2, z/N) \leq \delta\},\end{aligned}$$

allow the following characterization of the set \mathcal{L}_I of inertia states:

$$\mathcal{L}_I = \{\alpha_L, \alpha_L + 1, \dots, \beta, \alpha, \dots, \beta_H - 1, \beta_H\}.$$

The number l_I of these inertia states is

$$l_I = \beta_H - \alpha_L + 1 = \beta_H - \beta_L = \alpha_H - \alpha_L.$$

And given that

$$z_H - z_L = \frac{2N\delta}{a_{11} - a_{21} + a_{22} - a_{12}},$$

it follows that l_I is nondecreasing in N , and in δ .

In the presence of switching costs, the radius of $D_\delta(N)$ is clearly

$$R_\delta(N) = N - \beta_H,$$

whilst the coradius is

$$CR_\delta(N) = \alpha_H.$$

Hence, $CR_\delta(N)$ is nondecreasing in the size of the switching cost δ . However, the new inertia states \mathcal{L}_I —each of which has a radius of 1—provide Ellison’s (2000) intermediate “steps” between the two extreme states, 0 and N . The *modified* coradius is thus

$$\begin{aligned}CR_\delta^*(N) &= \alpha_H - \sum_{L \in \mathcal{L}_I} R_\delta(L) \\ &= \alpha_H - l_I = \alpha_L.\end{aligned}$$

Due to the symmetry of the expected-payoff functions about the mixed equilibrium z^* then,

$$R(N) - R_\delta(N) = CR(N) - CR_\delta^*(N)$$

(absent integer problems) as long as $\delta < a_{22} - a_{12}$. And since $R(N) > CR(N)$, $R_\delta(N)$ must be greater than $CR_\delta^*(N)$ —so that the risk-dominant equilibrium continues to be selected—in the presence of a switching cost $\delta < a_{22} - a_{12}$. This remains true for $a_{22} - a_{12} \leq \delta$ (so that $\{0\} \in \mathcal{L}_I$ and $CR_\delta^*(N) = 1$) as long as $R_\delta(N) > 1$. Once $R_\delta(N) = 1$ (i.e. once $\{N - 1\} \in \mathcal{L}_I$), however, N

ceases to be the unique long-run equilibrium; all states are stochastically stable.¹³ As $N \rightarrow \infty$, this occurs for $\delta \geq a_{11} - a_{21}$; hence, it is assumed in what follows that $\delta < a_{11} - a_{21}$.

Now, since $R_\delta(N) > CR_\delta^*(N)$, part 2 of Lemma 1 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 $W_\delta(z, N, \varepsilon) = O(\varepsilon^{-\alpha_L})$ as $\varepsilon \rightarrow 0$. The introduction of switching costs thus reduces the expected waiting time from $O(\varepsilon^{-\alpha})$ to $O(\varepsilon^{-\alpha_L})$. Intuitively, the new inertia states \mathcal{L}_I 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. Moreover, since

$$\frac{\partial z_L}{\partial \delta} = \frac{-N}{a_{11} - a_{21} + a_{22} - a_{12}} < 0,$$

it follows that the order of $W_\delta(z, N, \varepsilon)$ is a nonincreasing function of the switching cost δ . This makes sense, given the (weakly) larger number of inertia-state “steps” present under a higher switching cost.

The presence of switching costs also serves to attenuate the (weakly) positive effect of population size on transition times, since

$$\frac{\partial z_L}{\partial N} = \frac{a_{22} - a_{12} - \delta}{a_{11} - a_{21} + a_{22} - a_{12}}.$$

The intuition for this is again based on the accelerated “step-by-step” evolution allowed by the new inertia states. As the size of the population N grows, the number of inertia states 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 inertia states \mathcal{L}_I . Moreover, the more inertia states there are (i.e. the higher the switching cost δ), the stronger this attenuating effect will be.

5 “Step-by-Step” Evolution under Inertia

Turning to the general case, the step-by-step evolution made possible by inertia states in the previous section is obviously not dependent on the 2×2 setting. However, in order to guarantee that this serves to reduce the modified coradius of—and hence the order of the expected wait to reach—the long-run equilibrium in larger games, we do need to place some conditions on the stage game Γ . In this section, we will assume that Γ is a *pure coordination game*—i.e. that $\pi(s_i, s_j) > 0$ if and only if $i = j$.¹⁴ Further, we will call such a game *generic* if, for any $\Omega \subseteq \mathcal{L}$, $R(\Omega) > CR^*(\Omega)$ implies that $\arg \max_{L \in \mathcal{L} \setminus \Omega} \min_{(z_1, \dots, z_T) \in \rho(L, \Omega)} c^*(z_1, z_2, \dots, z_T)$ is a singleton.¹⁵ In fact, the main

¹³Even in this case, of course, transition times between limit sets are much lower than without switching costs, but the long-run equilibrium analysis loses its bite.

¹⁴In terms of payoff matrices, such games are those with zero off-diagonal payoffs.

¹⁵This assumption ensures that the modified coradius is attained only by states within one limit set’s basin of attraction, and it is violated only in knife-edge cases.

result on waiting times holds for a larger class of games, the details of which are provided in the Appendix.

The first step in analyzing the effect of introducing switching costs is to characterize the model's limit sets. Since $[P_{zz} = 1] \Leftrightarrow [BR(z) = z]$ under the best-response dynamic, absorbing states correspond to strict pure-strategy Nash equilibria of Γ , to which the unperturbed system (Z, P) converges for a pure coordination game. Since $BR(z) = z \Rightarrow BR_\delta(z) = z$, these states will remain limit sets in the presence of switching costs; intuitively, if a state cannot be escaped under the best-response dynamic in the absence of switching costs, this should not be altered when strategy change becomes costly.¹⁶

The key effect of introducing switching costs is to create a new set of limit sets born of system-wide inertia.

Definition 5 *A state $z \in Z$ is an inertia state if $\pi(B(z/N), z/N) - \pi(s, z/N) \in (0, \delta]$ for all $s \in C(z/N)$.*

An inertia state is clearly absorbing, with all agents finding inertia (i.e. keeping their current strategy rather than switching to a best response) optimal.

Proposition 1 *Suppose that $(Z, P_\delta, P_\delta(\varepsilon))$ is a model of evolution under inertia with noise, and that Γ is a pure coordination game. Then there is a $\delta_I > 0$ such that, if $\delta \geq \delta_I$, there exists a nonempty set \mathcal{L}_I of inertia states. Furthermore, $\delta_I \rightarrow 0$ as $N \rightarrow \infty$.*

Proof. All pure coordination games have at least one mixed-strategy equilibrium. Under random matching, there is a neighborhood $\nu_\delta(\sigma^*)$ of any mixed-strategy equilibrium σ^* such that $\pi(B(\sigma), \sigma) - \pi(s, \sigma) \leq \delta$ for all $\sigma \in \nu_\delta(\sigma^*)$ and all $s \in C(\sigma)$. Any z such that $\sigma = z/N$ belongs to $\nu_\delta(\sigma^*)$ for some σ^* is thus an absorbing state of (Z, P_δ) . For this set to be nonempty, δ must exceed some threshold δ_I , determined by the population size N (and Γ), such that at least one (integer-valued) state belongs to such a $\nu_\delta(\sigma^*)$. And since any σN approaches an (integer-valued) state as $N \rightarrow \infty$, this δ_I tends to 0 in the limit. ■

To characterize inertia states, the following concepts—essentially those of Samuelson (1994, 1997)—will be useful.

Definition 6 *Two states z and z' are adjacent under KMRY noise if $c(z, z') = 1$.*

In words, it takes a single mutation to move between two adjacent states.

Definition 7 *A set $\mathcal{L}' \subseteq \mathcal{L}$ of limit sets is a mutation-connected component of absorbing states if it is maximal with respect to the properties that L is an absorbing state for all $L \in \mathcal{L}'$ and that, for any two absorbing states L and L' in \mathcal{L}' , there is a path from L to L' that is contained in \mathcal{L}' and consists of single-mutation transitions.*

¹⁶Nonsingleton limit sets of $(Z, P, P(\varepsilon))$, on the other hand, would be vulnerable in the presence of sufficiently high switching costs. Such limit sets do not arise in pure coordination games.

One can thus move from any absorbing state in a mutation-connected component to any other absorbing state in the component, without leaving the component, by moving between adjacent states. Clearly any inertia state must belong to a mutation-connected component of inertia states by definition under random matching, and this will be important in the proof of the main result. Depending on Γ , there may be multiple mutation-connected components of inertia states (if there are multiple components of mixed-strategy equilibria of Γ).

However, inertia states are not the only effect of switching costs in general; outside the 2×2 case, it is possible that a proper subset of the agents find inertia optimal in a given state, whilst the remainder still have sufficient incentive to switch to a best response.

Definition 8 *A state $z \in Z$ is a semi-inertia state if there exist $s, s' \in C(z/N)$ such that $\pi(B(z/N), z/N) - \pi(s, z/N) \leq \delta$ and $\pi(B(z/N), z/N) - \pi(s', z/N) > \delta$.*

These states are obviously not absorbing, but serve to alter the unperturbed dynamic in the sense that $P_{\delta zz'} \neq P_{zz'}$ is now possible for a state z that is not an inertia state. This constitutes one of the main complications of the $m \times m$ case over the 2×2 example of Section 4, where such states cannot arise.

Nonetheless, since the limit sets \mathcal{L} of $(Z, P, P(\varepsilon))$ remain limit sets of $(Z, P_\delta, P_\delta(\varepsilon))$, and the inertia states \mathcal{L}_I are the only new limit sets under switching costs, it follows that $\mathcal{L}_\delta = \mathcal{L} \cup \mathcal{L}_I$. Of course, for sufficiently high δ , $\mathcal{L} \cap \mathcal{L}_I \neq \emptyset$, and ultimately $\mathcal{L} \subset \mathcal{L}_I$; but then $R_\delta(L) = 1$ for all $L \in \mathcal{L}_\delta$, at which point stochastic stability will no longer make a unique equilibrium selection under KMRY noise. Hence, we focus on the $\mathcal{L}_\delta \setminus \mathcal{L}_I \neq \emptyset$ case. For each $L \in \mathcal{L}_I$ in a pure coordination game $R_\delta(L) = 1$, so that inertia states are unlikely candidates for long-run equilibrium, being rather easy to escape.

As we have seen, however, the new inertia states \mathcal{L}_I are not inconsequential; they can provide Ellison's (2000) intermediate "steps" through which evolution occurs more quickly.

Theorem 1 *Suppose that $(Z, P_\delta, P_\delta(\varepsilon))$ is a model of evolution under inertia with KMRY noise and Γ a generic pure coordination game. If there is an $\Omega \subseteq \mathcal{L}$ such that $R(\Omega) > CR^*(\Omega)$ and $R_\delta(\Omega) > CR_\delta^*(\Omega)$, then for any $y \notin \Omega$:*

- $W_\delta(y, \Omega, \varepsilon)$ is at most of the order of $W(y, \Omega, \varepsilon)$;
- $W_\delta(y, \Omega, \varepsilon)$ has strictly lower order than $W(y, \Omega, \varepsilon)$ for δ exceeding a threshold $\underline{\delta}$; and
- $\partial O(W_\delta(y, \Omega, \varepsilon))/\partial \delta \leq 0$;

as $\varepsilon \rightarrow 0$.

The essence of the proof—which is relegated to the appendix—is that, for any given path leading to long-run equilibrium absent switching costs (including that of the modified coradius), we can construct a new such path with (almost) the same modified cost under switching costs by taking

advantage of step-by-step evolution through inertia states. And given sufficiently high switching costs, we can find such a path with a strictly lower modified cost, again through step-by-step evolution. We can thus shorten the long-run equilibrium's modified coradius in the presence of switching costs, reducing the order of the expected waiting time to reach it by part 2 of Lemma 1.

The same logic applies to any increase in switching costs (not just from 0 to δ), subject to continued satisfaction of the theorem's conditions; for such an increase serves to create a weakly larger set of inertia states, which—for the same reasons as before—cannot slow, and may speed, the system's evolution. At some level of switching costs, of course, the condition $R_\delta(\Omega) > CR_\delta^*(\Omega)$ —and thus the application of Lemma 1—will fail, so that Ω need no longer be the unique stochastically stable set. At this point, transition times have fallen further in the sense that short-run equilibrium is disturbed more frequently, but there need no longer be a unique long-run equilibrium to which the system is transiting.

A related point to note is that the effects of step-by-step evolution also work in reverse: the *uninterrupted* time spent in the long-run equilibrium Ω is of (weakly) lower order in expectation when switching costs are increased. This is a consequence of Ellison's result that the expected waiting time for leaving any given limit set L 's basin of attraction is $O(\varepsilon^{-R(L)})$, and the fact that the radius of the basin of attraction of any limit set $L \in \mathcal{L}$ is (weakly) lower in the presence of increased switching costs. Thus, under switching costs, the punctuated equilibrium of the noisy process becomes more frequently punctuated, if you like. Of course, as noise vanishes, the system still spends arbitrarily longer in the long-run equilibrium than in any other state.

Inertia states lie within a neighborhood of mixed-equilibrium states (or of states where all agents are playing dominated strategies), so that the step-by-step evolution they allow will only benefit paths running near such states. But whilst in general the threshold $\underline{\delta}$ required for a strict reduction in the order of transition times may thus exceed the δ_I required for the existence of inertia states, it too should be considered “small” in a similar sense; modified coradii for $\Gamma \in \mathcal{G}_I$ will in general pass through the states closest to mixed-strategy equilibria in order to economize on mutations, and these states require only “small” (*qua* vanishing in the large-population limit) δ to become inertia states.

Meanwhile, the assumption that $R(\Omega) > CR^*(\Omega)$ and $R_\delta(\Omega) > CR_\delta^*(\Omega)$ —meaning that Ω is the long-run equilibrium of both $(Z, P, P(\varepsilon))$ and $(Z, P_\delta, P_\delta(\varepsilon))$ —is also quite reasonable, since any reductions in the radius of $D(\Omega)$ owing to the presence of inertia states will in many cases imply a symmetric (about the appropriate mixed-strategy equilibrium) reduction in its modified coradius, as in the 2×2 example of Section 4. However, in general, although both $R_\delta(\Omega) \leq R(\Omega)$ and $CR_\delta^*(\Omega) \leq CR^*(\Omega)$ under the maintained assumptions, $R(\Omega)$ could fall by more than $CR^*(\Omega)$ in the presence of switching costs, possibly overthrowing Ω as the long-run equilibrium (if $R_\delta(\Omega) \leq CR_\delta^*(\Omega)$). Moreover, even absent this eventuality, once δ is high enough that $\mathcal{L} \subset \mathcal{L}_I$ (or more precisely, once a state within one mutation of Ω is an inertia state), $R_\delta(\Omega) = 1 = CR_\delta^*(\Omega)$ and the application of Lemma 1 fails. Given that δ_I and $\underline{\delta}$ are both “small” though, the range of δ for which Theorem 1 holds will be nonempty given a sufficiently large population.

Finally, the more general class of games for which Theorem 1 holds is somewhat intricate in the $m \times m$ case, and is thus discussed in the appendix. The 2×2 example of Section 4 is straightforward, thanks to the impossibility of semi-inertia states in that setting; waiting times fall for *any* 2×2 coordination game, as a consequence of the strategic equivalence of any such game to a generic pure coordination game. In the $m \times m$ case, the changes to the dynamic caused by semi-inertia states mean that there exist examples of (non-pure) coordination games for which Theorem 1 does not hold. Note, however, that we have been exploring sufficient, not necessary, conditions; even in cases where semi-inertia states serve to increase the cost of a path, this may be outweighed by the reductions possible from step-by-step evolution through inertia states.

6 How Noise Matters

That inertia should serve to speed evolution is a surprising result: the intuitive effect of inertia is to make change less likely, and hence less rapid. The logic of “step-by-step” evolution provides the explanation of course, but there is still something missing from the story. For whilst the unperturbed model of Section 5 incorporates the effect of switching costs through optimal behavior under P_δ , the mutation mechanism by contrast is unchanged under inertia. Hence, whilst the “probable switches” of the unperturbed model (Z, P_δ) are altered by changes in the switching cost δ , the mutations—or “improbable switches”—of the perturbed model $(Z, P_\delta, P_\delta(\varepsilon))$ are independent of δ . This is a consequence of the assumed fixed mutation rate ε of KMRY noise, which has been criticized by Bergin and Lipman (1996) given its importance for long-run selection results, but defended by Blume (2003) due to its representativeness of the large class of “skew-symmetric” noise models. Ellison’s model is framed using the “ ε -cost” language of KMRY noise, but—as he notes—can easily accommodate Bergin and Lipman-style state-dependent mutation rates with unbounded likelihood ratios, along with many other non-KMRY noise specifications. Some simplicity is sacrificed in the process, however, since the cost function $c(\cdot)$ loses its intuitive “mutation-counting” role in order to accommodate the variable mutation rates.

Suppose, for instance, that instead of a single switching cost for all agents and a fixed mutation rate, the model had heterogeneous switching costs drawn independently each period from a Normal distribution. This would give an ergodic process with a stochastic state-dependent rate of strategy-switching each period, and we could analyze the limit where the switching-cost distribution collapses to a point mass on some δ . This limit is analogous to the $\varepsilon \rightarrow 0$ case with a fixed mutation rate, and indeed approaches the same unperturbed model (Z, P_δ) . Moreover, this model of state-dependent mutations driven by stochastic switching costs can be accommodated within Ellison’s setup, although $c(\cdot)$ then ceases to have its “mutation-counting” interpretation and simply measures the probabilistic cost of any given transition. This probabilistic cost will be higher (i.e. a transition will have a lower order of probability) when the average switching cost δ is higher, due to the tail properties of the Normal distribution.

The lower-order waiting times of Theorem 1 become somewhat vulnerable for such non-KMRY

noise, but nonetheless survive under a reasonable condition. To see this, let $\overline{CR}_\delta^*(\Omega)$ be the fixed-mutation-rate modified coradius of Section 5 and rewrite the waiting time for the state-dependent case thus:

$$\begin{aligned} W_\delta(y, \Omega, \varepsilon(\delta)) &= O(\varepsilon^{-CR_\delta^*(\Omega)}), \\ &= O\left(\left(\varepsilon^{CR_\delta^*(\Omega)/\overline{CR}_\delta^*(\Omega)}\right)^{-\overline{CR}_\delta^*(\Omega)}\right) \\ &= O\left(\varepsilon(\delta)^{-\overline{CR}_\delta^*(\Omega)}\right), \quad y \notin \Omega, \quad \text{as } \varepsilon \rightarrow 0, \end{aligned}$$

where $\varepsilon(\delta)$ is a function of, *inter alia*, the switching cost. The result of Theorem 1 is now qualified. The “step-by-step” effect is still present; a rise in the switching cost δ still (weakly) reduces the fixed-mutation-rate modified coradius $\overline{CR}_\delta^*(\Omega)$ in the manner described in Section 5. However, if $\varepsilon(\delta)$ falls, the waiting time $W_\delta(y, \Omega, \varepsilon(\delta))$ may increase or decrease in order, according to the relative size of the changes in $\varepsilon(\delta)$ and $\overline{CR}_\delta^*(\Omega)$. Nonetheless, faster evolution will still result if $\varepsilon(\delta)$ does not decrease too sharply with δ . For

$$\begin{aligned} \frac{\partial O(W_\delta(y, \Omega, \varepsilon(\delta)))}{\partial \delta} &= O(W_\delta(y, \Omega, \varepsilon(\delta))) \frac{\partial \log O(W_\delta(y, \Omega, \varepsilon(\delta)))}{\partial \delta} \\ &= \varepsilon(\delta)^{-\overline{CR}_\delta^*(\Omega)} \left(-\frac{\partial \overline{CR}_\delta^*(\Omega)}{\partial \delta} \log \varepsilon(\delta) - \frac{\varepsilon'(\delta)}{\varepsilon(\delta)} \overline{CR}_\delta^*(\Omega) \right) \end{aligned}$$

is nonpositive if and only if

$$-\frac{\partial \overline{CR}_\delta^*(\Omega)/\partial \delta}{\overline{CR}_\delta^*(\Omega)} \geq \frac{\varepsilon'(\delta)}{\varepsilon(\delta) \log \varepsilon(\delta)}$$

(and is strictly negative if and only if the condition holds with strict inequality). The left-hand side of this condition is nonnegative by the arguments in the proof of Theorem 1, and is strictly positive if δ exceeds $\underline{\delta}$. The right-hand side, meanwhile, may be positive (if $\varepsilon'(\delta) < 0$), but vanishes to 0 as $\varepsilon \rightarrow 0$ if and only if $\varepsilon'(\delta)$ is of lower order than $\varepsilon(\delta) \log \varepsilon(\delta)$.

Hence, lower-order waiting times under switching costs are preserved with state-dependent mutations if the effect of switching costs on the probability of mutation is of lower order than $\varepsilon(\delta) \log \varepsilon(\delta)$, which is the order of the “entropy” (Shannon 1948) of the mutation mechanism—a measure of the degree of randomness in the process. Nor is this upper bound on $\varepsilon'(\delta)$ particularly low; after all, $\varepsilon^m = o(\varepsilon^n) = o(\varepsilon \log \varepsilon)$, for all $m > n \geq 1$.

Appendix

Whilst Theorem 1 is stated for pure coordination games, it in fact holds for a more general class of stage games. Specifically, it holds for the class \mathcal{G}_I of generic stage games satisfying the following conditions:

Strong acyclicity For all $z \in Z$, there exists an absorbing state $z' \in \mathcal{L}$ such that $P_{zz'} > 0$.

Increasing best-response differences $\pi(B(z/N), z/N) - \pi(s, z/N) \leq \pi(B(z/N), z'/N) - \pi(s, z'/N)$ for all s, z, z' such that $z'(B(z/N)) \geq z(B(z/N))$ and $[z'(s') = z(s'), \forall s' \notin \{s, B(z/N)\}]$.

Recall that the stage game Γ is “weakly acyclic” (Young 1993) if, from any state z , there exists a positive-probability path to some absorbing state in the unperturbed model (Z, P) . Strong acyclicity is a mild strengthening of weak acyclicity under the best-response dynamic, essentially saying that no state requires more than one round of strategy revisions for every agent to be playing a best response; all coordination games satisfy this requirement. Increasing best-response differences is somewhat more restrictive; it requires that other agents switching from any given strategy to the current best response should not diminish the incentive to do likewise. It thus gives something akin to strategic complementarity at the level of the population.¹⁷ $m \times m$ pure coordination games satisfy this restriction, as does any 2×2 coordination game, but some larger coordination games are excluded (as we discuss following the proof). The role of this property in the proof comes from its implication that, for all $L \in \mathcal{L}$ and all $z \in D(L)$, $BR_\delta(z) \subseteq D(L)$; no positive-probability transition under P_δ leads to a different zero-switching-cost basin of attraction $D(\cdot)$. Roughly speaking, any sequence of individual best responses starting in a given state must ultimately lead to the same limit set.

Proof of Theorem. Consider the modified coradius of $D(\Omega)$, defined in equation (1) on page 6. By Proposition 1, the introduction of a switching cost $\delta \geq \delta_I$ into $(Z, P, P(\varepsilon))$ creates a new set of “inertia states” \mathcal{L}_I , through which any given path may or may not pass. For any given $x \notin \Omega$, suppose that the original (zero- δ) path of minimum modified cost to Ω , $\rho^{\min}(x, \Omega) := \arg \min_{(z_1, \dots, z_T) \in \rho(x, \Omega)} c^*(z_1, z_2, \dots, z_T)$, is such that

$$c_\delta^*(\rho^{\min}(x, \Omega)) > c^*(\rho^{\min}(x, \Omega)).$$

Then there must exist a transition $(z_\tau, z_{\tau+1}) \subseteq \rho^{\min}(x, \Omega)$ with $c_\delta^*(z_\tau, z_{\tau+1}) > c^*(z_\tau, z_{\tau+1})$, which can occur only if z_τ is either a semi-inertia state or an inertia state in the presence of δ .

Case 1: z_τ is a semi-inertia state There must be a limit set L of $(Z, P, P(\varepsilon))$ such that $z_\tau \in D(L)$, so that $c^*(z_\tau, L) = 0 < c_\delta^*(z_\tau, L)$.¹⁸ But then $c^*(L, z_{\tau+1}) = c_\delta^*(L, z_{\tau+1}) = c^*(z_\tau, z_{\tau+1})$

¹⁷In terms of the stage game, this means that the increasing differences occur at the level of mixed strategies rather than the more usual pure strategies.

¹⁸This is where strong acyclicity of the stage game is used, giving us $c(z_\tau, L) = 0$ —as opposed to $c(z_\tau, \dots, L) = 0$ under weak acyclicity.

under the best-response dynamic, and we claim that a path of zero modified cost can be constructed between z_τ and L . To see this, note that the transition from z_τ effected by letting all agents behave optimally (either inertia or best response) has zero modified cost. For all $\Gamma \in \mathcal{G}_I$, we know that this transition leads to a state $z' \in D(L)$ by increasing best-response differences. We then allow optimal switches from z' to create another transition of zero modified cost that leads to a state in $D(L)$. We continue such transitions until the process arrives either in L or in an inertia state. If it is the former, we are done; if the latter, then we move on to Case 2.

Case 2: z_τ is an inertia state z_τ must again belong to the basin of attraction $D(L)$ of some limit set L of $(Z, P, P(\varepsilon))$, so that $c^*(z_\tau, L) = 0 < c_\delta^*(z_\tau, L)$. But then $c^*(L, z_{\tau+1}) = c_\delta^*(L, z_{\tau+1}) = c^*(z_\tau, z_{\tau+1})$ under the best-response dynamic, and we claim that a path of modified cost 1 can be constructed between $z_\tau = L_I$ and L in the presence of the switching cost δ , by taking advantage of L_I 's mutation-connected component \mathcal{L}'_I . To see this, note that $D(L)$ is mutation-connected under random matching; hence, there must exist a path of single-mutation best-response transitions (i.e. transitions where exactly one agent switches to $s \in B(z/N)$) from $z_\tau = L_I \in D(L)$ to some $z'' \in D(L) \setminus \mathcal{L}'_I$ that does not leave \mathcal{L}'_I (until the final transition). Such a path has modified cost 1, since $R(L_i) = 1$ for each intermediate inertia state $L_i \in \mathcal{L}'_I$ between z_τ and $D(L) \setminus \mathcal{L}'_I$. There is then a path of zero modified cost from z'' to L , if necessary through repeated application of Cases 1 and 2 (which must eventually terminate in L given the increasing number of agents best-responding).

If $\tau \neq 1$, then performing these operations gives a new path from x to Ω with modified cost equal to the original (zero-switching-cost) $c^*(\rho^{\min}(x, \Omega))$ —since if z_τ is an inertia state, then it is also an intermediate limit set in the new path and hence $R(z_\tau) = 1$ can be subtracted from its cost. If $\tau = 1$, meanwhile, so that $z_\tau = x$, then the new path has modified cost less than or equal to $c^*(z_1, z_2, \dots, z_T) + 1$. For this to increase the modified coradius of $D(\Omega)$, x would have to both belong to $D(L)$ for some $L \in \arg \max_{x \notin \Omega} \min_{(z_1, \dots, z_T) \in \rho(x, \Omega)} c^*(z_1, z_2, \dots, z_T)$, and become an inertia state $L'_I \in \mathcal{L}'_I$ in the presence of switching costs. But in this case, there must exist another limit set $L' \in \mathcal{L}$ such that $\mathcal{L}'_I \cup D(L')$ is mutation-connected;¹⁹ hence, there is a path of modified cost 1 from x to L' under switching costs, as above. It follows that $\min_{(z_1, \dots, z_T) \in \rho(x, \Omega)} c_\delta^*(z_1, z_2, \dots, z_T)$ for such an x is less than or equal to $\min_{(z_1, \dots, z_T) \in \rho(L', \Omega)} c^*(z_1, z_2, \dots, z_T) + 1$, which cannot increase the modified coradius of $D(\Omega)$ under the genericity assumption that $\arg \max_{L \in \mathcal{L} \setminus \Omega} \min_{(z_1, \dots, z_T) \in \rho(L, \Omega)} c^*(z_1, z_2, \dots, z_T)$ is a singleton.²⁰

¹⁹To see this, note that \mathcal{L}'_I is in a neighborhood of some mixed-strategy equilibrium, and mixed-strategy equilibria lie at the boundary of basins of attraction for all $\Gamma \in \mathcal{G}_I$ under random matching.

²⁰Note that genericity also ensures (through nonemptiness of $\mathcal{L} \setminus \Omega$): (a) that Γ has multiple strict pure-strategy Nash equilibria, so that the stochastic-stability analysis is nontrivial; and (b) that Γ has at least one mixed-strategy equilibrium (in the sense of an equilibrium in non-pure strategies).

Therefore,

$$\begin{aligned} CR_\delta^*(\Omega) &= \max_{x \notin \Omega} \min_{(z_1, \dots, z_T) \in \rho(x, \Omega)} c_\delta^*(z_1, z_2, \dots, z_T) \\ &\leq \max_{x \notin \Omega} \min_{(z_1, \dots, z_T) \in \rho(x, \Omega)} c^*(z_1, z_2, \dots, z_T) = CR^*(\Omega). \end{aligned}$$

And in fact, this inequality will be strict if an intermediate $CR^*(\Omega)$ transition $(z_\tau, z_{\tau+1})$, $\tau \neq 1$, that is costly absent switching costs is rooted in an inertia state when $\delta \geq \delta_I$; in this case, at least $R_\delta(z_\tau) = 1$ can be subtracted from $D(\Omega)$'s modified coradius under switching costs.²¹ Moreover, further reductions in the modified coradius arise if any $CR^*(\Omega)$ transition $(z_\tau, z_{\tau+1})$ can be replaced by a path through an intermediate inertia state(s) L_I'' such that $c_\delta^*(z_\tau, \dots, L_I'', \dots, z_{\tau+1}) < c^*(z_\tau, z_{\tau+1})$. Indeed, both of these eventualities will occur for sufficiently high δ , since they merely require the existence of appropriately located inertia states.

Since $CR_\delta^*(\Omega) \leq CR^*(\Omega)$, the order of $W_\delta(y, \Omega, \varepsilon)$ as $\varepsilon \rightarrow 0$ is less than or equal to that of $W(y, \Omega, \varepsilon)$ by part 2 of Lemma 1. And of course, it will be strictly lower if δ is high enough—exceeding $\underline{\delta}$ say—to create inertia states that can reduce the maximin modified cost in the step-by-step manner.

Finally, the same arguments apply when comparing waiting times under a switching cost δ with those under a higher switching cost δ' —provided the conditions of the theorem remain satisfied—since a higher switching cost simply creates a (weakly) larger number of inertia states. Hence, $O(W_\delta(y, \Omega, \varepsilon))$ is locally nonincreasing in δ , in the sense that $\partial O(W_\delta(y, \Omega, \varepsilon))/\partial \delta \leq 0$. ■

The role of strong acyclicity of the stage game is in part to aid expositional clarity, easing the construction of step-by-step paths to limit sets via inertia states that is required in the proof. In particular, for a given positive-probability path $(z_1, \dots, z_T) \in \rho(z, L)$, $z \in Z$, $L \in \mathcal{L}$, that becomes costly (i.e. vanishingly probable) given $\delta \geq \delta_I$, it is much easier to construct an alternative step-by-step path of (almost) zero modified cost if we only have to worry about reaching L , rather than reaching L via z_2, \dots, z_{T-1} —i.e. if $T = 2$. However, whilst we could still construct such paths in the absence of strong acyclicity, we would need to assume instead a sufficiently *low* δ that z_2, \dots, z_{T-1} *not* become (semi-)inertia states in its presence, so that each step of this sequence would remain zero-cost.

Increasing best-response differences, meanwhile, safeguards the construction of the same paths from the possible changes to the process caused by semi-inertia states. In particular, if $(z_1, \dots, z_T) \in \rho(z, L)$, $z \in Z$, $L \in \mathcal{L}$, passes through a semi-inertia state z' given $\delta \geq \delta_I$ and thus becomes costly, increasing best-response differences ensure that the new zero-cost transition away from z' (where all agents choose optimally between inertia and $s \in B(z'/N)$) does not lead the process out of $D(L)$. Intuitively, absent increasing best-response differences, the possibility of some agents remaining inert whilst others best-respond might lead the system far enough from L

²¹The same applies if an intermediate $CR^*(\Omega)$ transition $(z_\tau, z_{\tau+1})$, $\tau \neq 1$, is rooted in an intermediate limit set $L'' \in \mathcal{L}$ whose radius is lower in the presence of switching costs.




Figure 3: 3×3 game violating condition 2 of the class \mathcal{G}_I

that it is costly (probabilistically) to return, as illustrated in the following example.

Consider the 3×3 coordination game with payoff matrix

$$\begin{bmatrix} 1 & a & 0 \\ 0 & 1 & a \\ a & 0 & 1 \end{bmatrix},$$

where $a < 1$. The best-response structure for the case $a = 0.9$ is illustrated in the Z -space simplex of Figure 3, with the equilibria (three pure; four mixed) marked by heavy dots and $N_1 := (N, 0, 0)$, $N_2 := (0, N, 0)$, $N_3 := (0, 0, N)$. The hollow dot illustrates a state $z' \in D(N_3)$ with a transition leading to $D(N_2)$ that has positive probability under P_δ for a range of values of δ —in particular, for δ high enough that agents playing s_2 at z' find it optimal to remain inert, but small enough that those playing s_1 switch to $B(z'/N) = s_3$.²² Hence, if δ falls in this range, the assumption of increasing best-response differences rules out $a = 0.9$. Moreover, there will be similar violations of the condition for all $a \in (0.5, 1)$ (and for $a < -1$). However, all $m \times m$ pure coordination games satisfy increasing best-response differences, along with “nearby” coordination games. This is illustrated for the 3×3 case in Figure 4, where any game with basins of attraction whose

²²The only positive-probability transition away from z' under P , meanwhile, is of course that leading to N_3 .




Figure 4: 3×3 games belonging to the class \mathcal{G}_I

boundaries lie in the appropriate shaded regions belongs to \mathcal{G}_I ; the dashed lines depict the case of pure coordination.

Without assuming increasing best-response differences, we cannot be sure that switching costs will reduce the minimum modified cost of *arbitrary* paths to long-run equilibrium, as Figure 3 illustrates. Of course, waiting times may still be lower without the assumption, since the only path that really matters is the modified-coradius path, but without knowing where this path lies in general we cannot exploit this fact. There is, in any case, another good reason to focus on games with increasing best-response differences; namely, that the changes to the unperturbed dynamic in semi-inertia states can alter the long-run equilibrium without the assumption. Effectively, increasing best-response differences ensure that, under switching costs, all states are either inertia states or they belong to the same basin of attraction as they did absent switching costs. This is the appropriate setting to isolate the effects of switching costs on the speed of evolution.

References

- 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. E. (2003): “How Noise Matters,” *Games and Economic Behavior*, 44, 251–271.
- BÖRGERS, T., AND R. SARIN (1997): “Learning through Reinforcement and Replicator Dynamics,” *Journal of Economic Theory*, 77, 1–14.
- 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.
- FOSTER, D. P., AND H. P. YOUNG (1990): “Stochastic Evolutionary Game Dynamics,” *Theoretical Population Biology*, 38, 219–232.
- 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.
- KLEMPERER, P. (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.
- LIPMAN, B. L., AND R. WANG (2000): “Switching Costs in Frequently Repeated Games,” *Journal of Economic Theory*, 93, 149–190.
- ROBSON, A., AND F. VEGA-REDONDO (1996): “Efficient Equilibrium Selection in Evolutionary Games with Random Matching,” *Journal of Economic Theory*, 70, 65–92.
- SAMUELSON, L. (1994): “Stochastic Stability in Games with Alternative Best Replies,” *Journal of Economic Theory*, 64, 35–65.
- (1997): *Evolutionary Games and Equilibrium Selection*. The MIT Press, Cambridge, Massachusetts.

- SETHI, R. (1998): “Strategy-Specific Barriers to Learning and Nonmonotonic Selection Dynamics,” *Games and Economic Behavior*, 23, 284–304.
- SHANNON, C. E. (1948): “A Mathematical Theory of Communication,” *Bell System Technical Journal*, 27, 379–423.
- VAN DAMME, E., AND J. W. WEIBULL (2002): “Evolution in Games with Endogenous Mistake Probabilities,” *Journal of Economic Theory*, 106, 296–315.
- YOUNG, H. P. (1993): “The Evolution of Conventions,” *Econometrica*, 61, 57–84.