

Evolutionary Drift and Equilibrium Selection

KEN BINMORE

University College London

and

LARRY SAMUELSON

University of Wisconsin

This paper develops an approach to equilibrium selection in game theory based on studying the learning process through which equilibrium is achieved. The differential equations derived from models of interactive learning typically have stationary states that are not isolated. Instead, Nash equilibria that specify different out-of-equilibrium behaviour appear in connected components of stationary states. The stability properties of these components can depend critically on the perturbations to which the system is subjected. We argue that it is then important to incorporate such *drift* into the model. A sufficient condition is provided for drift to create stationary states, with strong stability properties, near a component of equilibria. Applications to questions of forward and backward induction are developed.

1. INTRODUCTION

Backward induction performs badly in predicting the outcome of laboratory experiments in some games, but well in others. Forward induction is similarly erratic.¹ In the face of such mounting evidence from carefully run experiments, game theorists are reconsidering their theories of equilibrium selection. Why do the principles of forward and backward induction sometimes work and sometimes fail? We need a theory that explains when and why experienced and well-motivated players honour or ignore such principles. In this paper, we argue that players' behaviour is often consistent with game theory, even when backward and forward induction fail—provided we recognize that people must *learn* to play games *in an imperfect world*.

In Binmore and Samuelson (1994) and Binmore, Gale and Samuelson (1995), we showed that simple models of learning can direct players to Nash equilibria in the Ultimatum Game in which player 2 receives a significant share of the surplus—a result closer to experimental observations than to the subgame-perfect prediction. This paper broadens the scope of our study of perturbed learning processes. We show that our results for the Ultimatum Game are not pathological by analysing several equally well-known games.

Our methodology requires examining the stability properties of the Nash equilibria of games, paying close attention to the perturbations of the learning process that are intended to stand in for the many imperfections of the world that are idealized away when formulating an abstract model. Until recently it was taken for granted that such perturbations could be relied upon to eliminate weakly dominated strategies. But Binmore *et al.* (1994, 1995) show to the contrary that vanishingly small perturbations can actually

1. Davis and Holt (1993) and Kagel and Roth (1995) survey the experimental literature.

play a key role in *stabilizing* Nash equilibria in weakly dominated strategies that would be refined away if forward or backward induction criteria were applied.

We use the term *drift* to summarize the perturbations to be studied and address the following questions. When does such drift matter? What can be said about equilibrium selection when it does?

The landscape metaphor. Our answers to these questions are easiest to understand when expressed in terms of the landscape metaphor used by biologists in discussing evolutionary dynamics. It is commonplace to think of a stable equilibrium as lying at the bottom of a pit that represents its basin of attraction. The dynamic system under study is then envisaged as a ball that rolls down the sides of the pit, eventually coming to rest at the equilibrium at the bottom. Evolutionary drift becomes important when the pit is replaced by a valley with a flat floor of equilibria. Once the ball reaches the floor of such a valley, its movement is determined by tiny shocks which can be neglected when the system is not close to equilibrium (because they are then vanishingly small compared with the selection pressures that power the evolutionary dynamics). These tiny shocks represent the perturbations we call drift.

The importance of genetic drift in biology has been emphasized by many authors. Kimura (1983) argues that drift may account for the bulk of genetic variation. Our concern is with drift through landscapes that contain *hanging* valleys, with one end of a hanging valley cutting into the wall of an adjoining, deeper valley. Hanging valleys appear to have received little attention in biology or elsewhere. In single-person decision problems, even a flat-bottomed valley is a sufficiently exceptional occurrence as to be safely ignored. But flat-bottomed, hanging valleys are common in landscapes corresponding to the strategic situations modeled by games, arising out of the same freedom to choose actions at unreached information sets that gives rise to the equilibrium refinements literature.

When a ball reaches the flat floor of a hanging valley, even very small amounts of drift can be large compared to the weak selection pressures that surround the equilibria lying on the floor. Small amounts of drift can then have a large effect. The outcome depends on the overall direction of the drift that operates in the valley. If the drift propels the ball away from unstable equilibria that can sit on the edge of a precipice where the hanging valley cuts into the wall of an adjoining, deeper valley, then the system will be stabilized in the hanging valley. Alternatively, the drift may propel the ball over the cliff-edge, whereupon strong selection pressures reappear and the ball plunges to the floor of the deeper valley. But the story may not end here, for the deeper valley may itself hang above a further valley. It may even hang above the original valley. The slopes transversed by one player in a strategic interaction are themselves moving in response to the changing play of other players, allowing the landscapes of evolutionary dynamics to resemble those of Escher, in which one may walk consistently down stairs, but find oneself on a higher story.

Refinements. The different equilibria found on the flat floor of a hanging valley, in the landscape metaphor, typically correspond to differences in play off the equilibrium path. The ability to adopt various plans of action for out-of-equilibrium contingencies is reflected in a variety of alternative best replies to the strategies of the equilibrium profile. Until recently, the equilibrium refinements literature sought to address the problem of how rational players might choose between alternative best replies by introducing explicit

or implicit “trembles.” Our approach replaces such trembles by a drift term built into an explicit equilibrating process.

When drift stabilizes the system in a hanging valley, plans of action for dealing with out-of-equilibrium contingencies that are condemned as “irrational” by one refinement concept or another may not be cleansed from the equilibrium strategies. The standard assumptions of refinement theories will then fail. The evolutionary process provides pressures that inhibit such irrationalities, but the limiting outcome of the evolutionary process depends on the relative *rates* at which different kinds of irrationality are eliminated. Very frequently, the selection process is so adept at stamping out certain very deleterious irrationalities that it reaches an equilibrium before it has had a chance to wipe out other types of irrationality. Refinement criteria that deny the possibility of the latter type of irrationality then fail to be honoured in the limit.

Outline. The model is presented in Section 2. Our point of departure is a system of continuous, deterministic differential equations describing how the proportions of the player-populations attached to each strategy adjust over time. We view this system as a model of the selection or learning forces that guide players’ decisions. Section 2 explains the sense in which these differential equations can serve as an approximation of an underlying, stochastic strategy-adjustment process governing the behaviour of the players in the game. In terms of the landscape metaphor, the differential equations describe the motion of the ball rolling over a landscape.

In some cases, the approximation provided by the selection dynamics will suffice for an analysis of long-run behaviour. This may not be the case, however, if these dynamics lead to a valley with a flat floor. If the valley completely encloses a floor of equivalent points, then we need proceed no further, but we must carefully consider the behaviour of the system in the valley if one end hangs above another valley.

We respond by adding a drift term to the selection dynamics. We view this as constructing a more detailed approximation of the underlying stochastic process, incorporating additional considerations into the model that were thought to be insignificant, and hence were excluded, when constructing the selection dynamic, but which have turned out to be important. Sections 3–4 present an example in which drift is crucial in assessing the outcomes of the learning process. The different points on the floor of the hanging valley in this example correspond to different actions at an information set that is not reached in the course of equilibrium play. If drift pushes the players toward actions at this information set that are compatible with the equilibrium, and hence pushes the system toward states that are contained within the equilibrium component, then drift can stabilize the equilibrium. Alternatively, if drift pushes players toward actions that are inconsistent with the equilibrium, then the system can fall off the edge of the equilibrium component and selection pressures can arise that lead to other equilibria.

Section 5 generalizes this intuition to establish a sufficient condition for a component of equilibria to be stable in the face of drift. Any pure-strategy equilibrium, with the property that a strict best response is played at every information set that is reached under equilibrium play, can potentially be stabilized by drift. We can break our sufficient condition into two parts, one addressing drift within the actions corresponding to particular information sets and one addressing relative rates of drift across information sets. At a given unreached information set, the drift must be compatible with the equilibrium, in the sense of pushing the system into the interior of the set of actions which cause the

equilibrium strategies to be best replies. Across information sets, drift at unreached information sets must be sufficiently strong relative to drift at information sets that are actually reached.

The former drift condition is more likely to be satisfied the larger is the component of states corresponding to an equilibrium. Hence, drift is unlikely to stabilize equilibria that impose quite specific, extreme requirements about play at unreached information sets. In the extreme, drift has virtually no chance of stabilizing an equilibrium with an unreached information set at which the set of actions compatible with the equilibrium has no interior.

Our willingness to work with the latter drift condition reflects a belief that players in a game are likely to devote considerable resources to analysing and making decisions that matter, but pay little attention to decisions that do not matter. At information sets that are reached, the selection dynamics are thus likely to be a good description of behaviour and drift is likely to be unimportant, while behaviour will be much more susceptible to drift at unreached information sets. The extent to which drift must be relatively larger at unreached information sets, in order to stabilize an equilibrium, varies from game to game. Our theoretical analysis avoids this issue by allowing relative drift levels to become arbitrarily sensitive to potential payoffs, but our results are clearly less likely to hold in situations where drift is generated by background noise that has nothing to do with the game.

The finding that drift can be important might appear to be a death knell for empirical applications of game theory. How can we hope to make use of a theory whose implications depend upon the details of an arbitrarily small drift process? Section 6 offers several examples that illustrate the forward and backward induction implications of drift. Building on our examples, Section 7 suggests that we might exploit the properties of drift to derive testable predictions from game-theoretic models. Far from being the end of empirical applications, we have hopes that an understanding of drift may provide the key to such work.

2. THE MODEL

We consider an n -player extensive game G of perfect recall. Simultaneous-move games are an important special case. Members of n subpopulations are repeatedly matched to play the game. We speak of “players” when referring to the game G , and “agents” when referring to the members of the populations in the evolutionary model.

Each agent is characterized by a pure strategy in the pure reduced normal form (Mailath, *et al.* (1993)) of G . A *population state* identifies the fraction of agents playing each of the pure strategies available to each subpopulation.

We shall identify a strategy in terms of the actions it specifies at each information set. Let H_i be the set of player i 's information sets and $A(h)$ the actions available at information set h . If no other information set for player i precedes $h \in H_i$, then the state z_h is a nonnegative vector of dimension $|A(h)|$ whose elements identify the fraction of agents in subpopulation i playing each of the $|A(h)|$ pure actions available at h (with the elements of z_h summing to one).² A state z_i for subpopulation i is a vector summarizing

2. If $h \in H_i$ is preceded by an information set for player i , then z_h is a nonnegative vector of dimension $|A(h)|$ whose elements identify the fraction of agents in subpopulation i whose strategies do not preclude reaching h and who play each of the $|A(h)|$ pure actions available at h (with the sum of z_h 's elements being less than or equal to one, where the residual is the proportion of subpopulation i playing strategies that preclude h). We confine the treatment of this case to footnotes to avoid obscuring the important issues.

the state at each of player i 's information sets. A population state z of the entire system is a vector summarizing the state of each information set for each subpopulation. Let Z_h , Z_i , and Z denote the set of states in each case.

The deterministic dynamics. Let $z(t)$ be the population state at time t . As the game is played, agents adjust their strategies. A formal model of this process will necessarily abstract away a host of perturbing influences to obtain an underlying selection process that drives strategy revisions. It is common to model the selection process as a deterministic Markov process represented by the differential equation

$$\dot{z} = \frac{dz}{dt} = f(z). \quad (1)$$

We call the vector-valued function f the *selection function* and refer to (1) as the *selection dynamic* to emphasize the origins of our approach in biological models of natural selection. In the models of Young (1993) and Kandori, Mailath and Rob (1993), f models a best-response learning process. Fudenberg and Levine (1995) examine a model in which f is a smoothed form of fictitious play.

The trajectories of the selection dynamic (1) are studied in the hope that they capture the important features of the strategy-selection process sufficiently well that they approximate the trajectories of the true process from which the selection model is abstracted. To study the extent to which this approximation is successful, we introduce a new differential equation

$$\dot{z} = \frac{dz}{dt} = f(z) + g(z). \quad (2)$$

The trajectories of this *perturbed selection dynamic* provide a better approximation of the true trajectories, because of the inclusion of a *drift function* g capturing small perturbations of the strategy-selection process excluded from the unperturbed dynamic (1). In biology, the vector g summarizes the numerous sources of genetic variation. In Kandori, Mailath and Rob (1993) and Young (1993), g models random alterations in strategies. In Fudenberg and Levine (1995), g captures perturbations of the payoffs.

If the unperturbed selection process (1) is to be useful in predicting the asymptotic behaviour of the true adjustment process, the trajectories of (1) must exhibit the same asymptotic behaviour as those of the perturbed process (2), for the case in which g is small. We shall find that this criterion often fails to be satisfied even if g is vanishingly small. We accordingly conclude that *drift sometimes matters*.

Drift. How can g matter, even though it may be arbitrarily small and we allow no discontinuities in the model? To discuss this issue, it is useful to introduce a distinction between time spans borrowed from the theory of the firm. As in Binmore, Samuelson and Vaughan (1995), the *short run* refers to a length of time too short for a movement from the initial condition to be perceptible. The *medium run* refers to a length of time sufficient for the forces of selection to operate, and hence for $z(t)$ to depart from $z(0)$, but not long enough for $z(t)$ to provide much indication of the asymptotic behaviour of the system. To simplify discussion of the long run, we assume that the deterministic processes currently under discussion converge. The *long run* then refers to a time large enough for $z(t)$ to serve as an approximation to a rest point of the system.

We compare the unperturbed and perturbed selection processes by computing their respective trajectories $z^1(t)$ and $z^2(t)$, beginning with the same initial condition $z^1(0) = z^2(0) = z^0$. Using the unperturbed trajectory $z^1(t)$ to approximate the perturbed trajectory $z^2(t)$ is certainly valid in the short run. It is a well-known continuity property of differential equations that $z^1(t)$ also approximates $z^2(t)$ for any fixed t when g is sufficiently small (Sánchez (1968)). Small levels of drift therefore never matter in the medium run. Problems arise only in the long run.

The study of the asymptotics of $z^1(t)$ and $z^2(t)$ as t becomes arbitrarily large and g becomes vanishingly small, requires taking two limits. Asymptotic predictions obtained from the unperturbed process (1) require that g be allowed to vanish *before* studying the limit $t \rightarrow \infty$. To argue that such a prediction will serve also for the perturbed process (2), given that g is sufficiently small, and hence that $z^1(t)$ can be used to approximate $z^2(t)$ in the long run, is to claim that it is innocent to reverse the order of limits. Drift matters precisely when these limits fail to commute.

Stochastic dynamics. Section 1's informational discussion of evolutionary dynamics using the landscape metaphor was couched in terms of small stochastic shocks to the system. We have elsewhere emphasized the importance of being realistic about building stochastic noise into selection processes (Binmore and Samuelson (1993), Binmore *et al.* (1993)). How can we then proceed as if a real-world stochastic process can be studied with a deterministic model like (2), which we then hope to approximate by (1)?

As an example of what we have in mind, consider a stochastic process in which the game is repeatedly played at discrete points in time, by agents drawn from large but *finite* subpopulations, each of size N . We are then led to a Markov process of the form³

$$z(t + \tau) = H(z(t)), \quad (3)$$

where $z(t)$ is a random variable representing the population state at time t , τ is the interval between successive periods, and H is a random function that depends on N , the procedure by which agents are matched, the payoffs in the game, the rules by which players choose their strategies, and all the perturbations that one would like to be negligible when N is large and τ small.

Biologists standardly reduce (3) to a deterministic differential equation like (2) simply by taking expectations, appealing to the law of large numbers to justify the procedure when N is large. Binmore, Samuelson and Vaughan (1995) study the circumstances under which this informal argument can be made to work in the context of a simple one-dimensional biological model. Appendix II confirms that the extension to the more general case of this paper creates no difficulties.⁴

In this section, we discuss how the deterministic, continuous-time dynamics (2) can provide a useful approximation of the discrete, stochastic process (3), as long as we are interested in large subpopulations N (allowing a deterministic approximation) and short time periods τ (allowing a continuous-time approximation). We confine our attention to a class of Markov processes satisfying (3) that is described in Appendix II, which proves the following. For any *fixed* time t , the solution $z^3(t)$ of a stochastic process in this class

3. The Markov assumption is strong, potentially calling for players to abandon large amounts of previously-collected information in order to concentrate on their current observation, but we leave its relaxation for future research.

4. Appendix II corrects an error in the proof of Binmore *et al.* (1995).

is approximated with probability p arbitrarily close to 1 by the solution $z^2(t)$ of a deterministic process of the form (2), provided that $(N, \tau) \rightarrow (\infty, 0)$ so that $N^2\tau \rightarrow 0$. It suffices for $N^2\tau \rightarrow 0$ if the limit $\tau \rightarrow 0$ is taken first and $N \rightarrow \infty$ second.

Because this result sounds like the continuity property we quoted to justify ignoring drift in the medium run, it is important to be clear on how it helps when studying the stochastic process (3). First solve the deterministic differential equation (2). Next, determine the asymptotic behaviour of the solution $z^2(t)$. If $z^2(t) \rightarrow z^*$ as $t \rightarrow \infty$, then the state z^* serves as a high-probability prediction of the first Nash equilibrium of the game that will be approached by the random variable $z^3(t)$, provided N is large and τ and $N^2\tau$ are small.⁵ The asymptotics of the perturbed selection dynamic (2) thus predict the *first* equilibrium visited by the Markov process (3) with high probability. However, the asymptotic behaviour of (2) does not predict the *asymptotic* behaviour of (3). To make this distinction clear, we introduce a further time span—the ultralong run.

The long run is enough time for a Markov process (3), with appropriately large N and small τ , to approach a rest point z^* of (2) with high probability.⁶ The system will then linger near z^* for a long time. But unlike the deterministic dynamic given by (2), the Markov process will not remain near z^* forever. Given long enough, some perturbing shock that is highly unlikely in any small time interval will eventually occur. If it occurs at time t' , it will bounce the system to a state $z^3(t')$ that is not contained in the basin of attraction of z^* . We can then solve the deterministic dynamic (2) with initial condition $z^2(0) = z^3(t')$ to obtain a high-probability prediction of the subsequent behaviour of the Markov process (3), including the next equilibrium w^* in whose neighbourhood $z^3(t)$ will linger. Given enough time, further such shocks will occur, causing $z^3(t)$ to visit all of the population states in Z , pausing at rest points like z^* and w^* , and rushing past other states in its passage between them. In the *ultralong run*, a probability distribution over Z will be established that describes the likelihood that $z^3(t)$ will be found in any particular region of Z at a large enough time. This limiting distribution assigns a negligible probability to regions not containing a rest point of (2).⁷

In brief, $z^2(t)$ summarizes the high-probability deterministic predictions that can be made about $z^3(t)$ given the initial condition z^0 . Because events that occur only with low probability in short time periods must eventually occur with high probability before some large enough time t' , $z^2(t)$ cannot successfully predict the behaviour of $z^3(t)$ after time t' . Hence, $z^2(t)$ only predicts the behaviour of the Markov process (3) over a *finite* time period. The *asymptotic* behaviour of the deterministic dynamic (2) is used to predict the behaviour of the Markov process (3) in an interval $[t'', t''']$, where t'' is sufficiently large to ensure that $z^2(t)$ is close to a rest point z^* , and t''' (with $t'' < t''' < t'$) is chosen to ensure that only small perturbations occur with significant probability in $[t'', t''']$. As $N \rightarrow \infty$, t'' remains bounded while $t''' \rightarrow \infty$.

Equilibrium selection criteria. We now have *two* equilibrium selection criteria—the history-independent distribution to which the Markov process (3) converges in the

5. To see this, choose t , N and τ such that $\|z^2(t) - z^*\| < \frac{1}{2}\epsilon$ and $\|z^3(t) - z^2(t)\| < \frac{1}{2}\epsilon$ with probability at least $1 - \epsilon$. Then $\|z^3(t) - z^*\| < \epsilon$ with probability at least $1 - \epsilon$.

6. It is important to remember that we are restricting attention to processes whose approximating deterministic dynamics converge.

7. The achievement of Kandori, Mailath and Rob (1993) and Young (1993) was to find cases in which high probability is attached only to regions that contain a *particular* equilibrium \hat{z} . They say that \hat{z} is selected in the long run. We say instead that \hat{z} is selected in the ultralong run.

ultralong run, and the history-dependent state z^* to which the approximating deterministic differential equation converges.⁸ Which is relevant is a matter of the time available for the system to converge.

Binmore, Samuelson and Vaughan (1995) study the expected waiting time to reach the ultralong-run distribution in a simple 2×2 game. Using the best-response dynamics of Kandori, Mailath and Rob (1993) and Young (1993), the expected waiting times are stupendously large for plausible parameter values.⁹ Models in which agents' choices are stochastic offer the system more opportunities to climb out of basins of attraction that are not selected in the long run, and reduce the expected waiting time to several thousand periods, though the waiting time still grows explosively in the population size. Young (1997) has followed Ellison (1993) in exploiting the speeding-up effects of local interaction effects in his model. Araki and Low (1997) have done the same for the model of Binmore, Samuelson and Vaughan (1995). Their simulations yield expected waiting times that might realistically be achieved in laboratories. We therefore do not discount the possibility that ultralong-run selection may be important in some practical contexts. But we believe that long-run equilibrium selection results will frequently be more relevant in applied work.

Properties of the dynamics. If we are to examine the long-run, when can we ignore drift and work with the simpler unperturbed selection dynamic (1) rather than the perturbed dynamic (2)? We begin by making four assumptions about the unperturbed and perturbed deterministic dynamic processes (1) and (2).

Assumption 1. The selection and drift functions f and g are Lipschitz continuous.

This ensures that the processes (1) and (2) each have a unique solution $z(t, z^0)$ for each time $t \geq 0$ and each initial condition z^0 [Hale (1969), p. 18].¹⁰ Lipschitz continuity is not necessary for this conclusion, but it is a standard assumption that is satisfied by common examples like the replicator dynamics. Our basic results continue to hold without this assumption as long as (1) and (2) have unique solutions, though some of the statements and arguments become more tedious. However, we would like to stress that the (ordinary) continuity of f and g plays an important role in our analysis, apart from issues of existence. Adjustment processes like the pure best-response dynamics, that can respond dramatically even to arbitrarily small payoff differences, are not continuous and hence are excluded, but we consider continuous processes to be more realistic.

Assumption 2. For each $t \geq 0$, the solutions to the unperturbed and perturbed selection dynamics are population states, meaning that the proportions of a subpopulation playing the available strategies are nonnegative and sum to one.

Let $\pi_{hk}(z) - \pi_{hk'}(z)$ denote the expected change in payoff if a randomly chosen member of subpopulation i , whose strategy allows information set $h \in H_i$ to be reached, chooses action $k \in A(h)$ instead of $k' \in A(h)$ at information set $h \in H_i$, given the behavior specified by state z (including the agent's choices at other information sets). We assume that the

8. We say that z^* is history-dependent because it depends on the initial condition z^0 .

9. Sandholm and Pauzner (1997) show that even moderate amounts of population growth make these long waiting times infinite. Robles (1996) reaches a similar conclusion.

10. If f and g satisfy Lipschitz continuity or the weaker condition of local Lipschitz continuity, then it is straightforward (because Z is convex) to extend f and g to locally Lipschitz functions on an open set containing Z . The result then follows from Theorems 1.1, 2.1, and 3.1 of Hale [(1969), ch. 1].

selection mechanism, restricted to each information set, is regular and monotonic (*cf.* Samuelson and Zhang (1992)):

Assumption 3.

(3.1) **Regularity.** For each player i , information set $h \in H_i$ and action $k \in A(h)$, the growth rate f_{hk}/z_{hk} is continuous on the state space Z .¹¹

(3.2) **Monotonicity.** For each player i and information set $h \in H_i$, growth rates are monotonic in the sense that, for actions k and k' in $A(h)$ and state $z \in Z$,

$$\pi_{hk}(z) - \pi_{hk'}(z) \geq 0 \iff \frac{f_{hk}(z)}{z_{hk}} \geq \frac{f_{hk'}(z)}{z_{hk'}}.$$

Monotonicity ensures that the proportion of a population playing a relatively high-payoff action grows faster than the proportion playing a relatively low-payoff action.¹²

In many biological models of mutation as well as in the models of Young (1993) and Kandori, Mailath and Rob (1993), the drift function g does not depend on payoffs at all. We think it more appropriate in economic contexts to follow Myerson's (1978) proper equilibrium in modeling more costly mistakes as being less likely to be made. However, we do not view agents as making detailed calculations of the costs of making various mistakes, as such agents would have then calculated enough to avoid making mistakes at all. Instead, we believe that people ordinarily make many decisions simultaneously, only a fraction of which can be analysed carefully. Most decisions are made by applying rules-of-thumb that agents have become accustomed to using in what appear to be similar circumstances, but which may not be suited to the problem at hand. These appear in our model as drift. The larger are the payoff consequences of a decision, the more likely is the decision to command an agent's analytical resources, and hence the less important will be drift.¹³ To capture this idea, we assume that players' behaviour is related to a measure Δ_h of the *potential* cost of making a mistake at the information set h .¹⁴ As in Binmore, Gale and Samuelson [9], the results would remain unchanged if we used other measures of payoff dispersion, but we concentrate on the difference between the largest and smallest expected payoffs because this seems simplest

$$\Delta_h(z) = \pi_{hk}(z) - \pi_{hk'}(z),$$

where $k \in A(h)$ and $k' \in A(h)$ maximize $\pi_{hk}(z) - \pi_{hk'}(z)$. For an information set $h \in H_i$ preceded by no other information set for player i , we then assume:¹⁵

Assumption 4. The vector g_h of drift terms associated with the actions $k \in A(h)$ satisfies

$$g_h(z) = \eta(\Delta_h(z))(\theta_h - z_h), \tag{4}$$

where θ_h is a fixed state in the interior of Z_h and $\eta: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is decreasing and Lipschitz continuous.

11. In particular, finite limits of the growth rates exist at the boundaries of the state space, where some population proportions are zero.

12. The rate at which the proportion playing action k at information set $h \in H_i$ grows may depend upon the choices made by agents in population i at other information sets in H_i , but the direction of the learning at h must reflect payoff differences at h .

13. McKelvey and Palfrey (1992) similarly suggest that players will be more likely to make mistakes or experiment when payoff implications are small.

14. Our debt to Myerson lies in making Δ_h a function of the vector $\pi_h(z)$ of payoffs to each action available at h in state z .

15. If h is preceded by another of player i 's information sets, then $g_h(z) = \eta(\Delta_h(z))(I_h(z)(\theta_h - z_h))$, where $I_h(z)$ is the proportion of subpopulation i whose actions do not preclude h .

The state θ_h is to be interpreted as a parameter determined by the characteristics of the subpopulations from which the agents are drawn. When the agents are distracted from the game at hand, or make a mistake for some other reason, θ_h is the expectation of the distribution of proportions of each subpopulation that takes a particular action at information set h . For example, agents may be accustomed to resolving bargaining problems by coordinating on the fifty/fifty outcome. When faced with the Ultimatum Game, they may not immediately take proper account of its unusual strategic structure, but simply apply the fifty/fifty rule of thumb instead. Such a tendency would be reflected in the specification of θ_h .

The assumption that θ_h lies in the interior of Z_h ensures that drift can always inject strategies that are currently not played in the population, and hence tends to carry the state z away from the boundary. This formulation is justified if there is always some chance z that each strategy might be played, and hence each state in Z_h realized, as a consequence of agents making mistakes.

We do not argue that Assumption 4 is plausible as a qualitative description of the drift likely to occur in *all* contexts. On the contrary, different drift conditions are likely to operate in different circumstances and one must tailor the assumptions made about drift to the application in hand. For example, potential entrants in the Chain-Store Game of Section 3 might be inclined to experiment more than incumbents, or might devote more attention to their decision, leading them to experience either more or less drift. We would capture this by generalizing Assumption 4 to allow different η functions for different subpopulations. But as long as drift levels are sufficiently sensitive to payoffs, our results continue to hold.


3. WHY DRIFT MATTERS

Although small drift terms can have no significant effect in the medium run, they matter in the long run whenever (1) and (2) have different asymptotics. This section uses the Chain-Store Game to illustrate this possibility. We choose this game because textbooks commonly use the same example to justify subgame-perfect equilibria.


The Chain-Store Game. Figure 1.1 shows the extensive form of Selten's (1978) Chain-Store Game. Player 1 moves first, choosing to enter (E) a market or not (N). If player 1 enters, player 2 can acquiesce (A) or resist entry (R). The payoffs satisfy the inequalities $a > e > c$, so that the entrant prefers to enter if the chain store acquiesces but prefers to stay out if the chain store resists. They also satisfy $b > d$, so that the chain store prefers acquiescing to resisting.

Figure 1.2 shows a phase diagram for an unperturbed dynamic. Any regular, monotonic dynamic gives a phase diagram with the same qualitative features. The horizontal axis measures the proportion of agents in population 2 playing R while the vertical axis measures the proportion of population 1 agents playing N . We can see two types of equilibria in Figure 1.2. There is a subgame-perfect equilibrium (denoted by S) in which player 1 enters and 2 acquiesces. There is also a component of Nash equilibria (denoted by \mathcal{N}) that are not subgame perfect, in which player 1 does not enter and player 2 resists entry with probability at least $(a - e)/(a - c)$. Depending on the initial state, the dynamics will converge either to the subgame-perfect equilibrium, or to the Nash equilibrium component \mathcal{N} . The subgame-perfect equilibrium is asymptotically stable, but the Nash equilibria are not. In terms of the landscape metaphor, \mathcal{N} is the bottom of a hanging valley, the left end of which opens into a pit with S at the bottom.


1.1:




1.2:




1.3:



1.4:



1.5:



1.6:




FIGURE 1
Chain-Store Game

How does \mathcal{N} fare when faced with perturbations caused by factors excluded from the model that led to the unperturbed selection dynamic? In the long run, these perturbations appear in the form of drift. Figure 1 shows two specifications of drift, and the corresponding phase diagrams for the Chain-Store Game. In each case, drift causes agents to choose randomly between their two strategies, with equal probability attached to each strategy (*i.e.* θ attaches probability $\frac{1}{2}$ to each strategy at each information set). Figures 1.3–1.4 illustrate the case when η of Assumption 4 is constant. Agents are then equally prone to drift in every state. In this case, the addition of drift yields a system that has a unique, asymptotically stable state that attracts the entire space and which is approximately the subgame-perfect equilibrium.¹⁶

Near the top of the state space, the payoffs to player 2's two strategies are virtually identical. We would accordingly expect player 2 to be more subject to drift near the top of the state space. Figure 1.5 shows the state space for such a drift process, which corresponds to making η of Assumption 4 *strictly* decreasing, with the corresponding phase diagram shown in Figure 1.6.¹⁷ There are now two asymptotically stable states. One of these (denoted by ξ) is approximately a Nash equilibrium that is not subgame perfect. In this case, drift stabilizes the system in the middle of a hanging valley. Different specifications of drift can thus give rise to different long-run behavior.

Drift and the medium run. Section 2 discusses how the *long-run* behaviour of the stochastic process (2) is approximated with high probability by the asymptotics of the perturbed differential equation (2). The asymptotics of the unperturbed deterministic process (1) sometimes provide an analogous *medium-run* approximation.

To illustrate this point, consider a specification of drift in the Chain-Store Game that *fails* to stabilize the perturbed dynamic in the hanging valley \mathcal{N} , as in Figures 1.3–1.4. The subgame-perfect equilibrium S will then be our long-run prediction for $z^3(t)$, the underlying stochastic process. Suppose, however, that the initial condition z^0 lies in the basin of attraction, relative to the *unperturbed* dynamic (given by (1)), of an equilibrium $v^* \in \mathcal{N}$. Then the unperturbed dynamic, $z^1(t)$, approaches v^* as $t \rightarrow \infty$. If the drift term is sufficiently small, the perturbed dynamic, $z^2(t)$, will also initially approach v^* . As $z^2(t)$ approaches \mathcal{N} , the selection term $f(z)$ becomes small, since this selection term is zero on the set of equilibria \mathcal{N} . The forces of selection will then eventually be dominated by those of drift. We have assumed that in the long run, drift will push the system away from \mathcal{N} to the subgame perfect equilibrium S . But if drift is small, $z^2(t)$ will move away from v^* very slowly. The same is therefore true with high probability of $z^3(t)$, and the medium-run behaviour of $z^3(t)$ thus calls for the system to approach v^* and then slowly drift away. In some cases, $z^3(t)$ may linger so long in the vicinity of v^* that only the medium-run behaviour is of practical relevance, and this medium-run behaviour is captured by $z^1(t)$ as well as $z^2(t)$, allowing us to work with the unperturbed dynamic. In other cases, it will be essential to examine long-run behaviour, and hence to work with the perturbed process $z^2(t)$ that incorporates drift.

Do we need dynamics? Our analysis is capable of stabilizing Nash equilibria in the hanging valley \mathcal{N} if the drift has appropriate properties. These equilibria are not subgame

16. "Approximately" is needed here because the inward-pointing drift prevents exact convergence to the subgame-perfect equilibrium.

17. Because player 2 is relatively prone to drift near the top of the state space, the drift trajectories near \mathcal{N} have flattened, causing relatively rapid movements in player 2's strategies toward the fifty/fifty mixture specified by θ . The trajectories are quite steep near those states at which R is played with probability $\frac{1}{3}$, where the payoff to player 1's strategies are virtually identical and hence player 1 is relatively subject to drift.

perfect because they involve the use of the strategy R , which is an inferior reply at the incumbent's information set. How does this result compare with the orthodox game-theoretic approach, in which the stability of equilibria is studied by introducing perturbations in the payoffs or strategies of the game? With the exception of Fudenberg, Kreps and Levine (1988)¹⁸, such orthodox analyses destabilize equilibria, like those in the component containing (N, R) , that sit in a hanging valley which calls for an inferior reply to be chosen at an unreached information set. Which stability analysis should command our attention? Samuelson (1994) examines a setting in which the trembles of orthodox stability analysis need to be infinitely large compared with perturbations in the equilibrating process before their effect becomes significant.

4. WHEN CAN DRIFT BE IGNORED?

Having shown that drift can matter, the question arises: when can the problems raised by drift safely be neglected?

If z^* is a *hyperbolic* stationary state of the unperturbed selection process (1), then the perturbed selection process (2) has a rest point close to z^* .¹⁹ The latter rest point moves arbitrarily close to z^* as the drift term becomes arbitrarily small, and has the same stability properties as z^* (Hirsch and Smale [(1974), Theorems 1–2, p. 305]). Hence, if we are working with hyperbolic rest points of $\dot{z} = f(z)$, then we can ignore drift. The rest points of $\dot{z} = f(z)$ provide approximate information about the rest points of $\dot{z} = f + g$ that lie nearby and are of the same type. The approximation becomes arbitrarily sharp as g gets small. Nonhyperbolic rest points, however, do not have this “structural stability.” The addition of arbitrarily small drift terms to the equation $\dot{z} = f(z)$ can completely change the nature of a nonhyperbolic rest point.

It is often said that almost all dynamic systems have only hyperbolic rest points.²⁰ Many mathematicians and physicists therefore ignore the possibility of nonhyperbolic rest points when working with similar dynamic systems. However, the economics of the applications to which learning models are applied usually force us to confront nonhyperbolic rest points. For example, the equilibria in the component \mathcal{N} of the Chain-Store Game are not hyperbolic. This is not exceptional. Any Nash equilibrium that specifies a path which does not reach every information set (excluding some games featuring fortuitous payoff ties) fails to be isolated, and hence fails to be hyperbolic, under all of the familiar selection dynamics.

We might try to respond by shifting the focus of our analysis to components of rest points, seeking components that satisfy a set-valued notion of asymptotic stability such as that offered by Bhatia and Szegő [(1970), Def. 1.5, p. 58].²¹ Since the Nash equilibria

18. They stabilize *all* Nash equilibria by allowing the players' knowledge of the rules of the game or their own payoffs to tremble.

19. A stationary state or rest point of a differential equation is *hyperbolic* if the Jacobian matrix of the differential equation, evaluated at the rest point, exists and has no eigenvalues with zero real parts. Hyperbolic rest points are isolated and are either sources, saddles or sinks (Hirsch and Smale [(1974), ch. 9]). Nonhyperbolic rest points need not be isolated and isolated rest points need not be hyperbolic.

20. The Peixoto theorem (Hirsch and Smale [(1974), p. 314]) shows that for two-dimensional systems, there is a precise sense in which “almost all” dynamic systems have only hyperbolic rest points. A similar result holds in higher dimensions for certain classes of dynamic systems, such as linear and gradient systems (Hirsch and Smale, pp. 313–315).

21. Ritzberger and Weibull (1995), Schlag (1993) and Swinkels (1993) examine components that satisfy a set version of asymptotic stability. A similar motivation but different techniques appear in Ritzberger (1993), who introduces the idea of an essential component, where (very roughly) a component is essential if all nearby games have nearby equilibria.

in a component like \mathcal{N} in the Chain-Store Game often specify the same path through the tree, differing only in the details of what happens off the equilibrium path, there is often no loss of information in treating components as a single entity. This approach is adequate when the component of rest points constitutes the flat floor of a pit in the evolutionary landscape. Drift will then be irrelevant, since we do not care where it takes us on the floor of the pit. But the Chain-Store Game of Figure 1 shows that we must often expect our components of rest points to form the floor of a hanging valley whose equilibria can be stabilized by some types of drift and destabilized by others. Drift then may or may not push the system over the cliff, and cannot be ignored.

5. WHEN DRIFT MATTERS

Section 3 shows that drift can stabilize a Nash equilibrium in the hanging valley \mathcal{N} of the Chain-Store Game. If z^* is a Nash equilibrium in a general game, this section defines the notion of z^* -compatible drift to provide a criterion that ensures stability for some Nash equilibrium w^* that lies on the floor of the same valley as z^* .

A Nash equilibrium z^* is a *strict-path* equilibrium if, at every information set reached with positive probability under z^* , the action prescribed by z^* is a strict best response. Notice that the strategies that support such an equilibrium need not be pure, since mixtures are allowed off the equilibrium path, though no mixtures can occur on the equilibrium path.²² For generic extensive-form games, all pure-strategy equilibria are strict-path equilibria. Figure 5 below contains a game with a pure-strategy but not a strict-path equilibrium.

We could apply similar techniques to Nash equilibria that allow mixtures along the equilibrium path, but postpone such an inquiry for future work. Mixed equilibria raise new analytical complications because many monotonic selection processes cannot converge to mixed equilibria. For example, mixed equilibria are centers for common versions of the replicator dynamics.²³

If z^* is a strict-path equilibrium, let $\zeta(z^*) \subset Z$ be the set of Nash equilibria specifying the same path of play as z^* .²⁴ We refer to $\zeta(z^*)$ as a *common-path set of equilibria* and refer to the path of play itself as a *strict path*. Because the equilibria in $\zeta(z^*)$ have the same path of play, they have the same collection of information sets that are reached with probability zero.

Let z^* be a strict-path equilibrium. We say that the specification of drift is z^* -compatible if we can find an equilibrium w^* , in the relative interior of the common-path set of equilibria $\zeta(z^*)$ and hence giving the same path of play as z^* , with the property that w^* attaches the action θ_h of Assumption 4 to each information set h that is reached with zero probability under equilibrium z^* , but is not precluded by previous choices of the player

22. This concept differs from simply applying quasi-strictness in the pure reduced normal form in that a strict-path equilibrium allows no mixtures along the equilibrium path, unlike a quasi-strict equilibrium. However, a quasi-strict equilibrium induces mixtures off the equilibrium path that have full support on the set of available actions, in order to satisfy the condition that the equilibrium contains every pure-strategy best response within its support, while a strict-path equilibrium does not impose such a requirement.

23. Fudenberg and Levine (1993a, b) direct attention to *self-confirming* equilibria in extensive-form games, the study of which we also postpone.

24. As in the less formal introductory sections we speak interchangeably of Nash equilibrium strategy profiles and the corresponding population states.

who moves at h .²⁵ If a z^* -compatible drift process were the only force acting at an information set h that is unreached under equilibrium z^* , then it would converge to the actions given by θ_h at information set h , which would support the outcome z^* as a Nash equilibrium.

The Chain-Store Game shows that we must pay attention to the sensitivity of drift levels to payoff differences. To examine this question, we parameterize the function η of Assumption 4. Replacing η by η_n for $n = 1, 2, \dots$, we can then study comparative statics questions by observing the result of varying n in the corresponding drift function g_n . In addition to the requirements of Assumption 4, we require:

Assumption 5.

$$\eta_n(0) = \eta_1(0), \quad n = 1, 2, \dots$$

$$\lim_{n \rightarrow \infty} \frac{\eta_n(\Delta)}{\eta_n(\Delta')} = 0 \quad \text{if } \Delta > \Delta'.$$

For a fixed θ , higher values of n then correspond to drift that is more payoff-sensitive. The function $\eta_n(\Delta) = e^{-n\Delta}$ satisfies these assumptions.

A drift function g_n is z^* -compatible if, and only if, the same is true of g_1 . It is therefore meaningful to speak of drift being z^* -compatible without specifying the value of the parameter n . In addition, if drift is compatible with one element of a common-path set of equilibria corresponding to a single outcome, then it is compatible with every element of that set.

It does not follow from z^* -compatibility that z^* itself is the limit of rest points of the perturbed dynamic as $n \rightarrow \infty$, but the following proposition, proved in Appendix I, indicates that something similar must be true for at least one equilibrium w^* with the same strict path.

Definition 1. An equilibrium w^* is stabilized by drift if, for any neighbourhood V containing w^* , there exists an integer $n(V)$ such that for any specification of drift that is at least as sensitive to payoffs as $g_{n(V)}$, i.e. for any $n > n(V)$, V contains two further neighbourhoods B and U_n , such that,

- (1.1) any trajectory of the perturbed dynamic that begins in B remains in V , eventually enters U_n , and never subsequently leaves U_n ;
- (1.2) $w^* \in U_n \subset B \subset V$;
- (1.3) U_n shrinks to w^* as $n \rightarrow \infty$.

Proposition 1. *Let drift be compatible with a strict-path equilibrium z^* . Then there exists an equilibrium w^* with the same strict path as z^* that is stabilized by drift.*

Proposition 1 shows that z^* -compatibility provides a criterion for drift to create a stable rest point on the floor of a hanging valley when drift is sufficiently payoff-sensitive. We shall often find it convenient to say that the payoff path corresponding to z^* and w^* is stabilized by such drift. In all the examples we have computed, the neighbourhood U_n contains only one rest point of the perturbed dynamics, which is asymptotically stable.

25. The equilibrium w^* is in the relative interior of $\zeta(z^*)$ if the actions prescribed by w^* at each nonprecluded information set reached with probability zero under z^* and w^* lie in the interior (relative to the product of the Z_h 's corresponding to these information sets) of the set of actions consistent with z^* and w^* being Nash equilibria.

Proposition 1 holds for arbitrarily small drift levels, since we can replace g_n by λg_n , for any sufficiently small λ , without affecting the conclusion. What is crucial is that n be sufficiently large, so that drift levels are sufficiently sensitive to payoff differences.

6. EXAMPLES


This section uses the results of Proposition 1 to examine several games that have become standard in thinking about backward and forward induction. Our conclusion is that the perturbed dynamic process may stabilize outcomes that fail to satisfy backward and forward induction properties. In each example, the common-path set of equilibria is a component of equilibria in the standard sense.

Backward induction. We first return to the Chain-Store Game of Figure 1. Let z^* be a Nash equilibrium in which the entrant does not enter, so that z^* is an element of the component \mathcal{N} . Then z^* must specify that the incumbent resist entry with probability at least $(a-e)/(a-c) = r^*$. Drift is z^* -compatible if, and only if, $\theta_r > r^*$, where θ_r is the probability with which drift induces a member of population 2 to resist entry. For example, if $r^* < \frac{1}{2}$ and $\theta_r = \frac{1}{2}$, so that drift prompts members of population 2 to choose between their two strategies with equal likelihood, while entry is unprofitable if only a minority of incumbents resist entry, then drift is z^* -compatible. This is the case illustrated in Figure 1. Alternatively, if $r^* > \frac{1}{2}$, so entry is unprofitable only if a majority of incumbents resist, then drift is z^* -compatible only if drift incorporates a bias in favour of resisting. Notice that compatibility depends only on the actions of agent 2.


If drift is z^* -compatible, then we can proceed to the next step. The perturbed process will have stable outcomes near \mathcal{N} if the drift process is sufficiently sensitive to payoffs. In Figure 1.3, drift is not sensitive to payoffs at all and there is no such stable state. In Figures 1.5–1.6, drift is sufficiently sensitive to payoffs, and the component \mathcal{N} is stabilized.

Why are the relative drift levels induced by payoff differences so important? Figure 2 shows a detail of the region of the Chain-Store Game's phase diagram near the component \mathcal{N} . The arrows represent the direction of the selection process and the direction of the combined drift of the agents in the two subpopulations. Figure 2.1 corresponds to Figures 1.3–1.4, where drift is compatible with \mathcal{N} but is not sensitive to payoff levels. The key here is that the trajectories of the drift process near \mathcal{N} create a force pushing the state to the right that is weaker than the force of the learning dynamics pushing the state to the left. The state is then pushed ever leftward until it reaches the cliff-edge at the end of the component \mathcal{N} and falls into the basin of attraction of the subgame-perfect equilibrium. In terms of our landscape metaphor, drift causes the state to roll to the end of the hanging valley corresponding to \mathcal{N} and fall off a precipice into a pit whose bottom is S .

In Figure 2.2, which corresponds to Figures 1.5–1.6, drift pressures pushing the state back into the basin of attraction of \mathcal{N} (*i.e.* to the right) are more powerful than learning pressures as long as the state is not too close to the edge of \mathcal{N} . Drift now stabilizes the state in the middle of the hanging valley \mathcal{N} . The heart of the proof of Proposition 1 involves showing that this relationship will hold whenever drift is compatible with an equilibrium and sufficiently sensitive to payoffs. Payoff sensitivity ensures that the drift near a component \mathcal{N} will be shaped primarily by the agent whose unreached information set creates the equilibrium indifference that gives rise to the component \mathcal{N} . Compatibility



2.1



2.2

FIGURE 2

Detail of phase diagram, Chain-Store Game

ensures that the result will be a force pushing the system into the basin of attraction of the component \mathcal{N} .²⁶

Because the ability of drift to stabilize equilibria arises out of the relative directions of the drift and learning forces, there is a trade-off between the properties that the drift and learning processes must have in order for drift to be effective. We have assumed relatively little about the learning process, and then have made strong assumptions about relative drift rates at reached and unreached information sets in order to ensure that the drift and learning dynamics sometimes have the relationship shown in Figure 2.2. We could impose fewer restrictions on the drift process if we were made stronger assumptions about the learning process, requiring learning to be relatively fast at reached compared to unreached information sets.

The first step in evaluating the Chain-Store Game was to check for compatibility of the drift process with component \mathcal{N} . Notice that for any fixed specification of drift, compatibility is more likely to be satisfied the larger is the component \mathcal{N} . This suggests that the perturbed process should be more likely to lead to outcomes near \mathcal{N} when payoffs are chosen so that \mathcal{N} is relatively large. This observation forms the basis for a comparative statics investigation discussed in Section 7.

Outside options. The shape of the extensive form of the game in Figure 3 suggests that it be called the “Dalek Game” (Binmore (1987–88)).

26. We could formulate a geometric version of Proposition 1 based on these ideas. Intuitively, suppose that, given a component \mathcal{N} , we could find a subset \mathcal{N}^* with the property that the drift process on \mathcal{N}^* points into the basin of attraction of \mathcal{N}^* under the selection dynamics. Then under the perturbed process, all trajectories beginning in a neighbourhood of \mathcal{N}^* would approach \mathcal{N}^* . This result holds for any drift process and hence replaces the compatibility and payoff sensitivity of drift in Proposition 1 with the geometric “pointing into” condition. Proposition 1 allows us to formulate our result while working directly with the game rather than with the phase diagram of the dynamics and directs attention to the role of drift at unreached information sets.




FIGURE 3
Dalek Game

The Dalek Game has two components of Nash equilibria, including a strict Nash equilibrium given by (M, L) with payoffs $(9, 3)$ and a component \mathcal{N} of equilibria with payoffs $(7, 4)$, in which player 1 takes the outside option (plays T) and player 2 plays R with probability at least $2/9$. The former is a (hyperbolic) sink under regular, monotonic dynamics while the stationary states in the latter component are not hyperbolic.

It is common to argue that forward induction forces us to restrict attention to the equilibrium (M, L) in this game.²⁷ How does this forward induction argument fare in our terms? For drift to be compatible with Nash equilibria in the component \mathcal{N} , drift must induce player 2 to play R with probability greater than $2/9$. If it does, then perturbed dynamics for which drift is sufficiently sensitive to payoffs will yield stable outcomes in which player 1 takes the outside option.

It is interesting to note that Balkenborg's (1994) careful experimental investigation of the Dalek Game finds that the outside option is virtually always chosen. This experimental result is one of a number which suggest that the forward induction criterion has little predictive power. From our point of view, its results are mildly encouraging because they are consistent with the type of drift analysed in Proposition 1. Section 7 proposes a more telling experiment capable of refuting the relevance of this type of drift to the learning behaviour of human experimental subjects.

Burning money. We consider another common forward-induction example, the general form of which is due to van Damme (1987, 1989) and Ben Porath and Dekel (1992). The Battle-of-the-Sexes game has two pure-strategy Nash equilibria and one mixed-strategy equilibrium. It is common to dismiss the mixed-strategy equilibrium. How do we

27. Kohlberg and Mertens (1986), early and forceful advocates of forward induction, introduced a version of the Dalek game. To apply forward induction, we might appeal to the iterated elimination of weakly dominated strategies. B is strictly dominated for player 1. Removing B causes R to be weakly dominated for player 2, the removal of which causes T to be weakly dominated for player 1, leaving (M, L) . Alternatively, we could appeal to the forward induction reasoning of van Damme (1987, 1989) or to the normal form variant of this reasoning given in Mailath, Samuelson and Swinkels (1993).




FIGURE 4
Burning-Money Game

choose between the two pure-strategy equilibria, given that the players have opposing preferences over these equilibria? Both equilibria are hyperbolic stationary states, so that appealing to drift is no help in this game.

Suppose that before the game is played, player 1 has the option of burning two dollars. The payoffs in the larger signaling game that results are shown in Figure 4. Notice that if the money is burned, then 2 is subtracted from player 1's payoffs. The iterated elimination of weakly dominated strategies leads to a unique outcome for this game of (*Not*, *T*; *LL*), giving player 1 her preferred payoff.

The assessment of this game again turns on the issue of compatibility. Consider the equilibrium in which player 1 burns the money and (*T*, *L*) is then played, yielding the payoff pair (2, 1). Compatibility requires that, at player 2's right information set, drift induces agents to play *R* with a probability exceeding $\frac{1}{2}$. A drift process that introduces strategies in equal proportions then cannot stabilize an equilibrium in which money is burned.

Consider the equilibrium in which no money is burned but (*B*, *R*) is played, yielding the payoff pair (1, 4). Compatibility requires that, at player 2's left information set, *R* is played with a probability exceeding $\frac{1}{4}$. In this case, compatibility appears to be plausible. The perturbed dynamics may then yield outcomes in which player 1 burns no money, but in which the forward induction argument carries no force and payoffs are (1, 4).

Incompatible drift. Not all Nash equilibria yield outcomes that can be stabilized by drift. To see that this is the case, consider the Nash equilibrium (*L*, *L*) in Figure 5. This component cannot be stabilized because there is no interior to the set of actions at player 2's information set that are consistent with the equilibrium, and hence no possibility for inward-pointing drift to be compatible with this equilibrium.

Cheap talk. One of the apparent successes of evolutionary game theory has been to use refinements of the evolutionarily stable strategy concept to examine issues of cheap talk. Blume, Kim and Sobel (1993), Kim and Sobel (1992), Matsui (1991), Sobel (1993),




FIGURE 5

Game with Nash equilibrium that cannot be stabilized

and Wärneryd (1990) establish conditions under which the evolutionary process, operating on the cheap talk game, selects efficient equilibria of the underlying game. To see why such a result might be expected, consider an outcome in which everyone plays an inefficient equilibrium. Let a strategy appear in which some agents send the currently unused message α and in which all agents play the efficient equilibrium whenever at least one agent sends message α . The resulting dynamics will lead to an outcome with only the efficient equilibrium being played. Two steps are important in making this argument. The first is to establish that an unused message exists. The second, upon which we shall focus, is to ensure that agents react to this message by playing the efficient equilibrium. The validity of both steps can be seen as raising questions of drift.²⁸


To consider cheap talk, we begin with the Stag Hunt Game shown in Figure 6.1. This game has two Nash equilibria, (A, A) and (B, B) , with the former being risk-dominant and the latter payoff-dominant.

Now suppose that, before playing the game, player 1 has an opportunity to announce either A or B . We interpret this as an announcement of the strategy that the agent claims she will play, but the announcement is cheap talk, in the sense that it imposes no restriction on the action that the player actually takes. The game with cheap talk is then given in Figure 6.2. One component of pure-strategy Nash equilibria gives payoffs $(10, 10)$. This component is asymptotically stable, a reflection of the fact that 10 is the largest payoff available in the game.


Consider an equilibrium in which player 1 plays AA (announce A and play A) and payoffs $(8, 8)$ are received. The conventional cheap talk argument now observes that if player 2's strategy is AB (respond to announcement A by playing A and to B by playing B), then it is a better reply for player 1 to choose BB , leading to a payoff of $(10, 10)$. The analysis is then completed with an argument as to why we would expect player 2 to be playing AB rather than AA . But the difference between AB and AA appears only off the equilibrium path. The plausibility of AB rather than AA therefore depends on what is assumed about drift.

Let \mathcal{N} be the component of equilibria in which player 1 chooses AA . Drift at the unreached information set corresponding to 1's announcement of B is compatible with Nash equilibria in \mathcal{N} if it causes player 2 to choose A at the left information set with

28. Banerjee and Weibull (1993) note that inefficient outcomes can persist in their evolutionary model with "discriminating" players for reasons analogous to those that allow inefficient outcomes to persist when there are no unused messages.



6.1



6.2

FIGURE 6
Stag hunt and Cheap Talk Games

probability at least $\frac{1}{5}$. The perturbed dynamics may therefore give stable outcomes near (A, A) . Because the announcements are cheap talk, drift can similarly lead to stable outcomes in which player 1 announces B but plays A , with player 2 playing A in response to the announcement of B and playing A at the unreached information set following an announcement of A with probability at least $\frac{1}{5}$.

We therefore consider it unsafe to conclude that evolutionary processes will necessarily select efficient outcomes in cheap talk games. This view is to be contrasted with much of the literature on evolutionary processes in cheap talk games. For example, Sobel (1993) rejects a component of equilibria if there is *any* realization of the underlying stochastic drift process that leads away from the component. This may be an appropriate notion for an ultralong-run analysis, since the ultralong run is a period of time long enough that any realization of the process that *can* happen *will* happen. For a long-run analysis, however, drift may yield quite different results.

Cooper, DeJong, Forsythe and Ross (1992) conduct an experimental investigation of the Stag Hunt Game and a Cheap Talk Game shown in Figure 6. They find that in the Stag Hunt Game, virtually all of the players chose action A . When cheap talk was allowed, virtually all of the announcements were strategy B , but strategy A was still often played, appearing slightly more than 30% of the time. The latter result is consistent with our observation that outcomes in which A is played can be stabilized by drift, though a more likely explanation appears to be that the system has simply not yet settled on an equilibrium.²⁹

Cooper, DeJong, Forsythe and Ross also investigate a game in which the two players simultaneously announce either A or B , and then play the Stag Hunt Game of Figure 6. Figure 7 presents the normal form of this game. A strategy in such a game is now a triple,

29. For example, most of the cases in which A was played involved an opponent who played B , suggesting disequilibrium rather than an equilibrium in which players have coordinated on A .

	<i>AAA</i>	<i>AAB</i>	<i>ABA</i>	<i>ABB</i>	<i>BAA</i>	<i>BAB</i>	<i>BBA</i>	<i>BBB</i>
<i>AAA</i>	8, 8	8, 8	8, 0	8, 0	8, 8	8, 8	8, 0	8, 0
<i>AAB</i>	8, 8	8, 8	8, 0	8, 0	0, 8	0, 8	10, 10	10, 10
<i>ABA</i>	0, 8	0, 8	10, 10	10, 10	8, 8	8, 8	8, 0	8, 0
<i>ABB</i>	0, 8	0, 8	10, 10	10, 10	0, 8	0, 8	10, 10	10, 10
<i>BAA</i>	8, 8	8, 0	8, 8	8, 0	8, 8	8, 0	8, 8	8, 0
<i>BAB</i>	8, 8	8, 0	8, 8	8, 0	0, 8	10, 10	0, 8	10, 10
<i>BBA</i>	0, 8	10, 10	0, 8	10, 10	8, 8	8, 0	8, 8	8, 0
<i>BBB</i>	0, 8	10, 10	0, 8	10, 10	0, 8	10, 10	0, 8	10, 10

FIGURE 7

Two-sided Cheap Talk Game

such as *ABA*, interpreted as announcing *A*, playing *B* if the opponent announces *A*, and playing *A* if the opponent announces *B*. Behaviour in this game differs markedly from the case of one-sided cheap talk, with the strategy *BBB* now being virtually always played by the experimental subjects.

This game has a component of equilibria that gives outcome (B, B) , for payoffs $(10, 10)$, that includes every pure strategy combination giving payoff $(10, 10)$ and which is asymptotically stable under a deterministic, monotonic dynamic. This stability arises out of the fact that 10 is the largest possible payoff in the game. The game also has two pure-strategy Nash equilibria which produce the outcome (A, A) , one in which both players choose *AAA* and one in which both choose *BAA*. In both cases, announcements are ignored. Neither of these equilibria is a strict-path equilibrium and neither is stable under a deterministic, monotonic dynamic (much like the equilibrium (L, L) of Figure 5). If strategy *AAA* is played, for example, then it takes only a trace of strategy *BAB* to trigger dynamics that lead to everyone playing strategy *BAB* for outcome (B, B) . No small amount of drift can then stabilize pure-strategy equilibria in which action *A* is chosen, making it no surprise that action *B* is commonly observed in the experimental outcome.³⁰

7. PREDICTIONS

Despite the Two-Sided Cheap Talk Game and Figure 5, we will typically find components with the property that some specifications of drift will stabilize them but other specifications of drift will not. The interesting question concerns the likelihood with which a component will be stable. This appears to be an impossible question to answer because it depends upon the specification of a drift process about which we are likely to have little information. On the other hand, progress on this question is essential if we aspire to a theory that will be useful in studying behaviour, whether in the laboratory or in the field. We accordingly turn to the question of how an understanding of the role of drift can provide the foundation for testable predictions. In particular, we suggest an experiment designed to examine the role of drift in equilibrium selection. The experiment proceeds in three stages.

The basic intuition for this experiment emerged from our discussion of the Chain-Store Game, where we noted that drift is more likely to be compatible with an equilibrium that allows a large set of out-of-equilibrium behaviours than an equilibrium with very

30. There remains the mixed, babbling equilibrium in which players randomize between *AAA* and *BAA*.

stringent out-of-equilibrium requirements (as in Figure 5). Our task now is to make this intuition usefully precise. Since the problem is inherently an exercise in comparative statics, the experiment must involve comparisons of outcomes of games that are similar but involve different payoffs.

The suggested experiment centers around a version of the Dalek Game studied experimentally by Balkenborg (1994) (Figure 3). Consider the version of this game shown in Figure 8. The component of equilibria supporting the outcome (T, R) is relatively large when x is small, and shrinks as x grows. Our model of drift then leads to the prediction that we should see the outcome (T, R) when x is small (say $x = 0$ or 1) and the outcome (M, L) when x is large (say $x = 6$).




FIGURE 8
Game 8—Modified Dalek Game

The equilibria (T, R) and (M, L) in Figure 8 are both subgame perfect. We are especially interested in the ability of drift to stabilize Nash equilibria that are not subgame perfect. In light of this, consider the game in Figure 9. Provided that $x > 0$, this game has a unique subgame-perfect equilibrium (M, L) , though (T, R) remains a Nash equilibrium. Our prediction is again that the latter equilibrium will appear when x is relatively small (such as $x = 1$) and (M, L) when x is large (say $x = 6$). Our suggested initial experiment thus involves running Games 8 and 9 while varying the value of x through the suggested range of values.

These predictions are based on the observation that, when x is large, the component supporting equilibrium (T, R) is small, and it is accordingly less likely that drift will be compatible with this component. Experimental outcomes may match our predictions, and yet drift still be compatible with (T, R) when x is large (contrary to our suggested explanation), simply because large values of x prompt experimental subjects to adopt initial play that happens to be in the basin of attraction of (M, L) , while small values yield initial play in the basin of attraction of (T, R) .

Our suggestion for addressing this possibility is to construct a second experiment with Games 8–9 in which different groups of subjects begin by playing a *common* game whose payoffs are gradually allowed to diverge, during the course of repeated play, until one group plays Game 8 and the other Game 9. This would ensure that the players cannot




FIGURE 9

Game 9—Modified Dalek Game

condition their initial behaviour (or their behaviour in the initial few periods) on the value of x and hence cannot begin with systematically different initial play. Any differences in outcomes could then reasonably be attributed to drift influencing the learning process in different ways.

To gain further insight into the role of drift, and pose a more demanding challenge for the theory, we suggest a third experiment in which the human subjects are initially matched not against other subjects but against a computer whose play is programmed to direct one group of subjects to equilibrium (M, L) and one group of (T, R) , in each of Games 8 and 9.³¹ The computerized players would then be eliminated and the subjects matched against each other. When x is small, both (M, L) and (T, R) are stable outcomes of the perturbed dynamics, in each of Games 8–9, and each group of subjects should continue to play the outcome established by its initial conditioning. When x is large, however, we expect only (M, L) to be stable, and expect (T, R) to be displaced by (M, L) in those groups of subjects who have initially been conditioned to play (T, R) .

8. CONCLUSION

The ideas behind this paper are simple: The criterion for a model to be successful is that it include important factors and exclude unimportant ones. But how do we know what is important and what is not? In the case of evolutionary games, the model itself provides some of the answers. If the model produces stationary states that are not hyperbolic and do not occur in components that satisfy some variation of asymptotic stability, then important factors have been excluded from the model and the latter should be expanded.

The factors to be added to the model are important, in the sense that they can have a significant impact on the behaviour of the dynamic system, but they also may be arbitrarily small in magnitude. It is presumably because they are small that they are excluded from the model in the first analysis. How can a model whose behaviour is shaped by

31. Binmore *et al.* (1993) and Winter and Zamir (1996) find a similar conditioning process to be quite effective.

arbitrarily small factors be of any use in applications? One conclusion of this paper is that, while the factors themselves may be small, their existence can nevertheless be used to derive experimentally testable predictions that do not depend upon observing arbitrarily small magnitudes.

APPENDIX

I. Proof of Proposition 1

Let z^* be a strict-path equilibrium. Construct the equilibrium w^* by letting w^* take the same action as z^* at any information set reached with positive probability under z^* , and let w^* take the mixed action θ_h at any information set h reached with zero probability under z^* , but not precluded by previous choices of the player moving at h . Notice that because θ is z^* -compatible, w^* is a Nash equilibrium contained in the relative interior of $\zeta(z^*)$. Let H_+ be the set of information sets reached with positive probability under equilibrium w^* and let H_- be the set of information sets reached with probability zero under w^* but not precluded by previous choices of the player moving at the information set.

Let $V(\delta, \epsilon)$ be the set of states such that $\|z_{h_i} - w_{h_i}^*\| \leq \delta$ for all $h_i \in H_+$ and $\|z_{h_i} - w_{h_i}^*\| \leq \epsilon$ for all $h_i \in H_-$, where $\|\cdot\|$ is the max norm.³² Let V be a neighbourhood of w^* .

Step 1. Choose $\epsilon^* > 0$ sufficiently small to ensure:

- (i) $V(\epsilon^*, \epsilon^*) \subset V$. This can be done because V is a neighbourhood of w^* .
- (ii) $V(0, \epsilon^*)$ is a set of strict-path Nash equilibria giving the same outcome as w^* and contained in the relative interior of $\zeta(z^*)$. This can be done because w^* lies in the relative interior of $\zeta(z^*)$.
- (iii) There exist $\bar{\beta} > \beta > 0$ such that the payoff to the action specified by w^* at any information set $h \in H_+$ exceeds the payoff to every other action available at h by at least $\bar{\beta}$; while the payoffs to two actions at an information set $h \in H_-$ differ by at most β . This can be done because w^* specifies strict best responses at information sets in H_+ , while information sets in H_- are unreached in equilibrium.
- (iv) There exists $k_0 > 0$ such that at every information set in $h \in H_+$ and for every state in $V(\epsilon^*, \epsilon^*)$, the total of the subpopulation proportions playing each of the nonequilibrium actions declines at a rate greater than or equal to k_0 . Condition (iii) and the monotonicity of f ensure that this can be done.

Step 2. Choose $\delta < \epsilon^*/4$ sufficiently small that, if there were no drift, then any trajectory beginning in the set $V(\delta, \delta)$ would converge to a state in $V(0, \epsilon^*/4)$ without leaving $V(\epsilon, \epsilon)$. We take $V(\delta, \delta)$ to be the set B of the theorem. To verify its existence, consider a trajectory whose initial condition lies in $V(\delta, \delta) \subset V(\epsilon^*, \epsilon^*)$. Let $\alpha(t)$ be the largest (over the information sets in H_+) proportion of a subpopulation at time t that is not playing the equilibrium action specified by w^* . Then $\alpha(0) \leq \delta$ by definition, and $\alpha(t)$ decreases to zero at an exponential rate that is at least k_0 (from (iv)). In addition, there is a constant k_1 such that the maximum payoff difference between actions at any information set $h \in H_-$ is less than $k_1\alpha(t)$. The monotonicity and Lipschitz continuity of f then ensures that there is a constant k_2 such that the proportion of the subpopulation playing action k at any information set in $h \in H_-$ satisfies $|\dot{z}_{hk}(t)| < k_2\alpha(t)$. We then need only choose $d < \epsilon^*/4$ sufficiently small that $\int_0^\infty \dot{z}_{hk}(t)dt \leq \int_0^\infty k_2\alpha(t)dt \leq k_2 \int_0^\infty \delta e^{-k_0 t} dt = \delta k_2/k_0 < \epsilon^*/4 - \delta$. This ensures that as actions at information sets in H_+ converge to the equilibrium actions specified by w^* , the proportion playing an action at an information set in H_- cannot change by more than $\epsilon^*/4 - \delta$, and hence a proportion beginning within δ of w^* must end within $\epsilon^*/4$ of w^* .

Step 3. For each sufficiently large n , there exists $\gamma_n \in [0, \epsilon^*)$ such that for any state in $V(\epsilon^*, \epsilon^*) \setminus V(\gamma_n, \epsilon^*)$, the total of the subpopulation proportions playing each of the nonequilibrium actions at any information set in H_+ , under the perturbed dynamic, declines at a rate that is at least $k_0/2$. To verify this, fix an information set $h \in H_+$ and let k^* be the equilibrium action at h . Then we need to find a value $\gamma < \epsilon^*$ such that the final inequality in the following holds (where “ $k \neq k^*$ ” is a shorthand for “ $k \in A(h) \setminus \{k^*\}$ ”):

$$\begin{aligned} \frac{\sum_{k \neq k^*} \dot{z}_{hk}}{\sum_{k \neq k^*} z_{hk}} &= \frac{\sum_{k \neq k^*} (f_{hk}(z) + \eta_n(\Delta_n)(\theta_{hk} - z_{hk}))}{\sum_{k \neq k^*} z_{hk}} \\ &\leq -k_0 + \frac{\eta_n(\bar{\beta})(\sum_{k \neq k^*} \theta_{hk} - \gamma)}{\gamma} < -\frac{k_0}{2}, \end{aligned} \tag{5}$$

32. The max norm $\|x\|$ is defined by $\|x\| = \max_i \{|x_i|\}$, where $|\cdot|$ is absolute value.

where γ is the total probability attached to nonequilibrium actions at h . By letting n be large, and hence $\eta_n(\beta)$ be small, we can find a value $\gamma < \varepsilon^*$ for which (5) holds. Let γ_n be the smallest such γ given n . Then $\lim_{n \rightarrow \infty} \gamma_n = 0$ (because $\lim_{n \rightarrow \infty} \eta_n(\beta) = 0$).

Step 4. For sufficiently large n , any trajectory beginning in $V(\delta, \delta)$ enters $V(\gamma_n, \varepsilon^*)$. To verify this, notice that for all states in the set $V(\varepsilon^*, \varepsilon^*) \setminus V(\gamma_n, \varepsilon^*)$, the subpopulation proportion attached to nonequilibrium actions at each information set in H_+ declines at rate at least $k_0/2$. If there were no drift, the absolute value of the derivative in the subpopulation proportion playing any action at an information set in H_- would be at most k_2 (from Step 2). We can then choose n sufficiently large that the absolute value of each such derivative with drift is at most $2k_2$. A sufficient condition for trajectories beginning in $V(\delta, \delta)$ to enter $V(\gamma_n, \varepsilon^*)$ is then $2k_2 \int_0^\infty \delta e^{-k/2} dt = 4\delta k_2/k_0 < \varepsilon^* - \delta$, which follows from our earlier choice of δ to ensure $\delta k_2/k_0 < \varepsilon^*/4 - \delta$.

Step 5. For all sufficiently large n , the set $V(\gamma_n, \varepsilon^*)$ is forward invariant. In addition, let $\{\lambda_m\}_{m=0}^\infty$ be a sequence of nonnegative numbers with $\lim_{m \rightarrow \infty} \lambda_m = 0$. For each λ_m , there is a smallest value of n , denoted by $n(m)$, such that any trajectory beginning in $V(\gamma_{n(m)}, \varepsilon^*)$ enters and subsequently remains in $V(\gamma_{n(m)}, \theta_m)$. The proof is then completed by taking the latter to be our set U_n . To verify these claims, it suffices to examine actions at information sets in H_- , since by definition the proportion of each subpopulation playing the equilibrium action at each information set in H_+ is growing at each state in $V(\gamma_n, \varepsilon^*)$ in which this proportion is only $1 - \gamma_n$. To check information sets in H_- , we note that the change in the subpopulation proportion playing an action k at an information set $h \in H_-$ is given by:

$$\dot{z}_{hk} = f_{hk}(z) + \eta_n(\Delta_h)(\theta_{hk} - z_{hk}).$$

The result then follows from noting that $\eta_H(\Delta_h) \geq \eta_n(\beta)$, $f_{hk}(z) < k_2 \gamma_n$, and

$$\lim_{n \rightarrow \infty} \gamma_n / \eta_n(\beta) = 0. \tag{6}$$

In particular, this ensures that for sufficiently large n , \dot{z}_{hk} takes the sign of $\theta_{hk} - z_{hk}$ whenever the absolute value of the latter difference exceeds $\lambda_m > 0$, no matter how small we take λ_m , which yields the result. To verify (6), notice that $\eta(\tilde{\beta})/\eta(\beta)$ approaches zero as n grows (because $\tilde{\beta} > \beta$). It then suffices to show that $\gamma_n/\eta(\tilde{\beta})$ is bounded above as n grows. If not, then $\eta(\tilde{\beta})/\gamma_n$ approaches zero. But then γ_n is not the smallest value of γ for which (5) holds, which is a contradiction. ||

II. Foundations

This section presents the stochastic foundations of (2). Our point of departure is a model in which each of the n populations contains a finite number of players N . We consider a discrete-time model, with periods of length τ , so that players are matched to play the game at times $\{0, \tau, 2\tau, \dots\}$. The state space is a finite subset $Z_N \subset Z$, where Z_N contains only those proportions that can be created by assigning the N agents of each population to various pure strategies. We let $z_{h,k}$ be the proportion of the i th population playing action $k \in A(h_i)$ at information set h_i . We often abbreviate this to simply z_k .

We shall be interested in the *expected* state at time $t + \tau$ given state $z(t)$. Let $v = 1/N$. Suppose there exists a function on $h(z, v)$ on $Z \times \mathbb{R}_+$ such that

$$\mathcal{E}\{z(t + \tau) | z(t)\} = z(t) + \tau[h(z(t), v)] + O(\tau^2 N^2) \tag{7}$$

where

$$h(z, v) = h(z, 0) + O(v). \tag{8}$$

The first “ τ ” that appears on the right side of (7) captures the fact that we expect only small changes in the state to occur in small intervals of time. Binmore and Samuelson (1993) and Binmore, Samuelson and Vaughan (1995) provide examples in which (7) is derived from explicit models of how agents change their strategies. The key feature of these models is that in each period of length τ , each agent takes an independent draw that causes them to retain their current strategy with probability $1 - \tau$ and consider changing to a new strategy with probability τ .³³ The “ $O(v)$ ” in (8) allows us to think of $h(z, 0)$ as the change that would be expected in an infinite

33. It suffices that the probability of changing strategies is proportional to τ . The probability that exactly one agent considers changing strategies is $\tau N + O(\tau^2 N^2)$ while the probability that more than one agent considers changing is $O(\tau^2 N^2)$. Given that one agent changes strategies, the result is to move z by at most $1/N$ in a direction described by $h(z, v)$, giving an expected movement that is described by (7).

population, with the expected change in a finite population of size N , given by $h(z, v)$, departing from this as a result of finite-sampling effects by an amount that decreases with the population size.

We now consider the differential equation

$$\dot{y} = h(z, 0). \tag{9}$$

We can rearrange (7) to give

$$\frac{\mathcal{E}\{z(t + \tau) | z(t)\} - z(t)}{\tau} = h(z(t), v) + O(\tau N^2). \tag{10}$$

We now direct attention to the case in which the populations are large and the interval of time τ is short. Letting the population grow allows us to smooth some of the randomness in the system, while letting τ shrink gives us a model in which agents' strategy revisions occur at idiosyncratic, uncoordinated times. As we take the limits $\tau \rightarrow 0$ and $N \rightarrow \infty$, with $\tau N^2 \rightarrow 0$, the left side of (10) appears to become the derivative dz/dt , causing (10) to converge to (9). The only apparent sticking point is the removal of the expectation on the left side of (10). This removal is often justified with an informal law-of-large-numbers argument.³⁴

When can the link between (10) and (9) be established formally? Binmore, Samuelson and Vaughan [(1995), Theorem 1], using techniques introduced by Börgers and Sarin (1995) and Boylan (1997, 1995), establish such a link for a model with one population and a one-dimensional state space. The following proposition extends this result to the current model. We assume that there exist a finite S such that, for any information set i and action $k \in A(h_i)$,

$$\mathcal{E}\{(z_{ik}(t + \tau) - z_{ik}(t))^2 | z(t)\} \leq \tau v S + O(\tau^2 N^2). \tag{11}$$

This condition is satisfied in the models developed in Binmore and Samuelson (1993) and Binmore *et al.* (1995). To see why we expect it to hold, notice that if h is continuous, then (10) ensures that there is S' with

$$\mathcal{E}\{z(t + \tau) - z(t) | z(t)\} \leq \tau S' + O(\tau^2 N^2).$$

But the key difference between $z_{i+\tau} - z_i$ and $(z_{i+\tau} - z_i)^2$ is that an agent who switches strategies appears as a term whose magnitude is $1/N$ in the former and $1/N^2$ in the latter, and it is this extra “ N ” in the denominator that accounts for our condition (11) on $\mathcal{E}\{(z_{i+\tau} - z_i)^2 | z_i\}$.

Proposition 2. *Let (7) and (11) hold, and let $h(z, v)$ be Lipschitz continuous on Z (in the max norm on Z), with Lipschitz constant C holding for all sufficiently small v . Then for any time T and any ϵ , we can choose a sufficiently large N and sufficiently small τ (so that τN^2 is sufficiently small) that the realization of the underlying stochastic strategy adjustment model at any time $t \in [0, T]$ is within ϵ (in terms of the max norm) of the expected value given by (9) with probability at least $1 - \epsilon$.*

Proof. Step 1. Fix a value $T > 0$ and fix a rational initial condition $z(0)$. Unless otherwise stated, t will be assumed to be admissible, *i.e.* to be of the form $t = k\tau$ for some integer k . We assume that N is always such that $z(0) \in Z_N$. We let $y(t)$ be the solution to the differential equation $\dot{y} = h(y, 0)$ given initial condition $y(0) = z(0)$. Our task is to show that, for any admissible $t \leq T$,

$$\text{prob}\{\|y(t) - z(t)\| \geq \epsilon\} < \epsilon. \tag{12}$$

Let $Y(t)$ solve $\dot{Y} = h(Y, v)$. Then it suffices to show, for any admissible $t \leq T$, that

$$\|Y(t) - y(t)\| < \frac{1}{2}\epsilon \tag{13}$$

$$\text{prob}\{\|Y(t) - z(t)\| \geq \frac{1}{2}\epsilon\} < \epsilon. \tag{14}$$

Step 2. We first establish (13). Because $h(y, 0)$ is Lipschitz continuous on Z and $h(y, v) = h(y, 0) + O(v)$, there exists K such that, for any pair of states x and y ,

$$\|h(x, 0) - h(y, 0)\| \leq K\|x - y\|$$

$$\|h(y, 0) - h(y, v)\| \leq Kv.$$

34. See, for example, van Damme [(1984), ch. 9.4] and Hofbauer and Sigmund [(1988), ch. 16.1].

We then have:

$$\begin{aligned} \|y(t) - Y(t)\| &\leq \int_0^t \|h(y(s), 0) - h(Y(s), v)\| ds \\ &\leq \int_0^t (\|h(y(s), 0) - h(Y(s), 0)\| + \|h(Y(s), 0) - h(Y(s), v)\|) ds \\ &\leq K \left[\int_0^t \|y(s) - Y(s)\| ds + vt \right]. \end{aligned}$$

For any $u \leq t$, we then also have

$$\|y(u) - Y(u)\| \leq K \left[\int_0^u \|y(s) - Y(s)\| ds + vt \right],$$

where $\|y(u) - Y(u)\|$ is nonnegative and continuous. We can then invoke Gronwall's lemma (Hirsch and Smale [(1974), p. 169]) to conclude:

$$\|y(t) - Y(t)\| \leq Kvt e^{Kt} < \varepsilon,$$

where the final inequality holds if N is sufficiently large (and hence v small).

Step 3. We now turn to (14). Our first step is to notice that

$$Y(t) - z(0) = \int_0^t h(Y(s), v) ds. \quad (15)$$

The next step is to find a corresponding expression for $z(t) - z(0)$. Define:

$$m(k\tau) = z(k\tau) - z(0) - \sum_{j=1}^k \mathcal{E} \{ z(j\tau) - z(j\tau - \tau) | z(j\tau - \tau) \}. \quad (16)$$

On rearranging, we get, for $t = k\tau$,

$$\begin{aligned} z(t) - z(0) &= m(t) + \sum_{j=1}^k \mathcal{E} \{ z(j\tau) - z(j\tau - \tau) | z(j\tau - \tau) \} \\ &= m(t) + \sum_{j=1}^k h(z(j\tau - \tau), v) + kO(\tau^2 N^2) \\ &= m(t) + \int_0^t h(z([s/\tau]\tau), v) ds + O(\tau N^2), \end{aligned} \quad (17)$$

where $[x]$ denotes the integer part of x and $kO(\tau^2 N^2) = O(k\tau^2 N^2) = O(t\tau N^2) = O(\tau N^2)$. Now we subtract (17) from (15) to get (where $|\cdot|$ is a vector of absolute values):

$$|Y(t) - z(t)| \leq |m(t)| + \int_0^t |h(Y(s), v) - h(z([s/\tau]\tau), v)| ds + O(\tau N^2).$$

This implies that

$$\begin{aligned} \|Y(t) - z(t)\| &\leq \|m(t)\| + \int_0^t \|h(Y(s), v) - h(z([s/\tau]\tau), v)\| ds + O(\tau N^2) \\ &\leq \|m(t)\| + C \int_0^t \|Y(s) - z([s/\tau]\tau)\| ds + O(\tau N^2), \end{aligned}$$

where $h(z, v)$ is Lipschitz continuous on Z with Lipschitz C for all sufficiently small v . This in turn implies that, for all $u \leq t$, we have

$$\|Y(u) - z([u/\tau]\tau)\| \leq M(t) + C \int_0^u \|Y(s) - z([s/\tau]\tau)\| ds, \quad (18)$$

where

$$M(t) = \sup_{0 \leq s \leq t} \|m([s/\tau]\tau)\| + O(\tau N^2) + O(\tau),$$

where the final $O(\tau)$ is added so that (18) holds for all $u \leq t$ and not just admissible u (i.e. u of the form $k\tau$ for some k). Again noting that $\|Y(u) - z([u/\tau]\tau)\|$ is nonnegative and continuous at all but finitely many points, we

can apply Gronwall's lemma to (18) to obtain

$$\|Y(t) - z([t/\tau]\tau)\| \leq M(t)e^{Ct}$$

for all t with $0 \leq t \leq T$.³⁵ We use this to conclude that, for all admissible t ,

$$\text{prob} \{ \|Y(t) - z(t)\| \geq \frac{1}{2}\epsilon \} \leq \text{prob} \{ \|M(t)\| \geq \frac{1}{2}\epsilon e^{-Ct} \}.$$

It then suffices to show, for admissible t ,

$$\text{prob} \left\{ \max_{0 \leq s \leq t} \|m(s)\| \geq \frac{1}{4}\epsilon e^{-Ct} \right\} \leq \epsilon,$$

where τ and τN^2 are taken to be sufficiently small that the error terms in the definition of $M(t)$ are less than $\epsilon e^{-Ct}/4$. Hence, letting $m_i(t)$ be an element of $m(t)$, it suffices to show that, for large N and small τ and τN^2 (where $|Z|$ denotes the dimension of Z),

$$\text{prob} \left\{ \max_{0 \leq s \leq t} |m_i(s)| \geq \frac{1}{4}\epsilon e^{-Ct} \right\} \leq \frac{\epsilon}{|Z|}, \tag{19}$$

since the probability that $\|m_i(s)\|$ exceeds $\frac{1}{4}\epsilon e^{-Ct}$ in the max norm is less than the sum of the probabilities that $|m_i(s)|$ exceed $\frac{1}{4}\epsilon e^{-Ct}$ and there are $|Z|$ such probabilities.

Step 4. We now use inequality (19). From Kolmogorov's inequality, we have (hereafter deleting the "i" subscript)³⁶

$$\text{prob} \left\{ \max_{0 \leq s \leq t} |m(s)| \geq \frac{1}{4}\epsilon e^{-Ct} \right\} \leq \frac{16e^{2Ct}}{\epsilon^2} \text{var} \{m(t)\}. \tag{20}$$

Hence, since the expected value of $m(t)$ is zero, it suffices for (19) to show that, for all $t \leq T$,

$$\mathcal{E} \{m(t)\}^2 \leq \frac{\epsilon^3}{16|Z|} e^{-2Ct}. \tag{21}$$

To do this, we define (hereafter deleting the "i" subscript on m)

$$\Delta_j = m(j\tau) - m(j\tau - \tau).$$

Then, if $j > l$, we have

$$\mathcal{E} \{\Delta_j \Delta_l\} = \mathcal{E} \{\Delta_l \mathcal{E} \{\Delta_j | z(l\tau)\}\} = 0,$$

since $\mathcal{E} \{\Delta_j | z(l\tau)\} = \mathcal{E} \{m(j\tau) | z(l\tau)\} - \mathcal{E} \{m(j\tau - \tau) | z(l\tau)\} = m(l\tau) - m(l\tau) = 0$. Then we have

$$\mathcal{E} \{m(k\tau)\}^2 = \mathcal{E} \{\sum_{j=1}^k \Delta_j\}^2 = \mathcal{E} \{\sum_{j=1}^k \sum_{l=1}^k \Delta_j \Delta_l\} = \sum_{j=1}^k \mathcal{E} \{\Delta_j\}^2.$$

Hence, it suffices for (21) that

$$\sum_{j=1}^k \mathcal{E} \{\Delta_j\}^2 \leq \frac{\epsilon^3}{16|Z|} e^{-2Ck\tau}. \tag{22}$$

Since $\mathcal{E} \{\Delta_j\} = \mathcal{E} \{m(j\tau) - m(j\tau - \tau)\} = 0$, we have

$$\begin{aligned} \mathcal{E} \{\Delta_j\}^2 &= \text{var} \Delta_j = \mathcal{E} \{(z(j\tau) - z(j\tau - \tau))^2 | z(j\tau - \tau)\} \\ &\leq \tau vS + O(\tau^2 N^2). \end{aligned}$$

35. Gronwall's lemma is typically stated with the requirement that the integrand be continuous. However, the proof requires only that the derivative of the integral be integrable, which will hold if the integrand is continuous almost everywhere.

36. If the increments $m(j\tau) - m(j\tau - \tau)$ were independent, (20) would be a statement of a simple form of Kolmogorov's inequality for sums of independent random variables (Billingsley [(1986), p. 296]). In our case, these increments are not independent, since the distribution of $m(j\tau)$ may depend on the realized value of $m(j\tau - \tau)$. However, Kolmogorov's inequality requires only that $m(t)$ is a martingale (Chung [(1974), p. 331]).

Equation (22) then becomes

$$\begin{aligned}\sum_{j=1}^k \mathcal{E}\{\Delta_j\}^2 &\leq k(\tau vS + O(\tau^2 N^2)) \\ &\leq T vS + O(\tau N^2) \\ &\leq \frac{\varepsilon^3}{16|Z|} e^{-2cT},\end{aligned}$$

where the final inequality holds as long as N is large and τN^2 is small. \parallel

Acknowledgements. Financial support from National Science Foundation grants SES-9122176 and SBR-9320678, the Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 303 at the University of Bonn, and the British Economic and Social Research Council is gratefully acknowledged. We thank Drew Fudenberg, Tilman Börgers, Klaus Ritzberger, Karl Schlag, Hyun Song Shin, Jörgen Weibull, and three anonymous referees for helpful comments and discussions. We are grateful to the Department of Economics at the University of Bonn and the Institute for Advanced Studies at the Hebrew University of Jerusalem, where part of this work was done, for their hospitality.

REFERENCES

- ARAKI, K. and LOW, H. (1997), "Long Run Equilibria: Noisy Learning in a World of Local Interaction" (Mimeo, University College London).
- BALKENBORG, D. (1994), "An Experiment on Forward vs. Backward Induction" (SFB Discussion Paper B-268, University of Bonn).
- BANERJEE, A. and WEIBULL, J. W. (1993), "Evolutionary Selection with Discriminating Players" (Research Paper in Economics, University of Stockholm).
- BEN-PORATH, E. and DEKEL, E. (1992), "Signaling Future Actions and the Potential for Sacrifice" *Journal of Economic Theory*, **57**, 36–51.
- BHATIA, N. P. and SZEGÖ, G. P. (1970) *Stability Theory of Dynamical Systems* (New York: Springer-Verlag).
- BILLINGSLEY, P. (1986) *Probability and Measure* (New York: John Wiley and Sons).
- BINMORE, K., HSU, S., SWIERZBINSKI, J. and PROULX, C. (1993), "Focal Points and Bargaining", *International Journal of Game Theory*, **22**, 381–409.
- BINMORE, K. (1987–88), "Modelling Rational Players: I and II", *Economics and Philosophy*, **3** and **4**, 179–214 and 9–55.
- BINMORE, K., GALE, J. and SAMUELSON, L. (1995), "Learning to be Imperfect: The Ultimatum Game", *Games and Economic Behaviour*, **8**, 56–90.
- BINMORE, K. and SAMUELSON, L. (1993), "Muddling Through: Noisy Equilibrium Selection" (SSRI Working Paper 9410, University of Wisconsin).
- BINMORE, K. and SAMUELSON, L. (1994), "An Economist's Perspective on the Evolution of Norms", *Journal of Institutional and Theoretical Economics*, **150**, 45–63.
- BINMORE, K. and SAMUELSON, L. (1994), "Evolutionary Drift and Equilibrium Selection" (SSRI Working Paper 9529, University of Wisconsin).
- BINMORE, K., SAMUELSON, L. and VAUGHAN, R. (1995), "Musical Chairs: Modelling Noisy Evolution", *Games and Economic Behavior*, **11**, 1–35.
- BLUME, A., KIM, Y.-G. and SOBEL, J. (1993), "Evolutionary Stability in Games of Communication", *Games and Economic Behavior*, **5**, 547–575.
- BÖRGERS, T. and SARIN, R. (1995), "Learning Through Reinforcement and Replicator Dynamics", (Mimeo, University College London).
- BOYLAN, R. T. (1992), "Laws of Large Numbers for Dynamical Systems with Randomly Matched Individuals", *Journal of Economic Theory*, **57**, 473–504.
- BOYLAN, R. T. (1995), "Continuous Approximation of Dynamical Systems with Randomly Matched Individuals", *Journal of Economic Theory*, **66**, 615–625.
- CHUNG, K. (1974) *A Course in Probability Theory* (2nd Ed.) (Boston: Academic Press).
- COOPER, R., DEJONG, D. V., FORSYTHE, R. and ROSS, T. W. (1992), "Communication in Coordination Games", *Quarterly Journal of Economics*, **107**, 739–773.
- DAVIS, D. D. and HOLT, C. A. (1993) *Experimental Economics*. (Princeton: Princeton University Press).
- ELLISON, G. (1993), "Learning, Local Interaction, and Coordination", *Econometrica*, **61**, 1047–1072.
- FUDENBERG, D., KREPS, D. M. and LEVINE, D. K. (1988), "On the Robustness of Equilibrium Refinements", *Journal of Economic Theory*, **44**, 354–380.
- FUDENBERG, D. and LEVINE, D. K. (1993), "Self-Confirming Equilibrium", *Econometrica*, **61**, 523–546.
- FUDENBERG, D. and LEVINE, D. K. (1993), "Steady State Learning and Nash Equilibrium", *Econometrica*, **61**, 547–574.
- FUDENBERG, D. and LEVINE, D. K. (1995), "Theory of Learning in Games" (Mimeo, Harvard University and UCLA).

- HALE, J. (1969) *Ordinary Differential Equations* (New York: John Wiley and Sons).
- HIRSCH, M. W. and SMALE, S. (1974) *Differential Equations, Dynamical Systems, and Linear Algebra* (New York: Academic Press).
- HOFBAUER, J. and SIGMUND, K. (1988) *The Theory of Evolution and Dynamical Systems* (Cambridge: Cambridge University Press).
- KAGEL, J. H. and ROTH, A. E. (1995) *Handbook of Experimental Economics* (Princeton: Princeton University Press).
- KANDORI, M., MAILATH, G. J. and ROB, R. (1993), "Learning, Mutation, and Long Run Equilibria in Games", *Econometrica*, **61**, 29–56.
- KIM, Y.-G. and SOBEL, J. (1992), "An Evolutionary Approach to Pre-Play Communication", *Econometrica*, **63**, 1181–1194.
- KIMURA, M. (1983) *The Neutral Theory of Molecular Evolution* (Cambridge: Cambridge University Press).
- KOHLBERG, E. and MERTENS, J.-F. (1986), "On the Strategic Stability of Equilibria", *Econometrica*, **54**, 1003–1038.
- MAILATH, G., SAMUELSON, L. and SWINKELS, J. (1993), "Extensive Form Reasoning in Normal Games", *Econometrica*, **61**, 273–302.
- MATSUI, A. (1991), "Cheap-Talk and Cooperation in Society", *Journal of Economic Theory*, **54**, 245–258.
- McKELVEY, R. D. and PALFREY, T. R. (1992), "An Experimental Study of the Centipede Game", *Econometrica*, **60**, 803–836.
- MYERSON, R. B. (1978), "Proper Equilibria", *International Journal of Game Theory*, **7**, 73–80.
- RITZBERGER, K. (1993), "The Theory of Normal Form Games from the Differentiable Viewpoint", *International Journal of Game Theory*, **23**, 207–236.
- RITZBERGER, K. and WEIBULL, J. (1995), "Evolutionary Selection in Normal Form Games", *Econometrica*, **63**, 1371–1400.
- ROBLES, J. (1996), "Evolution and Lock in with Nonstationary Transition Probabilities", (Mimeo, University of Colorado).
- SAMUELSON, L. (1994), "Stochastic Stability in Games with Alternative Best Replies", *Journal of Economic Theory*, **64**, 35–64.
- SAMUELSON, L. and ZHANG, J. (1992), "Evolutionary Stability in Asymmetric Games", *Journal of Economic Theory*, **57**, 363–391.
- SÁNCHEZ, D. A. (1968) *Ordinary Differential Equations and Stability Theory: An Introduction* (San Francisco: W. H. Freeman and Company).
- SANDHOLM, W. H. and PAUZNER, A. (1997), "Evolution, Population Growth, and History Dependence" (Mimeo, Northwestern University and Tel Aviv University).
- SCHLAG, K. H. (1993), "Dynamic Stability in the Repeated Prisoners' Dilemma Played by Finite Automata" (SFB Discussion Paper B-243, University of Bonn).
- SELTEN, R. (1978), "The Chain-Store Paradox", *Theory and Decision*, **9**, 127–159.
- SOBEL, J. (1993), "Evolutionary Stability in Communication Games", *Economics Letters*, **42**, 301–312.
- SWINKELS, J. (1993), "Adjustment Dynamics and Rational Play in Games", *Games and Economic Behavior*, **5**, 455–484.
- VAN DAMME, E. (1984), "A Relation Between Perfect Equilibria in Extensive Form Games and Proper Equilibria in Normal Form Games", *International Journal of Game Theory*, **13**, 1–13.
- VAN DAMME, E. (1987), "Stable Equilibria and Forward Induction" (SFB Discussion Paper A-128, University of Bonn).
- VAN DAMME, E. (1989), "Stable Equilibria and Forward Induction", *Journal of Economic Theory*, **48**, 476–596.
- WÄRNERYD, K. (1990), "Cheap talk, Coordination, and Evolutionary Stability", *Games and Economic Behavior*, **5**, 532–546.
- WINTER, E. and ZAMIR, S. (1996), "Experimenting Ultimatum Bargaining with Virtual Players" (Mimeo, Hebrew University of Jerusalem).
- YOUNG, P. (1993), "The Evolution of Conventions", *Econometrica*, **61**, 57–84.
- YOUNG, P. (1997) *Microstrategy and Macrostructure* (Princeton: Princeton University Press).