Scand. J. of Economics 106(*), ∗–∗, 2004
DOI: 10.1111/j.1467-9442.2004.000376.x

Noisy Directional Learning and the Logit Equilibrium*

Simon P. Anderson

University of Virginia, Charlottesville, VA 22903-3328, USA
sa9w@virginia.edu

Jacob K. Goeree

California Institute of Technology, Pasadena, CA 91125, USA
jkg@hss.caltech.edu

Charles A. Holt

University of Virginia, Charlottesville, VA 22903-3328, USA
cah2k@virginia.edu

Abstract

We specify a dynamic model in which agents adjust their decisions toward higher payoffs, subject to normal error. This process generates a probability distribution of players' decisions that evolves over time according to the Fokker–Planck equation. The dynamic process is stable for all potential games, a class of payoff structures that includes several widely studied games. In equilibrium, the distributions that determine expected payoffs correspond to the distributions that arise from the logit function applied to those expected payoffs. This "logit equilibrium" forms a stochastic generalization of the Nash equilibrium and provides a possible explanation of anomalous laboratory data.

Keywords: Bounded rationality; noisy directional learning; Fokker–Planck equation; potential games; logit equilibrium

JEL classification: *C*62; *C*73

I. Introduction

Small errors and shocks may have offsetting effects in some economic contexts, in which case there is not much to be gained from an explicit analysis of stochastic elements. In other contexts, a small amount of randomness can

* We gratefully acknowledge financial support from the National Science Foundation (SBR-9818683 and SBR-0094800), the Alfred P. Sloan Foundation and the Dutch National Science Foundation (NWO-VICI 453.03.606). We wish to thank (without implicating) Kenneth Arrow, Harmen Bussemaker, Peter Hammond, Michihiro Kandori, Regis Renault, Robert Rosenthal, Thomas Sargent and Arthur Schram for helpful comments and suggestions on earlier drafts.

have a large effect on equilibrium behavior.[1] Regardless of whether "errors" or "trembles" are due to random preference shocks, experimentation or actual mistakes in judgment, their effect can be particularly important when players' payoffs are quite sensitive to others' decisions, e.g. when payoffs are discontinuous as in auctions, or highly interrelated as in coordination games. Nor do errors cancel out when the Nash equilibrium is near a boundary of the set of feasible actions and noise pushes actions towards the interior, as in a public goods contribution game where the Nash equilibrium is at zero contributions (full free riding). Errors are more likely when payoff differences across alternatives are small, so the consequences of mistakes are minor. For example, when managers are weakly motivated by profits to owners, they may not exert much effort to find the optimal action.

Stochastic elements have been incorporated successfully into a wide array of economic theories. These stochastic elements have been typically assumed to be driven by exogenous shocks.[2] Noise that is endogenous to the system can arise from decision errors, which are endogenously affected by the costs of making errors in that a more costly mistake is less likely to be made. Despite Simon's (1957) early work on modeling bounded rationality, the incorporation of noise in the analysis of economic games is relatively recent. Rosenthal (1989) and McKelvey and Palfrey (1995) propose noisy generalizations of the standard Nash equilibrium.[3] McKelvey and Palfrey's "quantal response equilibrium" allows a wide class of probabilistic choice rules to be substituted for perfect maximizing behavior in an equilibrium context. Other economists have introduced noise into models of learning and evolutionary adjustment; see for instance Foster and Young (1990), Fudenberg and Harris (1992), Kandori *et al*. (1993), Binmore, Samuelson and Vaughan (1995), and Chen, Friedman and Thisse (1997). In particular, Foster and Young (1990) and Fudenberg and Harris (1992) use a Brownian motion process, similar to the one specified in Section II.

Our goal in this paper is to provide a unified approach to equilibrium and evolutionary dynamics for a class of models with continuous decisions. The dynamic model is based on an assumption that decisions are changed locally in the direction of increasing payoff, subject to some randomness. Specifically, we propose a model of noisy adjustment to current conditions that, in equilibrium, yields a steady-state probability distribution of decisions for each player. Our modeling approach is inspired by two strands of thought,

[1] For example, in evolutionary models of coordination a small mutation rate may prevent the system from getting stuck in an equilibrium that is risk dominated; see e.g. Kandori, Mailath and Rob (1993) and Young (1993). Similarly, a small amount of noise or "trembles" can be used to rule out certain Nash equilibria; see Selten (1975).

[2] For instance, real business cycle models and much econometric work make this assumption.

[3] See Smith and Walker (1993) and Smith (1997) for an alternative approach.

directional adaptive behavior and randomness, both of which are grounded in early writings on bounded rationality.

Selten and Buchta's (1998) "learning direction theory" postulates that players are more likely to shift decisions in the direction of a best response to recent conditions. They show that such behavior was observed in an experimental trial of a first-price auction. However, Selten and Buchta (1998) do not expressly model the rate of adaption. One contribution of this paper is to operationalize learning direction theory by specifying an adjustment process. Our model is also linked to the literature on evolutionary game theory in which strategies with higher payoffs become more widely used. Such evolution can be driven by increased survival and fitness arguments with direct biological parallels, as in e.g. Foster and Young (1990), or by more cognitive models in which agents learn to use strategies that have worked better for themselves, as in e.g. Roth and Erev (1995) and Erev and Roth (1998), or in which they imitate successful strategies used by others, as in Vega-Redondo (1997) and Rhode and Stegeman (2001). An alternative to imitation and adaptation has been to assume that agents move in the direction of best responses to others' decisions. This is the approach we take.[4]

In addition to "survival of the fittest", biological evolution is driven by mutation of existing types, which is the second element that motivates our work. In the economics literature, evolutionary mutation is often specified as a fixed "epsilon" probability of switching to a new decision that is chosen randomly from the entire feasible set; see the discussion in Kandori (1997) and the references therein. Instead of mutation via new types entering a population, we allow existing individuals to make mistakes with the probability of a mistake being inversely related to its severity; see also Blume (1993, 1997), Young (1998) and Hofbauer and Sandholm (2002). The assumption of error-prone behavior can be justified by the apparent noisiness of decisions made in laboratory experiments with financially motivated subjects. To combine the two strands of thought, we analyze a model of noisy adjustment in the direction of higher payoffs. The payoff component is more important when the payoff gradient is steep, while the noise component is more important when the payoff gradient is relatively flat. The main intuition behind this approach can be illustrated in the simple case of only two decisions, 1 and 2, with associated payoffs π_1 and π_2, where the probability

[4] Models of imitation and reinforcement-learning are probably more likely to yield good predictions in noisy, complex situations where players do not have a clear understanding of how payoffs are determined, but rather can see clearly their own and others' payoffs and decisions. Best-response and more forward-looking behavior may be more likely in situations where the nature of the payoff functions is clearly understood. For example, in a Bertrand game in which the lowest-priced firm makes all sales, it is implausible that firms would be content merely to copy the most successful (low) price.

of switching from 1 to 2 is increasing in the payoff difference, $\pi_2 - \pi_1$. An approach based on payoff differences (or gradients in the continuous case) has the property that adding a constant to all payoffs has no effect, while scaling up the payoffs increases the speed at which decisions migrate toward high-payoff decisions. The belief that "noise" is reduced when payoffs are scaled up is supported by some laboratory experiments, e.g. the reduction in the rate of "irrational rejections" of low offers in ultimatum games as the amount of money being divided rises from \$20 to \$400; see List and Cherry (2000).

The next step in the analysis is to translate this noisy directional adjustment into an operational description of the dynamics of strategic choice. For this step, we use a classic result from theoretical physics, namely the Fokker–Planck equation that describes the evolution of a macroscopic system that is subject to microscopic fluctuations (e.g. the dispersion of heat in some medium). The state of the system in our model is a vector of the individual players' probability distributions over possible decisions. The Fokker–Planck equation shows how the details of the noisy directional adjustment rule determine the evolution of this vector of probability distributions. These equations thus describe behavioral adjustment in a stochastic game, in which the relative importance of stochastic elements is endogenously determined by payoff derivatives.

The prime interest in the dynamical system concerns its stability and steady state (a vector of players' decision distributions that does not change over time). The adjustment rule is particularly interesting in that it yields a steady state in which the distributions that determine expected payoffs are those that are generated by applying a logit probabilistic choice rule to these expected payoffs. Our approach derives this *logit equilibrium*, as in McKelvey and Palfrey (1995), from a completely different perspective than its usual roots. We prove stability of the adjustment rule for an important class of games, i.e. "potential games" for which the Nash equilibrium can be found by maximizing some function of all players' decisions. In particular, the Liapunov function that is maximized in the steady state of our model is the expected value of the potential function plus the standard measure of entropy in the system, weighted by an error parameter.

The dynamic model and its steady state are presented in Section II. Section III contains an analysis of global stability for an interesting class of games, i.e., potential games, which include public goods, oligopoly and two-person matrix games. Section IV concludes.

II. Evolution and Equilibrium with Stochastic Errors

We specify a stochastic model in continuous time to describe the interaction of a finite number of players. In our model, players tend to move towards decisions with higher expected payoffs, but such movements are subject to random shocks. At any point in time, the state of the system is characterized by probability distributions of players' decisions. The steady-state equilibrium is a fixed point at which the distributions that determine expected payoffs have converged to distributions of decisions that are based on those expected payoffs. This stochastic equilibrium reflects bounded rationality in that the optimal decision is not always selected, although decisions with higher payoffs are more likely to be chosen. The degree of imperfection in rationality is parameterized in a manner that yields the standard Nash equilibrium as a limiting case. The specific evolutionary process we consider shows an intuitive relationship between the nature of the adjustment and the probabilistic choice structure used in the equilibrium. In particular, with adjustments that are proportional to marginal payoffs plus normal noise, the steady state has a logit structure.

There are $n \geq 2$ players that make decisions in continuous time. At time t, player $i = 1, \ldots, n$ selects an action $x_i(t) \in (x_L, x_H)$. Since actions will be subject to random shocks, behavior will be characterized by probability distributions. Let $F_i(x, t)$ be the probability that player i chooses an action less than or equal to x at time t. Similarly, let the vector of the $n-1$ other players' decisions and probability distributions be denoted by $x_{-i}(t)$ and $F_{-i}(x_{-i}, t)$, respectively. The instantaneous expected payoff for player i at time t depends on the action taken and on the distributions of others' decisions:

$$\pi^e{}_i(x_i(t), t) = \int \pi_i(x_i(t), x_{-i})\, dF_{-i}(x_{-i}, t), \quad i = 1, \ldots, n. \tag{1}$$

We assume that payoffs, and hence expected payoffs, are bounded from above. In addition, we assume that expected payoffs are differentiable in $x_i(t)$ when the distribution functions are. The latter condition is ensured when the payoffs $\pi_i(x_i, x_{-i})$ are continuous.[5]

In a standard evolutionary model with replicator dynamics, the assumption is that strategies that do better than the population average against the distribution of decisions become more frequent in the population. The idea behind such a "population game" is that the usefulness of a strategy is evaluated in terms of how it performs against a distribution of strategies in the population of other players. We use the population game paradigm in a similar manner by assuming that the attractiveness of a pure strategy is based on its expected payoff given the distribution of others' decisions in the

[5] Continuity of the payoffs is sufficient but not necessary. For instance, in a first-price auction with prize value V, payoffs are discontinuous, but expected payoffs, $(V - x_i)\pi_{j \neq i} F_j(x_i)$, are twice differentiable when the F_j are twice differentiable. More generally, the expected payoff function will be twice differentiable even when the payoffs $\pi_i(x_i, x_{-i})$ are only piece-wise continuous.

population. To capture the idea of local adjustment to better outcomes, we assume that players move in the *direction* of increasing expected payoff, with the rate at which players change increasing in the marginal benefit of making that change.[6] This marginal benefit is denoted by $\pi_i^{e\prime}(x_i(t), t)$, where the prime denotes the partial derivative with respect to $x_i(t)$. However, individuals may make mistakes in the calculation of expected payoff, or they may be influenced by non-payoff factors. Therefore, we assume that the directional adjustments are subject to error, which we model as an additive disturbance, $w_i(t)$, weighted by a variance parameter σ_i:[7]

$$dx_i(t) = \pi_i^{e\prime}(x_i(t), t)dt + \sigma_i \, dw_i(t), \quad i = 1, \ldots, n. \tag{2}$$

Here $w_i(t)$ is a standard Wiener (or white noise) process that is assumed to be independent across players and time. Essentially, dx_i/dt equals the slope of the individual's expected payoff function plus a normal error with zero mean and unit variance.

The deterministic part of the local adjustment rule (2) indicates a "weak" form of feedback in the sense that players react to the distributions of others' actions (through the expected payoff function), rather than to the actions themselves. This formulation is motivated by laboratory experiments that use a random matching protocol. Random matching causes players' observations of others' actions to keep changing even when behavior has stabilized. When players gain experience they will take this random matching effect into account and react to the "average observed decision" or the distribution of decisions rather than to the decision of their latest opponent.

The stochastic part of the local adjustment rule in (2) captures the idea that such adaptation is imperfect and that decisions are subject to error. It is motivated by observed noise in laboratory data where adjustments are often unpredictable, and subjects sometimes experiment with alternative decisions. In particular, "errors" or "trembles" may occur because current conditions are not known precisely, expected payoffs are only estimated, or decisions are affected by factors beyond the scope of current expected payoffs, e.g. emotions like curiosity, boredom, inertia or desire to change. The random shocks in (2) capture the idea that players may use heuristics or "rules of thumb" to respond to current payoff conditions. We assume that these responses are, on average, proportional to the correct expected payoff gradients, but that calculation errors, extraneous factors and imperfect

[6] Friedman and Yellin (1997) show that when adjustment costs are quadratic in the speed of adjustment, it is optimal for players to alter their actions partially and in proportion to the gradient of expected payoff.
[7] See Basov (2001) for a multi-dimensional generalization of (2) and a careful discussion of the boundary conditions needed to ensure that no probability mass escapes (x_L, x_H).

information require that a stochastic term be appended to the deterministic part of (2). Taken together, the two terms in (2) simply imply that a change in the direction of increasing expected payoff is more likely, and that the magnitude of the change is positively correlated with the expected payoff gradient.

The adjustment rule (2) translates into a differential equation for the distribution function of decisions, $F_i(x, t)$. This equation will depend on the density $f_i(x, t)$ corresponding to $F_i(x, t)$, and on the slope, $\pi_i^{e\prime}(x, t)$, of the expected payoff function. It is a well-known result from theoretical physics that the stochastic adjustment rule (2) yields the Fokker–Planck equation for the distribution function.[8]

Proposition 1. *The noisy directional adjustment process (2) yields the Fokker–Planck equation for the evolution of the distributions of decisions*:

$$\frac{\partial F_i(x,t)}{\partial t} = -\pi_i^{e\prime}(x,t) f_i(x,t) + \nu_i f_i'(x,t), \quad i = 1, \ldots, n, \tag{3}$$

where $\nu_i = \sigma_i^2/2$.

Binmore *et al*. (1995) use the Fokker–Planck equation to model the evolution of choice probabilities in 2×2 matrix games. Instead of using the expected-payoff derivative as we do in (2), they use a non-linear genetic-drift function. Friedman and Yellin (1997) consider a one-population model in which all players get the same payoff from a given vector of actions, which they call "games of common interest". (This is a subset of the class of potential games discussed in the next section.) They start out with the assumption that the distribution evolves according to (3), but without the error term (i.e., $v_i = 0$). This deterministic version of Fokker–Planck is used to show that behavior converges to a (local) Nash equilibrium in such games.

A derivation of the Fokker–Planck equation is shown in the Appendix. Existence of a (twice differentiable) solution to the Fokker–Planck equation is demonstrated in most textbooks on stochastic processes; see e.g. Smoller (1994) and Gihman and Skohorod (1972). Notice that there is a separate equation for each player $i = 1, \ldots, n$, and that the individual Fokker–Planck equations are interdependent only through the expected payoff functions. In contrast, replacing the expected payoff in (2) by the instantaneous payoff, $\pi(x_1(t), \ldots, x_n(t))$, results in a *single* Fokker–Planck equation

[8] This result has been derived independently by a number of physicists, including Einstein (1905), and the mathematician Kolmogorov (1931). The first term on the RHS of (3) is known as a drift term, and the second term is a diffusion term. The standard example of pure diffusion without drift is a small particle in a suspension of water; in the absence of external forces the particle's motion is completely determined by random collisions with water molecules (Brownian motion). A drift term is introduced, for instance, when the particle is charged and influenced by an electric field.

that describes the evolution of the joint density of $x_1(t), \ldots, x_n(t)$. This formulation might be relevant when the same players repeatedly interact as in experiments with a *fixed*-matching protocol. Most experiments, however, employ a random-matching protocol in which case the population-game approach discussed here is more natural.

The Fokker–Planck equation (3) has a very intuitive economic interpretation. First, players' decisions tend to move in the direction of greater payoff, and a larger payoff derivative induces faster movement. In particular, when payoff is increasing at some point x, lower decisions become less likely, decreasing $F_i(x, t)$. The rate at which probability mass crosses over at x depends on the density at x, which explains the $-\pi_i^{e\prime}(x,t)f_i(x,t)$ term on the RHS of (3). The second term, $\nu_i f_i'$, reflects aggregate noise in the system (due to intrinsic errors in decision-making), which causes the density to "flatten out". Locally, if the density has a positive slope at x, then flattening moves mass toward lower values of x, increasing $F_i(x, t)$, and vice versa, as indicated by the second term on the RHS of equation (3).

Since $\nu_i = \sigma_i^2/2$, the variance coefficient ν_i in (3) determines the importance of errors relative to payoff-seeking behavior for individual i. First consider the limiting case $\nu_i = 0$. If behavior in (3) converges, it must be the case that $\pi_i^{e\prime}(x)f_i(x) = 0$, which is the necessary condition for an interior Nash equilibrium: either the necessary condition for payoff maximization is satisfied at x, or else the density of decisions is zero at x. As ν_i goes to infinity in (3), the noise effect dominates and the Fokker–Planck equation tends to $\partial F_i/\partial t = \nu_i \partial^2 F_i/\partial x^2$, which is equivalent to the "heat equation" that describes how heat spreads out uniformly in some medium.[9] In this limit, the steady state of (3) is a uniform density with $f_i' = 0$.

In a steady state of the process in (3), the RHS is identically zero, which yields the equilibrium conditions:

$$f_i'(x) = \pi_i^{e\prime}(x) f_i(x)/\nu_i, \quad i = 1, \ldots, n, \tag{4}$$

where the t arguments have been dropped since these equations pertain to a steady state. These equations can be simplified by dividing both sides by $f_i(x)$ and integrating, to obtain:

$$f_i(x) = \frac{\exp(\pi_i^e(x)/\nu_i)}{\int_{x_L}^{x_H} \exp(\pi_i^e(s)/\nu_i)ds}, \quad i = 1, \ldots, n, \tag{5}$$

where the integral in the denominator is a constant, independent of x, which ensures that the density integrates to one.

[9] The heat equation $\partial f_i/\partial t = \nu_i \partial^2 f_i/\partial x^2$ follows by differentiating both sides with respect to x.

The formula in (5) is a continuous analogue to the logit probabilistic choice rule. Since the expected payoffs on the RHS depend on the distributions of the other players' actions (see (1)), the equations in (5) are not explicit solutions. Instead, these equations constitute *equilibrium conditions* for the steady-state distribution: the probability distributions that determine expected payoffs must match the choice distributions determined by the logit formula in (5). In the steady-state equilibrium these conditions are simultaneously satisfied. The steady-state equilibrium is a continuous version of the quantal response equilibrium proposed by McKelvey and Palfrey (1995).[10] Thus we generate a logit equilibrium as a steady-state from a more primitive formulation of noisy directional learning, instead of imposing the logit form as a model of decision error. To summarize:

Proposition 2. *When players adjust their actions in the direction of higher payoff, but are subject to normal error as in (2), then any steady state of the Fokker–Planck equation (3) constitutes a logit equilibrium as defined by (5).*

This derivation of the logit model is very different from the usual derivations. Luce (1959) uses an axiomatic approach to tie down the form of choice probabilities.[11] In econometrics, the logit model is typically derived from a "random-utility" approach.[12] Both of these derivations are static in

[10] Rosenthal (1989) proposed a similar equilibrium with endogenously determined distributions of decisions, although he used a linear probabilistic choice rule instead of the logit rule. McKelvey and Palfrey (1995) consider a more general class of probabilistic choice rules, which includes the logit formula as a special case. Our model with continuous decisions is similar to the approach taken in Lopez (1995).

[11] Luce (1959) postulated that decisions satisfy a "choice axiom", which implies that the ratio of the choice probabilities for two decisions is independent of the overall choice set containing those two choices (the independence of irrelevant alternatives property). In that case, he shows that there exist "scale values" u_i such that the probability of choosing decision i is $u_i/\Sigma_j u_j$. The logit model follows when $u_i = \exp(\pi_i/v)$.

[12] This footnote presents the random-utility derivation of the logit choice rule for a finite number of decisions. Suppose there are m decisions, with expected payoffs $u_1, \ldots, u_m$. A probabilistic discrete choice model stipulates that a person chooses decision k if: $u_k + \varepsilon_k > u_i + \varepsilon_i$, for all $i \neq k$, where the ε_i are random variables. The errors allow the possibility that the decision with the highest payoff will not be selected, and the probability of such a mistake depends on both the magnitude of the difference in the expected payoffs and on the "spread" in the error distribution. The logit model results from the assumption that the errors are i.i.d. and double-exponentially distributed. The probability of choosing decision k is then $\exp(u_k/v)/\Sigma_i \exp(u_i/v)$, where v is proportional to the standard deviation of the error distribution. There are two alternative interpretations of the ε_i errors: they can either represent mistakes in the calculation or perception of expected payoffs, or they can represent unobservable preference shocks. These two interpretations are formally equivalent, although one embodies bounded rationality and the other implies rational behavior with respect to unobserved preferences. See Anderson, de Palma and Thisse (1992, Ch. 2) for further discussion and other derivations of the logit model.

nature. Here the logit model results from the behavioral assumption of directional adjustment with normal error.

Some properties of the equilibrium distributions can be determined from the structure of (4) or (5), independent of the specific game being considered. Equation (5) specifies the choice density to be proportional to an exponential function of expected payoff, so that actions with higher payoffs are more likely to be chosen, and the local maxima and minima of the equilibrium density will correspond to local maxima and minima of the expected payoff function. The error parameter determines how sensitive the density is to variations in expected payoffs. As the error parameter goes to infinity, the slope of the density in (4) goes to zero, and so the density in (5) becomes uniform, i.e., totally random and unaffected by payoff considerations. Conversely, as the error parameter becomes small, the density in (5) will place more and more mass on decisions with high expected payoffs. In the literature on stochastic evolution, it is common to proceed directly to the limiting case as the amount of noise goes to zero.[13] This limit is not our primary interest, for two reasons. First, econometric analysis of data from laboratory experiments yields error parameter estimates that are significantly different from zero, which is the null hypothesis corresponding to a Nash equilibrium. Second, the limiting case of perfect rationality is generally a Nash equilibrium, and our theoretical analysis was originally motivated as an attempt to explain data patterns that are consistent with economic intuition but which are not predicted by a Nash equilibrium. As we have shown elsewhere, the (static) logit model (5) yields comparative static results that conform with both economic intuition and data patterns from laboratory experiments, but are not predicted by the standard Nash equilibrium; see Anderson, Goeree and Holt (1998a, b, 2001) and Capra, Goeree, Gomez and Holt (1999). The dynamic adjustment model presented here gives a theoretical justification for using the logit equilibrium to describe decisions when behavior has stabilized, e.g. in the final periods of laboratory experiments.

To summarize the main result of this section, the steady-state distributions of decisions that follow from the adjustment rule (2) satisfy the conditions that define a logit equilibrium. Therefore, when the dynamical system described by (3) is stable, the logit equilibrium results in the long run when players adjust their actions in the direction of higher payoff (directional learning), but are subject to error. In the next section, we use Liapunov function methods to prove stability and existence for the class of potential games.

[13] One exception is Binmore and Samuelson (1997), who consider an evolutionary model in which the mistakes made by agents (referred to as "muddlers") are not negligible. At the aggregate level, however, the effect of noise is washed out when considering the limit of an infinite population.

III. Stability Analysis

So far, we have shown that any steady state of the Fokker–Planck equation (3) is a logit equilibrium. We now consider the dynamics of the system (3) and characterize sufficient conditions for a steady state to be attained in the long run. Specifically, we use Liapunov methods to prove stability for a class of games that includes some widely studied special cases. A Liapunov function is non-decreasing over time and has a zero time derivative only when the system has reached an equilibrium steady state. The system is (locally) stable when such a function exists.[14]

Although our primary concern is the effect of endogenous noise, it is instructive to begin with the special case in which there is no decision error and all players use pure strategies. Then it is natural to search for a function of all players' decisions that will be maximized (at least locally) in a Nash equilibrium. In particular, consider a function, $V(x_1, \ldots, x_n)$, with the property $\partial V/\partial x_i = \partial \pi_i/\partial x_i$ for $i = 1, \ldots, n$. When such a function exists, Nash equilibria can be found by maximizing V. The $V(\cdot)$ function is called the potential function, and games for which such a function exists are known as potential games; see Monderer and Shapley (1996).[15]

The usefulness of the potential function is not just that it is (locally) maximized at a Nash equilibrium. It also provides a direct tool to prove equilibrium stability under the directional adjustment hypothesis in (2). Indeed, in the absence of noise, the potential function itself is a Liapunov function; see also Slade (1994). This can be expressed as:

$$\frac{dV}{dt} = \sum_{i=1}^{n} \frac{\partial V}{\partial x_i}\frac{dx_i}{dt} = \sum_{i=1}^{n} \frac{\partial \pi_i}{\partial x_i}\frac{dx_i}{dt} = \sum_{i=1}^{n} (\partial \pi_i / \partial x_i)^2 \geq 0, \qquad (6)$$

where the final equality follows from the directional adjustment rule (2) with no noise, i.e., $\sigma_i = 0$.[16] Thus the value of the potential function is strictly increasing over time unless all payoff derivatives are zero, which is a necessary condition for an interior Nash equilibrium. The condition that $dV/dt = 0$ need not generate a Nash equilibrium: the process might come to rest at a local maximum of the potential function that corresponds to a *local* Nash equilibrium from which large unilateral deviations may still be profitable.

[14] See e.g. Kogan and Soravia (2002) for an application of Liapunov function techniques to infinite dimensional dynamical systems.
[15] Rosenthal (1973) first used a potential function to prove the existence of a pure-strategy Nash equilibrium in congestion games.
[16] This type of deterministic gradient-based adjustment has a long history; see Arrow and Hurwicz (1960).

Our primary interest concerns noisy decisions, so we work with the expected value of the potential function. It follows from (1) that the partial derivatives of the expected value of the potential function correspond to the partial derivatives of the expected payoff functions:

$$\pi_i^{e\prime}(x_i, t) = \frac{\partial}{\partial x_i} \int V(x_i, x_{-i})\, dF_{-i}(x_{-i}, t), \quad i = 1, \ldots, n. \tag{7}$$

Again, the intuitive idea is to use something that is maximized at a logit equilibrium to construct a Liapunov function, i.e., a function whose time derivative is non-negative and only equal to zero at a steady state. When $v_i > 0$ for at least one player i, then the steady state is not generally a Nash equilibrium, and the potential function must be augmented to generate an appropriate Liapunov function. Look again at the Fokker–Planck equation (3); the first term on the RHS is zero at an interior maximum of expected payoff, and the $f_i'(x, t)$ term is zero for a uniform distribution. Therefore, we want to augment the Liapunov function with a term that is maximized by a uniform distribution. Consider the standard measure of noise in a stochastic system, entropy, which is defined as $-\sum_{i=1}^{n} \int f_i \log(f_i)$. It can be shown that this measure is maximized by a uniform distribution, and that entropy is reduced as the distribution becomes more concentrated. The Liapunov function we seek is constructed by adding entropy to the expected value of the potential function:

$$\begin{aligned} L = & \int_{x_L}^{x_H} \cdots \int_{x_L}^{x_H} V(x_1, \ldots, x_n) f_1(x_1, t) \ldots f_n(x_n, t)\, dx_1 \ldots dx_n \\ & - \sum_{i=1}^{n} \nu_i \int_{x_L}^{x_H} f_i(x_i, t) \log(f_i(x_i, t))\, dx_i. \end{aligned} \tag{8}$$

The v_i parameters determine the relative importance of the entropy terms in (8), which is not surprising given that v_i is proportional to the variance of the Wiener process in player i's directional adjustment rule (2). Since entropy is maximized by a uniform distribution (i.e., purely random decision-making), it follows that decision distributions that concentrate probability mass on higher-payoff actions will have lower entropy. Therefore, one interpretation of the role of the entropy term in (8) is that, if the v_i parameters are large, then entropy places a high "cost" of concentrating probability on high-payoff decisions.[17]

[17] The connection between entropy and the logit choice probabilities is well established in physics and economics. For example, Anderson *et al.* (1992) showed that logit demands are generated from a representative consumer with a utility function that has an entropic form.

We prove that the dynamical system described by (3) converges to a logit equilibrium, by showing that the Liapunov function (8) is non-decreasing over time.[18]

Proposition 3. *For the class of potential games, behavior converges to a logit equilibrium when players adjust their actions in the direction of higher payoff, subject to normal error as in (2).*

Proof: In the Appendix we show that the Liapunov function is non-decreasing over time; by taking the time derivative of the Liapunov function, partially integrating, and using the Fokker–Planck equation, we can express this time derivative in a form that is analogous to (6)

$$\frac{dL}{dt} = \sum_{i=1}^{n} \int_{x_L}^{x_H} \frac{(\partial F_i(x_i, t)/\partial t)^2}{f_i(x_i, t)} dx_i \geq 0. \tag{9}$$

The entropy term in (8) is maximized by the uniform densities $f_i(x, t) = 1/(x_H - x_L)$, $i = 1, \ldots, n$. It follows from this observation that the maximum entropy is given by $\log(x_H - x_L)\ \Sigma_i v_i$, which is finite. The expected value of the potential function is bounded from above since, by assumption, expected payoffs are. Therefore, the Liapunov function, which is the sum of expected potential and entropy, is bounded from above. Since L is non-decreasing over time for any potential game, we must have $dL/dt \to 0$ as $t \to \infty$, so $dF_i/dt \to 0$ in this limit. By (3) this yields the logit equilibrium conditions in (4). The solutions to these equilibrium conditions are the logit equilibria defined by (5). Q.E.D.

When there are multiple logit equilibria, the equilibrium attained under the dynamical process (3) is determined by the initial distributions $F_i(x, 0)$. In other words, the dynamical process (3) is not ergodic. This follows because, with multiple equilibria, the Liapunov function (8) has multiple local maxima and minima, and since the Liapunov function cannot decrease over time *any* of these extrema are necessarily rest points of the dynamical process.

We now show that (local) maxima of the Liapunov function correspond to (locally) stable logit equilibria; see Anderson *et al.* (2001) and Hofbauer and
1 Sandholm (2002, 2003).

[18] The notion of convergence used here is "weak convergence" or "convergence in distribution": the random variable $x(t)$ weakly converges to the random variable X if $\lim_{t\to\infty}\text{Prob}[x(t) \leq x] = \text{Prob}[X \leq x]$ for all x. Proposition 3 thus implies that the random variable $x_i(t)$ defined in (2) weakly converges to a random variable that is distributed according to a logit equilibrium distribution, for any starting point $x_i(0)$.

Proposition 4. *A logit equilibrium is locally (asymptotically) stable under the process (3) if and only if it corresponds to a strict local maximum of the Liapunov function in (8). When the logit equilibrium is unique, it is globally stable.*

Proof: We first show that strict local maxima of the Liapunov function are locally (asymptotically) stable logit equilibria. Let $F^*(x)$ denote a vector of distributions that constitutes a logit equilibrium which corresponds to a strict local maximum of the Liapunov function. Suppose that at F^* the Liapunov function attains the value L^*. Furthermore, let U be the set of distributions in the neighborhood of F^* for which $L \geq L^* - \varepsilon$, where $\varepsilon > 0$ is small. Since ε can be made arbitrarily small, we may assume that U contains no other stationary points of L. Note from (9) that L is non-decreasing over time, so no trajectory starting in U will ever leave it. Moreover, since F^* is the only stationary point of L in U, Proposition 3 implies that all trajectories starting in U necessarily converge to F^* in the limit $t \to \infty$, i.e., F^* is locally (asymptotically) stable. Hence, strict local maxima of L are locally stable logit equilibria.

Next, we prove that any locally (asymptotically) stable logit equilibrium, F^*, is a strict local maximum of L. Since F^* is locally (asymptotically) stable, there exists a local neighborhood U of F^* that is invariant under the process (3), and whose elements converge to F^*. The Liapunov function is strictly increasing along a trajectory starting from any distribution in U (other than F^* itself), so L necessarily attains a strict local maximum at F^*. Finally, when the logit equilibrium is unique, it corresponds to the unique stationary point of L. Proposition 3 holds for any initial distribution, so the logit equilibrium is globally stable. Q.E.D.

It follows from (9) that $dF_i/dt = 0$ when the Liapunov function is (locally) maximized, which, by (3) and (4), implies that a logit equilibrium is necessarily reached. Recall that, in the absence of noise, a local maximum of the Liapunov function does *not* necessarily correspond to a Nash equilibrium; the system may come to rest at a point where "large" unilateral deviations are still profitable; see Friedman and Yellin (1997). In contrast, with noise, local maxima of the Liapunov function always produce a logit equilibrium in which decisions with higher expected payoffs are more likely to be made. In fact, even (local) *minima* of the Liapunov function correspond to such equilibria, although they are unstable steady states of the dynamical system.

Propositions 3 and 4 do not preclude the existence of multiple locally stable equilibria. In such cases, the initial conditions determine which equilibrium will be selected. As shown in the proof of Proposition 4, if the initial distributions are "close" to those of a particular logit equilibrium, then that equilibrium will be attained under the dynamic process (3). The 2×2 example presented at the end of this section illustrates the possibility of multiple stable equilibria.

Since the existence of potential functions is crucial to the results of Proposition 3, we next discuss conditions under which such functions can be found. A necessary condition for the existence of a potential function is that $\partial^2\pi_i/\partial x_j\partial x_i = \partial^2\pi_j/\partial x_i\partial x_j$ for all i, j, since both sides are equal to $\partial^2 V/\partial x_i\partial x_j$. Hence, the existence of a potential function requires $\partial^2[\pi_i - \pi_j]/\partial x_j\partial x_i = 0$ for all i, j. Moreover, these "integrability" conditions are also sufficient to guarantee existence of a potential function. It is straightforward to show that payoffs satisfy the integrability conditions if and only if: $\pi_i(x_1, \ldots, x_n) = \pi_c(x_1, \ldots, x_n) + \theta_i(x_i) + \varphi_i(x_{-i})$ for $i = 1, \ldots, n$, where π_c is the same for players, hence it has no i subscript. To see that this class of payoffs solves the integrability condition, note that the common part, π_c, cancels when taking the difference of π_i and π_j, and the player-specific parts, θ_i and φ_i, vanish upon differentiation. If we define $V(x_1, \ldots, x_n) = \pi_c(x_1, \ldots, x_n) + \sum_{i=1}^{n} \theta_i(x_i)$, we can write the above payoffs π_i as the sum of two components: a common component and a component that only depends on others' decisions

$$\pi_i(x_1, \ldots, x_n) = V(x_1, \ldots, x_n) + \alpha_i(x_{-i}), i = 1, \ldots, n. \tag{10}$$

where we have defined $\alpha_i(x_{-i}) = \varphi_i(x_{-i}) - \sum_{j \neq i} \mathrm{h}_j(x_j)$. The common part, V, has no i subscript, and is the same function for all players, although it is not necessarily symmetric in the x_i. The individual part, $\alpha_i(x_{-i})$, may differ across players. The common part includes benefits or costs that are determined by one's own decision, e.g. effort costs. The $\alpha_i(x_{-i})$ term in (10) does not affect the Nash equilibrium since it is independent of one's own decision, e.g. others' effort costs or gifts received from others. It follows from this observation that the partial derivative of $V(x_1, \ldots, x_n)$ with respect to x_i is the same as the partial derivative of $\pi_i(x_1, \ldots, x_n)$ with respect to x_i for $i = 1, \ldots, n$, so $V(\cdot)$ is a potential function for this class of payoffs. Proposition 3 then implies that behavior converges to a logit equilibrium for this class of games.

The payoff structure in (10) covers a number of important games. For instance, consider a linear public goods game in which individuals are given an endowment, ω. If an amount x_i is contributed to a public good, the player earns $\omega - x_i$ for the part of the endowment that is kept. In addition, every player receives a constant (positive) fraction m of the total amount contributed to the public good. Therefore, the payoff to player i is: $\pi_i = \omega - x_i + mX$, where X is the sum of all contributions including those of player i. The potential for this game is: $V(x) = \omega + mX - \Sigma_i x_i$, and $\alpha_i(x_{-i}) = \Sigma_{j \neq i} x_j$. Another example is the minimum-effort coordination game, as in e.g. Bryant (1983), for which: $\pi_i = \min_{j=1\ldots N}\{x_j\} - cx_i$, where effort costs $c \in [0, 1]$. Here, $V(x) = \min_{j=1\ldots N}\{x_j\} - \Sigma_i cx_i$ (see also Section IV). In both of these applications the common part represents a symmetric production function, included once, minus the sum of all players' effort costs. In previous work on public goods and coordination games, we showed that the logit equilibrium is unique; see Anderson *et al.* (1998b, 2001). Therefore, the directional adjustment process studied here is globally stable for these games.

It is also straightforward to construct potential functions for many oligopoly models. Consider a Cournot oligopoly with n firms and linear demand, so that $\pi_i = (a - bX)\, x_i - c_i(x_i)$, where X is the sum of all outputs and $c_i(x_i)$ is firm i's cost function. Since the derivative of firm i's profit with respect to its own output is given by $\partial\pi_i/\partial x_i = a - bX - bx_i - c_i{}'$, the potential function is easily derived as: $V = aX - b/2X^2 - b/2\sum_i x_i^2 - \sum_i c_i(x_i)$. Some non-linear demand specifications can also be incorporated.

As a final example, consider the class of symmetric two-player matrix games with two decisions:

		Player 2 D_1	Player 2 D_2
Player 1	D_1	a, a	b, c
	D_2	c, b	d, d

Player i is characterized by a probability x_i of choosing decision D_1.[19] Thus the payoff to player i is linear in the probability x_i:

$$\pi_i(x_i, x_{-i}) = d + (a - b - c + d)\, x_i\, x_{-i} + (b - d)\, x_i + (c - d)\, x_{-i}, i = 1, 2. \tag{11}$$

It is straightforward to show that for this payoff structure the potential function is given by $V = (a - b - c + d)x_1x_2 + (b - d)(x_1 + x_2)$. The potential for asymmetric two-player games can be constructed along similar lines.[20] Hence, the choice probabilities converge to those of the logit equilibrium for the whole class of these commonly considered games.

[19] This formulation corresponds to the setting in some laboratory experiments when subjects are required to select probabilities rather than actions, with the experimenter performing the randomization according to the selected probabilities. This method is used when the focus is on the extent to which behavior conforms to a mixed-strategy Nash equilibrium. Ochs (1995) used this approach in a series of matching-pennies games. Ochs reports that choices are sensitive to players' own payoffs, contrary to the mixed-strategy Nash prediction, and he finds some empirical support for the logit equilibrium.

[20] In an asymmetric game, the letters representing payoffs in (11) would have i subscripts, $i = 1, 2$. Asymmetries in the constant or final two terms pose no problems for the construction of a potential function, so the only difficulty is to make the $(a_i - b_i - c_i + d_i)$ coefficient of the interaction terms match for the two players. This can be accomplished by a simple rescaling of all four payoffs for one of the players, which does not affect the stability proof of Proposition 3.

IV. Conclusion

Models of bounded rationality are appealing because the calculations required for optimal decision-making are often quite complex, especially when optimal decisions depend on what others are expected to do. This paper begins with an assumption that decisions are adjusted locally toward increasing payoffs. These adjustments are sensitive to stochastic disturbances. When the process settles down, systematic adjustments no longer occur, although behavior remains noisy. The result is an equilibrium probability distribution of decisions, with errors in the sense that optimal decisions are not always selected, although more profitable decisions are more likely to be chosen. The first contribution of this paper is to use a simple model of noisy directional adjustments to derive an equilibrium model of behavior with endogenous decision errors that corresponds to the stochastic generalization of Nash equilibrium proposed by Rosenthal (1989) and McKelvey and Palfrey (1995). The central technical step in the analysis is to show that directional adjustments subject to normal noise yield a Fokker–Planck equation, with a steady state that corresponds to a "logit equilibrium". This equilibrium is described by a logit probabilistic choice function coupled to a Nash-like consistency condition.

The second contribution of this paper is to prove stability of the logit equilibrium for all potential games. We use Liapunov methods to show that the dynamic system is stable for a class of interesting payoff functions, i.e., those for potential games. This class includes minimum-effort coordination games, linear/quadratic public goods and oligopoly games, and two-person 2×2 matrix games in which players select mixed strategies. The process model of directional changes adds plausibility to the equilibrium analysis, and an understanding of stability is useful in deciding which equilibria are more likely to be observed.

Models of bounded rationality are of interest because they can explain behavior of human decision-makers in complex, changing situations. The stochastic logit equilibrium provides an explanation of data patterns in laboratory experiments that are consistent with economic intuition but which are not explained by a Nash equilibrium analysis; see McKelvey and Palfrey (1995) and Anderson *et al*. (1998a, b, 2001). The presence of decision errors is important when the Nash equilibrium is near the boundary of the set of feasible decisions, so that errors are biased toward the interior. In addition, errors have non-symmetric effects when payoff functions are sensitive to noise in others' behavior. In the presence of noise, equilibrium behavior is not necessarily centered around the Nash prediction; errors that push one player's decision away from a Nash decision may make it safer for others to deviate. In some parameterizations of a "traveler's dilemma" game, for example, the Nash equilibrium is at the lower end of the feasible set, whereas behavior in laboratory experiments conforms more closely to a logit equilibrium with a unimodal density located at the upper end; see Capra *et al*. (1999).

The stochastic elements in our model are intended to capture a variety of factors, such as errors, trembles, experimentation and non-payoff factors such as emotions. In some contexts, behavior may be largely driven by a specific bias, like inequality aversion in bargaining situations, as in Fehr and Schmidt
2 (1999) and altruism in public goods games, as in Andreoni (1987). These factors can be used to specify more general payoff functions that incorporate social preferences. The analysis of this paper would still apply in the sense that the gradient of the enriched payoff function would determine the direction of adjustments, and the resulting steady state would correspond to a logit equilibrium that incorporates these other regarding preferences. In summary, adding an error term to a gradient adjustment rule yields a tractable model with a steady-state equilibrium that has appealing theoretical and empirical properties.

Appendix

Derivation of the Fokker–Planck Equation

Recall that the directional adjustments are stochastic: $dx(t) = \pi^{e\prime}(x(t), t)dt + \sigma\, dw(t)$ (see (2)), where we have dropped the player-specific subscripts for brevity. Note that the payoff derivative $\pi^{e\prime}$ depends on time through the decision x and through other players' distribution functions. After a small time change, Δt, the change in a player's decision can be expressed as:

$$\Delta x(t) \equiv x(t + \Delta t) - x(t) = \pi^{e\prime}(x, t)\Delta t + \sigma \Delta w(t) + o(\Delta t). \tag{A1}$$

where $\sigma\Delta w(t)$ is a normal random variable with mean zero and variance $\sigma^2\Delta t$, and $o(\Delta t)$ indicate terms that go to zero faster than Δt (i.e., K is of $o(\Delta t)$ when $K/\Delta t \to 0$ as $\Delta t \to 0$). A player's decision, therefore, is a random variable $x(t)$ that has a time-dependent density $f(x, t)$. Let $h(x)$ be an arbitrary twice-differentiable function that vanishes at the boundaries, as does its derivative. At time $t + \Delta t$, the expected value of $h(x)$ can be expressed directly as:

$$E\{h(x(t + \Delta t))\} = \int_{x_L}^{x_H} h(x)f(x, t + \Delta t)dx. \tag{A2}$$

The directional adjustment rule in (A1) can be used to obtain an alternative expression for the expected value of $h(x)$ at time $t + \Delta t$:

$$E\{h(x(t + \Delta t))\} = E\{h(x(t) + \Delta x(t))\} \approx E\{h(x(t)) + \pi^{e\prime}(x, t)\Delta t + \sigma\Delta w(t))\}. \tag{A3}$$

where we neglected terms of $o(\Delta t)$. The rest of the proof is based on a comparison of the expected values in (A2) and (A3). A Taylor expansion of (A3) will involve $h'(x)$ and $h''(x)$ terms, that can be partially integrated to convert them to expressions in $h(x)$. Since

$h(\cdot)$ is arbitrary, one can equate equivalent parts of the expected values in (A2) and (A3), which yields the Fokker–Planck equation in the limit as Δt goes to zero.

Let $g(y)$ be the density of $\sigma\Delta w(t)$, i.e., a normal density with mean zero and variance $\sigma^2\Delta t$. The expectation in (A3) can be written as an integral over the relevant densities:

$$E\{h(x(t+\Delta t))\} = \int_{-\infty}^{\infty}\int_{x_L}^{x_H} h(x+\pi^{e\prime}(x,t)\Delta t+\sigma y)f(x,t)g(y)dxdy. \qquad \text{(A4)}$$

A Taylor expansion of the RHS of (A4) yields:

$$\int_{-\infty}^{\infty}\int_{x_L}^{x_H} \{h(x)+h'(x)[\pi^{e\prime}(x,t)\Delta t+\sigma y]+\frac{1}{2}h''(x)[\pi^{e\prime}(x,t)\Delta t+\sigma y]^2+\cdots f(x,t)g(y)dxdy,$$

where the dots indicate terms of $o(\Delta t)$. Integration over y eliminates the terms that are linear in y, since it has mean zero. In addition, the expected value of y^2 is $\sigma^2\Delta t$, so the result of expanding and integrating the above expression is:

$$\int_{x_L}^{x_H} h(x)f(x,t)dx + \Delta t\int_{x_L}^{x_H} h'(x)\pi^{e\prime}(x,t)f(x,t)dx + \Delta t\frac{\sigma^2}{2}\int_{x_L}^{x_H} h''(x)f(x,t)dx + o(\Delta t).$$

The integrals containing the h' and h'' term can be integrated by parts to obtain integrals in $h(x)$:

$$\int_{x_L}^{x_H} h(x)f(x,t)dx - \Delta t\int_{x_L}^{x_H} h(x)(\pi^{e\prime}(x,t)f(x,t))'dx + \Delta t\frac{\sigma^2}{2}\int_{x_L}^{x_H} h(x)f''(x,t)dx, \qquad \text{(A5)}$$

where a prime indicates a partial derivative with respect to x, and we used the fact that h and its derivative vanish at the boundaries. Since (A5) is an approximation for (A2) when Δt is small, take their difference to obtain:

$$\int_{x_L}^{x_H} h(x)[f(x,t+\Delta t)-f(x,t)]dx = \Delta t\int_{x_L}^{x_H} h(x)\left[-(\pi^{e\prime}(x,t)f(x,t))'+(\sigma^2/2)f''(x,t)\right]dx. \qquad \text{(A6)}$$

The terms in square brackets on each side must be equal at all values of x, since the choice of the $h(x)$ function is arbitrary. Dividing both sides by Δt, taking the limit $\Delta t \to 0$ to obtain the time derivative of $f(x, t)$, and equating the terms in square brackets yields:

$$\frac{\partial f(x,t)}{\partial t} = -(\pi^{e\prime}(x,t)\, f(x,t))' + \frac{\sigma^2}{2}f''(x,t). \qquad \text{(A7)}$$

Since the primes indicate partial derivatives with respect to x, we can integrate both sides of (A7) with respect to x to obtain the Fokker–Planck equation in (3).

Derivation of Equation (9)

The Liapunov function in (8) depends on time only through the density functions, since the x's are variables of integration. Hence the time derivative is:

$$\begin{aligned}\frac{dL}{dt} &= \sum_{i=1}^{n}\int_{x_L}^{x_H}\cdots\int_{x_L}^{x_H} V(x_1,\ldots,x_n)\prod_{j\neq i} f_j(x_j,t)\,\frac{\partial f_i(x_i,t)}{\partial t}\,dx_1\ldots dx_n \\ &\quad -\sum_{i=1}^{n}\nu_i\int_{x_L}^{x_H}(1+\log(f_i(x_i,t)))\,\frac{\partial f_i(x_i,t)}{\partial t}\,dx_i\,. \end{aligned} \tag{A8}$$

The next step is to integrate each of the expressions in the sums in (A8) by parts. First note that $\partial f_i/\partial t = \partial^2 F_i/\partial t\partial x_i$ and that the anti-derivative of this expression is $\partial F_i/\partial t$. Moreover, the boundary terms that result from partial integration vanish because
3 $F_i(0,\ t) = 0$ and $F_i(,t) = 1$ for all t, i.e., $\partial F_i/\partial t = 0$ at both boundaries. It follows that partial integration of (A8) yields:

$$\begin{aligned}\frac{dL}{dt} &= -\sum_{i=1}^{n}\int_{x_L}^{x_H}\cdots\int_{x_L}^{x_H}\frac{\partial V(x_1,\ldots,x_n)}{\partial x_i}\prod_{j\neq i} f_j(x_j,t)\,\frac{\partial F_i(x_i,t)}{\partial t}\,dx_1\ldots dx_n \\ &\quad +\sum_{i=1}^{n}\nu_i\int_{x_L}^{x_H}\frac{f_i'(x_i,t)}{f_i(x_i,t)}\,\frac{\partial F_i(x_i,t)}{\partial t}\,dx_i. \end{aligned} \tag{A9}$$

Equation (7) can be used to replace $\int \partial V/\partial x_i\ dF_{-i}$ with $\pi_i^{e\prime}$, and then the integrals in (A9) can be combined as:

$$\begin{aligned}\frac{dL}{dt} &= \sum_{i=1}^{n}\int_{x_L}^{x_H}\left\{-\pi^{e\prime}(x_i,t)+\nu_i\frac{f_i'(x_i,t)}{f_i(x_i,t)}\right\}\frac{\partial F_i(x_i,t)}{\partial t}\,dx_i \\ &= \sum_{i=1}^{n}\int_{x_L}^{x_H}\frac{(\partial F_i(x_i,t)/\partial t)^2}{f_i(x_i,t)}\,dx_i, \end{aligned} \tag{A10}$$

where the final equation follows from (3). Note that the RHS of (A10) is strictly positive unless $\partial F_i/\partial t = 0$ for $i = 1,\ldots,n$, i.e., when the logit conditions in (4) are satisfied.

References

Anderson, S. P., de Palma, A. and Thisse, J.-F. (1992), *Discrete Choice Theory of Product Differentiation*, MIT Press, Cambridge, MA.

Anderson, S. P., Goeree, J. K. and Holt, C. A. (1998a), Rent Seeking with Bounded Rationality: An Analysis of the All-pay Auction, *Journal of Political Economy 106*, 828–853.

Anderson, S. P., Goeree, J. K. and Holt, C. A. (1998b), A Theoretical Analysis of Altruism and Decision Error in Public Goods Games, *Journal of Public Economics 70*, 297–323.

Anderson, S. P., Goeree, J. K. and Holt, C. A. (2001), Minimum-effort Coordination Games: Stochastic Potential and the Logit Equilibrium, *Games and Economic Behavior 34*, 177–199.

Andreoni, J. (1995), Cooperation in Public Goods: Kindness or Confusion? *American Economic Review 85*, 891–904.

Arrow, K. J. and Hurwicz, L. (1960), Stability of the Gradient Process in *n*-Person Games, *Journal of the Society of Industrial and Applied Mathematics 8*, 280–294.

Basov, S. (2001), A Noisy Model of Individual Behavior, Working Paper, University of Melbourne.

Binmore, K. and Samuelson, L. (1997), Muddling Through: Noisy Equilibrium Selection, *Journal of Economic Theory 74*, 235–265.

Binmore, K., Samuelson, L. and Vaughan, R. (1995), Musical Chairs: Modeling Noisy Evolution, *Games and Economic Behavior 11*, 1–35.

Blume, L. E. (1993), The Statistical Mechanics of Strategic Interaction, *Games and Economic Behavior 5*, 387–424.

Blume, L. E. (1997), Population Games, in W. B. Arthur, S. N. Durlauf and D. A. Lane (eds.), *The Economy as a Complex Evolving System II*, Addison-Wesley, Reading, MA.

Bryant, J. (1983), A Simple Rational Expectations Keynes-type Model, *Quarterly Journal of Economics 98*, 525–528.

Capra, C. M., Goeree, J. K., Gomez, R. and Holt, C. A. (1999), Anomalous Behavior in a Traveler's Dilemma? *American Economic Review 89* (3), 678–690.

Chen, H.-C., Friedman, J. W. and Thisse, J.-F. (1997), Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach, *Games and Economic Behavior 18*, 32–54.

4 Cyert, R. M. and March, J. G. (1963), *A Behavioral Theory of the Firm*, Prentice-Hall Inc, Englewood Cliffs, NJ.

Einstein, A. (1905), *Annalen der. Physik 17* (4), 549.

Erev, I. and Roth, A. E. (1998), Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria, *American Economic Review 88* (4), 848–881.

Fehr, E. and Schmidt, K. M. (1999), A Theory of Fairness, Competition, and Cooperation, *Quarterly Journal of Economics 114*, 769–816.

Foster, D. and Young, P. (1990), Stochastic Evolutionary Game Dynamics, *Theoretical Population Biology 38*, 219–232.

Friedman, D. and Yellin, J. (1997), Evolving Landscapes for Population Games, Working Paper, University of California, Santa Cruz.

Fudenberg, D. and Harris, C. (1992), Evolutionary Dynamics with Aggregate Shocks, *Journal of Economic Theory 57* (2), 420–441.

Gihman, I. I. and Skorohod, A. V. (1972), *Stochastic Differential Equations*, Springer-Verlag, Berlin.

5 Goeree, J. K. and Holt, C. A. (2004), An Experimental Study of Costly Coordination, forthcoming in *Games and Economic Behavior.*

Hofbauer, J. and Sandholm, W. (2002), On the Global Convergence of Stochastic Fictitous Play, *Econometrica 70*, 2265–2294.

Kandori, M. (1997), Evolutionary Game Theory in Economics, in D. Kreps and K. Wallis (eds.), *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress*, Vol. *1*, Cambridge University Press, Cambridge.

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

Kogan, M. and Soravia, P. (2002), Lyapunov Functions for Infinite-Dimensional Systems, *Journal of Functional Analysis 192*, 342–363.

Kolmogorov, A. N. (1931), *Mathematische Annalen 104*, 415.

List, J. A. and Cherry, T. L. (2000), Learning to Accept in Ultimatum Games: Evidence from an Experimental Design that Generates Low Offers, *Experimental Economics 3*, 11–29.

Lopez, G. (1995), Quantal Response Equilibria for Models of Price Competition, Ph.D. dissertation, University of Virginia.

Luce, R. D. (1959), *Individual Choice Behavior*, John Wiley, New York.

McKelvey, R. D. and Palfrey, T. R. (1995), Quantal Response Equilibria for Normal Form Games, *Games and Economic Behavior 10*, 6–38.

Monderer, D. and Shapley, L. S. (1996), Potential Games, *Games and Economic Behavior 14*, 124–143.

Ochs, J. (1995), Games with Unique Mixed Strategy Equilibria: An Experimental Study, *Games and Economic Behavior 10*, 202–217.

Rhode, P. and Stegeman, M. (2001), Evolution through Imitation: The Case of Duopoly, *International Journal of Industrial Organization 19*, 415–454.

Rosenthal, R. W. (1973), A Class of Games Possessing Pure-strategy Nash Equilibria, *International Journal of Game Theory 2*, 65–67.

Rosenthal, R. W. (1989), A Bounded Rationality Approach to the Study of Noncooperative Games, *International Journal of Game Theory 18*, 273–292.

Roth, A. E. and Erev, I. (1995), Learning in Extensive-form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term, *Games and Economic Behavior 8*, 164–212.

Selten, R. (1975), Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Form Games, *International Journal of Game Theory 4*, 25–55.

Selten, R. and Buchta, J. (1998), Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions, in I. E. D. Budescu, I. Erev and R. Zwick, (eds.), *Games and Human Behavior: Essays in Honor of Amnon Rapoport*, Lawrence Erlbaum Associates, Mahwah, NJ.

Simon, H. A. (1957), *Models of Man*, John Wiley, New York.

Slade, M. E. (1994), What Does an Oligopoly Maximize?, *Journal of Industrial Economics 58*, 45–61.

Smith, V. L. (1997), Monetary Rewards and Decision Cost in Experimental Economics: An Extension, Working Paper, University of Arizona.

Smith, V. L. and Walker, J. M. (1993), Monetary Rewards and Decision Cost in Experimental Economics, *Economic Inquiry 31*, 245–261.

Smoller, J. (1994), *Shock Waves and Reaction–Diffusion Equations*, Springer-Verlag, Berlin.

Vega-Redondo, F. (1997), The Evolution of Walrasian Behavior, *Econometrica 65*, 375–384.

Young, P. (1993), The Evolution of Conventions, *Econometrica 61*, 57–84.

Young, P. (1998), *Individual Strategy and Social Structure*, Princeton University Press, Princeton, NJ.

Author Query Form

Journal: The Scandinavian Journal of Economics

Article : 011

Dear Author,

During the copy-editing of your paper, the following queries arose. Please respond to these by marking up your proofs with the necessary changes/additions. Please write your answers on the query sheet if there is insufficient space on the page proofs. Please write clearly and follow the conventions shown on the attached corrections sheet. If returning the proof by fax do not write too close to the paper's edge. Please remember that illegible mark-ups may delay publication.
Many thanks for your assistance.

Query Refs.	Query	Remarks
1	Hofbauer and Sandholm 2003 not in the list. Please supply.	
2	Andreoni 1987 is 1995 in the list. Which is correct?	
3	The text falling between (A8) and (A9) contains "$F_i(, t) = 1$". What comes ahead of the comma, if anything?	
4	Are the following cited? Cyert and March 1963	
5	Goeree and Holt 2004 (if retained, update if possible) If not, remove from the list?	