Adaptive Learning in Financial Markets

Bryan R. Routledge
Carnegie Mellon University

We investigate adaptive or evolutionary learning in a repeated version of the Grossman and Stiglitz (1980) model. We demonstrate that any process that is a monotonic selection dynamic will converge to the rational expectations asset demands if the proportion of informed traders is fixed. We also show that these learning processes have a unique asymptotically stable fixed point at the Grossman–Stiglitz (GS) equilibrium. The robustness of learning to noisy experimentation is studied using Binmore and Samuelson's (1999) deterministic drift approximation. Conditions on economic and learning process parameters for adaptive learning to lead to the GS rational expectations equilibrium are presented.

An important role of financial markets is to aggregate and transmit information. Radner (1979) and others [see Jordan and Radner (1982)] address the information processing role by characterizing rational expectations equilibria where individuals use the information contained in the market clearing price. These models, however, do not address how individuals acquire sufficient knowledge about economic structure, parameters, others' objectives, and others' beliefs to make appropriate inferences from a market-clearing price. In addition, rational expectation models typically assume an extremely high degree of individual rationality. It is important, therefore, to explore the robustness of rational expectations models to assumptions about knowledge and rationality. In this article we explore how individuals can discover and use an endogenous relationship through adaptive or evolutionary learning. We relax knowledge and rationality assumptions by modelling behavior as resulting from a process of imitation and experimentation rather than explicit optimization.

We investigate a repeated economy version of the Grossman and Stiglitz (1980) model. The Grossman–Stiglitz (GS) model provides a good framework for considering adaptive learning since it is the standard model of endogenous information acquisition and is the basis for several other learning models like those of Bray (1982) and Bray and Kreps (1987). In the GS model, traders in a one-period economy can choose to acquire a costly signal

For helpful comments I thank Franklin Allen, Jonathan Berk, Murray Carlson, Richard Green, Alan Kraus, Bart Lipman, Christine Parlour, Uday Rajan, Chester Spatt, and participants of the 1996 Western Finance Association Conference and the Fields Institute Conference on "New Directions in the Theory of Markets and Games" as well as three anonymous referees. Remaining errors are, of course, mine. I am grateful for the financial support of the Social Science and Humanities Research Council of Canada which funded a portion of this research. Address correspondence to Bryan Routledge, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA, 15213-3890, or email: rout@andrew.cmu.edu.

The Review of Financial Studies Winter 1999 Vol. 12, No. 5, pp. 1165–1202

of the risky asset's terminal dividend. Those traders choosing not to acquire the signal make an inference about the signal from the market-clearing asset price. In our model, both the information choice and the inference are determined by adaptive learning rather than optimization. Behavioral rules of thumb are evolved by copying other traders and from experimentation. We demonstrate conditions under which adaptive learning leads to behavior that is similar to that observed in the GS rational expectations equilibrium.

Learning about equilibrium relationships is complicated because it involves learning about other agents in the economy. Since other agents are also learning, agents are learning about others' learning and so on. Even though exogenous variables are independent and identically distributed, learning causes the endogenous price-signal relation to be nonstationary. There are two approaches to modeling learning in this environment. The first is Bayesian learning where beliefs about the endogenous parameters are updated with each realization of the economy via Bayes' rule. These models are quite complex since an individual's Bayesian updating needs to be consistent with the learning of all others. In the second approach, people use past realizations to estimate the endogenous price-signal relationship using some econometric technique. While the econometric technique is reasonable, it is misspecified due to the nonstationarity induced by the learning. For example, in least-squares learning, individuals treat the endogenous relationship as if it were exogenous (and fixed) and estimate it using OLS.[1] Adaptive learning most closely resembles the second approach since it is non-Bayesian and involves copying past successful behavior which, due to the nonstationarity, may not be successful in the current period.

There are several characteristics of adaptive learning that are appealing. First, adaptive learning does not require assumptions of unbounded rationality or that individuals possess a complete and correct knowledge of the economy. In fact, strong rationality assumptions are often motivated by appeals to adaptation or evolution [e.g., Friedman (1953) or Lucas (1986)]. By comparing adaptive learning to the GS rational expectations equilibrium, we explore the conditions under which such a motivation is appropriate. In addition, by studying alternatives to complete rationality we can explore the link between individual behavior and the aggregate behavior of markets. This is in contrast to "behavioral finance" research that explains market

[1] Examples of Bayesian learning are in Blume, Bray and Easley (1982) (an overview), Blume and Easley (1993) and Vives (1993, 1995). Bray and Kreps (1987) models Bayesian learning in a GS-type model. In their model, traders do not know a parameter of the economy (which is known in the GS model). They demonstrate that Bayesian learning does convere to the GS rational expectations equilibrium. However, as the authors point out, this result is not robust to slight perturbations of the model. In general, Bray and Kreps can only conclude that beliefs updated using Bayes' rule converge and put positive weight on the true parameters. Examples of least-squares learning include Marcet and Sargent (1988, 1989). Evans and Honkapohja (1998) consider more general econometric techniques. Bray (1982) considers least-squares learning in a GS-type model. She determines the parametric restrictions for learning to converge to the GS rational expectations equilibrium.

phenomenon in terms of the cognitive psychology of individuals without a formal aggregation link.[2] Second, adaptive learning facilitates exploring how an economy reaches equilibrium. As Bray and Kreps (1987) point out, Bayesian learning can only address learning within an equilibrium and not learning about an equilibrium. Third, adaptive learning also allows us to easily model the information acquisition decision of traders in the GS model. This would be difficult in either a least-squares or Bayesian learning context. Finally, using techniques developed by Binmore and Samuelson (1999), we can tractably model adaptive learning that incorporates both imitation and experimentation in the GS economy.

Adaptive learning demands less of traders' rationality. However, in order to abandon unsuccessful behavior and imitate successful behavior, traders need to have some measure of success (or fitness) of their behavior, as well as have some information about others' behavior. Consider the following example of adaptive learning. Traders choose an initial behavior at random. In this article, the behavior determines the trader's information about the dividend and parameters for her risky asset demand. Their collective behavior determines market-clearing prices and, in turn, individual utility or fitness. Adaptive learning describes how traders update their behavior. Assume that individuals update their behavior periodically with lower fitness individuals updating more frequently. People may be very unsophisticated in their updating behavior. For example, if a trader is unsatisfied with her behavior, she can simply copy the behavior of the first individual she meets. In particular, she need not seek out high fitness individuals. This simple adaptive learning process will produce the GS rational expectations asset demands if the proportion of informed traders is held constant and when the traders can also choose their information, the process has a unique asymptotically stable fixed point at the GS equilibrium. The key assumption in this example is that lower fitness behavior is re-evaluated more frequently. This is sufficient to ensure that the adaptive learning process is monotonic. Monotonicity in the learning process is that growth rates in the proportions of agents using the various possible behaviors are ordered by fitness. This embodies the intuition that successful behavior is copied and/or unsuccessful behavior is abandoned. The results we develop in the first portion of the article are for any monotonic learning process.

In the previous example, it is important that traders can observe their own fitness and the behavior of others. Financial markets provide an interesting setting to investigate adaptive learning. Since feedback on returns or

[2] See Shiller (1997) for a survey of this large area of research. Examples of using individual behavior to explain market phenomena include using prospect theory to understand the equity premium puzzle [Benartzi and Thaler (1995)] and overweighting recent evidence (the extrapolating fallacy or the recency effect) to explain the apparent success of contrarian investment strategies [Lakonishok, Shleifer, and Vishny (1994)].

wealth is typically quite frequent, traders typically have a good measure of their own fitness. Information on other traders' fitness is also available. For example, the information about the fitness of others can be inferred from information on the distribution of aggregate wealth or, in contexts like mutual fund management, from ranking services. Note that in the example, information on others' fitness is not crucial. The difficult assumption for adaptive learning is that traders can observe other's strategy. In some financial contexts like floor trading, behavior is public and observable. However, since most financial trade is typically occurs through a large and anonymous market, directly observing behavior may be difficult. In this situation, one might assume that information about the strategy of others is obtained through hiring senior personnel, mergers, or trade publications.

Given one is willing to assume that traders can observe their fitness and the behavior of others, the specific details of how the imitation occurs are not important. The results developed in the first part of the article are quite robust: adaptive learning that is a monotone selection dynamic produces to the GS equilibrium. The second portion of the article investigates the robustness of imitation-driven monotonic adaptive learning processes to noise from random experimentation. Here, we find that adaptive learning is less robust. In particular, we model imitation and experimentation using the techniques of Binmore and Samuelson (1999). They use deterministic drift as a proxy for small amounts of random experimentation. We demonstrate that for any given economy, there are learning processes with drift that yield behavior almost identical to the GS rational expectations equilibrium. More interestingly, we show that for any learning process with drift, there exist economies such that the limiting behavior of the adaptive learning process is distinct from that of the GS equilibrium. In these situations, all traders become informed. The development of these results illustrates the important economic and learning process parameters. In particular, for the existence of the GS rational expectations equilibrium, there must be some noise in the risky asset supply. In order for adaptive learning with drift to lead to the GS rational expectations equilibrium, it is necessary that the noise in the risky asset supply be large relative to the level of experimentation in the learning process.

The next section describes the repeated GS model. It describes a single period of the repeated economy and traders' rules of thumbs. Section 2 characterizes the convergence properties of monotonic learning processes driven by imitation. Section 3 uses the Binmore and Samuelson (1999) technique to consider the robustness of the imitation convergence results to noisy experimentation. All proofs are in the appendix.

1. Model

The economy is the Grossman and Stiglitz (1980) one-period endowment economy where traders can choose whether or not to purchase a costly signal of the terminal dividend (risky asset payoff). Since we wish to investigate adaptive learning, the one-period model will be repeated. Individuals learn from previous generations, but the economy each generation faces consists only of a single period. In standard economic models, once preferences, information, and endowments are specified, behavior is fully determined as the solution to an optimization problem. In contrast, adaptive learning individuals are specified by rules of thumb and a learning process specifies how these rules evolve. This section characterizes the repeated GS environment and provides the basis for modeling adaptive learning.

1.1 Repeated GS Model

At each date t a generation consisting of an unit-measure continuum of one-period lived individuals trades a risk-free and a risky asset in a perfectly competitive market. At the end of each period the payoffs on the securities are realized and fully consumed. No wealth can be transferred between generations. This preserves the one-period nature of the GS model.[3] Traders have constant absolute risk aversion (CARA) preferences, with risk aversion a, over their end-of-period wealth. The exogenous end-of-period payoff on the risky asset, denoted d_t, is given by

$$d_t = y_t + z_t, \tag{1}$$

where y_t and z_t are random variables. Since no wealth is transferred between periods, the payoff and price of the risk-free asset are normalized to one. The market clearing price of the risky asset is denoted P_t. The supply of the risky asset is e_t and is a random variable. The assumption of a stochastic risky asset supply is required in the original GS model to obtain existence of a rational expectations equilibrium. The random variables, y_t, z_t, and e_t are independently and identically distributed as uncorrelated, mean-zero, and jointly normal with strictly positive variances, σ_y^2, σ_z^2, and σ_e^2. The z_t and e_t are not observed at t. However, traders can choose to observe the common signal y_t by paying cost c before trading.

Traders must choose whether or not they wish to become informed and their demand for the risky asset, x_t^n. Preferences are represented as CARA utility over end-of-period wealth. Recall that there is no wealth transferred between periods. Let W_{t0}^n and W_{t1}^n indicate initial and terminal wealth for

[3] A repeated one-period model is suggested in Grossman and Stiglitz (1980). However, it is important to emphasise that although the economy is repeated, it is not a dynamic model like Wang (1993). A dynamic models with adaptive learning are complicated since a component of returns, the future price, depends on the state of learning in the future. LeBaron, Arthur, and Palmer (1997) simulate adaptive learning in a dynamic market but restricts attention to myopic behavior. The repeated one-period model we study here is similar to Bray (1982) and Bray and Kreps (1987). An alternative interpretation is one of infinitely lived traders who are forced to consume before trading can begin again (trade at t and dividends are realized and completely consumed at $t + 1/2$).

trader n in period t. The utility of trader n is defined as

$$U(W_{t1}^n) = -\exp\left(-aW_{t1}^n\right). \tag{2}$$

End-of-period wealth is determined by the information choice, demands for the risky asset, x^n, and a budget constraint with initial wealth W_{t0}^n. That is, $W_{t1}^n = W_{t0}^n + x_t^n(d_t - P_t) - c1_{\{I\}}$. Note the trader pays the cost of information only if she is informed (the $1_{\{\}}$ is an indicator function). If the trader chooses to be informed, risky asset demands may depend on the signal y_t as well as the asset price P_t. If the trader chooses to remain uninformed asset demands may only depend on the price.

1.2 Individual Behavior

Traders choose whether or not they wish to purchase information and how they wish to formulate their asset demands. We summarize their choice as an element from the finite set with $N+1$ elements, $L^n \in \{I\} \cup L^U$, where $L^U = \{\ell_0^1, \dots, \ell_0^{N0}\} \times \{\ell_1^1, \dots, \ell_1^{N1}\}$. If $L^n = I$, then the trader is informed. All informed traders have the same asset demands. If $L^n = \ell^n = (\ell_0^{n0}, \ell_1^{n1})$, then the trader has chosen not to be informed and her asset demands are determined using the two parameters ℓ_0^{n0} and ℓ_1^{n1}. How these parameters are used is described below. At generation t, the vector ϕ_t describes the proportion of the agents using each of the strategies. $\phi_t \in S^{N+1}$, the $N+1$ simplex. Let $\lambda_t = \phi_t^0$ represent the proportion of agents who are informed (i.e., using I). $1 - \lambda_t = \sum_{n=1}^{N} \phi^n$ is the proportion of uninformed traders. The finiteness of L^U simplifies the specification of the learning processes considered in the next section. However, in Section 3, the assumption can be relaxed as we focus directly on the average demand parameters.

The asset demands of an informed trader that maximize expected utility follow directly from (conditional) normality and CARA utility. Since all the informed traders observe the same signal, y_t, there is no additional information available for them in the market clearing price. Informed traders need not learn the endogenous signal-price relation. Therefore, conditional on being informed, the asset demand schedule is fixed by utility maximization.[4] Asset demand schedule for an informed trader ($L^n = I$) is

$$x_t^I = \frac{E[d|y] - P}{aV[d|y]} = \gamma^I(y_t - P_t), \tag{3}$$

where $\gamma^I = (a\sigma_z^2)^{-1}$.

The asset demands for an uninformed trader are more difficult to specify since uninformed traders must make an inference about the signal from

[4] Routledge (1995) considers a more general case where all agents, including informed traders, must also learn the exogenous signal-dividend relation. This more general environment is investigated by simulating an adaptive learning process using a genetic algorithm.

the equilibrium asset price. If the price-signal relation is known (as in the original GS model) then linear asset demands maximize expected utility.[5] However, in this model the price-signal relation is not ex ante known and must be learned. One approach would be to specify prior beliefs about the price-signal relation, updated via Bayes' rule, and calculate utility maximizing asset demands [see Bray and Kreps (1987)]. However, I wish to investigate adaptive learning in this article. Therefore we use an approach similar to that used in least squares learning [see Bray (1982)]. Here the functional form of the demands is specified and the parameters are learned. Therefore, by assumption, the asset demand schedule for trader n, conditional on being uninformed, is

$$x^n = \gamma^U \left(\ell_0^{n0} + \ell_1^{n1} P_t - P_t\right), \tag{4}$$

where $\ell^n = (\ell_0^{n0}, \ell_1^{n1}) \in L^U$ are the asset demand parameters and $\gamma^U > 0$ is constant across traders and across time.

Focusing on linear demands is a binding restriction. Bayesian demands which maximize subjective expected utility are not linear. However, with adaptive learning, behavior is modified based on success or fitness. As will be demonstrated in Lemma 4, the linear asset demands are not inconsistent with maximimal fitness. In addition, the linear demands are compatible with demands in the GS model. This will allow comparison of adaptive learning to the static GS rational expectations equilibrium.

1.3 Market Clearing

The linearity of the GS model allows the vector ϕ_t to be summarized by the proportion of informed traders λ_t and average demand parameters ℓ_t. Market clearing price, conditional distributions, and utility depend only on λ_t and ℓ_t. Define the average demand parameters $\ell_t = (\ell_{0t}, \ell_{1t})$ as

$$\ell_t = \frac{1}{1-\lambda_t} \sum_{\ell^n \in L^U} \phi_t^n \ell^n. \tag{5}$$

Not surprisingly, the linearity of the demands implies a linear market clearing price relation.

Lemma 1. *Given the proportion of informed traders λ_t and the average demand parameters of the uninformed ℓ_t, the market clearing price is linear in signal y_t and asset supply e_t as*

$$P_t = \frac{(1-\lambda_t)\gamma^U \ell_{0t} + \lambda_t \gamma^I y_t - e_t}{\lambda_t \gamma^I + (1-\lambda_t)\gamma^U (1-\ell_{1t})}. \tag{6}$$

[5] This requires that the asset demands of all agents are linear. Given this linearity, the conditional distribution of y on observing P is normal. Optimal linear demands follow from this normality and the CARA preferences. This is described in detail in Grossman and Stiglitz (1980).

From the market clearing price, an inference about the signal y_t can be formed. In order to make the correct inference, one must know ϕ_t, which can be summarized by λ_t and ℓ_t. This is more knowledge than our adaptive learning traders possess. However, calculating the distribution of the signal conditional on observing asset price, λ_t, and ℓ_t is useful for the analysis that follows.

Lemma 2. *Given ϕ, the distribution of signal y conditional on observing price P is normal with conditional mean $E[y \mid P, \phi] = E[y \mid P, \lambda, \ell] = \ell_0^*(\lambda, \ell) + \ell_1^*(\lambda, \ell)P$, where*

$$\begin{aligned} \ell_0^*(\lambda, \ell) &= (1 + k(\lambda))\, \ell_0^{**}(\lambda) - k(\lambda)\ell_0 \\ \ell_1^*(\lambda, \ell) &= (1 + k(\lambda))\, \ell_1^{**}(\lambda) - k(\lambda)\ell_1 \end{aligned} \tag{7}$$

$$k(\lambda) = \frac{\lambda(1-\lambda)\gamma^I \gamma^U \sigma_y^2}{(\lambda\gamma^I)^2\sigma_y^2 + \sigma_e^2}, \tag{8}$$

$$\ell_0^{**}(\lambda) = 0, \qquad \ell_1^{**}(\lambda) = \frac{\lambda\gamma^I \left(\lambda\gamma^I + (1-\lambda)\gamma^U\right)\sigma_y^2}{\lambda\gamma^I \left(\lambda\gamma^I + (1-\lambda)\gamma^U\right)\sigma_y^2 + \sigma_e^2} \tag{9}$$

and conditional variance

$$V[y \mid P, \phi] = V[y \mid P, \lambda] = \frac{\sigma_e^2\sigma_y^2}{\left(\lambda\gamma^I\right)^2 \sigma_y^2 + \sigma_e^2} \tag{10}$$

1.4 Fitness

The intuition for adaptive learning, formally defined in Section 2, is that successful behavior is imitated and (or) unsuccessful behavior is abandoned. In order to model adaptive learning, we need to measure success. This is done by calculating the fitness of the learning states using the CARA preferences for end-of-period wealth. In this article, fitness is defined as

$$f^n(\phi_t) = E\left[U\left(W_1^n\right) \mid \phi_t\right]. \tag{11}$$

Fitness depends on the population state ϕ_t since it determines the equilibrium price-signal relation. Calculating fitness conditional on ϕ_t is appropriate since it is the actual price-signal relation, which depends on ϕ_t, that determines the success of a behavior. In other words, in an adaptive learning model, once an individual's behavior is known, beliefs are not important. Fitness is not based on subjective expected utility, but on realized expected utility.[6]

[6] In Blume and Easley (1992), the distinction between fitness and subjective expected utility is central to their analysis of the evolution of portfolio strategies. A similar issue is raised in Biais and Shadur (1998) where incorrect beliefs improve bargaining power, but fitness is not calculated with respect to those beliefs.

It is an immediate consequence of Lemmas 1 and 2 that the λ_t and ℓ_t contain all the relevant information for calculating the conditional expectations; that is $E[U(W) \mid \phi_t] = E[U(W) \mid \lambda_t, \ell_t]$. The class of learning dynamics we consider in this article are invariant to monotonic transformations of the utility. Therefore we normalize the fitness of the informed traders to be one. Define $f^I(\phi) = f^I(\lambda_t, \ell_t) = 1$ as the relative fitness of an informed trader $(L^n = I)$ and $f^n(\phi) = f^n(\lambda_t, \ell_t)$ as the relative fitness of an uninformed trader using demand parameters ℓ^n given the population characterized by ϕ_t or λ_t and ℓ_t.

Lemma 3. *Normalizing the fitness of the informed trader as* $f^I(\phi) = f^I(\lambda, \ell) \equiv 1$, *uninformed trader fitness is*

$$f^n(\phi) = f^n(\lambda, \ell) = \exp(ac)\left(\frac{\sigma_z^2}{\sigma_z^2 + V[y|P, \phi]}\right)^{1/2} \xi\left(\ell^n, \lambda, \ell\right), \quad (12)$$

where $\xi \colon \mathbb{R}^5 \to [0, 1]$ *and is jointly continuous except at* 0.

Before discussing the fitness expression, it is helpful to calculate the asset demands that maximize uninformed traders' fitness. The following lemma establishes that the fitness maximizing asset demands are linear as in Equation (4) with parameters $\hat{\ell}_0^*(\phi)$ and $\hat{\ell}_1^*(\phi)$.

Lemma 4. *Asset demands,* $x(P)$, *that maximize uninformed fitness are linear as in Equation (4) with parameters* $\ell^n = \hat{\ell}^*(\phi) = \hat{\ell}^*(\lambda, \ell)$, *where*

$$\hat{\ell}_0^*(\lambda, \ell) = \ell_0^*(\lambda, \ell)\frac{\gamma^*(\lambda)}{\gamma^U} \qquad \hat{\ell}_1^*(\lambda, \ell) = \ell_1^*(\lambda, \ell)\frac{\gamma^*(\lambda)}{\gamma^U} - \frac{\gamma^U - \gamma^*(\lambda)}{\gamma^U} \quad (13)$$

$$\gamma^*(\lambda) = \left(a\left(\text{var}[y \mid P, \phi] + \sigma_z^2\right)\right)^{-1}. \quad (14)$$

If the constant γ^U in Equation (4) happens to be $\gamma^*(\lambda)$ then the fitness maximizing parameters are the same parameters used to calculate the $E[y \mid P, \lambda, \ell]$ inference. In this case, learning to maximize fitness is equivalent to learning how to properly infer the signal from the market clearing price. If $\gamma^U \neq \gamma^*(\lambda)$, then the inference parameters in Equation (7) are adjusted for the conditional variance of the dividend, which depends on the proportion of informed traders.

Looking back at Equation (12), there are three elements that contribute to the fitness of the uninformed trader. The first is the cost of the information that the uninformed traders do not bear. The second is the informativeness of the price about the signal. This is determined by the size of the conditional variance of y given P. Note that $V[y \mid P, \phi]$ in Equation (10) is decreasing in the proportion of informed traders. The information cost and price informativeness play the same role in the GS model. The final term is the function $\xi(\ell^n, \lambda, \ell)$. It captures any inference errors made by the uninformed trader. Inference errors are when $\ell^n \neq \hat{\ell}^*(\phi)$, which, by Lemma 4, always reduces fitness; that is $\xi(\ell^n, \lambda, \ell) = 1$ only at $\ell^n = \hat{\ell}^*(\phi)$.[7]

1.5 The GS Rational Expectations Equilibrium

In order to compare the outcome of a learning process to the GS rational expectations equilibrium, it is helpful to state the GS equilibrium in the notation used thus far. In particular, the population state ϕ^*, with proportion λ^* of informed traders, is the GS rational expectations equilibrium if it has two features. The first is that uninformed traders make correct inferences (i.e., no inference errors). For a given λ, this is the fixed point of Equation (13) and is denoted $\hat{\ell}^{**}(\lambda)$.[8] The rational expectations equilibrium is where $\phi^{n^*} = 1 - \lambda^*$, for $\ell^{n^*} = \hat{\ell}^{**}(\lambda^*)$. The second condition is that traders have no incentive to change their information choice. This condition is expressed as $f^{n^*}(\phi^*) = 1$. In this article, I will consider the economies where the GS equilibrium is such that $\lambda^* \in (0, 1)$.

In order for this economy to replicate the GS rational expectations equilibria, the set of possible demand parameters for the uninformed must be rich enough to include the GS equilibrium. Therefore we make the assumption that the rational expectations solution is an element of the set of possible demand parameters; that is, $\hat{\ell}^{**}(\lambda^*) \in L^U$. In addition, we assume that for any fixed $\lambda \in [0, 1]$, the rational expectation demand parameters are contained in the convex hull of the grid; that is, $\hat{\ell}^{**}(\lambda) \in C(L^U)$. As long as the grid is sufficiently fine (large enough N), it is not necessary that these parameters lie exactly on the grid.

2. Adaptive Learning

One period of the GS economy is characterized by the vector ϕ_t containing the proportions of the various behaviors. We now focus on the adaptive learning process that governs the evolution of the ϕ_t. In this section, I will focus on learning driven by imitation. The intuition of adaptation is that successful behavior is copied and (or) unsuccessful behavior is abandoned, resulting in a higher growth in the proportion of the population following successful behavior. I will consider the convergence properties of adaptive learning and give a simple example of an adaptive learning model. In

---

[7] From the proof of Lemma 3, note that absolute utility is the expectation of a squared normal random variable and it need not be finite. Therefore the function ξ is strictly positive and continuous as long as the slope inference errors are not too large. As the slope-based inference errors become large the absolute level of the uninformed expected utility approaches negative infinity and the ratio of informed to uninformed utility approaches zero.

[8] The existence and uniqueness of the fixed point is established in Grossman and Stiglitz (1980). In this article, the existence and uniqueness follows directly from $k(\lambda) \geq 0$ and the linearity in Equations (7) and (13).

this section I will assume ϕ_t evolves according to a *monotone selection dynamic*, g.[9]

Definition. *Let* g: $S^{N+1} \rightarrow \mathbb{R}^{N+1}$.

$$\frac{d\phi_t^n}{dt} = \phi_t^n g^n(\phi_t) \tag{15}$$

is a monotone selection dynamic if it satisfies for all $\phi \in S^{N+1}$

(1) *Continuity:* g *is Lipschitz continuous*
(2) *Simplex restriction:* $\Sigma \phi^n g^n = 0$
(3) *Regular:* $g(\phi)$ *is bounded for all* $\phi \in S^{N+1}$
(4) *Monotonic: for all* $L^n, L^{n'} \in \{I\} \times L^U$

$$f^n(\phi) > (=) f^{n'}(\phi) \qquad \Rightarrow \qquad g^n(\phi) > (=) g^{n'}(\phi) \tag{16}$$

The continuity ensures that the differential equation in Equation (15) has a unique solution. The second condition ensures that ϕ_t remains on the simplex S^{N+1}. The assumption that growth rates are finite implies that any behavior is never literally abandoned and novel behavior is never introduced. Therefore we will restrict our attention to initial populations where all behavior is represented [i.e., $\phi_0 \in interior(S^{N+1})$]. These three conditions are technical. The crucial economic assumption is the fourth condition. The monotonicity links the fitness to growth rates and relates the dynamic to learning. The crucial assumption is that better strategies are imitated more frequently and (or) bad strategies are abandoned, resulting in a higher growth rate for better strategies. How traders may implement this algorithm in a financial market is discussed after presenting the results and an example.

2.1 Convergence Properties of Monotone Selection Dynamics

In order to understand the properties of learning, we consider the choice between being an informed and uninformed trader separate from the choice of inference parameters. For a moment, fix the proportion of informed agents at some arbitrary level, λ. Consider the selection dynamic, g, that is monotointic for all $n \geq 1$ with the exception that $g^0 = g^I = 0$ (i.e., λ_t is constant). The following proposition establishes a sufficient condition for the parameters of the economy that implies adaptive learning will converge on the rational expectations parameters, $\hat{\ell}^{**}(\lambda)$. These parameters produce the same demands as in the standard GS rational expectations equilibrium with

[9] Since the repeated GS model is discrete, the differential equations are the limit of difference equations. It is more convenient to use differential equations. See Weibull (1995) for a discussion of difference equations in evolutionary game theory.

a fixed and exogenously specified proportion of informed traders. Define $\hat{k}(\lambda)$ using Equation (8) as

$$\hat{k}(\lambda) = \frac{\gamma^*(\lambda)}{\gamma^U} k(\lambda). \tag{17}$$

Proposition 1. *For a monotone selection dynamic, with fine enough grid of strategies (large N), holding fixed the proportion of informed agents and for any initial population: if $\hat{k}(\lambda) < 1$ the adaptive learning will converge such that all uninformed traders are using a strategy arbitrarily close to the GS rational expectations demand parameters.*

The proof of Proposition 1 relies on the fact that the fitness maximizing demand parameters, $\hat{\ell}^*(\phi)$, are a linear function of the average demand parameters. In particular, since $\hat{k} < 1$, extreme demand parameters (i.e., on the boundary of the convex hull of L^U) are never fitness maximizing and their proportionate use in the population will decline under a monotone selection dynamic.[10] Once the use of the extreme parameters declines, they have a vanishing influence on the average demand parameters. Now, the next-to-extreme parameters are never fitness maximizing. In this way, a monotone selection dynamic hones in on the rational expectations demands.[11] Finally, $\lambda \to 1$ implies $\hat{k} < 1$. Thus, for any economy, one can always fix a proportion of informed traders, λ, high enough such that the uninformed traders' demands converge to the rational expectations demands.

The next proposition builds on this result to demonstrate that the GS rational expectations equilibrium, where we do not constraint the proportion of informed agents, is the unique asymptotically stable fixed point of any monotonic selection dynamic.[12]

Proposition 2. *For any monotonic selection dynamic, the GS equilibrium ϕ^* is asymptotically stable where the proportion of informed traders con-*

---

[10] The condition $\hat{k} < 1$ is identical to the condition in Bray (1982) for the convergence of intermittent least squares learning. In Bray (1982), traders update their demands by using the parameters from an OLS regression on the dividend-price observations. In the intermittent form of this learning, traders have a long (infinite) series of price-dividend observations before adjusting their behavior. In this case, Bray can appeal to the consistency properties of OLS estimates. In an incremental form of OLS where agents update their parameters each period, the convergence results no longer require $\hat{k} < 1$. Analogously, in our model, it may be possible to relax the requirement that $\hat{k} < 1$ if fitness was defined based on the realized utility and not expected utility. Finally, from Equations (8) and (17), note that changing the arbitrary constant γ^U does not change $\hat{k}$ and has no affect on convergence.

[11] This proof is analogous to the proof of Samuelson and Zhang (1992) showing that monotone selection dynamic eliminates iteratively dominated strategies.

[12] The standard definitions for dynamic systems are used. See Hirsch and Smale (1974) for details. ϕ is a *stationary point or dynamic equilibrium* of Equation (15) if $\phi^n g^n(\phi) = 0$ for all n. ϕ is *asymptotically stable* if there exists a neighborhood Q of ϕ such that all trajectories of Equation (15) that originate in Q converge to ϕ.

verges to the GS equilibrium from above ($\lambda_t \to \lambda^*$ from above). In addition, if the grid is fine enough (large enough N), ϕ^* is the only such point.

If $\lambda < \lambda^*$, uninformed traders always have lower fitness than informed. If $\lambda > \lambda^*$ and uninformed traders are using appropriate inference parameters [i.e., from Equation (13)], then uninformed traders have a higher fitness than the informed traders. This is due to the conditional variance of the signal conditional on price decreasing in λ. Since growth rates are related to fitness under a monotonic selection dynamic, λ moves toward λ^*. The monotonic selection dynamic also drives inference errors to zero. Close to the GS equilibrium, $\hat{\ell}^{**}(\lambda^*)$ are the best demand parameters and are eventually adopted by all uninformed traders. Since the proposition only establishes asymptotic stability and not global attraction as in Proposition 1, it is not necessary to assume that $\hat{k} < 1$. Finally, the GS equilibrium is the unique stable point since it is the only situation where there does not exist some fitness-improving behavior. For example, $\lambda = 1$ is not asymptotically stable since close to $\lambda = 1$, uninformed traders using demand parameters close to $\hat{\ell}^{**}(1)$ will have higher fitness and therefore the proportion using these parameters will grow.

The reason that λ approaches λ^* from above is that inference errors (which reduce the uninformed fitness) are nonzero before learning has converged. Informed and uninformed fitness is equal only when λ is strictly above λ^*, which compensates the uninformed traders for the nonzero inference errors. This trajectory for λ is consistent with the experimental asset market result in Sunder (1992). In that article, the proportion of informed traders was held constant and the cost of information was determined in equilibrium. In Sunder's experiment, the equilibrium cost of information approached the rational expectations equilibrium cost from above. This is consistent with λ approaching λ^* from above given the constant cost of information in our model. Oddly, Sunder observes that when the cost of information was fixed, the λ failed to converge to its rational expectations level.

2.2 An Example of a Monotone Selection Dynamic

We can use Björnerstedt and Weibull (1996) to construct an example of a monotonic learning process to highlight the important behavioral assumptions of adaptive learning in a financial market. Assume individuals currently using behavior L^n choose to revise their strategy with probability $\rho^n(\phi)$. Given they choose to reevaluate their behavior, they will choose strategy L^m with a probability given by $\pi^n_m(\phi)$. Both of these probabilities can depend on the current population ϕ_t (and therefore depend on fitness levels, average inference parameters etc.).[13] Since there is a continuum of traders,

[13] We need to assume that the functions ρ and π are Lipshitz continuous.

if we assume enough independence in individual updating, we can appeal to the law of large numbers to produce the following selection dynamic:

$$\frac{d\phi^n}{dt} = \sum_m \phi^m \rho^m(\phi)\pi_n^m(\phi) - \rho^n(\phi)\phi^n. \tag{18}$$

A monotone selection dynamic follows from specific assumptions about ρ and π. For example, suppose that people with lower fitness are more likely to reevaluate their strategy. That is, let $\rho^n(\phi) = \rho(f^n)$, where $\rho : \mathbb{R} \to (0, 1)$ and is strictly decreasing in fitness. In addition, assume that a trader that is updating her behavior simply adopts one strategy from the population at random. That is, people copy the behavior of the first person they meet regardless of that person's fitness. Appealing again to the large number of traders, this implies that $\pi_m^n = \phi^m$. Putting these two assumptions into Equation (18) yields:

$$\frac{d\phi^n}{dt} = \phi^n \left(\sum_m \phi^m \rho(f^m) - \rho(f^n) \right). \tag{19}$$

Since ρ is a decreasing function of fitness, the dynamic is a monotonic selection dynamic and therefore has the properties used to derive Propositions 1 and 2.[14] The assumptions required to generate Equation (19) are mild. However, in order for traders to implement adaptive learning, they must observe their fitness and the behavior of others. How individuals might have access to this information in a financial market setting is discussed below.

First, it is crucial for learning that behavior revisions be related to fitness. For example, if the probability of revising one's strategy does not relate to fitness [i.e., $\rho(f^n) = 0$ in Equation (19)], then the population never changes $(d\phi^n/dt = 0)$. Therefore, to justify the assumption of a monotone selection dynamic, it is important traders have some information about their fitness. In our model, each period, wealth (dividend) is realized and consumed. This produces a realized utility according to Equation (2). Since it is not necessary that a trader updates her strategy each period (in the above example, ρ can be strictly less than 1), she will have several realizations of utility. This provides a noisy measure of one's own absolute fitness. This is sufficient to implement the example embodied in Equation (18).[15]

While having an indication of one's own fitness is crucial, the example of a monotone selection dynamic presented in Equation (19) does not require the traders to know other traders' fitness. The crucial assumption is that the

[14] As an aside, note that if the function ρ is linear, Equation (19) produces the commonly used replicator dynamics introduced by Taylor and Jonker (1978).

[15] The normalization by the fitness of the informed trader (see Lemma 3) simplifies the calculations in the article. However, it is straightforward to construct an example of Equation (19) using the absolute fitness measure of Equation (11).

function $\rho(f^n)$ is monotonic decreasing. The assumption that lower fitness leads to a higher probability of strategy revision does not require traders to know the fitness of other individuals. However, to calibrate the function ρ, one must have some feel for the likely range for fitness. However, this is less demanding than knowing the specific fitness of a specific trader.[16]

A concern with using realizations of the economy as an estimator of fitness is that the observations about fitness are not drawn from a stationary environment due to the learning of other traders. By defining fitness based on expected utility, our model is closely related to Bray's (1982) model of intermittent OLS learning. In her model, she assumes traders observe a long series from the economy before anyone updates. The consistency properties of OLS estimation simplify the analysis by eliminating estimation error (see note 10). The difficulty with using a small sample of realizations to infer fitness is that estimation errors across traders will be correlated. In particular, the dividend realization will influence the learning dynamic, making the analysis much more complicated. In dynamic models, the link between economic shocks and learning is interesting. For example, in dynamic settings, Bayesian learning can induce ARCH patterns in returns [see Veronesi (1999)]. Formally, in the context of this example, one can ensure that traders have a sufficiently precise estimate of their fitness by choosing a low frequency of strategy revision [small values for $\rho(f)$].

The second requirement to implementing adaptive learning is that traders have some observation about the behavior of other traders. In the example presented, traders do not need to observe the fitness of another trader. However, they do need to observe the strategy of a trader selected at random. Alternatively, one can assume that traders observe the strategies of more successful traders.[17] In some financial market settings, such as with floor traders, it is feasible that traders can observe some noisy signal of other's behavior. Presumably this is part of the reason why a floor trader's training often involves assisting a more senior trader and gaining hands-on experience with limited exposure. In other financial market contexts, such as proprietary trading, other traders' behavior is not directly observable. Imitative learning in this context may involve cross-firm hiring or acquisitions.

[16] If one wanted to construct another example of a monotone selection dynamic where traders compared their fitness to the fitness of the population (e.g., average fitness), traders could construct a noisy measure of population fitness using information about consumption (terminal wealth) in the economy. For example, one can estimate average fitness by observing average terminal wealth in the economy in a period. Of course, since utility/fitness is concave, a better estimate of average fitness could be constructed from observing more information about the wealth distribution (quartiles, deciles, etc.). Using information about aggregate consumption to evaluate oneself and make decisions appears in the class of preferences referred to as "catching up with the Joneses" [Abel (1990)]. In other contexts such as mutual funds, there is explicit publicly available information on relative performance that one could use to implement adaptive learning.

[17] For example, if $\pi_m^n = \phi^m \varepsilon^m$, then a sufficient condition to generate a monotone selection dynamic is that $\varepsilon^m \geq \varepsilon^n$, i.f.f. $f^m \geq f^n$. In the example in Equation (19), $\varepsilon^m = 1$ for all m.

Finally, some imitative learning can occur through publication and discussion of strategies in various trade and popular media.[18]

3. Adaptive Learning and Experimentation

Propositions 1 and 2 follow from the monotonicity of imitative learning. The exact specification of how traders perform the learning by imitation is not important. However, the previous example depends heavily on the assumption that updating and copying is independent enough to eliminate randomness. Since this assumption is perhaps unrealistic, it is necessary to consider the robustness of the results to learning processes that have a small amount of noise. Understanding the role of random experimentation is necessary if we wish to use adaptive learning to understand experimental asset market studies like Sunder (1992), which has a small number of traders. More generally, understanding the robustness of rational expectations equilibria to small amounts of behavior randomness is an important aspect of the current debate about behavioral finance. This section extends the previous results to capture the effect of noise from experimentation or noise in the imitation process.

Ideally one would capture the stochastic elements of learning directly. Unfortunately the Markov process generated by a stochastic adaptive learning algorithm is typically too large for direct analysis. Several articles approach this type of model by simulating learning using a genetic algorithm or other random search algorithms [e.g., Routledge (1995) or LeBaron, Arthur and Palmer (1998)]. While these simulations offer insight, it can be difficult to determine which of the many simulation parameters are crucial. In our model, we are able to use techniques developed in Binmore and Samuelson (1999) to capture the important characteristics of noisy adaptive learning in a deterministic differential equation. Instead of focusing on the learning of individual traders, we can characterize the behavior of the proportion of informed traders, λ_t, and the average inference parameters of the uninformed traders, ℓ_t. As noted previously, these parameters are sufficient to calculate the equilibrium asset price, the correct inference parameters, and fitness.

Binmore and Samuelson (1999) construct a model to characterize the evolution of the average behavior using a deterministic differential equation.[19] In their analysis, the noise from experimentation is replaced by determin-

---

[18] For example, Lappen (1998) describes the strategies of money managers (in this case, all are well-known academics) and offers an example of how strategies can be (partially) observed. Such articles raise two interesting questions not addressed in our model. First, in our model there is no difference in the skill of traders. In practice, traders can imitate the strategies discussed in Lappen with different degrees of success. Second, there is some strategic component to disclosing strategies in this manner in that such exposure increases the flow into the funds discussed.

[19] See Binmore and Samuelson (1999) or Samuelson (1997) for additional details. Similar techniques are also used in Gale, Binmore, and Samuelson (1995), Sethi and Somanathan (1996), Somanathan (1997).

istic drift. The authors establish that equilibria of the deterministic process are arbitrarily close to the expected state of the stochastic process for some long but finite time period for small levels of noise and drift. To understand the average behavior of a learning process with a small degree of experimentation, the exact form of the noise does not matter. Hence the effects of random noise can be captured by the deterministic drift. Unfortunately their result does not necessarily apply for arbitrarily long time periods. If the deterministic system (with drift) has multiple stationary points, shocks (often called mutations) in the stochastic process, however improbable, can move the system between the multiple basins of attraction present in the deterministic system. In such a case, it is the relative likelihood of shocks that determines the expected limiting behavior of the stochastic process and the exact specification of the stochastic process matters and deterministic drift is a less useful tool. However, if the deterministic system has a unique stable equilibrium, then the specification of shocks is not important since, on average, the system will remain in the neighborhood of the stable deterministic equilibrium. In the GS model presented here, deterministic learning processes are constructed which have unique stable equilibria.

In this section I focus on the evolution of the proportion of informed traders, λ_t, and the average inference parameters, ℓ_t. Note that the average inference parameter is a continuous variable and the assumption regarding the discrete grid of individual demand parameters is no longer binding. We continue to use $f^I(\lambda_t, \ell_t) = 1$ as the fitness of the informed trader. Let $f^U(\lambda_t, \ell_t)$ be the fitness of an uninformed trader using the average inference parameters [calculated in Equation (12) where $\ell^n = \ell_t$]. The important parameters of the economy are denoted $E = (\sigma_e^2, \lambda^*)$. In the analysis we will hold all parameters of the economy constant except the level of noise in the risky asset, σ_e^2 and the cost of information, c, implied by the GS equilibrium proportion of informed traders, λ^*.

Consider the adaptive learning process, $\mathcal{L}$, for the GS model defined by the following set of deterministic differential equations for the state variables (λ_t, ℓ_t).

$$\frac{d\lambda_t}{dt} = g_\lambda(\lambda_t, \ell_t) + \eta_\lambda m_\lambda(\lambda_t, \ell_t) \tag{20}$$

$$\frac{d\ell_t}{dt} = g_\ell(\lambda_t, \ell_t) + \eta_\ell m_\ell(\lambda_t, \ell_t) \tag{21}$$

where $g_\lambda, m_\lambda\colon [0, 1] \times \mathbb{R}^2 \to \mathbb{R}$, $g_\ell, m_\ell\colon [0, 1] \times \mathbb{R}^2 \to \mathbb{R}^2$, and η_λ and η_l are positive constants. The g functions describe the imitative portion of the learning process. The m functions are drift and represent the effects of experimentation. The η are small positive constants that determine the amount of drift. We make the following assumptions about $\mathcal{L}$. They are discussed below.

0. *Continuity*
   The functions, g_λ, g_ℓ, m_λ, and m_ℓ are Lipshitz continuous.
1. *Monotonicity*
   (a) For $\lambda \in (0, 1)$, $g_\lambda(\lambda, \ell) \leq (<) 0 \Leftrightarrow f^U(\lambda, \ell) \geq (>) 1$.
   (b) For a fixed $\lambda \in (0, 1)$, the solution to $d\ell/dt = g_\ell(\lambda, \ell)$ implies $\lim_{t\to\infty} \ell_t = \hat{\ell}^{**}(\lambda)$.
2. *Regularity*
   (a) $\lim_{\lambda\to 0} g_\lambda(\lambda, \ell) \geq 0$ and $\lim_{\lambda\to 1} g_\lambda(\lambda, \ell) \leq 0$.
   (b) $\lim_{\lambda\to 1} \|g_\ell(\lambda, \ell)\| = 0$.
3. *Drift*
   (a) For all ℓ, $\lim_{\lambda\to 0} m_\lambda(\lambda, \ell) > 0$ and $\lim_{\lambda\to 1} m_\lambda(\lambda, \ell) < 0$.
   (b) For all λ, the solution to $d\ell/dt = m_\ell(\lambda, \ell)$ implies $\lim_{t\to\infty} \ell_t = \ell_m$, such that $\ell_m \neq \hat{\ell}^{**}(\lambda)$ for any λ.

The assumptions on the g functions are consistent with the more primitive behavioral assumption in the monotonic selection dynamic used in the previous section. The monotonicity assumptions are natural analogs to Propositions 1 and 2. The regularity assumption [Assumption 2(a)] is similar to the simplex restriction of the monotonic selection dynamic and implies λ_t remains in the interval [0,1]. Assumption 2(b) is consistent with the bounded growth rates assumed in the previous section. It implies that the convergence to the rational expectations inference parameters slows when there are fewer uninformed traders (i.e., as $\lambda \to 1$). Assumption 3 is discussed below in Section 3.2.

3.1 Imitation Only

In the case where $\eta_\lambda = \eta_\ell = 0$, the stable dynamic equilibria reflect the two types of behavior. The first is the GS equilibrium, $(\lambda^*, \hat{\ell}^{**}(\lambda^*))$. The second type is where $\lambda = 1$ and inference errors remain large.

Proposition 3. *Given an economy, E, with $\sigma_e^2 > 0$, the stable dynamic equilibria for the learning process in Equations (20) and (21) with $\eta_\lambda = \eta_\ell = 0$ are the GS equilibrium, $\{(\lambda^*, \hat{\ell}^{**}(\lambda^*))\}$ and where all traders are informed $(\lambda = 1)$ such that $\{(1, \ell) \notin G(E)\}$, where $G(E) = \{(\lambda, \ell) | f^U(\lambda, \ell) \geq 1\}$.*

Figure 1 characterizes this result. The proportion of informed traders, λ, is represented on the horizontal axis and one dimension of the uninformed traders' inference error, $\ell - \hat{\ell}^*(\lambda, \ell)$, is represented on the vertical axis. The set $G(E)$ is shown and, by Assumption 1(a), represents the region where $d\lambda/dt \leq 0$. The important properties of $G(E)$ are that it is nonempty, closed, and contains only points $\lambda > \lambda^*$ (see Lemma A2 in the appendix). When $\lambda > \lambda^*$, traders can make small inference errors and still have a higher fitness than the informed traders. The line H identifies the set where inference

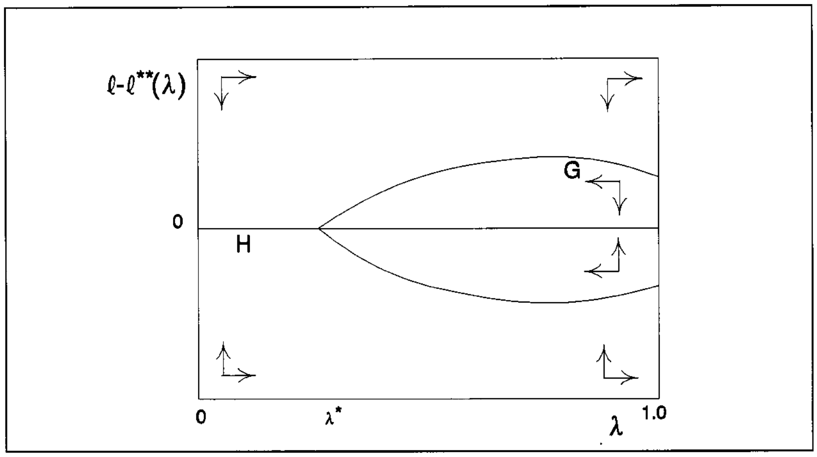

Figure 1
Learning process with no drift
The arrows are for a typical learning dynamic with $\eta_\lambda = \eta_\ell = 0$. The equilibria are at the GS values $(\lambda^*, 0)$ or at $\lambda = 1$.

errors are zero and, by Assumption 1(b), $d\ell/dt = 0$. The arrows represent directions of movement for a learning dynamic. Figure 2 represents phase trajectories for specific g_λ and g_ℓ.[20] Each line on the diagram represents an adaptive learning trajectory in the economy for a different initial λ and ℓ. The two types of stable dynamic equilibria are where the phase lines converge (the GS equilibrium) or reach the boundary ($\lambda = 1$) outside the set G.

As in Proposition 2, an asymptotically stable dynamic equilibria occurs at the GS rational expectations equilibrium. However, unlike in the previous section, we can identify stable equilibria where all traders are informed ($\lambda = 1$) and inference errors are large enough such that an uninformed trader has lower utility than the informed trader. Once (almost) all traders are informed, inference parameters are (almost) not updated. This is due to the fact that inference parameter learning slows as the proportion of uninformed traders decreases [Assumption 2(b)]. The reduced form equations of Equations (20) and (21) are consistent with the monotonic selection dynamic considered in the previous section, since the $\lambda = 1$ stable points are not asymptotically stable. One can perturb the system away from these stable points and λ will converge back to $\lambda = 1$. However, the inference parameters will move closer to the rational expectations parameters. Although these points do not

[20] These functions are the replicator dynamics for the evolution of λ and a "best response" dynamic for the updating of inference parameters. These functions are chosen only to illustrate the more general points of this section. The inference parameter space (to avoid higher dimensional diagrams) is for the intercept parameter ℓ_0. In these simulations, inference errors on the slope parameter ℓ_1, are held at zero.

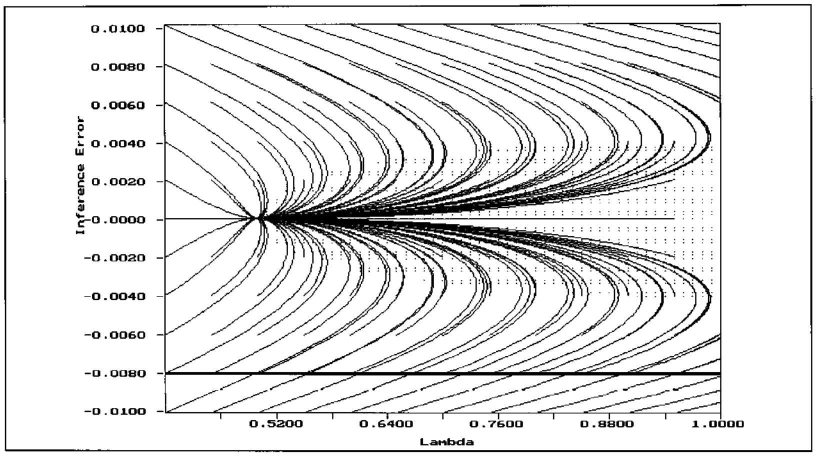

Figure 2
Phase plots in learning process without drift
Each line represents a different initial condition. Note that all lines converge to GS equilibrium ($\lambda^* = 0.5$) or the $\lambda = 1$ boundary.

persist in the monotonic selection of Proposition 2, they play an important role when we consider noise proxied by drift.

3.2 Adaptation and Drift

In our model, experimentation has two important effects. First, it will ensure that some traders are always experimenting with being uninformed. It pushes the system away from the $\lambda = 1$ stationary point. However, experimentation also makes it more difficult for the uninformed trader to learn the correct inference parameters since each variation in behavior alters the $\hat{\ell}^*(\lambda, \ell)$. Proposition 3 shows that purely imitative learning leads the economy either to the GS equilibrium or to the situation where everyone is informed. Which one of these outcomes is more robust to experimentation, and hence more likely, depends on the relative depths of the basins of attraction. Using the Binmore and Samuelson (1999) approach, this salient feature of random experimentation can be captured with deterministic drift. Locally, pushing the system away from the GS equilibrium or from the everyone-informed stationary points with deterministic drift has the same effect as stochastic experimentation.

The key features of experimentation are captured in the drift functions, m. Assumption 3(a) keeps λ away from the boundaries [the inequality in Assumption 3(a) is strict]. Assumption 3(b) pulls the inference parameters toward some arbitrary and incorrect value. The η_λ and η_ℓ parameters scale the drift and ensure that both of these effects are small. Small drift means that the direction and stationary points of Equations (20) and (21) are similar to direction and zero points of g_λ and g_ℓ everywhere except near $\lambda = 1$. To help describe this, define $G'(E)$ as the points where λ_t is decreasing and $h(\lambda; E)$ as the zero points of Equation (21) for a given λ.[21] That is,

$$G' = \{(\lambda, \ell) \mid g_\lambda(\lambda, \ell) + \eta_\lambda m_\lambda(\lambda, \ell) \leq 0\} \tag{22}$$

$$h(\lambda) = \{\ell \mid g_\ell(\lambda, \ell) + \eta_\ell m_\ell(\lambda, \ell) = 0\} \tag{23}$$

The following lemma shows that we can specify a learning process with drift using three parameters: λ'', Δ, and ℓ_m. We can choose a learning process $\mathcal{L}(\lambda'', \Delta, \ell_m)$ with drift levels $\eta_\lambda > 0$ and $\eta_\ell > 0$ such that for all $\lambda < \lambda''$, $G'(E)$ is within Δ of G(E) and $h(\lambda; E)$ is within Δ of $\hat{\ell}^{**}(\lambda)$. The final part of the lemma establishes that drift matters most when most traders are informed (as $\lambda_t \to 1$).

Lemma 5. *Given $\Delta > 0$ and $\lambda'' < 1$, there exists $\eta > 0$ such that for all $\eta_\lambda < \eta$ and $\eta_\ell < \eta$*

(a) for all $(\lambda', \ell') \in G'$ [$(\lambda, \ell) \in G$] with $\lambda' < \lambda''$, there exists $(\lambda, \ell) \in G$ [or $(\lambda', \ell') \in G'$] such that $\|(\lambda, \ell) - (\lambda', \ell')\| < \Delta$,
(b) for all $\lambda \leq \lambda''$, $\|h(\lambda) - \hat{\ell}^{**}(\lambda)\| < \Delta$, and
(c) there exists an $\bar{\varepsilon}$ such that $\varepsilon < \bar{\varepsilon}$ implies $(1 - \varepsilon, \ell) \in G'$ for all ℓ and $\|h(1 - \varepsilon) - \ell_m\| < \Delta$.

Dynamic equilibria for a learning process, $\mathcal{L}$, occur where $h(\lambda)$ intersect the boundary of G'. The following proposition establishes that at least one stable dynamic equilibrium exists.

Proposition 4. *For a given economy E, there exists at least one stable dynamic equilibrium to the learning process, $\mathcal{L}(\lambda'', \Delta, \ell_m)$.*

The static GS rational expectations equilibrium is not a dynamic equilibrium of an adaptive learning process $\mathcal{L}(\lambda'', \Delta, \ell_m)$ due to the small amount of drift. However, it is possible that all stable dynamic equilibria of the adaptive learning process lie very close to static GS equilibrium. In such a case the static GS rational expectations model and the adaptive learning model with drift make almost identical predictions about the behavior of an economy. The following proposition shows, for a given economy, one can always find learning processes such that the dynamic equilibria lie arbitrarily close to the static GS equilibria.

Proposition 5. *Given $\Delta > 0$ and economy $E = (\sigma_e^2, \lambda^*)$, there exist λ'', Δ'', and ℓ''_m such that for all $\lambda' > \lambda''$, $\Delta' < \Delta''$ and $\|\ell'_m - \hat{\ell}^{**}(1)\| < \|\ell''_m - \hat{\ell}^{**}(1)\|$, all the dynamic equilibria of the learning process $\mathcal{L}(\lambda', \Delta', \ell'_m)$ are within distance Δ of the GS equilibrium $(\lambda^*, \hat{\ell}^{**}(\lambda^*))$.*

[21] Note that from Assumption 1(b) and the implicit function theorem, $h(\lambda; E)$ is a function (and not just a correspondence).

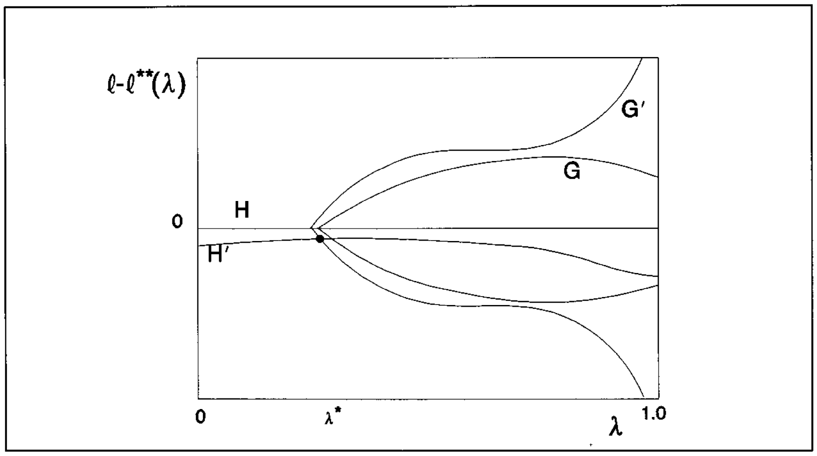

Figure 3
Learning process with drift
The unique dynamic equilibrium is near the GS values (λ^*, 0) (denoted by the dot).

This proposition is illustrated in Figures 3 and 4. $h(\lambda)$ (denoted H' in the figure) and G' cross only once and this crossing occurs near to the GS equilibrium. All trajectories in Figure 4, including those which begin with $\lambda = 1$, converge to a point just below the GS equilibrium.

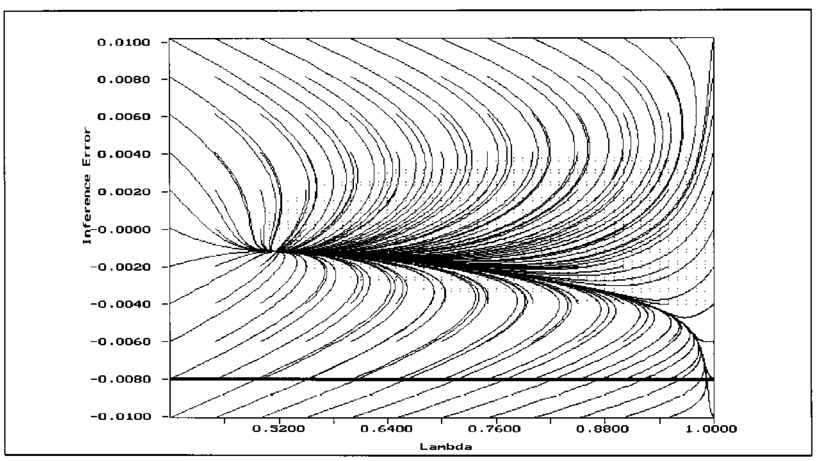

Figure 4
Phase plots in learning process with drift
Each line represents a different initial condition. Note that all lines converge to a point which is close to the GS equilibrium.

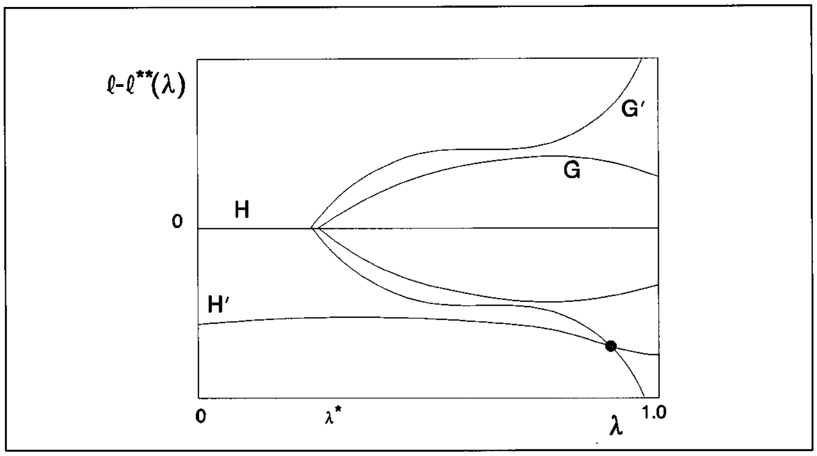

Figure 5
Learning process with drift
The unique dynamic equilibrium (denoted with the dot) is far from the GS values.

Proposition 5 demonstrates that for a given economy, a learning process with arbitrarily small drift converges near to the GS equilibrium. However, the converse does not hold. For a given learning process, economies exist such that the learning process does not converge near to the GS equilibrium. Figures 5 and 6 describe an example where the unique dynamic equilibrium of the learning process is near $\lambda = 1$, away from the GS rational expectations

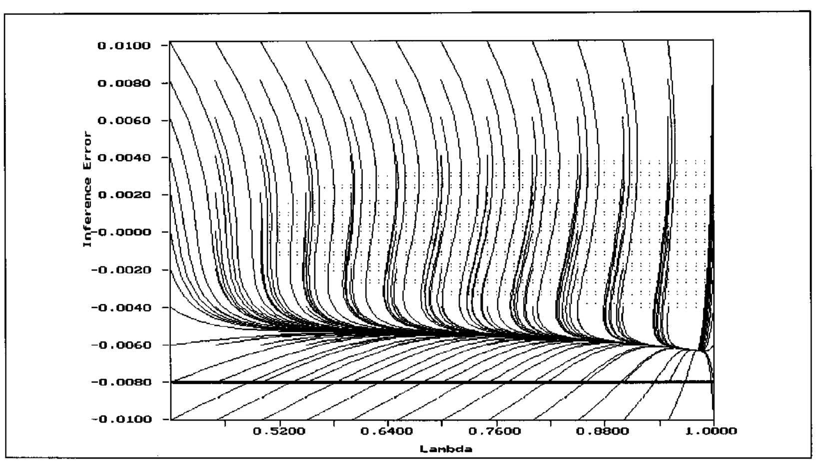

Figure 6
Phase plots in learning process with drift
Each line represents a different initial condition. Note that all lines converge to a point close to $\lambda = 1$.

equilibrium. In this case, all trajectories, including those initiating at the GS equilibrium, converge to this point.

The key for the example in Figure 5 is that the drift in the inference parameters is large relative to the size of the set $G(E)$. If uninformed traders make no inference errors ($\ell = \hat{\ell}^{**}(\lambda)$), when $\lambda > \lambda^*$ it is more economical to infer the signal from the asset price than to be informed and bear the cost of information. This is the case as long as inference errors are small. As the level of noise in the risky asset supply decreases, holding the GS equilibrium λ^* constant, the size of the set $G(E)$ shrinks as does the tolerable inference error. Proposition 6 demonstrates that for a fixed learning process, even one with a small amount of drift, there exist economies such that all dynamic equilibria of the learning process are far from the static GS rational expectations equilibrium.

Proposition 6. *For a given learning process* $\mathcal{L}(\lambda'', \Delta, \ell_m)$ *with* $0 < \lambda^* < \lambda'' < 1$ *such that*

$$\inf_{\lambda<\lambda''} \left\{ \left\| h(\lambda) - \ell^{**}(\lambda) \right\| \right\} > \sup_{\lambda<\lambda''} \left\{ \left\| (\lambda, \ell) - (\lambda, \ell') \right\| \mid (\lambda, \ell) \in G, (\lambda, \ell') \in G' \right\} \quad (24)$$

there exists $\sigma_e^{2''}$ *such that for all economies* $E = (\sigma_e^{2'}, \ell^*)$ *with* $\sigma_e^{2'} < \sigma_e^{2''}$, *all the dynamic equilibria* (λ', ℓ') *have* $\lambda' > \lambda''$.

The proof demonstrates that the set G(E) shrinks to a line (permitting no inference errors) as the asset supply noise vanishes. The dynamic equilibria of the learning process occurs when $h(\lambda)$ intersects the boundary of G', not G. The sufficient condition in Equation (24) ensures that the amount of drift in the inference parameters is roughly the same size as the drift in the information choice so that $G'(E)$ shrinks enough as asset supply variance shrinks.

Propositions 5 and 6 demonstrate that for adaptive learning to produce behavior similar to that observed in the static GS rational expectations equilibrium, it is necessary for drift (a proxy for experimentation) to be small relative to the amount variance in the risky asset supply. This provides an important link between the structure of learning and the structure of the economy. The intuition for this result can be demonstrated by tracing a trajectory of the adaptive learning process beginning at the GS equilibrium. At the GS equilibrium, informed and uninformed traders have identical fitness and uninformed traders make no inference errors. Drift, representing the underlying experimentation process, induces inference errors and thus lower fitness for the uninformed. Adaptive learning causes the proportion of informed traders to increase. If the asset supply variance, σ_e^2, is high, the increased λ has a significant impact on the informativeness of price about the signal (i.e., $V[y \mid P, \lambda, \ell]$ decreases substantially). Therefore, uninformed traders' fitness improves enough to offset inference errors. As uninformed traders learn more about the correct inference parameters and their fitness increases, the proportion of uninformed traders decreases. In this case, the adaptive learning process remains close to the GS equilibrium. In contrast, if asset supply noise is low, then the initial increase in λ has only a small effect on the informativeness of price and uninformed fitness remains below that of the informed traders and most traders become informed. Once most traders are informed, the learning rate of the few remaining uninformed traders slows (since there are fewer uninformed people to copy) and drift in the inference parameters has a larger effect ($\lambda > \lambda''$ in Lemma 5c). In this case, the dynamic equilibrium remains well away from the GS equilibrium.[22]

4. Conclusions

Adaptive learning is a relatively unsophisticated behavior of imitating successful behavior, abandoning unsuccessful behavior, and occasional experimentation. In other fields of research, genetic algorithms and genetic programming, both examples of adaptive learning, have proven very successful in solving complex optimization problems.[23] Whether or not adaptive learning is a useful description of economic behavior is an open question. In this article I have shown that the details of imitative learning are less important. In particular, imitative learning that can be represented with a monotone selection dynamic will converge on rational expectations demands and has a unique asymptotically stable point at the GS equilibrium. However, the result is less robust to including the effects of experimentation. This is unfortunate because this implies that the important parameters of learning may be difficult to observe.

Whether or not the limiting behavior of the adaptive learning process is similar to that in the GS rational expectations equilibrium depends on the level of noise in the asset supply in the economy and the size of experimentation (captured by drift) in the learning process. There are two interpretations for the uncertainty in the asset supply. One view is that this noise is caused by irrational traders trading with biased views or based on irrelevant data. In our model, this view of supply noise is not justified since we have tried to incorporate "irrational" behavior in a learning process that is adaptive and not Bayesian. An alternative view of asset supply noise is that it is a reduced form for trades which arise endogenously in a more complex, but

[22] There can be multiple stationary points for the learning dynamic. In general, $h(\lambda)$ will intersect the G' boundary of $q \geq 1$ times, where q is odd and $(q+1)/2$ are the stable dynamic equilibria. In these cases, the Binmore–Samuelson techniques are less able to capture the effects of stochastic noise with deterministic drift.

[23] See Holland (1975) or Goldberg (1989) for details on genetic algorithms and Koza (1992) for genetic programming. Sargent (1993) provides an overview of genetic algorithms as models of learning in economic settings.

rational, economic model.[24] Since the asset supply noise parameter turns out to be crucial, the details of more complex models underlying supply noise are important to empirically test our adaptive learning model.

The adaptive learning model in the GS framework is well suited to laboratory experiments. Experimental asset markets have proven a useful tool for investigating rational expectations equilibria.[25] By controlling the information subjects observe about the behavior and success of others as well as the economic parameter of asset supply noise, insight can be gained into the appropriateness of modeling economies using adaptive learning and perhaps allow one to calibrate the learning process. This article focuses primarily on the limiting behavior of adaptive learning processes. However, we are able to identify some properties of the dynamic path. In particular we have shown that the proportion of informed traders will approach its equilibrium level from above. If one is willing to make additional assumptions about the learning process, more properties of the dynamic path can be developed.

I have focused on the situation where all traders follow the same learning algorithm. One interesting conjecture suggested by the model is that individuals who learn faster are more likely to be uninformed. This runs counter to the common view that sophisticated investors process more information. In a model of adaptive learning, more sophisticated traders are better able to make inferences from prices and therefore spend less resources on acquiring information. In order to fully investigate such a conjecture, a more detailed model of how individual adaptive learning algorithms differ is required. This article is a first step to exploring such questions by characterizing adaptive learning in a well-known rational expectations model.

Appendix

Proof of Lemma 1. There is a continumn of traders but only a finite number of types. Aggregate their demands in Equations (3) and (4) according to the proportions in ϕ_t and equate with the aggregate supply.

$$\lambda_t \left(\gamma^I (y_t - P_t)\right) + \sum_{\ell^n \in L^U} \phi_t^n \left(\gamma^U (\ell_0^n + \ell_1^n P_t - P_t)\right) = e_t. \tag{A.1}$$

Since the equation is linear, the market clearing price depends only on λ_t and the average inference parameters, $\ell_t = (\ell_{0t}, \ell_{1t})$ defined in Equation (5), and not the complete ϕ_t vector. Solving for P_t gives Equation (6). ■

Proof of Lemma 2. This formulation follows Bray (1982). By inspection of Equation (6), (λ, ℓ) is a sufficient statistic for ϕ. Joint normality of the random variables and the linear

[24] Black (1986) or De Long et al. (1990) view noise arising from irrational traders. Jackson (1991), Wang (1993), Bernardo and Judd (1995), and Paul (1995) take the opposing view.

[25] See Sunder (1995) for a survey of experimental asset markets.

price relation in Equation (6) imply that y, conditional on P, λ, and ℓ, is normal with $E[y \mid P, \lambda, \ell]$ and $V[y \mid P, \lambda, \ell]$ as follows [see Morrison (1990)]:

$$\begin{aligned} E[y \mid P, \lambda, \ell] &= E[y] + \frac{\sigma_{Py}}{\sigma_P^2}\,(P - E[P]) \\ &= -\frac{\lambda(1-\lambda)\gamma^I\gamma^U\sigma_y^2}{\left(\lambda\gamma^I\right)^2\sigma_y^2 + \sigma_e^2}\ell_0 \\ &\quad + \frac{\left(\lambda\gamma^I + (1-\lambda)\gamma^U(1-\ell_1)\right)\lambda\gamma^I\sigma_y^2}{\left(\lambda\gamma^I\right)^2\sigma_y^2 + \sigma_e^2}P. \end{aligned} \tag{A.2}$$

Equation (7) follows using the definitions in Equation (8) and simplifying the slope and intercept expressions. The conditional variance is

$$V[y \mid P, \lambda, \ell] = \sigma_y^2 - \frac{\left(\sigma_{Py}\right)^2}{\sigma_P^2} = \frac{\sigma_e^2\sigma_y^2}{\left(\lambda\gamma^I\right)^2\sigma_y^2 + \sigma_e^2}. \tag{A.3}$$

Preliminary to Lemma 3

Fitness is calculated by taking expectations iteratively. For informed traders $(L^n = I)$, fitness is

$$F^I = E\left[E\left[E\left[-\exp(-aW_1^n) \mid y, P, \phi\right] \mid P, \phi\right] \mid \phi\right]. \tag{A.4}$$

To calculate fitness of the uninformed traders conditional on the population ϕ, consider an agent with asset demands as in Equation (4) with arbitrary parameters $\ell^n = (\ell_0^n, \ell_1^n)$ (not necessarily on the finite grid L^U) and any arbitrary (positive) constant γ^U,

$$F(\ell^n, \gamma^U) = E\left[E\left[-\exp(-aW_1^n) \mid P, \phi\right] |\phi\right]. \tag{A.5}$$

Since the asset demands in Equation (4) are linear, we can calculate fitness for any constant γ^U as long as the ℓ^n parameters are appropriately adjusted for each n. From Equation (4), note that for all P,

$$x^n = \gamma^U\left(\ell_0^n + \ell_1^n P - P\right) = \hat{\gamma}^U\left(\hat{\ell}_0^n + \hat{\ell}_1^n P - P\right), \tag{A.6}$$

as long as for all n,

$$\hat{\ell}_0^n = \ell_0^n\frac{\gamma^U}{\hat{\gamma}^U} \quad \hat{\ell}_1^n = \ell_1^n\frac{\gamma^U}{\hat{\gamma}^U} - \frac{\gamma^U - \hat{\gamma}^U}{\hat{\gamma}^U}. \tag{A.7}$$

Since the demands are identical, there is no change in the economy from the normalization by a different γ^U. In particular,

$$F(\ell^n, \gamma^U) = F(\hat{\ell}^n, \hat{\gamma}^U). \tag{A.8}$$

Therefore I will calculate fitness for a conveniently chosen γ^U and then use Equations (A.7) and (A.8) to calculate fitness for the parameters in L^U and the given γ^U. Define $\gamma^*(\lambda)$ using the conditional variance relation in Equation (10):

$$\gamma^*(\lambda) = \left(a(\text{var}[y \mid P, \phi] + \sigma_z^2)\right)^{-1}. \tag{A.9}$$

Note that this expression depends on λ through the conditional variance of y. This is allowable since λ is part of the conditioning information in ϕ. Initially we calculate the fitnesses conditional on observing price. These are denoted $F^I(P)$ and $F^U(\ell^n, \gamma^*; P)$ for the informed and uninformed traders.

Lemma A.1. *Given a population state ϕ, the fitnesses conditional on P are as follows:*

$$F^I(P) = \exp(a\,c)\left(\frac{\sigma_z^2}{\sigma_z^2 + V[y \mid P]}\right)^{\frac{1}{2}} F^U(\ell^*, \gamma^* \mid P) \tag{A.10}$$

$$F^U(\ell^n, \gamma^*; P) = \exp\left(\frac{\zeta^n(P)^2}{2(\sigma_z^2 + V[y \mid P])}\right) F^U(\ell^*, \gamma^* \mid P), \tag{A.11}$$

where $\zeta^n(P) = (\ell_0^n - \ell_0^*) + (\ell_1^n - \ell_1^*)P$ [and ℓ^* are defined in Equation (7)].

Proof of Lemma A.1. The proof of the proposition proceeds by calculating three quantities: $F^U(\ell^n, \gamma^*; P)$, $F^U(\ell^*, \gamma^*; P)$, and $F^I(P)$. All expectations and variances are conditioning on ϕ.

(i) $F^U(\ell^n, \gamma^*; P)$
The risky asset demands of an uninformed trader from Equation (4) and the parameters ℓ^n and γ^* are

$$\begin{aligned} x^n &= \gamma^*\left(\ell_0^n + \ell_1^n P - P\right) \\ &= \frac{E[y \mid P] - P + \zeta^n(P)}{a(\sigma_z^2 + V[y \mid P])}. \end{aligned} \tag{A.12}$$

$E[y \mid P]$ and $V[y \mid P]$ are calculated in Lemma 2 and $\zeta^n(P)$ is defined above. End-of-period wealth for the uninformed trader is

$$W_1^n = W_0^n = \left(\frac{E[y \mid P] - P + \zeta^n(P)}{a(\sigma_z^2 + V[y \mid P])}\right)(y + z - P). \tag{A.13}$$

End-of-period wealth, conditional on P and ϕ, is normally distributed with conditional mean and variance of

$$E\left[W_1^n \mid P\right] = W_0^n + \frac{(E[y \mid P] - P + \zeta^n(P))(E[y \mid P] - P)}{a(\sigma_z^2 + V[y \mid P])} \tag{A.14}$$

$$V\left[W_1^n \mid P\right] = \frac{(E[y \mid P] - P + \zeta^n(P))^2}{a^2(\sigma_z^2 + V[y \mid P])} \tag{A.15}$$

These conditional moments determine the fitness (using the CARA preferences and normality).

$$\begin{aligned} F^U(\ell^n, \gamma^*; P) &= E\left[-\exp\left(-aW_1^n\right) \mid P\right] \\ &= -\exp\left(-a\left(E\left[W_1^n \mid P\right] - \frac{a}{2}V\left[W_1^n \mid P\right]\right)\right) \end{aligned}$$

$$= -\exp\left(-aW_0^n\right)\exp\left(\frac{(E[y \mid P]-P)^2}{2(\sigma_z^2+V[y \mid P])}\right) \times \exp\left(\frac{\zeta^n(P)^2}{2(\sigma_z^2+V[y \mid P])}\right) \quad \text{(A.16)}$$

(ii) $F^U(\ell^*, \gamma^*; P)$
When $\ell^n = \ell^*$ [defined in Equation (7)], $\zeta^n(P) = 0$ for all P. Therefore

$$F^U(\ell^*, \gamma^*; P) = -\exp\left(-aW_0^n\right)\exp\left(-\frac{(E[y \mid P]-P)^2}{2(\sigma_z^2+V[y \mid P])}\right) \quad \text{(A.17)}$$

and

$$F^U(\ell^n, \gamma^*; P) = \exp\left(\frac{\zeta^n(P)^2}{2(\sigma_z^2+V[y \mid P])}\right)F^U(\ell^*, \gamma^*; P) \quad \text{(A.18)}$$

(iii) $F^I(P)$
Following similar steps which lead to Equation (A.16), the informed trader's fitness conditional on $\{y, P, \phi\}$ is

$$E\left[-\exp\left(-aW_1^n\right) \mid y, P, \phi\right] = -\exp\left(-a\left(E\left[W_1^n \mid y, P,\right] - \frac{a}{2}V\left[W_1^n \mid y, P\right]\right)\right) = -\exp\left(-aW_0^n\right)\exp(a\,c)\exp\left(-\frac{(y-P)^2}{2\sigma_z^2}\right). \quad \text{(A.19)}$$

Taking expectations over y given P (and ϕ) yields $F^I(P)$:

$$F^I(P) = -\exp\left(-aW_0^n\right)\exp(a\,c)\left(\frac{\sigma_z^2}{\sigma_z^2+V[y \mid P]}\right)^{\frac{1}{2}} \times \exp\left(-\frac{(E[y \mid P, L]-P)^2}{2(\sigma_z^2+V[y \mid P])}\right) = \exp(a\,c)\left(\frac{\sigma_z^2}{\sigma_z^2+V[y \mid P]}\right)^{\frac{1}{2}}F^U(\ell^*, \gamma^*; P). \quad \text{(A.20)}$$

This completes the proof of Lemma A.1. ■

Proof of Lemma 3. The proof calculates the expectation over P of the expressions in Lemma A.1. First, by inspection of Equations (A.16) and (A.20), initial wealth, W_0, will have no impact on the relative utilities. Therefore, without loss of generality, set $W_0 = 0$. Define $X = \ell_0^* + (\ell_1^* - 1)P \sim \text{normal}(\mu_X, \sigma_X^2)$, and

$$A = \ell_0^n - \ell_0^* - \left(\frac{\ell_1^n - \ell_1^*}{\ell_1^* - 1}\right)\ell_0^*, \quad B = \frac{\ell_1^n - \ell_1^*}{\ell_1^* - 1}, \quad C = \frac{1}{\sigma_z^2 + V[y \mid P, \phi]},$$
$$D = \exp(a\,c)\left(C\sigma_z^2\right)^{\frac{1}{2}}. \quad \text{(A.21)}$$

A, B, C, and D do not depend on P. Write Equation (A.16) as

$$F^U(\ell^n, \gamma^*; P) = -\exp\left(-\frac{C}{2}\left(X^2 - (A + BX)^2\right)\right) \tag{A.22}$$

and Equation (A.20) as

$$F^I(P) = -D\exp\left(-\frac{C}{2}X^2\right). \tag{A.23}$$

X is a linear function of P and therefore normally distributed. Since P and X are perfectly correlated, $F^U(\ell^n, \gamma^*; P) = F^U(\ell^n, \gamma^*; X)$ and $F^I(P) = F^I(X)$. Calculating the expectation over X, yields

$$\begin{aligned} &E\left[-\exp\left(-\frac{C}{2}\left(X^2 - (A + BX)^2\right)\right)\right] \\ &\quad = -\left(1 + C\sigma_X^2(1 - B^2)\right)^{\frac{1}{2}} \\ &\qquad \times \exp\left(-\frac{C}{2}\left(\frac{-A^2 + \mu_X^2 - B^2\mu_X^2 - CA^2\sigma_x^2 - 2AB\mu_X}{1 + C\sigma_X^2(1 - B^2)}\right)\right). \end{aligned} \tag{A.24}$$

In calculating this expression, one can verify that the expectation is finite only if $1 + C(1 - B^2) > 0$. Equation (A.24) is $E[F^U(\ell^n, \gamma^*; X)]$. Multiplying Equation (A.24) by D and setting $A = B = 0$ yields $E[F^I(X)]$. Normalizing $F^I = 1$ and accounting for the fact that CARA utilities are negative gives

$$\begin{aligned} \hat{f}^U(\ell^n, \gamma^*) &= \frac{E[F^I(X)]}{E[f^U(\ell^n, \gamma^* \mid X)]} \\ &= \exp(ac)\left(\frac{\sigma_z^2}{\sigma_z^2 + V[y \mid P, \phi]}\right)^{\frac{1}{2}} \hat{\xi}(\ell^n, \lambda, \ell) \end{aligned} \tag{A.25}$$

$$\hat{\xi}(\ell^n, \lambda, \ell) = \begin{cases} \left(\frac{1 + C\sigma_X^2(1 - B^2)}{1 + C\sigma_X^2}\right)^{\frac{1}{2}} & \\ \quad \cdot \exp\left(-\frac{C[A(1 + C\sigma_X^2) + B\mu_X]^2}{2(1 + C\sigma_X^2)(1 + C\sigma_X^2(1 - B^2))}\right) & \text{if } 1 + C\sigma_X^2(1 - B^2) > 0 \\ 0 & \text{otherwise} \end{cases} \tag{A.26}$$

If $E[F^U(\ell^n, \gamma^*; X)]$ is not finite, then $\hat{f}^U(\ell^n, \gamma^*) = 0$. To complete the proof, we use the relationship in Equations (A.7) and (A.8) to calculate $\xi(\ell^n, \lambda, \ell) = \hat{\xi}(\hat{\ell}^n, \lambda, \ell)$ for $\ell_n \in L^U$ and the desired constant γ^U to get $f^n(\lambda, \ell)$. ■

Proof of Lemma 4. Linear asset demands of the other traders imply the linear market clearing price relation in Equation (6). Therefore, conditional on observing P and knowing the population state ϕ, end-of-period wealth W_1^n is normally distributed. Therefore the asset demands $x^*(P)$, which maximize the innermost expectation in Equation (A.5), are

$$x^*(P) = \frac{E[y \mid P, \phi] - P}{a(V[y \mid P, \phi] + \sigma_z^2)}, \tag{A.27}$$

where $E[y \mid P, \phi]$ and $V[y \mid P, \phi]$ are linear in P (see Lemma 2). Finally,

$$\underset{x(P)}{\operatorname{argmax}} \left(E\left[-\exp\left(-aW_1^n\right) \mid P, \phi\right]\right) = \underset{x(P)}{\operatorname{argmax}} \left(E\left[-\exp\left(-aW_1^n\right) \mid \phi\right]\right). \quad \text{(A.28)}$$

Therefore the fitness maximizing demand parameters are chosen to give Equation (A.27). Use ℓ^* from Equation (7) and γ^* from Equation (A.9) and the adjustment in Equation (A.7). ■

Proof of Proposition 1. Define $S = \{\phi \in S^{N+1} \mid \phi^0 = \lambda\}$ for a fixed λ. We need to show for any $\phi_0 \in interior(S)$, if $\hat{k}(\lambda) < 1$, then given an $\varepsilon > 0$, there exists an N' and T such that for all $N > N'$ (sufficiently fine grid) and $t > T$ (long enough horizon), $\phi_t^n < \varepsilon$ for all n such that $\|\ell^n - \hat{\ell}^{**}(\lambda)\| > \varepsilon$. To begin, define the following sequence of sets:

$$\begin{aligned} L_1 &= \left\{\ell^n \in L^U \mid \exists \phi \in S \ s.t. \ \|\ell^n - \hat{\ell}^*(\phi)\| < \varepsilon\right\} \\ L_{j+1} &= \left\{\ell^n \in L^U \mid \exists \phi \in S \ with \ \phi^m = 0 \right. \\ &\qquad \left. for \ all \ \ell^m \notin L_j \ s.t. \ \|\ell^n - \hat{\ell}^*(\phi)\| < \varepsilon\right\}. \end{aligned} \quad \text{(A.29)}$$

Since $\hat{\ell}^{**}(\lambda) \in C(L^U)$, there exists a sufficiently fine grid (large enough $N > N'$) such that $\|\ell^n - \hat{\ell}^{**}(\lambda)\| < \varepsilon$ for some $\ell^n \in L^U$. We can write Equation (13) as

$$\hat{\ell}^*(\lambda, \ell) = \left(1 + \hat{k}(\lambda)\right) \hat{\ell}^{**}(\lambda) - \hat{k}(\lambda)\ell, \quad \text{(A.30)}$$

where $\ell = \Sigma \phi^n \ell^n$ and $\hat{k} < 1$. Using this, it is easy to verify the following:

1. $\hat{\ell}^{**}(\lambda) \in C(L_j)$ for all j.
2. $L_{j+1} \subseteq L_j$ for all j.
3. There exists a finite J such that $L_{J+1} = L_J$.
4. $\ell^n \in L_J$ implies $\|\ell^n - \hat{\ell}^{**}(\lambda)\| < \varepsilon$.

To complete the proof we need to show that for any monotone selection dynamic, $\phi_t^n \to 0$ for all $\ell^n \notin L_J$. We use the arguments of Samuelson and Zhang (1992). Let M be the set of demand parameters that are not in L^J but do not vanish under a monotone selection dynamic; that is, $M = \{\ell^n \mid \lim_{t\to\infty} \phi_t^n > 0 \text{ and } \ell^n \notin L_J\}$. We assume $M \neq \varnothing$ and derive a contradiction. Consider the smallest j such that there exists $\ell^m \in M$ such that $\ell^m \in L_{j-1}$ and $\ell^m \notin L_j$. Let $\ell^{m'} = \operatorname{argmin}\{\|\ell^m - \ell^n\| \mid \ell^n \in L_j\}$. For all $\phi \in S$ such that $\phi^n = 0$ for all $\ell^n \notin L_j \cup \{\ell^m\}$, $f^{m'}(\phi) > f^m(\phi)$. By construction, $\lim_{t\to\infty} \phi_t^n = 0$ for all $\ell^n \notin L_j \cup \{\ell^m\}$. Therefore there exists a T such that for all $t > T$, $f^{m'}(\phi_t) > f^m(\phi_t)$ and by the monotone selection dynamic there is a $\delta > 0$ such that

$$\frac{d}{dt}\left[\frac{\phi^m}{\phi^{m'}}\right] = \left(g^m(\phi_t) - g^{m'}(\phi_t)\right) \frac{\phi^m}{\phi^{m'}} < -\delta \frac{\phi^m}{\phi^{m'}}. \quad \text{(A.31)}$$

This implies that $lim_{t\to\infty} \phi_t^m = 0$, which is a contradiction of the assumption that $\ell^m \in M$. Therefore $M = \varnothing$ and $\phi_t^n \to 0$ for all $\ell^n \notin L_J$. ■

Proof of Proposition 2.
(i) ϕ^* is asymptotically stable for any monotone selection dynamic:
Without loss of generality, reorder the set L^U so that $\phi^* = (\lambda^*, 1 - \lambda^*, 0, \ldots, 0)$; and

let $\ell^{n*} = \ell^{**}(\lambda^*)$. By construction of the GS equilibrium, $f^I(\phi^*) = f^{n*}(\phi^*) > f^n(\phi^*)$ for $n \neq n^*$. Monotonicty of the selection dynamic implies $g^I = g^{n*}$ and by the simplex restriction $g^I = g^{n*} = 0$. Thus ϕ^* is a stationary point of Equation (15). To show that it is also asymptotically stable, consider the neighborhood of ϕ^* defined by $S^* = \{(\lambda^* + \delta, 1 - \lambda^* - \delta - \varepsilon, \varepsilon^2, \ldots, \varepsilon^N) \mid \varepsilon^n > 0, \varepsilon = \Sigma\varepsilon^n, -\Delta < \delta < \Delta, \text{and } \Delta > 0\}$. Since $f^n(\phi)$ is continuous for all $\phi \in S^*$ (for $f^n(\phi) > 1$), $f^{n*}(\phi) > f^n(\phi)$, and by monotonicity $g^{n*}(\phi) > g^n(\phi)$. (Note that this does not require $\hat{k} < 1$). Therefore, if $\phi_t \in S^*$ for all t, then $\phi_t^n \to 0$. We now need to show that the monotonic selection dynamic implies $\lambda_t \in (\lambda^* - \Delta, \lambda^* + \Delta)$. Since ℓ^{n*} maximizes fitness (Lemma 4), we can apply the envelope theorem to calculate that $df^{n*}(\phi^*)/d\lambda > 0$ [using the fact that $dV[y \mid P, \phi]/d\lambda < 0$ from see (10)]. Therefore, for any $(\varepsilon^2, \ldots, \varepsilon^N)$, there exists a $\bar{\lambda} \in (\lambda^* - \Delta, \lambda^* + \Delta)$ such that for $\phi_t \in S^*$, $f^I(\phi) > f^{n*}(\phi)$ if and only if $\lambda_t < \bar{\lambda}$, and therefore $g^I(\phi) - g^{n*}(\phi) > 0$ if and only if $\lambda_t < \bar{\lambda}$ and therefore $\lambda_t \in (\lambda^* - \Delta, \lambda^* + \Delta)$. Finally, since ℓ^{n*} maximizes fitness and $f^I(\phi) = f^{n*}(\phi)$ only at $\phi = \phi^*$ (i.e., $\lambda = \lambda^*$), if $\varepsilon^n = 0$, then $\bar{\lambda} = \lambda^*$ and if $\varepsilon^n > 0$, then $\bar{\lambda} > \lambda^*$. Therefore ϕ^* is asymptotically stable and λ_t will converge to λ^* from above.

(ii) Uniqueness: $\phi \neq \phi^*$ is not asymptotically stable:
By construction, ϕ^* is the only point where $\lambda \in (0, 1)$ and $f^I(\phi) = f^n(\phi)$ for all n such that $\phi^n > 0$. Therefore all other stationary points of Equation (15) have $\lambda = 1$ or $\lambda = 0$. It is straightforward to show that each of these is not stable. Consider the case ϕ', where $\lambda' = 1$. Since $\hat{\ell}^{**}(\lambda) \in C(L^U)$ for all λ, as long as the grid of possible demand parameters is sufficiently fine (large enough N), there exists an ℓ^n close to $\hat{\ell}^{**}(1)$ such that $f^n(\phi') > f^I(\phi')$. By the continuity of f, there exists a neighborhood of ϕ', denoted S', such that for all $\phi \in S'$ $f^n(\phi) > f^I(\phi)$, which by monotonicity implies that $g^n(\phi) > g^I(\phi)$. If ϕ' were stable, then $\phi_t \in S'$ for all t. But in $\phi \in S'$ implies $g^n(\phi) > g^I(\phi)$ and $\lambda_t \to 0$, which contradicts ϕ' being stable. An identical argument rules out stationary points with $\lambda = 0$. ■

Preliminary to Proposition 3
Lemma A.2. *For a given economy, E, where $\sigma_e^2 > 0$, the set $G(E) = \{(\lambda, \ell) \mid f^U(\lambda, \ell) \geq 1\}$ has the following properties:*

- *(a) $G(E)$ is closed.*
- *(b) $G(E)$ is not empty.*
- *(c) For $\lambda > \lambda^*$, there is a neighborhood of $(\lambda, \hat{\ell}^{**}(\lambda))$ that is in $G(E)$.*
- *(d) The only $(\lambda, \ell) \in G(E)$ with $\lambda \leq \lambda^*$ is $(\lambda^*, \hat{\ell}^{**}(\lambda^*))$.*

Proof of Lemma A.2. (a) f^U is continuous in λ and ℓ in a neighborhood of $f^U = 1$ and the set $\{f^U \geq 1\}$ is closed. $G(E)$ is closed (i.e., inverse image of a closed set for a continuous function is closed).

(b) By the definition of the GS equilibrium, $(\lambda^*, \hat{\ell}^{**}(\lambda^*)) \in G(E)$. Therefore $G(E)$ is not empty.

(c) At $\hat{\ell}^{**}(\lambda)$, uninformed agents are making no inference errors (by definition) and since $\lambda > \lambda^*$, $f^U(\lambda, \hat{\ell}^{**}(\lambda)) > 1$. Since f^U is jointly continuous in λ and in the inference parameters, there exists $\bar{\varepsilon}$ such that for all $(\varepsilon_\lambda, \varepsilon_\ell) \in \mathbb{R}^3$, $f^U(\lambda + \varepsilon_\lambda, \hat{\ell}^{**}(\lambda) + \varepsilon_\ell) > 1$ for $\lambda > \lambda^*$ and $\|(\varepsilon_\lambda, \varepsilon_\ell)\| < \bar{\varepsilon}$.

(d) By construction of the GS equilibrium, $f^U(\lambda^*, \hat{\ell}^{**}(\lambda^*)) = 1$. Since inference errors reduce fitness (Lemma 4) for all ε (with $\|\varepsilon\| > 0$), $f^U(\lambda^*, \hat{\ell}^{**}(\lambda^*) + \varepsilon) < 1$. Again, by construction of the GS equilibrium, for $\lambda < \lambda^*$, $f^U(\lambda, \ell) \leq f^U(\lambda, \hat{\ell}^{**}(\lambda)) < 1$. ■

Proof of Proposition 3. (a) First consider possible dynamic equilibria such that $\lambda \in (0, 1)$. By definition, $g_\lambda \leq 0 \leftrightarrow (\lambda, \ell) \in G(E)$. For $\lambda \in (0, 1)$, $g_\lambda = 0$ if and only if (λ, ℓ) is on the border of G (denoted ∂G) and $g_\ell = 0 \leftrightarrow \ell = \hat{\ell}^{**}(\lambda)$. From Lemma A.2(c) and (d), the only intersection of the ∂G with $\hat{\ell}^{**}(\lambda)$ is at λ^*. Therefore, for $\lambda \in (0, 1)$, the unique dynamic equilibrium is $(\lambda^*, \hat{\ell}^{**}(\lambda^*))$. To verify that this equilibrium is stable, note that $g_\lambda > 0$ for $\lambda < \lambda^*$ (Lemma 5(d)) and $g_\lambda < 0$ in a neighborhood of $\hat{\ell}^{**}(\lambda)$ for $\lambda > \lambda^*$ (Lemma 5(c)). This, coupled with Assumption 1(b) on g_ℓ, establishes stability.

(b) Next consider $\lambda = 1$. By the assumptions about the behavior of g_λ and g_ℓ as $\lambda \to 1$, all points $(1, \ell) \notin G$ are fixed points of the learning dynamic. If $(1, \ell) \in G$ is a fixed point, it is not stable since $g_\lambda(1 - \varepsilon, \ell) < 0$ (for $\varepsilon > 0$). For $(1, \ell) \notin G$, $g_\lambda(1 - \varepsilon, \lambda) > 0$. Thus for small enough ε and $\lambda_0 = 1 - \varepsilon$, $\lim_{t\to\infty} \lambda_t = 1$. By assumption, $\|g_\ell\|$ is small for λ close to 1. Thus for $\ell_0 = \ell$, $\lim_{t\to\infty} \ell_t$ can be made arbitrarily close to ℓ. Thus $(1, \lambda) \notin G$ is stable. (Note the points $(1, \lambda) \notin G$ are not asymptotically stable since the equilibria $(1, \ell)$ are connected).

(c) Finally consider $\lambda = 0$. Since $g_\lambda(\varepsilon, \lambda) > 0$ (for small positive ε), any possible dynamic equilibrium with $\lambda = 0$ is not stable. ■

Proof of Lemma 5. (a) By definition, $(\lambda', \ell') \in G'$ implies that $g_\lambda(\lambda', \ell') + \eta_\lambda m_\lambda(\lambda', \ell') \leq 0$. If $g_\lambda(\lambda', \ell') \leq 0$ then $(\lambda', \ell') \in G$. On the other hand, if $g_\lambda(\lambda', \ell') > 0$ then $0 < g_\lambda(\lambda', \ell') < -\eta_\lambda m_\lambda(\lambda', \ell')$. By choosing small enough η, $g_\lambda(\lambda', \ell')$ is within Δ' of $g_\lambda(\lambda, \ell) = 0$. By the continuity of g_λ, (λ', ℓ') is with distance Δ of (λ, ℓ). Alternatively, $(\lambda, \ell) \in G$ implies that $g_\lambda(\lambda, \ell) \leq 0$. If $g_\lambda(\lambda, \ell) + \eta_\lambda m_\lambda(\lambda, \ell) \leq 0$, then $(\lambda, \ell) \in G'$. If not, then $0 < g_\lambda(\lambda, \ell) + \eta_\lambda m_\lambda(\lambda, \ell) < \eta_\lambda m_\lambda(\lambda, \ell)$. Thus for small enough η, by the continuity of g_λ and m_λ, there exists (λ', ℓ') within distance Δ of (λ, ℓ) such that $g_\lambda(\lambda', \ell') + \eta_\lambda m_\lambda(\lambda', \ell') \leq 0$.

(b) Define $M(\lambda) = \sup_{\{\ell\}}\{\|m_\ell(\lambda, \ell)\|\}$. Since m_ℓ is bounded, $M(\lambda) < \infty$. For any $\lambda \in (0, \lambda'')$, $g_\ell(\lambda, \hat{\ell}^{**}(\lambda)) = 0$ and $g_\ell(\lambda, h(\lambda)) + \eta_\ell m_\ell(\lambda, h(\lambda)) = 0$. Therefore $\|g_\ell(\lambda, \hat{\ell}^{**}) - g_\ell(\lambda, h(\lambda))\| = \|\eta_\ell m_\ell(\lambda, h(\lambda))\| < \eta_\ell M(\lambda)$. By choice of η, $\eta_\ell M(\lambda)$ can be made arbitrarily small. By the continuity of g_ℓ, $\|h(\lambda) - \hat{\ell}^{**}(\lambda)\| < \Delta$.

(c) For $\lambda < 1$, $g_\ell(\lambda, \ell) = 0$ if and only if $\ell = \hat{\ell}^{**}(\lambda)$ and $m_\ell(\lambda, \ell) = 0$ if and only if $\ell = \ell_m$. However, by assumption, $\ell_m \neq \ell^{**}(\lambda)$. Thus $h(\lambda) \neq \hat{\ell}^{**}(\lambda)$. $\lim_{\lambda\to 1} h(\lambda) = \ell_m$ follows from the assumptions that $\lim_{\lambda\to 1} g_\ell(\lambda, \ell) = 0$ for all ℓ and $m_\ell(\lambda, \ell) = 0$ only at $\ell = \ell_m$. ■

Prelude to Proof of Proposition 4

Since the learning dynamic in Equations (20) and (21) is an autonomous (i.e., time independent) system of equations, we can define an invariant set around a stable dynamic equilibrium. For $x \in \mathbb{R}^n$, let $dx/dt = g(x)$ (where $g(x)$ is Lipshitz continuous) and $L(x_0, t)$ be the unique continuously differentiable solution, the set V is an *invariant set* if all trajectories which begin in V remain in V. If x^* is a stable dynamic equilibrium, then for every neighborhood Q of x^*, there exists an $R \subset Q$ such that $V = \{L(x_0, t) \mid x_0 \in R, t \geq 0\} \subset Q$ and V is an invariant set (proof omitted).

Proof of Proposition 4. The proof shows the intersection of $h(\lambda)$ (where $d\ell/dt = 0$) and the boundary of G', denoted $\partial G'$, and where $d\lambda/dt = 0$ is nonempty and that at least one of the intersections is stable.

For small $\varepsilon > 0$, $(\varepsilon, h(\varepsilon)) \notin G'$ and $(1 - \varepsilon, h(1 - \varepsilon)) \in G'$. Therefore, the infemum of

$$(\lambda', \ell') = \inf_{\lambda}\{(\lambda, h(\lambda)) \in G'\} \tag{A.32}$$

has $0 < \lambda' < 1$. Since $h(\lambda)$ is continuous, $(\lambda', h(\lambda')) \in \partial G'$. Therefore (λ', ℓ') is a dynamic equilibrium. It remains to be shown that this point is stable.

To show that (λ', ℓ') is stable, consider any arbitrary neighborhood $Q' \in \mathbb{R}^3$ of (λ', ℓ'). The following will construct the neighborhood $R' \subset Q'$ with the property that all trajectories originating in R' must remain in Q'. To begin, without loss of generality, let $Q' = (\lambda' - \Delta, \lambda' + \Delta) \times Q$ where $\Delta > 0$ and $Q \in \mathbb{R}^2$ is a neighborhood of ℓ'. The construction of R' is such that $R' = \Lambda \times R$ where $\Lambda = (\lambda' - \varepsilon, \lambda' + \varepsilon)$ with $0 < \varepsilon < \frac{1}{2}\Delta$ and R is a neighborhood of ℓ' with $R \subset Q$.

Let the solution to $d\ell/dt = g_\ell(\lambda, \ell) + \eta_\ell m_\ell(\lambda, \ell)$ for a fixed and constant λ be denoted $\ell_t = L(\lambda, \ell_0, t)$. L is continuously differentiable in each of its parameters. Recall that the unique asymptotically stable fixed point of this system is given by the continuous function $h(\lambda)$. Since $h(\lambda)$ is continuous, there exists an $\varepsilon > 0$ such that $h(\lambda) \in R$ for all $\lambda \in \Lambda$. Furthermore, since each $h(\lambda)$ is also a stable dynamic equilibrium (when λ is held constant), there exist sets $U(\lambda)$, neighborhoods of $h(\lambda)$, such that $\ell_0 \in U(\lambda)$ implies that $L(\lambda, \ell_0, t) \in Q$ for all $t \geq 0$. Since each $U(\lambda)$ is open, $h(\lambda)$ is continuous, and $\varepsilon > 0$ can be chosen arbitrarily small, there exists an open set U [a neighborhood of $\ell' = h(\lambda')$] such that for all $\lambda \in \Lambda$, $h(\lambda) \in U$ and $\ell_0 \in U$ implies that $L(\lambda, \ell_0, t) \in R$ for all $t \geq 0$. For example, U could be the intersection of the appropriately chosen $U(\lambda)$.

For each $\lambda \in \Lambda$, an invariant set $V(\lambda) \subset R$ can be constructed as in Lemma A.3, $V(\lambda) = \{L(\lambda, \ell_0, t) \mid \ell_0 \in U, t \geq 0\}$. Consider the intersection between any two of these sets $V(\lambda) \cap V(\lambda + \varepsilon)$. The following will show that this intersection is an invariant set under both $L(\lambda, \cdot, \cdot)$ and $L(\lambda + \varepsilon, \cdot, \cdot)$.

Consider $\ell \in V(\lambda) \cap V(\lambda + \varepsilon)$. It follows:

(i) By definition $\ell \in V(\lambda)$ implies there exists $\ell_0 \in U$ and τ such that $\ell = L(\lambda, \ell_0, \tau)$.
(ii) Similarly, $\ell \in V(\lambda + \varepsilon)$ implies there exists $\ell_{0\varepsilon} \in U$ and τ_ε such that $\ell = L(\lambda + \varepsilon, \ell_{0\varepsilon}, \tau_\varepsilon)$.
(iii) Since $V(\lambda)$ is an invariant set, $L(\lambda, \ell_0, t) \in V(\lambda)$ for all $t \geq 0$.
(iv) Similarly, since $V(\lambda + \varepsilon)$ is an invariant set, $L(\lambda + \varepsilon, \ell_{0\varepsilon}, t) \in V(\lambda + \varepsilon)$ for all $t \geq 0$.
(v) Since L is (continuously) differentiable, the implicit function theorem can be applied[26] and there exists $\zeta(t)$ such that $L(\lambda, \ell_0, t) = L(\lambda + \varepsilon, \ell_0 + \zeta, t)$. By choice of small ε, $\|\zeta(t)\|$ can be made arbitrarily small.[27] Since U is open, $\ell + \zeta(t) \in U$. Since this is true for any t, $L(\lambda, \ell_0, t) \in V(\lambda + \varepsilon)$ for $t \geq 0$.
(vi) Similarly, there exists a small $\zeta_\varepsilon(t)$ such that $L(\lambda + \varepsilon, \ell_{0\varepsilon}, t) = L(\lambda, \ell_{0\varepsilon} + \zeta, t)$ and $\ell_{0\varepsilon} + \zeta_\varepsilon \in U$. Therefore $L(\lambda + \varepsilon, \ell_{0\varepsilon}, t) \in V(\lambda)$ for all $t \geq 0$.

Taken together, these points imply that $L(\lambda, \ell_0, t) \in V(\lambda) \cap V(\lambda + \varepsilon)$ and $L(\lambda + \varepsilon, \ell_{0\varepsilon}, t) \in V(\lambda) \cap V(\lambda + \varepsilon)$. The intersection of the invariant sets is also invariant. Let V be the intersection of all the invariant sets. That is,

$$V = \bigcap_{\lambda \in \Lambda} V(\lambda) \tag{A.33}$$

Since *(i)*–*(vi)* are true for the intersection of two arbitrary sets which differ in the parameter by ε, it follows that for all $\lambda \in \Lambda$, $\ell \in V$ implies that $L(\lambda, \ell, t) \in V$ for all $t \geq 0$.

[26] The implicit function theorem applies so long as the determinant Jacobean of L, $|\nabla L(\lambda, h(\lambda), t)|$ is nonzero. This follows from the fact that $h(\lambda)$ is the unique fixed point and that it is asymptotically stable.

[27] Note that since the $h(\lambda)$ are asymptotically stable, the set U can be chosen to ensure that trajectories converge. Therefore $\lim_{t \to \infty} \|\zeta(t)\| \approx \|h(\lambda) - h(\lambda + \varepsilon)\|$, which is small by the continuity of $h(\lambda)$.

We have constructed an invariant set V (such that $V \subset R \subset Q$) for the inference parameters. As long as $\lambda_t \in \Lambda$, then $\ell_t \in V$. It remains to be shown that R and Λ (i.e., choice of ε) can be chosen to ensure that $\lambda_t \in \Lambda$. By construction, the dynamic equilibrium (λ', ℓ') is on the boundary of the closed set G', such that $(\lambda' - \varepsilon, \ell') \notin G'$ and $(\lambda' + \varepsilon, \ell')$ is an interior point of G'. Therefore, for a small enough neighborhood R of ℓ', $\ell \in R$ and $\lambda' - \varepsilon < \lambda < \lambda'$ implies $(\lambda, \ell) \notin G'$ and $d\lambda/dt > 0$. Similarly, for $\ell \in R$ and $\lambda' \leq \lambda < \lambda' + \varepsilon$ implies $(\lambda, \ell) \in G'$ and $d\lambda/dt \leq 0$.

Therefore, for any neighborhood Q' of (λ', ℓ'), there exists the neighborhood $\Lambda \times U$ such that $(\lambda_0, \ell_0) \in \Lambda \times U$ implies $(\lambda_t, \ell_t) \in \Lambda \times V$ and by construction $\Lambda \times V \subset (\lambda' - \Delta, \lambda' + \Delta) \times Q = Q'$. Therefore (λ', ℓ') is a stable dynamic equilibrium of the learning process. ∎

Proof of Proposition 5. Consider the learning process $\mathcal{L}(\lambda'', \Delta'', \hat{\ell}''_m)$ with $\lambda'' = 1 - \Delta'' > \lambda^*$ and $\ell''_m = \hat{\ell}^{**}(1) - (\Delta'', \Delta'')$.[28] For small Δ'' and $\lambda \geq \lambda''$, $(\lambda, h(\lambda))$ is an interior point of G. Therefore there are no dynamic equilibrium of $\mathcal{L}$ with $\lambda \geq \lambda''$ (i.e., no intersections of $h(\lambda)$ and $\partial G'$). Applying Proposition 2, there exists at least one dynamic equilibrium with $\lambda < \lambda''$. Recall that the intersection of $\hat{\ell}^{**}(\lambda)$ and ∂G is unique and is the GS equilibrium (Proposition 3). Since for $\lambda < \lambda''$, $\|h(\lambda) - \hat{\ell}^{**}(\lambda)\| < \Delta''$ [see Lemma 5(b)], Δ'' can be chosen small enough to ensure that $\sup\{\|(\lambda, \ell) - (\lambda^*, \hat{\ell}^{**}(\lambda^*))\| \mid (\lambda, \ell) \in \partial G \cap h(\lambda)\} < \varepsilon$. Lemma 5(a) implies that for $\lambda < \lambda''$ and $(\lambda', \ell') \in \partial G'$, $\inf\{\|(\lambda', \ell') - (\lambda, \ell)\| \mid (\lambda, \ell) \in \partial G\} < \Delta''$ (i.e., the boundary of G is close to the boundary of G'). Therefore, by the triangle inequality, if (λ, ℓ) is a dynamic equilibrium (i.e., in $h(\lambda) \cap \partial G'$) then $\|(\lambda, \ell) - (\lambda^*, \hat{\ell}^{**}(\lambda^*))\| < \Delta'' + \varepsilon < \Delta$ for small enough Δ''. Therefore all the dynamic equilibria of $\mathcal{L}(\lambda'', \Delta'', \ell''_m)$ are within Δ of the GS equilibrium. The same is also true of learning processes with less drift, namely, $\lambda' \geq \lambda''$, $\Delta' \leq \Delta''$, and $\|\ell'_m - \lambda^{**}(1)\| \leq \|\ell''_m - \hat{\ell}^{**}(1)\|$. ∎

Prelude to Proof of Proposition 6

The following function measures the size of the inference errors that still leave the uninformed trader at least as well off as the informed trader. This function characterizes the size of $G(E)$. For $\lambda \geq \lambda^*$, define $\alpha_\lambda(G(E))$ as follows,

$$\alpha_\lambda(G(E)) = \sup_{\{\ell, \ell'\}} \left\{ \|\ell - \ell'\| \mid (\lambda, \ell), (\lambda, \ell') \in G(E) \right\}. \tag{A.34}$$

Lemma A.4. *For economies, $E = (\sigma_e^2, \lambda^*)$, with fixed λ^*, for all $\lambda \geq \lambda^*$,*

$$\lim_{\sigma_e^2 \to 0} \alpha_\lambda(G(E)) = 0. \tag{A.35}$$

Proof of Lemma A.4. For $\lambda \in [\lambda^*, 1]$, define the correspondence $G_\lambda(E)$ as

$$G_\lambda(E) = \{\ell \mid (\lambda, \ell) \in G(E)\} = \{\ell \mid f^u(\lambda, \ell; E) \geq 1\}. \tag{A.36}$$

By Lemma A.2(b), $G_\lambda(E)$ is nonempty. Since f^U is jointly continuous in σ_e^2 and c (which is implied by λ^*) and the set $\{f^U \geq 1\}$ is closed, the graph of the correspondence $G_\lambda(E)$ is closed. Therefore $G_\lambda(E)$ is upper hemicontinuous.

[28] This assumes that for all λ, $\hat{\ell}^{**}(\lambda) \neq \hat{\ell}^{**}(1) - (\Delta'', \Delta'')$. This is without loss of generality since choice of ℓ''_m can be perturbed slightly if required.

In the economy where $\sigma_e^2 = 0$, that is, $E = (0, \lambda^*)$, $f^U(\lambda, \ell; E) \leq 1$ and equals 1 only at $\ell = \hat{\ell}^{**}(\lambda)$. Therefore, $G_\lambda(0, \lambda^*)$ is the singleton $\{\hat{\ell}^{**}(\lambda)\}$. Consider a sequence E_t, with $\sigma_e^2(t) > 0$ and constant λ^*, such that $E_t \to (0, \lambda^*)$. Take any two sequences $\ell_t \to \ell$ and $\ell'_t \to \ell'$ such that $\ell_t, \ell'_t \in G_\lambda(E_t)$. Since G_λ is upper hemicontinuous, $\ell, \ell' \in G_\lambda(0, \lambda^*)$. But since $G_\lambda(0, \lambda^*)$ is a singleton, $\ell = \ell' = \hat{\ell}^{**}(\lambda)$ and $\|\ell_t - \hat{\ell}^{**}(\lambda)\| + \|\ell'_t - \hat{\ell}^{**}(\lambda)\| \to 0$. By the triangle inequality $\|\ell_t - \ell'_t\| \to 0$. Since this is true for any two sequences, $\alpha_\lambda(E_t) \to 0$. ■

Proof of Proposition 6. It is sufficient to show that for small σ_e^2, $\partial G'$ and $h(\lambda)$ do not intersect for $\lambda < \lambda''$. From Lemma A.4, $G(E)$ approaches the set $\{(\lambda, \hat{\ell}^{**}(\lambda)) \mid \lambda > \lambda^*\}$ as σ_e^2 goes to zero (holding λ^* constant). From Lemma 5(c), $h(\lambda) \neq \hat{\ell}^{**}(\lambda)$. Therefore there exists $\sigma_e^{2''}$ such that for all $\sigma_e^2 < \sigma_e^{2''}$, $h(\lambda)$ and ∂G do not intersect. The fact that $h(\lambda)$ and $\partial G'$ do not intersect for $\lambda < \lambda''$ in economies with $\sigma_e < \sigma_e''$ follows from the property that λ drift does not exceed inference drift. That is, the property that

$$\inf_{\lambda<\lambda''}\{\|h(\lambda)-\ell^{**}(\lambda)\|\} > \sup_{\lambda<\lambda''}\{\|(\lambda, \ell)-(\lambda, \ell')\| \mid (\lambda, \ell) \in \partial G, (\lambda, \ell') \in \partial G'\}. \quad \text{(A.37)}$$

Therefore there are no equilibria with $\lambda < \lambda''$. Any equilibria (at least one exists), must have the property that $\lambda > \lambda''$. ■

References

Abel, A. B., 1990, "Asset Prices Under Habit Formation and Catching Up With the Joneses," *American Economic Review*, 80(2), 38–42.

Benartzi, S., and R. H. Thaler, 1995, "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics*, 107, 797–817.

Bernardo, A. E., and K. L. Judd, 1995, "The Relationship Between Price Changes and Volume," working paper, UCLA.

Biais, B., and R. Shadur, 1998, "Darwinian Selection does not Eliminate Irrational Traders," working paper, Université de Toulouse; forthcoming in *European Economic Review*.

Binmore K. G., and L. Samuelson, 1999, "Evolutionary Drift and Equilibrium Selection," *Review of Economic Studies*, 66, 363–393.

Björnerstedt, J., and J. Weibull, 1996, "Nash Equilibrium and Evolution by Imitation," in K. J. Arrow, E. Colombatto, and M. Perlman (eds.), *The Rational Foundations of Economic Behaviour: Proceedings of the IEA Conference Held in Turin, Italy*, MacMillian Press, New York.

Black, F., 1986, "Noise," *Journal of Finance*, 41, 529–543.

Blume, L. E., M. Bray, and D. Easley, 1982, "Introduction to the Stability of Rational Expectations Equilibrium," *Journal of Economic Theory*, 26, 313–317.

Blume, L., and D. Easley, 1992, "Evolution and Market Behavior," *Journal of Economic Theory*, 58, 9–40.

Blume, L., and D. Easley, 1993, "Rational Expectations and Rational Learning," working paper, Cornell University.

Bray, M., 1982, "Learning, Estimation and the Stability of Rational Expectations," *Journal of Economic Theory*, 26, 318–339.

Bray, M., and D. Kreps, 1987, "Rational Learning and Rational Expectations," in G. Feiwel (ed.), *Arrow and the Accent of Modern Economic Theory*, New York University Press, New York.

De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann, 1990, "Noise Trader Risk in Financial Markets," *Journal of Political Economy*, 98, 703–737.

Evans, G. W., and S. Honkapohja, 1998, "Economic Dynamics with Learning: New Stability Results," *Review of Economic Studies*, 65, 23–44.

Friedman, M., 1953, *Essays in Positive Economics*, University of Chicago Press, Chicago.

Gale, J., K. G. Binmore, and L. Samuelson, 1995, "Learning to be Imperfect: The Ultimatum Game," *Games and Economic Behavior*, 8, 56–90.

Goldberg, D., 1989, *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, Reading, Mass.

Grossman, S. J., and J. E. Stiglitz, 1980, "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70, 393–408.

Hirsch M. W., and S. Smale, 1974, *Differential Equations, Dynamical Systems and Linear Algebra*, Academic Press, San Diego.

Holland, J. H., 1975, *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor (1992 edition, MIT Press, Cambridge, Mass).

Jackson, M. O., 1991, "Equilibrium Price Formation and the Value of Private Information," *Review of Financial Studies*, 4, 1–16.

Jordan and Radner, 1982, "Rational Expectations in Microeconomic Models: An Overview," *Journal of Economic Theory*, 26, 201–223.

Koza, J. R., 1992, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, MIT Press, Cambridge, Mass.

Lakonishok, J., A. Shleifer, and R. W. Vishny, 1994, "Contrarian Investment, Extrapolation and Risk," *Journal of Finance*, 49, 1541–1578.

Lappen, A.A., 1998, "Ivory-Tower Investing," *Institutional Investor*, 32(3), 43–55.

LeBaron, B., W. B. Arthur, and R. Palmer, 1998, "Time Series Properties of an Artificial Stock Market," working paper, University of Wisconsin; forthcoming in *Journal of Economic Dynamic Control*.

Lucas, R. E., 1986, "Adaptive Behavior and Economic Theory," in R. M. Hogarth and M. W. Reder (eds.), *Rational Choice: The Contrast between Economics and Psychology*, University of Chicago Press, Chicago.

Marcet, A., and T. Sargent, 1988, "The Fate of Systems with 'Adaptive' Expectations," *American Economic Review*, 78, 168–171.

Marcet, A., and T. Sargent, 1989, "Convergence of Least-Squares Leaning in Environments with Hidden State Variables and Private Information," *Journal of Political Economy*, 97, 1309–1322.

Morrison, D. F, 1990, *Multivariate Statistical Methods*, 3rd ed., McGraw-Hill, New York.

Paul, J., 1995, "Information Aggregation without Exogenous 'Liquidity' Traders," working paper, University of Michigan.

Radner, R., 1979, "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," *Econometrica*, 47, 655–678.

Routledge, B. R., 1995, "Artificial Selection: Genetic Algorithms and Learning in a Rational Expectations Model," working paper, Carnegie Mellon University.

Samuelson, L., 1997, *Evolutionary Games and Equilibrium Selection*, MIT Press, Cambridge, Mass.

Samuelson, L., and J. Zhang, 1992, "Evolutionary Stability in Asymmetric Games," *Journal of Economic Theory*, 57, 363–391.

Sargent, T. J., 1993, *Bounded Rationality in Macroeconomics*, Oxford University Press, Cambridge, Mass.

Sethi, R., and E. Somanathan, 1996, "The Evolution of Social Norms in Common Property Resource Use," *American Economic Review*, 86, 766–788.

Shiller, R. J., 1997, "Human Behavior and the Efficiency of the Financial System," working paper, Yale University; forthcoming in J. B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, North-Holland, Amsterdam.

Somanathan, E., 1997, "Evolutionary Stability of Pure-Strategy Equilibria in Finite Games," *Games and Economic Behavior*, 21, 253–265.

Sunder, S., 1992, "Market for Information: Experimental Evidence," *Econometrica*, 60, 667–695.

Sunder, S, 1995, "Experimental Asset Markets: A Survey," in J. H. Kagel and A. E. Roth (eds.), *The Handbook of Experimental Economics*, Princeton University Press, Princeton, N.J.

Taylor, P. D., and L. B. Jonker, 1978, "Evolutionary Stable Strategies and Game Dynamics," *Mathematical Biosciences*, 40, 145–156.

Veronesi, P., 1999, "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Model," *Review of Financial Studies*, 12, 975–1007.

Vives, X., 1993, "How Fast do Rational Agents Learn?" *Review of Economic Studies*, 60, 329–347.

Vives, X., 1995, "The Speed of Information Revelation in a Financial Market Mechanism," *Journal of Economic Theory*, 37, 178–204.

Wang, J., 1993, "A Model of Intertemporal Asset Prices Under Asymmetric Information," *Review of Economic Studies*, 60, 249–282.

Weibull, J. W., 1995, *Evolutionary Game Theory*, MIT Press, Cambridge, Mass.