LEARNING TO SIGNAL IN MARKETS[1]

Georg Nöldeke
Department of Economics
Princeton University
Princeton, New Jersey 08544 USA

Larry Samuelson
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, Wisconsin 53706 USA

October 21, 1994

---

[1] First draft January 13, 1993. Financial support from the National Science Foundation and the Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 303 at the University of Bonn, is gratefully acknowledged.

Abstract

We formulate a dynamic learning-and-adjustment model of a market in which sellers choose signals that potentially reveal their types. If the dynamic process selects a unique limiting outcome, then that outcome must be an undefeated equilibrium; though to be undefeated does not suffice to be the sole limiting outcome. If a Riley outcome exists that provides "high" type sellers with a higher utility than any other equilibrium outcome, then that outcome is the unique limiting outcome of our model. In the absence of a Riley outcome, or if high type workers obtain higher utility in a pooling equilibrium than in the Riley outcome, a unique limit outcome will only emerge under very stringent conditions. If these conditions fail, the market will cycle between various equilibria and, possibly, nonequilibrium outcomes.

Journal of Economic Literature Classification Numbers C70, C72, D82, D83.

LEARNING TO SIGNAL IN MARKETS

by Georg Nöldeke and Larry Samuelson

1 Introduction

Signaling models, pioneered by Spence [21, 22], are plagued by a common difficulty: very few interesting results hold for all of the sequential equilibrium outcomes. In response, the equilibrium refinements literature endeavors to restrict attention to "plausible" equilibria. Plausibility is defined by placing restrictions on out-of-equilibrium beliefs, where these restrictions are based on the incentives for various types of fully-rational informed players to deviate from equilibrium.[1]

In this paper, we take a different approach to equilibria in signaling games. Rather than speculate on the beliefs that rational players will adopt when faced with out-of-equilibrium actions, we model a process by which boundedly rational players learn to play signaling games.[2] We examine a market containing a large number of buyers and sellers who repeatedly (and anonymously) interact, with sellers sending signals and buyers bidding to make purchases from sellers. Buyers do not know whether a given seller is a "high" or "low" type, though they may infer these types from the sellers' signals.

Two processes shape strategies in this market. First, buyers and sellers regularly adjust their conjectures about others' behavior in light of the observed market outcomes (while always playing best replies against these conjectures.) We refer to this as the *learning dynamics*. Second, agents are occasionally replaced by entrants in the market. Entrants may hold different conjectures about play at currently unreached information sets, and hence may pursue different strategies, than their predecessors. By potentially providing the stimulus that tips the market from one equilibrium to another, these replacements play an important role in determining market behavior.

We study the limiting outcome (or stationary distribution) of this learning-and-replacement process, in the limit as the rate at which players are replaced by entrants approaches zero. Using results developed in Nöldeke and Samuelson [16], we find that we can learn much from the following "local stability" condition, which is necessary for an equilibrium outcome to be the unique limiting outcome of our dynamic process:[3]

---

[1] See van Damme [24] for a survey of the refinements literature. Before equilibrium refinements were unleashed in their full fury, Riley [17] and Wilson [25] proposed resolutions to the nonexistence-of-equilibrium problem discovered by Rothschild and Stiglitz [18] in screening models. The "Riley outcome" also appears as a sequential equilibrium outcome in signaling models and is selected by many refinement concepts.

[2] Stiglitz and Weiss [23] advocate a dynamic approach to equilibrium in signaling and screening models. Links between evolutionary models and equilibrium refinements have been examined in the context of cheap talk games. See, for example, Nöldeke and Samuelson [15] and Sobel [20].

[3] This intuitive account neglects some issues arising from the existence of non-singleton absorbing sets ("cycles") of the learning dynamics that are addressed in the course of the analysis.

> (I) Starting from the equilibrium outcome, the replacement of a single agent, causing a deviation to an "out of equilibrium" signal, cannot induce learning dynamics that lead away from the given outcome and *to another equilibrium outcome*.

The requirement that replacements should not induce learning dynamics that lead away from the equilibrium is reminiscent of equilibrium refinements like the intuitive criterion of Cho and Kreps [3], where an equilibrium is rejected if an out-of-equilibrium signal prompts belief revision that disrupts the equilibrium. However, conventional equilibrium refinements construct arguments about belief revision by specifying how players *should* interpret out-of-equilibrium signals. These arguments are often represented in the form of "speeches" that players are imagined to deliver when making out-of-equilibrium moves. In our model, out-of-equilibrium signals acquire meaning endogenously as part of the evolutionary process. As a result, the single-crossing condition, which often plays a central role in determining which type of the sender has most to gain from a particular deviation and hence what buyers should believe when faced with such a deviation, is irrelevant in our work.

The requirement in (I) that the learning dynamics be able to reach another equilibrium is reminiscent of the "Stiglitz critique" ([3], p.203): The mere fact that a deviation could initially destabilize an equilibrium does not suffice to exclude that equilibrium from consideration. Note, however, that in our model, an equilibrium fails to be the unique limiting outcome if it is *possible* for the learning dynamics to lead to another equilibrium; the possibility that the learning dynamics might also return to the original equilibrium is irrelevant.

Mailath, Okuno-Fujiwara and Postlewaite [13] have recently argued that in order to reject an equilibrium, an equilibrium refinement should not only identify an out-of-equilibrium signal and belief revisions leading away from the equilibrium; but should also identify an alternative equilibrium in which this signal is sent and in which the senders of this signal fare better than in the original equilibrium. This is similar to our requirement that the learning dynamics be able to reach another equilibrium. We exploit this similarity to show that a defeated equilibrium cannot be locally stable.

Being undefeated is a necessary but not sufficient condition for an equilibrium outcome to be locally stable. Instead, many undefeated outcomes are rejected by condition (I). The undefeated equilibrium concept rejects an equilibrium only if a belief revision process can lead *immediately* to an alternative equilibrium, while equilibria in our market can be destabilized by learning dynamics that only indirectly reach alternative equilibria. We exploit such indirect paths to show that an equilibrium outcome can only satisfy condition (I) if there does not exist an alternative equilibrium in which *high*-quality sellers send a signal that is not used in the given equilibrium and fare better than in the given equilibrium outcome. We call equilibrium outcomes satisfying this condition H-dominant.

Our necessary conditions (undefeated and H-dominant) for local stability imply that no market can have more than one locally stable pooling equilibrium and one locally stable (partially) separating equilibrium. The unique candidate for a locally stable separating equilibrium is (a slight modification of) the familiar Riley outcome, while the unique candidate for a locally stable pooling equilibrium is one which gives the high-quality seller the highest payoff among all equilibrium outcomes. When can we find sufficient conditions for these outcomes to be locally stable, and what is the relationship between local stability and the limiting outcome of our dynamic market process?

We first consider the case of a separating equilibrium. We show that if the Riley outcome is H-dominant and each type of seller strictly prefers his equilibrium signal to that of the other type, then the Riley outcome is locally stable in the sense of (I). If it is also the case that no other equilibrium gives high sellers a higher utility that the Riley outcome (i.e. the Riley outcome is the *unique* H-dominant equilibrium outcome), then condition II (introduced below) also holds and the Riley outcome is the unique limiting outcome of the model.

We next show that only in extraordinary cases will pooling equilibria be stable. If there exists an alternative sequential equilibrium outcome (satisfying an additional technical condition) and if there are *any* circumstances that would induce some type of seller to leave the pooling equilibrium signal, then we can construct a path as in condition (I) that leads from the pooling to the alternative equilibrium, precluding the local stability of the pooling equilibrium.

Our model thus selects a unique equilibrium outcome if a Riley outcome exists (in which sellers are not indifferent between equilibrium signals) and is the unique H-dominant equilibrium, i.e., is the best possible equilibrium outcome for high-quality sellers. On the one hand, the arguments supporting this result are reminiscent of the arguments advanced by many equilibrium refinements for choosing separating equilibria.[4] At the same time, our result requires the Riley outcome to be the unique H-dominant equilibrium. There is no conflict between separation and dominance arguments in such markets, and the Riley outcome appears to be the unique “plausible” equilibrium outcome. Suppose instead that a Riley outcome exists but is not H-dominant, with the unique H-dominant equilibrium being a pooling equilibrium.[5] It is then no longer obvious which outcome is plausible. Refinements such as the intuitive criterion, divinity and D1 still invoke separation arguments to select the Riley outcome; arguing that the pooling equilibrium will be destabilized by the high-quality seller’s incentive to deviate from the equilibrium in order to signal his type. Others invoke dominance considerations to argue that it is implausible to select the Riley outcome in the presence of a dominating pooling equilibrium.[6] The perfect se-

---

[4] For example, the intuitive criterion (Cho and Kreps [3]), divinity (Banks and Sobel [1]), and D1 (Cho and Sobel [4]).

[5] The equilibrium producing the Riley outcome and a pooling equilibrium could both be H-dominant, though only if the signal used in the pooling equilibrium is also used by low-quality sellers in the Riley outcome. In this case, the model selects either both equilibria or just the Riley outcome, selecting the latter if a single mutation can produce learning dynamics that lead away from the pooling equilibrium.

[6] See, for example Fudenberg and Tirole ([8], p. 458) and Mailath, Okuno-Fujiwara and Postlewaite [13].

quential equilibrium concept (Grossman and Perry [10] responds to this conflict by failing to exist.[7] Our dynamic model may select the pooling equilibrium in this case, but only under extremely stringent conditions. In general, multiple equilibrium outcomes appear in the limit, in effect replacing previous non-existence results with a set–valued solution concept.[8]

The reasoning driving our results shares some similarities with the intuition underlying refinement criteria, but there are also essential differences. In particular, to show that an "undominated" Riley outcome is the sole limiting outcome of our model, we must not only verify condition (I), but must also show that

> (II) By repeatedly constructing replacement-and-learning paths of the type appearing in step (I), *every* equilibrium can be linked to the undominated Riley outcome.

This requirement has no natural counterpart in the refinement literature, which focuses on the stability of single equilibrium outcomes. Condition (II) ensures that the limiting outcome cannot consist of a collection of equilibrium

outcomes that form a "stable set" in the sense that the types of paths constructed in step (I) allow the system to cycle

among the members of the set (but not to leave the set). The outcome in a market with both a Riley outcome and a pooling equilibrium that is H-dominant (i.e., yields a higher utility for high sellers than the Riley outcome and uses a signal that does not appear in the Riley outcome) consists of just such a set.

Our work departs from previous evolutionary models in three important ways. First, our model has a large number of agents, allowing individual agents to learn from observing the behavior of many other agents, but there is no random matching of agents. Instead, all agents interact simultaneously in a market game and the resulting market outcome provides the information agents use to adapt their behavior. We consider this to be the most appealing economic setting in which to invoke evolutionary arguments.

Second, as is many recent evolutionary models, the agents in our model are subject to continual "mutations", which take the form of new entrants into the market in our case. However, these mutations affect only agents' conjectures at unreached information sets, with strategies always being best replies to conjectures. Mutations then cannot move the market from every state to every other state. As a result, there may be multiple limiting outcomes, with the realized outcome depending on initial conditions and the realization of the random variable governing learning. This mutation process contrasts with that of

---

[7] The non-existence of equilibrium in the screening model of Rothschild and Stiglitz [18] is closely related. See van Damme [24].

[8] If a Riley outcome does not exist, then again we have multiple equilibrium outcomes appearing in the limit except under exceptional circumstances.

Kandori, Mailath and Rob [12], where mutations affect strategies and cause the probability of moving between any two states to be positive.[9]

Third, our formal model is a Markov process, which is again standard in evolutionary analyses. Because nonsingleton absorbing sets are hard to handle, the common practice when working with Markov models is to find conditions under which all absorbing sets are singletons. Young's [26] acyclicity condition is in this tradition. We do not restrict our analysis to the case in which absorbing sets are singletons, instead working directly with nonsingleton absorbing sets.

The following section presents the model and establishes some useful preliminary results. Section 3 investigates our learning dynamics. Section 4 examines the limiting behavior of the market when mutations can occur. Section 5 establishes necessary conditions for local stability, yielding the instability of defeated equilibria as a special result. Section 6 examines the stability of the Riley outcome. Section 7 derives the stringent conditions for pooling equilibria (and partially separating equilibria) to be stable outcomes. Section 8 concludes.

2 The Model

2.1 Signaling Games and Markets

Spence [21, 22] examined signaling in the context of a labor market containing a large number of potential workers, some of high productivity and some of low productivity. Firms could not observe workers' productivities, but could observe their choices of education levels. Wage rates for each education level were set by firms engaging in "Bertrand" price competition for the workers. Spence's key insight was that market equilibria could arise in which workers used education levels to signal productivities.

Cho and Kreps [3] provide a game-theoretic analysis of Spence's market signaling model. There is a single worker, who may be of high or low productivity, but whose type cannot be observed by the firms. The proportions of high and low types in Spence's model are replaced by prior probabilities of the single worker's type in Cho and Kreps. Two firms bid for the services of the worker. Spence's signaling market equilibrium is now a separating sequential equilibrium of the three-player signaling game of incomplete information.

We follow Cho and Kreps [3] in modeling the signaling market as a game. This reflects our quest for insight into equilibrium refinements for signaling games. However, our game has a large number of sellers (or workers) and buyers (or firms).[10] We believe that the best

---

[9] Kandori, Mailath and Rob [12] work with a normal form game where conjectures at unreached information sets play no role. Our mutation process would not allow us to differentiate between different strict equilibria in their model.

[10] In this respect, we follow Gale [9], though for quite different reasons. Gale [9] argues that it is restrictive to replace population proportions by prior probabilities and consider a game in which there is

hope for understanding equilibria in signaling games is to examine the learning process by which these equilibria are achieved. In turn, we think an important aspect of such a dynamic model is the possibility that agents may learn from observing the actions of others with whom they do not directly interact. A market consisting of large numbers of buyers and sellers provides a natural setting for such information transmission. In particular, observing the outcome in a market can provide agents with a substantial amount of payoff-relevant information, such as information about the wage offers that appear in response to various education levels.[11]

2.2 The Market

We consider a market with a finite number of sellers and buyers. Each seller is endowed with one unit of a commodity. These units can be of high or low quality. We let A denote the set of all agents, A_B the set of buyers, A_L the set of low-quality sellers, and A_H the set of high-quality sellers. Let

$$\phi^0 = \frac{|A_H|}{|A_H| + |A_L|} \in (0, 1)$$

be the fraction of high-quality sellers.

Each seller has a strategy set S containing a finite number of signals. Buyers quote prices p from a finite set $P \subset \mathbb{R}_+$. We let $\underline{p}$ and $\overline{p}$ be the largest and smallest prices in P, and assume that Δ (the smallest money unit) is the uniform step size of the price grid, so that P consists of prices of the form $\underline{p} + k\Delta$ for some integer k.[12]

A strategy for a seller is a signal in S. A strategy for a buyer is a price schedule mapping from S into P.[13] Throughout this paper we restrict attention to pure strategies. Since our game has many buyers and sellers, pure strategy combinations in which different agents

---

only one worker. For example, if there is a single worker and two firms, then the two firms must earn zero profits in equilibrium. In a market, one would expect the equilibrium profits of firms to depend upon the relative numbers of firms and workers of various types, possibly being positive if the number of workers is large relative to the number of firms.

[11] An alternative is to assume that triplets of players, consisting of one worker and two buyers, are repeatedly matched to play the Cho-Kreps signaling game. In order for current outcomes to drive strategy choices, such models commonly seek refuge in the assumption that all agents play a round–robin tournament in every period (Nöldeke and Samuelson [16]) or that an infinite number of games is played in each period (Kandori, Mailath and Rob [12]).

[12] We have modeled both buyers and sellers as choosing from finite sets. Here, we again follow Gale [9]. In Gale's case, the motivation for working with a discrete model is to employ the techniques of general equilibrium analysis for economies with finite numbers of goods. In our case, the motivation is to employ models of evolutionary systems as finite Markov chains.

[13] Buyers can condition prices on signals but cannot condition on the number or the identity of sellers sending each signal. This will be the case if, for example, buyers must determine price schedules before observing which signals are sent. If buyers could condition prices on the identity or number of sellers choosing a signal, strategic issues arise that have no counterpart in the usual signaling model. See Footnote 15.

of the same type play different pure strategies will look much like mixed strategy profiles in a Cho-Kreps game.

A seller who has chosen signal s sells his unit to one of the buyers who offers the **market price**, which is the maximum price (from the finite set P) that is chosen by some buyer for signal s. If there is more than one buyer offering the market price then the unit is allocated to every such buyer with equal probability.

A seller's payoff from choosing signal s and receiving price p is given by

$$p - c_L(s)$$

for L sellers and

$$p - c_H(s)$$

for H sellers; where $c_L(s) : S \to \mathbb{R}$ and $c_H(s) : S \to \mathbb{R}$ are functions identifying the cost of sending signal s for L and H sellers.

We assume that the marginal cost of switching from signal s to s' is different for different types of the seller:

Assumption 1 *For all $s \neq s'$:*

$$c_L(s) - c_L(s') \neq c_H(s) - c_H(s'). \tag{1}$$

In particular, this assumption implies that we are not examining cheap talk games. Assumption 1 is consistent with, but much weaker than, the single crossing property, which would require that signals can be ordered in such a way that H sellers have the lower marginal costs of switching from s to s' for any $s' > s$.

For buyers, the payoff *per unit purchased* depends on the quality of the unit purchased and may also depend on the signal used by the seller. This payoff is given by

$$v_T(s) - p$$

for $T \in \{L, H\}$, where $v_T(s) : S \to \mathbb{R}$ identifies the value of a unit offered by a T-quality seller who sends signal s. We assume that buyers can always ensure non-negative profits and prefer to trade with high-quality sellers:

$$\forall s \in S : v_H(s) > v_L(s) > \underline{p}. \tag{2}$$

All agents are risk neutral. Let

$$v_\phi(s) \equiv \phi v_H(s) + (1 - \phi) v_L(s)$$

denote the expected value of a unit that is high quality with probability ϕ. Note $v_H(s) = v_1(s)$ and $v_L(s) = v_0(s)$.

We have now modeled the signaling market as a game, and we refer to this game as G.

A **state** θ is a specification of a strategy (denoted $\sigma_a(\theta)$) and a **conjecture** (denoted $\sigma_{-a}(\theta)$) for every agent a. For $a \in A_L \cup A_H$, the strategy $\sigma_a(\theta)$ is an element of S. For $a \in A_B$, $\sigma_a(\theta) : S \to P$ is a function assigning prices to signals, with $\sigma_a(s, \theta)$ being the price assigned to state s. For $T \in \{L, H\}$ and $s \in S$, the conjecture $\sigma_{-a}(s, T, \theta) \geq 0$ specifies the number of T sellers (excluding agent a if $a \in A_T$) that agent a believes to be choosing signal s; for $p \in P$, $\sigma_{-a}(s, p, \theta) \geq 0$ specifies the number of buyers (excluding agent a if $a \in A_B$) that agent a believes to offer price p at signal s. We assume that all agents are informed about the number of other agents in the market, so for all a conjectures satisfy:

$$\begin{aligned} \sum_{s \in S} \sigma_{-a}(s, T, \theta) &= |A_T \setminus \{a\}|, \; T = L, H \\ \forall s : \sum_{p \in P} \sigma_{-a}(s, p, \theta) &= |A_B \setminus \{a\}|. \end{aligned}$$

We let $BR(\sigma_{-a})$ denote the set of best responses for agent a against the conjecture σ_{-a}. We let $\sigma^a \equiv (\sigma_a(\theta), \sigma_{-a}(\theta))$ describe agent a's strategy and conjecture in state θ, and suppress the θ whenever convenient.

For $T \in \{L, H\}$ let

$$S_T(\theta) = \{s | \exists a \in A_T : \sigma_a(\theta) = s\}.$$

Then $S_T(\theta)$ identifies the signals sent by type T sellers in state θ.

In our dynamic model, agents cannot observe price offers at signals that are not chosen by any seller. More formally, after the game has been played agents can observe what we call the **outcome** of a state θ, denoted $\rho(\theta)$, consisting of the distribution of sellers across signals and the distribution of prices offered at those signals that are chosen by some seller:

$$\begin{aligned} \forall s \in S, T \in \{L, H\} : \quad & \rho(s, T, \theta) = |\{a \in A_T | \sigma_a(\theta) = s\}| \\ \forall s \in S_L(\theta) \cup S_H(\theta) : \quad & \rho(s, p, \theta) = |\{a \in A_B | \sigma_a(s, \theta) = p\}|. \end{aligned}$$

To simplify notation we will often simply write ρ for $\rho(\theta)$. We let $\rho(\sigma^a)$ denote the outcome that would appear if agent a plays σ_a and all other agents' strategies match agent a's conjecture σ_{-a}.

2.3 Equilibrium

An equilibrium requires that all agents play best responses against their conjectures. Because agents can observe only outcomes and not strategies, an equilibrium does not require that conjectures agree with the actual strategy profile, but they must induce the same outcome:

Definition 1 *A state θ is an **equilibrium** if:*

$$\forall a : \quad \sigma_a \in BR(\sigma_{-a}) \tag{3}$$

$$\forall a : \quad \rho(\sigma^a) = \rho(\theta). \tag{4}$$

This definition of equilibrium imposes no restrictions on the prices offered by buyers at signals that are not used, and requires only that sellers' conjectures about those prices be such that the unused signals are not best replies.[14] Note that if θ is an equilibrium, then all sellers of the same type must obtain the same expected utility level, allowing us to denote the equilibrium utility level of a type T seller by $U_T(\theta)$.[15]

We will often find it convenient to refer to states that may not be equilibrium states, but in which buyers are choosing best replies and hold correct conjectures over the outcome:

Definition 2 *A state θ is **competitive** if the equilibrium conditions hold for all buyers:*

$$\forall a \in A_B : \qquad \sigma_a \in BR(\sigma_{-a}) \tag{5}$$

$$\forall a \in A_B : \qquad \rho(\sigma^a) = \rho(\theta). \tag{6}$$

In a model with continuum action sets, the competition by which buyers bid for units would ensure that in every competitive state the market price (at any signal s that is used) would equal the "competitive price", by which we would mean a price just equal to the expected value of a unit supplied at s. This feature is very attractive in a static equilibrium analysis. In spite of this, we eschew continuum actions sets because this same feature raises significant problems for a dynamic analysis. In the equilibrium of a continuum-action-set model every buyer is playing a weakly dominated strategy, earning a zero payoff and being indifferent between the equilibrium price and any lower price. As a result the pressures pushing the system toward the equilibrium outcome are weak. The system can then drift away from the equilibrium outcome, rendering this outcome unstable.

To define the concept of a competitive price with discrete action sets, let Φ be given by

$$\Phi \equiv \{\phi \in [0,1] : \phi = \frac{M}{N} \textit{ for some } |A_L| + |A_H| \geq N \geq 1 \textit{ and } |A_H| \geq M \geq 0\}.$$

Φ is the set of all possible conjectures concerning the proportion of sellers who are type H and choose a given signal. Then we have:

Definition 3 *The **competitive price** for signal s given probability $\phi \in \Phi$ of high quality is given by*

$$p^*_\phi(s) \equiv max\{p \in P : v_\phi(s) \geq p\}.$$

---

[14] This definition of equilibrium is akin to Kalai and Lehrer's [11] notion of "private belief equilibrium" and Fudenberg and Levine's [6] notion of "self-confirming equilibrium". Fudenberg and Kreps [5] and Fudenberg and Levine [6, 7] show that when working with dynamic processes for extensive form games, it is especially important to take account of the information that agents can obtain from playing the game; and show that doing so leads to self-confirming equilibrium.

[15] An equivalent claim would not be true if buyers could condition their responses on the number or identities of sellers choosing each signal. Disparities in seller payoffs could then be enforced, for example, by a strategy that attaches a relatively favorable response to signal s in equilibrium, but a much less favorable response if one additional seller chooses s.

We have assumed in (2) that the set P contains sufficiently low prices to make the purchase of a low-quality unit profitable at each signal s, which in turn ensures that $p^*_\phi(s)$ exists for all s and ϕ.

We assume the price grid is sufficiently fine for it to make a difference if prior information concerning a seller's type is replaced by the information that the seller is either type H or type L (recall $\phi^0 \in (0,1)$):

Assumption 2 *For all $s \in S$:*

$$p^*_1(s) > p^*_{\phi^0}(s) > p^*_0(s).$$

We cannot be sure that competitive bidding between buyers actually produces the competitive price for all ϕ in our discrete model, since the profits from the competitive price may be so close to zero that no buyer has an incentive to deviate to the competitive price if all other buyers are offering the next-lowest price. The following assumption and lemma eliminate this problem:

Assumption 3 *For all $s \in S$ and $\phi \in \Phi$:*

$$v_\phi(s) > p^*_\phi(s) + \frac{\Delta}{\mid A_B \mid -1}. \tag{7}$$

Note that for given A_H and A_L, Assumption 3 will hold for generic price grids (so that profits at the finite number of relevant competitive prices are nonzero) as long as the number of buyers is sufficiently large.[16] We use Assumption 3 to characterize competitive states as those in which buyers hold correct conjectures and offer competitive prices:[17]

Lemma 1 *A state θ is competitive if and only if all buyers hold correct conjectures over the outcome (i.e., $\rho(\sigma^a) = \rho(\theta)$ for all buyers) and*

$$\forall s \in S_L \cup S_H, \forall a \in A_B : \ \sigma_a(s) = p^*_{\phi(s,\theta)}(s),$$

where

$$\phi(s,\theta) = \frac{|\{a \in A_H | \sigma_a = s\}|}{|\{a \in A_L \cup A_H | \sigma_a = s\}|}.$$

Proof: See Appendix. □

We also impose a genericity assumption that simplifies the analysis:

---

[16] Because $v_\phi(s) - p^*_\phi(s) < \Delta$, at least three buyers are needed for (7) to hold.

[17] While seemingly obvious, results of this kind are not automatic in games with discrete choice sets and condition (7) is required in the proof. The result in Lemma 1 will also hold if we allow for mixed strategies for buyers. The restriction to pure strategies for sellers, however, is essential.

Assumption 4 *There do not exist signals* $s \neq s'$ *such that*

$$\underline{p} - c_H(s) = p_0^*(s') - c_H(s')$$

or

$$\underline{p} - c_H(s) = p_1^*(s') - c_H(s').$$

Hence, there are no signals s and s' such that the H seller is exactly indifferent between receiving the lowest possible price at s and receiving the competitive price for either L sellers or H sellers at s'.

Finally, it will be useful to define a notion of equilibrium that requires the prices set by buyers at unsent signals to be competitive prices against some conjecture about the types of sellers who send that signal, and hence corresponds more closely to a sequential equilibrium of the Cho-Kreps model. To simplify terminology we call such equilibrium states sequential.

Definition 4 *Let* $\tilde{\Theta}$ *be the set of states in which, for each signal* $s \in S$*, all buyers offer prices in* $[p_0^*(s), p_1^*(s)]$ *and all agents conjecture this to be the case. Equilibrium state* θ *is* ***sequential*** *if* $\theta \in \tilde{\Theta}$*. Outcome* ρ *is a sequential equilibrium outcome if there exists a sequential equilibrium* θ *such that* $\rho(\theta) = \rho$*.*

2.4 Market Dynamics

We now embed this model in a dynamic process. We view this process as a description of how boundedly rational agents might find their way to an equilibrium. By “boundedly rational”, we mean agents who play best replies against their conjectures, but whose conjectures about others’ play is shaped by experience. These conjectures may accordingly be incorrect in their description of play at information sets that have only rarely been reached.

Our dynamic process is a modified version of the evolutionary model of Nöldeke and Samuelson [16]. The most important difference is that the mutation process in [16] affects both the conjectures and strategies of agents. Mutations can then endow agents with strategies that are not best replies to their conjectures. In this paper, mutations affect only conjectures about the prices offered at signals not used in the market. This reflects our belief that the most natural interpretation for mutations in a market model is that of new entrants to the market, who use the information that can be obtained from the market outcome to form conjectures and choose best replies to those conjectures, but are left to form their own conjectures about behavior at unsent signals.

We assume that game G is played at time periods $t \in \{0, 1, 2, \cdots\}$. As in the previous section we restrict attention to pure states, i.e. states in which all agents conjecture that all other agents are playing pure strategies and all agents play pure strategies that are best responses to these conjectures. The state space of our system is given by

$$\Theta = \{\theta | \forall a : \sigma_a(\theta) \in BR(\sigma_{-a}(\theta))\}.$$

Let θ be the state at the beginning of period t. In period t, the game G is played according to the strategies $\sigma_a(\theta)$. Each agent then takes an independent random draw determining whether the agent continues to be active in the market in period $t + 1$. With probability $1 - \lambda$ the agent stays in the market and proceeds to the following learning stage; with probability λ the agent is replaced by a new agent *of the same type* (i.e., a buyer, type L seller or type H seller). Because our entering agents are the counterpart of mutations in other evolutionary models, we refer to λ both as the **replacement rate** and the **mutation rate**, and speak of a mutation occurring when an agent is replaced.

Now learning occurs. Each continuing agent (independently) learns with probability $\mu \in (0, 1)$ in each period.[18] Because they enter the game devoid of conjectures, all new agents learn.

If an agent does not learn, then his strategy and conjecture remain unchanged until the next period. If an agent learns he observes the *outcome* $\rho(\theta)$ from period t. He uses this information to update his conjecture σ_{-a} to match observed behavior. The agent receives no information concerning the part of the strategy profile that cannot be inferred from the market outcome, i.e., the choices of buyers at signals that were not chosen in state θ. If the agent is a continuing agent he retains his old conjecture about buyers' choices at these signals. If the agent is an entrant his conjecture for these signals is drawn from a probability distribution that depends only on the state θ and has full support on the set of all conjectures that result in the outcome $\rho(\theta)$.

Given his new conjecture the agent then adjusts his strategy. In particular, new agents adopt a best response against their conjecture, putting strictly positive probability on all possible best responses. If a continuing agent finds that his current strategy is a best response then the agent does not change strategies. If the current strategy is not a best response, then

- if the agent is a seller, he switches to a signal that is a best reply against his current conjecture, putting strictly positive probability on every possible best reply.
- if the agent is a buyer, he switches his price at every signal at which such a switch yields higher profits given his conjecture (i.e., at every signal sent by some seller in period t where the buyer is not currently playing a best response), putting positive probability on every best response that can be so achieved.[19]

These adjustment rules reflect a belief that decision-making is costly, either in physical or mental resources. Agents accordingly often proceed as they have previously played the game, without devoting attention to the actual market outcome. When agents learn, they

---

[18] We require there to be positive probability that each agent learns in each period as well as positive probability that the agent simply proceeds to the next period without any change in strategy or conjecture. Within these bounds, we take no position on whether learning occurs only occasionally or quite frequently.

[19] The buyer and seller cases are worded differently because the buyer must make choices for multiple information sets. For both the buyer and seller, the behavior strategy at an information set is changed if and only if it is not a best response to the agent's conjecture *at that information set.*

observe the current state and then change their behavior if and only if they conjecture that such a change will increase their payoff.

For every replacement rate $\lambda \in [0, 1]$ (fixing all other probabilities) this gives us a Markov process, which we refer to as $\Gamma(\lambda)$. The process without replacement, $\Gamma(0)$, will be called the **learning dynamics**.

For any state $\theta \in \Theta$ we let $B(\theta, \lambda)$ denote the basin of attraction of state θ under the process $\Gamma(\lambda)$.[20] A nonempty set $A \subset \Theta$ is **absorbing** for $\Gamma(\lambda)$ if A is minimal with respect to the following property:

$$\forall \theta \notin A : A \cap B(\theta, \lambda) = \emptyset.$$

Hence, the Markov process cannot leave A once it has entered it, but every state in A lies in the basin of attraction of every other state in A. If an absorbing set A is a singleton, then we refer to it as an **absorbing state**.

The market dynamics $\Gamma(\lambda)$ allow the process to reach states which are not contained in $\tilde{\Theta}$, i.e., in which some market prices (and conjectures about those market prices) do not correspond to competitive prices against *any* conjecture. First, the replacement of agents allows arbitrary drift of prices and conjectures about prices at unused signals. Second, learning (without replacements) can cause prices at a signal s to fall below $p_0^*(s)$ if that signal is currently used: If all buyers offer prices above $p_{\phi(s)}^*(s)$ in the current state θ, then any price below the current price is part of a best response at s. The learning dynamics may thus reach a state in which all buyers offer $\underline{p}$ at s. We refer to such an episode as a **crash** at s. If $\underline{p} < p_0(s)$, then a crash at s yields a state that is not in $\tilde{\Theta}$.

These complications could be avoided by simply restricting the dynamics (and the initial state) to $\tilde{\Theta}$, which is easy to achieve by restricting the support of the corresponding distributions in the learning and replacement stages of our process. Working with such a restricted dynamics would simplify some of our arguments and would enforce the sequentiality of limit outcomes (see Section 5.2), without affecting any of our main conclusions. However, we find it difficult to motivate such a restriction on the dynamics, and prefer to proceed with the larger state space Θ defined above.

3 Learning Dynamics

Our goal is to examine the limiting outcome or stationary distribution of the market dynamics $\Gamma(\lambda)$; for the case of very small replacement rates λ. We do this by examining the limit, as replacement rates go to zero, of the stationary distributions of $\Gamma(\lambda)$. This limit will be concentrated on absorbing sets of the learning dynamics $\Gamma(0)$. In this section, we accordingly examine the learning dynamics of our signaling market, with particular attention paid to characterizing the absorbing sets of this process.

---

[20] The basin of attraction of state θ is the set of states from which the Markov process can transit to θ in a finite number of steps with positive probability.

3.1 Singleton Absorbing Sets

The following result, which is almost immediate from the definitions, gives us a complete characterization of singleton absorbing sets.

Proposition 1 *A state θ is an absorbing state of the learning dynamics if and only if θ is a equilibrium state.*

Proof: Because agents always play best responses against their conjectures, condition (i) in the definition of an equilibrium state holds for all $\theta \in \Theta$. Suppose θ is absorbing. Then there can be no agent such that $\rho(\sigma^a) \neq \rho(\theta)$ since the conjecture of such an agent would change upon receiving a learn draw, preventing θ from being absorbing. Hence θ is an equilibrium. Conversely, suppose θ is an equilibrium. Then no agent changes his conjecture upon receiving a learn draw, since the conjecture is not falsified; nor does any agent change his strategy, since the the learning dynamics call for agents to change their strategies only when not playing a best response. Hence θ is absorbing. $\square$

Propositions of this type, establishing a connection between absorbing states and equilibria, are familiar in evolutionary models. Conventional evolutionary analyses follow such a result by either placing sufficient restrictions on the evolutionary process to ensure convergence to an absorbing state or establishing results conditional on such convergence. In contrast, we address the possibility of nonconvergence with our second proposition. It provides a characterization of nonsingleton absorbing sets, showing that each such set contains what we call a four-state cycle. Remark 1 shows that such cycles correspond to mixed-strategy equilibria of the underlying game. We thus have a close connection between limit outcomes and equilibria, even though our model does not have sufficient structure to ensure convergence.

3.2 Nonsingleton Absorbing Sets

To discuss non-singleton absorbing sets we need a definition:

Definition 5 *Let $\theta \in \Theta$. The **competitive successor** of θ, denoted $c(\theta)$, is the (unique) state in Θ that is competitive and satisfies:*

$$\forall a \in A_L \cup A_H : \quad \sigma^a(c(\theta)) = \sigma^a(\theta) \tag{8}$$

$$\forall a \in A_B : \quad \forall s \notin S_L(\theta) \cup S_H(\theta) : \sigma_{-a}(s, c(\theta)) = \sigma_{-a}(s, \theta) \tag{9}$$

$$\forall a \in A_B : \quad \forall s \notin S_L(\theta) \cup S_H(\theta) : \sigma_a(s, c(\theta)) = \sigma_a(s, \theta). \tag{10}$$

The competitive successor is then the state that emerges if sellers retain their current signals and conjectures (this is (8)), price competition among buyers pushes the price at all used signals and buyers' conjectures to the competitive price (i.e., $c(\theta)$ is competitive), and buyers' conjectures about and actions at unused signals remain unchanged ((9) and (10)).

Given a state θ, its competitive successor can be reached under the learning dynamics through a realization in which sellers' strategies and conjectures remain fixed while buyers hold a Bertrand auction at each signal to purchase the commodities from sellers sending that signal. More formally:

Lemma 2 *Every state θ is contained in the basin of attraction of its competitive successor under the learning dynamics:*

$$\forall\theta \in \Theta : \theta \in B(c(\theta), 0).$$

Proof: It suffices to show that with positive probability, the learning dynamics leads from θ to $c(\theta)$, i.e., there exists a finite number of periods, along with a configuration of learn draws and choices of best replies for learning agents in those periods, that leads from θ to $c(\theta)$. Consider the following sequence. First, a number of periods occur in which buyer a receives the learn draw if and only if there exists an $s \in S_L \cup S_H$ such that $\sigma_a(s) > p^*_{\phi(s)}(s)$ and $\sigma_a(s) = \max_{a\in A_B} \sigma_a(s)$. (No sellers receive the learn draw in any of the periods of the sequence. This ensures condition (8).) Such a buyer earns a negative payoff at signal s, and to achieve a best reply must switch to some lower price at signal s. This buyer may also adjust his price at other signals $s' \in S_L(\theta) \cup S_H(\theta)$, but must switch to a price which is strictly lower than the maximum of other buyers' offers, whenever this maximum exceeds the competitive price at signal s'. Hence, after a finite number of periods, this leads to a state in which no buyer offers a price at any signal that is higher than the competitive price at that signal. Second, there then follows a number of periods in which buyer a gets the learn draw if and only if there exists a signal $s \in S_L(\theta) \cup S_H(\theta)$ at which $\sigma_a(s) = \min_{a\in A_B} \sigma_a(s)$. If $\sigma_a(s) < p^*_{\phi(s)}(s)$, then, in order to play a best reply buyer a must increase his offer at signal s (by Assumption 3), but will never choose an offer above $p^*_{\phi(s)}$. This buyer may again adjust his price at other signals $s' \in S_L(\theta) \cup S_H(\theta)$ (at which he is neither playing a best response nor offering the minimal price), but it cannot be a best response for buyer a to lower his price to or below the minimum price at such a signal (cf. the proof of Lemma 1). Hence, after a finite number of periods, we reach a state which is competitive and satisfies (8). Since learning does not affect buyers' conjectures about (or actions at) signals not contained in $S_L(\theta) \cup S_H(\theta)$ this state must be $c(\theta)$. $\square$

We refer to the learning sequence described in this proof, with sellers not learning and buyers being driven to competitive prices, as a **competitive learn sequence**. The ability to construct competitive learn sequences implies that we can extract a sequence of states from the learning dynamics that resembles an alternating-best-response dynamics: starting from a given state buyers first undergo a competitive learn sequence and adjust to the competitive successor, then all sellers choose a best response[21] against the competitive

---

[21] In the formal definition we consider a particular selection from the best response correspondence if necessary, ensuring that the resulting dynamics is deterministic.

successor, then buyers adjust to the competitive successor, and so on. Call this sequence the **alternating-best-response dynamics**.

Since each of the transitions in the alternating-best-response dynamics occurs with positive probability under the learning dynamics, every non-singleton absorbing set of the learning dynamics must contain a cycle of the alternating-best-response dynamics (among many other transition possibilities). Hence, studying the relatively simple alternating-best-response dynamics allows us to obtain a partial characterization of non-singleton absorbing sets. In particular, we prove the following result by showing that whenever the alternating-best-response dynamics does not converge to a singleton, it must converge to the limit cycle described in the statement of the Lemma. We will refer to such a cycle as a **four-state cycle**.

Proposition 2 *Suppose $X \subset \Theta$ is a non-singleton absorbing set of $\Gamma(0)$. Then X contains a four-state cycle, i.e., a set of states $F = \{\theta', c(\theta'), \theta, c(\theta)\}$ such that:*

$$\exists \hat{s}_H \in S: \qquad S_L(\theta) = S_H(\theta) = S_H(\theta') = \{\hat{s}_H\} \tag{11}$$

$$\exists \hat{s}_L \neq \hat{s}_H \in S: \qquad S_L(\theta') = \{\hat{s}_L\} \tag{12}$$

$$\forall a \in A_L \cup A_H: \qquad \sigma_a(\theta) \in BR(\sigma_{-a}(c(\theta'))) \tag{13}$$

$$\forall a \in A_L \cup A_H: \qquad \sigma_a(\theta') \in BR(\sigma_{-a}(c(\theta))) \tag{14}$$

$$\forall a \in A_B: \qquad \sigma_a(\theta') = \sigma_a(c(\theta)) \tag{15}$$

$$\forall a \in A_B: \qquad \sigma_a(\theta) = \sigma_a(c(\theta')). \tag{16}$$

Proof: See Appendix. □

The four-state cycle is a sequence of four states with the property that H sellers send the same signal throughout the cycle, denoted $\hat{s}_H$ (see (11)), while L sellers alternate between two signals, $\hat{s}_L$ or $\hat{s}_H$ (see (11)–(12)). The four-state cycle takes the form of L sellers first sending the signal $\hat{s}_L$ (state θ'). Buyers then adjust strategies to reach the competitive successor $c(\theta')$ of state θ', which involves increasing the price at $\hat{s}_H$ to the competitive price $p_1^*(\hat{s}_H)$. This makes signal $\hat{s}_H$ attractive, prompting L sellers to switch to signal $\hat{s}_H$ (see (11) and (13)) while buyer actions remain unchanged (see (16)). This situation corresponds to state θ. Buyers again adjust to the competitive successor, $c(\theta)$, this time reducing the price at $\hat{s}_H$ to $p_{\phi^0}^*(\hat{s}_H)$. This prompts L sellers to switch back to $\hat{s}_L$ (see (14)) while buyer actions remain unchanged (see (15)), yielding state θ' and starting the cycle anew.

Given a four-state cycle F we let $s_L(F)$ denote the signal used only by L sellers and $s_H(F)$ the signal used by all H sellers in the cycle. Let $U_L(F) = p_0^*(s_L(F)) - c_L(s_L(F))$ denote the utility level type L sellers achieve from choosing $s_L(F)$.

From the construction of the four-state cycle it is clear that there must exist $p_F \in [p_{\phi^0}^*(s_L(F)), p_1^*(s_L(F))]$ such that

$$U_L(F) = p_F - c_L(s_h(F)). \tag{17}$$

This observation suggests that four-state cycles should correspond to equilibria of the underlying game G. This is indeed the case, though the equilibria in question are generically mixed:[22]

Remark 1 *Let $F = \{\theta', c(\theta'), \theta, c(\theta)\}$ be a four-state cycle. Then there exists a mixed strategy equilibrium of the game G, in which H sellers choose $s_H(F)$, buyers mix between prices so as to render L sellers indifferent between signals $s_L(F)$ and $s_H(F)$, and L sellers mix between these signals.*

4 Limit Sets and Stable Outcomes

We now investigate the perturbed learning dynamics. The absorbing sets of the perturbed dynamics $\Gamma(\lambda)$ are independent of λ (as long as $\lambda > 0$). We can thus write $\{A_1, \cdots A_k\}$ for the collection of absorbing sets of $\Gamma(\lambda)$ for $\lambda > 0$. We let $\{A_1(0), \cdots A_n(0)\}$ denote the collection of absorbing sets of the learning dynamics $\Gamma(0)$ and let $A(0) = \cup_i A_i(0)$. Notice that the process $\Gamma(\lambda)$ may have fewer absorbing sets than the learning dynamics, because positive replacement rates may allow the system $\Gamma(\lambda)$ to move from one absorbing set of the learning dynamic $\Gamma(0)$ to another.

The process $\Gamma(\lambda)$ may have many absorbing sets. In particular, our model places considerable structure on the process by which mutations alter agents' strategies, so that there may be transitions between states that cannot be accomplished via mutation. This is in contrast to previous work with perturbed dynamic systems, such as Kandori, Mailath and Rob [12], Young [26], Samuelson [19], and Nöldeke and Samuelson [16], where the mutation process is generally structured to ensure that $\Gamma(\lambda)$ has a single absorbing set containing all states.

For every A_j and $\lambda > 0$ there is a unique stationary distribution of the Markov process $\Gamma(\lambda)$, denoted $\zeta_j(\lambda)$, with support A_j. Since $\Gamma(\lambda)$ is aperiodic,[23] it follows that conditional on reaching the absorbing set A_j the Markov process $\Gamma(\lambda)$ converges to $\zeta_j(\lambda)$.

We are interested in the behavior of our evolutionary model as the probability of mutations becomes small. We accordingly focus on the **limit distributions** $\lim_{\lambda \to 0} \zeta_j(\lambda)$. We will be interested in characterizing the supports of the limit distributions, which we call limit sets:[24]

---

[22] The proof of the following result is somewhat involved and not directly relevant to the remainder of the paper, and so is omitted. The difficulty in proving this result is that generically (the genericity is with respect to Lebesgue measure in the Euclidean space that describes the finite number of payoffs in the model), p_F will fail to be in P and buyers will have to mix over price offers at $s_H(F)$ to make type L sellers indifferent between $s_L(F)$ and $s_H(F)$. Because commodities are sold only at the maximal price offered at a given signal, it is not completely trivial to show that the appropriate mixture can indeed arise as an equilibrium outcome.

[23] Aperiodicity follows from the fact that, for any state, there is positive probability that no agents are replaced and no agents receive the learn draw, causing the system to remain in its current state.

[24] We concentrate on the supports of the limiting distribution because, unlike the details of the limit distributions, limit sets are independent of the probability of receiving the learn draw μ and independent of

Definition 6 *A set $L^* \subset \Theta$ is a* ***limit set*** *if it is the support of a limit distribution.*

Each absorbing set A_j contains a unique limit set. Lemma 3, below, shows that limit sets will usually contain large numbers of states. These states may well produce a variety of outcomes. We will be especially interested in whether the evolutionary model selects a "stable" outcome of G in the sense that all states appearing in a limit set of G produce the same outcome:

Definition 7 *If there is an outcome ρ^* such that every state in a limit set L^* results in outcome ρ^*, then we say that ρ^* is a* ***stable outcome*** *for G or that ρ^* is* ***stable****.*

Even if every limit set yields a stable outcome, there may be multiple stable outcomes. Which stable outcome characterizes the limiting behavior of the evolutionary process will then depend on initial conditions and the realization of the random learning and mutation processes. We can avoid these complications if there is only one stable outcome:

Definition 8 *If every limit set yields the stable outcome ρ^*, then we say that ρ^* is the* ***global outcome*** *of G or that ρ^* is* ***globally stable****.*

We now turn to the question of which outcomes are stable outcomes or global outcomes. Our key result here is that a stable outcome must be a "locally stable" outcome. To define this latter concept, we first identify absorbing sets of the learning dynamics that can be linked together by single mutations. Let $M(\theta)$ be the set of all states that can be reached from state θ with a single mutation. For $X \subset \Theta$, let $M(X)$ denote the union of the sets $M(\theta)$ for $\theta \in X$. We then have the following definition of a locally stable component, which is equivalent to the one given in Nöldeke and Samuelson [16]:

Definition 9 *A non-empty set $C \subseteq A(0)$ is a* ***locally stable component*** *if it is minimal with respect to the property:*

$$\forall \theta' \in A(0) \setminus C : M(C) \cap B(\theta', 0) = \emptyset. \tag{18}$$

Outcome ρ^* is a ***locally stable outcome****, or ρ^* is* ***locally stable*** *if there exists a locally stable component C such that the outcome of every state in C is ρ^*.*

Condition (18) requires that a locally stable component be stable against single mutations in the following sense: from any state that can be reached from a locally stable component via a single mutation, the learning dynamics must lead back to a state in the locally stable component.[25] (Because $C \subseteq A(0)$, the learning dynamics can never cause the market

---

the probability distributions from which best responses and mutants' conjectures are selected.

[25] The minimality requirement in the definition of a locally stable component plays the same role as the minimality requirement in the definition of an absorbing set. It ensures that a locally stable component does not contain a subset of states from which it is impossible to reach the other states in the component with a single mutation.

to leave a locally stable component without mutations.) Every absorbing set A_j of the perturbed dynamics must contain at least one locally stable component: A_j must contain at least one absorbing set of the learning dynamic. Since it is not possible to leave A_j with any number of mutations it then follows that A_j must also contain a locally stable component.

The following result states that limit sets are unions of locally stable components.

Proposition 3 *Let θ be contained in a limit set L^*. Then θ is contained in a locally stable component C. Furthermore, all $\theta' \in C$ are contained in the limit set L^*.*

The proof of Proposition 3 is a straightforward modification of arguments developed by Young [26] and Samuelson [19], and is omitted. To see the intuition behind this result (which is explained in some detail in Nöldeke and Samuelson [16]), note that the system can move from one absorbing set of the learning dynamics to another only via mutations. As the probability of a mutation gets arbitrarily small, limit sets will consist of those absorbing sets of the learning dynamics that are the "easiest to reach" in the sense that it takes the fewest mutations to reach them from other absorbing sets of the learning dynamics. If states θ and θ' are both singleton absorbing sets, and if a single mutation (plus subsequent learning) suffices to move the system from θ to θ', then θ' must be at least as easy to reach as θ. The absorbing set θ will then appear in a limit set only if θ' does. A similar statement applies to nonsingleton absorbing sets. Hence, if one state of a locally stable component appears in a limit set, so do all states in the locally stable component. Furthermore, it takes fewer mutations to move from an absorbing set that is *not* part of a locally stable component to one that *is* contained in a locally stable component than it takes to move in the reverse direction, so only locally stable components appear in limit sets. Hence, limit sets are unions of locally stable components.

Proposition 3 leads to a simple necessary condition for the stability of an outcome:

Corollary 1 *Outcome ρ^* is a stable outcome only if it is a locally stable outcome. If every state in every locally stable component results in an outcome ρ^* then ρ^* is the global outcome of G.*

From Corollary 1, an obvious way to begin our search for stable and global outcomes is to investigate locally stable components. We will find the implications of local stability to be sufficiently strong that little further investigation is needed.

In verifying that a set of states is a locally stable component, we will often exploit the fact that if a locally stable component contains an equilibrium state it must contain all equilibria that differ only in the prices offered at unused signals and agents' conjectures about these prices:

Lemma 3 *Let C be a locally stable component and let $\theta \in C$ be an equilibrium. Then C contains all equilibria θ' such that*

$$a \in A_L \cup A_H \Rightarrow \sigma_a(\theta') = \sigma_a(\theta). \tag{19}$$

Proof: Consider equilibrium state θ in a locally stable component C. Let equilibrium state θ' satisfy (19). Then θ and θ' differ only in prices offered at unused signals and conjectures about these prices. Let a single mutation occur that for some agent a changes $\sigma^a(\theta)$ to $\sigma^a(\theta')$. This mutation does not affect the signal agent a sends (if a is a seller) or the prices that a offers at signals that are sent (if a is a buyer) and thus yields an equilibrium θ_1 with $\rho(\theta_1) = \rho(\theta)$. Hence, $\theta_1 \in M(\theta)$ and since θ_1 is an absorbing state of the learning dynamics, $\theta_1 \in C$. By a sequence of such transitions, one can move from θ to θ', giving $\theta' \in C$. $\square$

Finally, to be a locally stable outcome is a necessary condition for an outcome to be a stable outcome, but is not sufficient, because (a) a limit set may consist of multiple locally stable components yielding different outcomes and (b) there may be locally stable components that do not appear in limit sets. However, Corollary 1 indicates that if all locally stable components yield outcome ρ^*, then ρ^* is not only a stable outcome, but also the global outcome. A special case to which this corollary obviously applies is the following:

Remark 2 *Let there be a unique locally stable component. Then this component must be the unique limit set. If all states in the locally stable component give outcome* ρ^**, then* ρ^* *is the global outcome.*

5 Necessary Conditions for Stability

5.1 Minimal Equilibria

This section begins our investigation of stability by deriving a simple necessary condition for local stability of an equilibrium outcome, namely that the equilibrium outcome be *minimal*.

An equilibrium state θ is said to be **separating** if $S_L(\theta) \cap S_H(\theta) = \emptyset$. An equilibrium is said to be **quasi-separating** if $S_L(\theta) \not\subseteq S_H(\theta)$ and $S_H(\theta) \not\subseteq S_L(\theta)$, so that for each type of seller there is at least one signal sent by only that type. An equilibrium is said to be **pooling** if $S_L(\theta) = S_H(\theta)$ and is said to be **quasi-pooling** if $S_L(\theta) \subseteq S_H(\theta)$. If an equilibrium is neither quasi-pooling nor quasi-separating (i.e., $S_H \subset S_L$ but $S_L \not\subset S_H$), then we say that it is **partially separating**. We are led to these definitions by Proposition 4 below, which shows that quasi-separating and quasi-pooling equilibria are essentially disguised separating or pooling equilibria. The same is not true for partially separating equilibria. We will often devote attention to equilibria that are either separating or partially separating, which we will denote as **(partially) separating**.[26]

---

[26] If a market has no equilibrium, then it has no singleton absorbing sets (Proposition 1) and every limit set must be a union of nonsingleton absorbing sets of the learning dynamics. There is then no stable outcome. The most interesting games from our perspective are thus ones with an equilibrium outcome.

Neither $S_L(\theta)$ nor $S_H(\theta)$ needs to be a singleton in order for θ to be an equilibrium, though $S_L(\theta) \cap S_H(\theta)$ must be a singleton in every equilibrium (from Assumption 1). In particular, $S_L(\theta) = S_H(\theta)$ must be a singleton in any pooling equilibrium. In order for $\rho(\theta)$ to be a stable outcome, however, mutations must not allow indifferent agents to "drift" between signals. To make this idea precise, we first need a definition:

Definition 10 *A pure strategy equilibrium θ is* ***minimal*** *if θ is one of the following:*

> *(10.1) A separating or pooling equilibrium in which both $S_L(\theta)$ and $S_H(\theta)$ are singletons.*
>
> *(10.2) A partially separating equilibrium in which $S_H(\theta)$ is a singleton and $S_L(\theta)$ contains two elements (one of which must be the signal contained in $S_H(\theta)$).*

Notice that in any minimal equilibrium, there is only one signal used by H sellers, which we will denote $s_H(\theta)$. In a partially separating or pooling equilibria there will also be L sellers using $s_H(\theta)$. We will find it convenient to let $s_L(\theta)$ denote the signal (if it exists) used only by L sellers in a minimal equilibrium.

By Assumption (1), all pooling equilibria must be minimal. Generically, *all* pure strategy equilibria in the game G are minimal and satisfy (10.1). Additional pure strategy equilibria arise only out of coincidental payoff ties that render agents indifferent between signals. While it would be easy to rule out such indifferences by assumption, we prefer not to impose such an assumption since the partially separating equilibria described in (10.2) mimic mixed strategy equilibria in the Cho-Kreps game.

We now show that every component containing an equilibrium also contains a minimal equilibrium of the "same type":

Proposition 4 *Suppose θ is a quasi-separating (quasi-pooling, partially separating) equilibrium state. Then there exists a minimal separating (pooling, partially separating) equilibrium θ' such that*

$$\begin{aligned} S_T(\theta') &\subseteq S_T(\theta) &\quad (20)\\ U_T(\theta') &\geq U_T(\theta). &\quad (21)\end{aligned}$$

Furthermore, if θ is contained in a locally stable component C, then θ' is also contained in C.

Proof: See Appendix. □

To see why this result holds, consider first an equilibrium θ in which there are multiple signals that are sent by type $T \in \{L, H\}$ and only type T sellers. Then any mutation switching a type T seller from one to another of these signals does not affect utilities and gives a new equilibrium contained in the same component as θ. A sequence of such mutations yields an equilibrium in the component containing θ in which there is at most one signal sent by only type L sellers, at most one signal sent by only H types, and at most one signal sent by both types (from Assumption 1). Let this new equilibrium be called θ'. If θ is quasi-separating, then all three signals are sent in θ'. A further sequence of mutations and learning can then switch sellers away from the signal sent by both types to the other two signals, yielding a minimal separating equilibrium. If θ is quasi-pooling, then only the latter two signals are sent in θ'. A further sequence of mutations and learning can then switch sellers from the signal sent only by H types to the signal sent by both types, yielding a minimal pooling equilibrium. If θ is partially separating, then only the first and third signals are sent in θ'. In this case, θ' is already a minimal partially separating equilibrium. Note that it need not be the case that a partially separating equilibrium is contained in a component that also contains either a minimal separating or a minimal pooling equilibria.

Propositions 4 (for singleton absorbing sets) and 2 (for nonsingleton absorbing sets) imply that every locally stable component contains a state in which at most two signals are used. Such "minimal" states are the ones which are most vulnerable to mutations (which only affect unused signals) and thus provide the obvious starting point to test whether a given component is locally stable. In addition, the following is immediate from Proposition 4:

Corollary 2 *Every locally stable outcome is the outcome of a minimal equilibrium.*

5.2 Sequential Equilibria

Our definition of equilibrium imposes no constraints on choices at signals that are not used in an equilibrium state, and no constraints on conjectures about these choices, beyond the requirement that the unused signals not be conjectured to be best replies. By having sellers conjecture prices below p_0^* at unsent signals, we may then be able to construct equilibria which cannot be supported as sequential equilibrium outcomes (i.e., are not contained in $\tilde{\Theta}$).

It would appear as if such nonsequential equilibrium outcomes cannot be stable outcomes: There must be a single mutation by some seller, endowing him with a conjecture that a price at least as high as $p_0^*(s)$ is offered at each out-of-equilibrium signal s, that triggers a deviation; with the learning dynamics then apparently raising the prices at unused signals to at least p_0^* and hence leading to an absorbing set contained in $\tilde{\Theta}$. However, it may be the case that the only way to reach an absorbing set of the learning dynamics after an initial mutation is for a crash to reduce the price at some signal s below $p_0^*(s)$. The resulting absorbing set may then not be a sequential equilibrium and sequentiality is *not* a necessary condition for stability.[27]

---

[27] Both sequentiality of locally stable outcomes and Lemma 4 below automatically obtain if the learning dynamics is restricted to states in $\tilde{\Theta}$. With this restriction the game in Example 1 has a unique non-singleton

---

absorbing set containing its four-state cycle.

To see the issues involved here, consider the following example:

	$\underline{p}$	L	P	H	$\underline{p}$	L	P	H
U_L	-10	-2	1	8	-14	-6	-3	4
U_H	-10	-2	1	8	-6	2	5	12
		s_0				s_1		

Example 1

There are two signals, s_0 and s_1. At each signal there are four prices, including price $\underline{p}$, the competitive price when facing only L sellers, the competitive price when facing the pool of all sellers, and the competitive price when facing only H sellers. We label the latter three prices L, P, and H, and then need report only seller payoffs. The payoffs given are those of an L seller (top row) and H seller (bottom row) for each signal and price combination.

This market has no separating equilibria. There are pooling equilibria, in which signal s_0 is sent and price $\underline{p}$ must be conjectured at s_1 by all H sellers, so that the corresponding equilibrium outcome ρ^* is not a sequential equilibrium outcome. There is also a unique four-state cycle, corresponding to a mixed strategy equilibrium in which H sellers send s_1 and L sellers mix between s_0 and s_1. This four-state cycle is not contained in an absorbing set, instead lying in the basin of attraction of a pooling equilibrium giving outcome ρ^*. In particular, at that point in the four-state cycle in which L sellers have just switched from signal s_0 to s_1, there is a positive probability that the market price at signal s_1 falls to $\underline{p}$ and all sellers switch to s_0, leading to outcome ρ^*. By Proposition 2 this implies that all absorbing sets of the learning dynamics are singletons, resulting in the pooling outcome. Consequently, the non-sequential outcome ρ^* is the unique locally stable outcome of the game and also the global outcome of the game.[28]

The non-sequential pooling outcome ρ^* in Example 1 can be supported by conjectures for L (but not H) sellers which are concentrated on competitive prices. The following lemma shows that this is no coincidence: every locally stable component must contain a state in which L sellers obtain a utility level which is at least:

$$\bar{U}_L = \max_{s \in S} p_0^*(s) - c_L(s),$$

and hence which can be supported by "sequential conjectures" for L sellers.

Lemma 4 *Every locally stable component contains either an equilibrium θ satisfying*

$$U_L(\theta) \geq \bar{U}_L \tag{22}$$

or a four-state cycle F satisfying

$$U_L(F) = \bar{U}_L. \tag{23}$$

---

[28] One may conjecture that the non-sequential pooling equilibrium is a global outcome in Example 1 only because the market has no pure strategy sequential equilibrium outcome. However, there exist games (with more signals) in which a nonsequential separating equilibrium is the global outcome, even though a pure strategy (pooling) sequential equilibrium exists.

Proof: See Appendix. □

The inequality in (22) must be an equality if θ is (quasi- or partially) separating.

5.3 Undefeated Equilibria

Lemma 4 establishes a lower bound for the utility level of L sellers in a locally stable outcome. The results in this section are driven by considerations involving H sellers. In particular, it is shown that an equilibrium outcome θ fails to be a locally stable outcome if it is the case that a mutation could cause H sellers to deviate to their equilibrium signal in a different minimal equilibrium θ' in which they receive a (weakly) higher utility level than in θ. Formally, let us define:

Definition 11 *A minimal equilibrium θ' H-**dominates** a minimal equilibrium θ if*

$$U_H(\theta') \geq U_H(\theta) \text{ and } s_H(\theta') \notin S_L(\theta) \cup S_H(\theta). \tag{24}$$

*If there is no minimal equilibrium H-dominating θ then we say that θ is H-**dominant**.*

Note that the definition of H-dominance differs from more standard dominance notions in that the H sellers' equilibrium signal in θ' must be unused in θ. This condition is required so that mutations can work on $s_H(\theta')$ to lead the market from θ to the H-dominating equilibrium.[29]

If θ is not H-dominant, then $\rho(\theta)$ cannot be a stable outcome because a single mutation can cause an H seller to send the signal that he sends in the equilibrium θ' that H-dominates θ. Because this gives H sellers a higher payoff than equilibrium θ, learning dynamics ensue that can lead to θ', precluding the stability of $\rho(\theta)$. More precisely:

Proposition 5 *If $\rho(\theta)$ is a locally stable outcome then θ is H-dominant.*

Proof: See Appendix. □

Together, Proposition 4, Lemma 4, and Proposition 5 have strong implications for the number of locally stable outcomes:

Corollary 3 *No game G has more than two locally stable outcomes. If a game has two locally stable outcomes, ρ and ρ', then one of them is the outcome of a minimal pooling equilibrium θ' and the other is the outcome of a minimal (partially) separating equilibrium θ satisfying*

$$s_L(\theta) = s_H(\theta').$$

---

[29] The ability of unused signals to upset equilibria and the conditions under which such signals exist has been a major theme of the cheap talk literature (see, e.g., Nachbar [14], Blume, Kim and Sobel [2], and Sobel [20]).

Proof: We first note that there can be at most one locally stable pooling equilibrium outcome. In particular, if there were more than one such outcome, then one would H-dominate the other, contrary to Proposition 5.

We next show that there can be at most one locally stable (partially) separating equilibrium. Suppose θ is (partially) separating and $\rho(\theta)$ is locally stable. Then θ is minimal (Corollary 2), and H-dominant (Proposition 5). Let $\bar{S}_L = \{s : p_0^*(s) - c_L(s) = \bar{U}_L\}$. We argue that $s_L(\theta)$ is the unique signal in $\bar{S}_L$. First, Lemma 4 implies that $s_L(\theta) \in \bar{S}_L$. Second, by Assumption 2,

$$p^*_{\phi(\theta,s_H(\theta))}(s_H(\theta)) > p_0^*(s_H(\theta)),$$

and since $U_L(\theta) = \bar{U}_L \geq p^*_{\phi(\theta,s_H(\theta))}(s_H(\theta)) - c_L(s_H(\theta))$ (because θ is an equilibrium), we must then have $s_H(\theta) \notin \bar{S}_L$. So, if $s_L(\theta)$ were not the unique element in $\bar{S}_L$ then any other $s' \in \bar{S}_L$ would be unused in the minimal equilibrium θ. Suppose such a signal s' exists. Consider a single mutation that causes an L seller (by endowing him with the conjecture that a price $p_0^*(s')$ is offered at s') to switch from $s_L(\theta)$ to s'. Then let a competitive learn sequence occur. If

$$U_H(\theta) \geq p_0^*(s') - c_H(s'),$$

then a subsequent learn draw for all sellers establishes an equilibrium in which the L seller continues to send s', contradicting local stability of ρ. If this inequality is reversed, then there must exist a pooling equilibrium at signal s' that H-dominates θ, again contradicting the local stability of θ (Proposition 5). Hence, $s_L(\theta)$ is the unique element of $\bar{S}_L$.

Now suppose there exists a locally stable outcome $\rho(\theta') \neq \rho(\theta)$. If θ' is (partially) separating then the preceding argument shows $s_L(\theta') = s_L(\theta)$. Furthermore, $s_H(\theta) = s_H(\theta')$ since otherwise θ and θ' cannot both be H-dominant. This in turn implies

$$p^*_{\phi(\theta,s_H(\theta))}(s_H(\theta)) = p^*_{\phi(\theta',s_H(\theta))}(s_H(\theta)),$$

and it is then easy to see that θ and θ' must lie in the same locally stable component, contradicting the local stability of both $\rho(\theta)$ and $\rho(\theta')$. Hence, there can be at most one (partially) separating locally stable equilibrium outcome.

Finally, if θ is a (partially) separating, locally stable outcome and θ' is a partially stable, pooling equilibrium outcome, then θ' must be a pool at either $s_L(\theta)$ or $s_H(\theta)$, since otherwise either θ or θ' would fail to be H-dominant. If θ' were a pool at $s_H(\theta)$, then, from Assumption 2, $U_L(\theta') \leq U_L(\theta) = \bar{U}_L$. A strict inequality contradicts Lemma 4 for θ'. If equality holds, θ and θ' are contained in the same component, contradicting local stability of both $\rho(\theta)$ and $\rho(\theta')$. □

Notice that this proof establishes:

Remark 3 *No (partially) separating outcome is a locally stable outcome if there are multiple signals s at which* $\bar{U}_L = p_0^*(s) - c_L(s)$.

Example 3 below shows that markets exist in which there are two stable outcomes.

Mailath, Okuno-Fujiwara and Postlewaite [13] have recently introduced the notion of one equilibrium defeating another in a Cho-Kreps signaling game. The following definition is a slight modification of the one offered by Mailath, Okuno-Fujiwara and Postlewaite [13].[30]

Definition 12 *Let states θ and θ' be minimal equilibrium states in Θ. Then θ'* ***defeats*** *θ if there exists a signal s', unused in θ but used in θ' such that, for $T \in \{L, H\}$,*

$$s' \in S_T(\theta') \Rightarrow U_T(\theta') \geq U_T(\theta) \tag{25}$$

with at least one strict inequality. A minimal equilibrium θ is ***undefeated*** *is there is no minimal equilibrium defeating it.*

The intuition behind this definition is that, beginning with the defeated equilibrium, a seller can send an out-of-equilibrium signal that prompts belief revisions that not only lead away from the defeated equilibrium (as is the case in most refinements) but also leads to the defeating equilibrium. While Mailath, Okuno-Fujiwara, and Postlewaite develop this intuition in a model with rational agents, the following result shows that our evolutionary approach provides an alternative justification for focusing on undefeated equilibria.

Corollary 4 *If $\rho(\theta)$ is a locally stable outcome, then equilibrium θ is undefeated.*

Proof: Let $\rho(\theta)$ be a locally stable outcome and let θ be defeated by θ'. From Lemma 4, we have $U_L(\theta) \geq \bar{U}_L$. Then θ' can defeat θ only if $U_H(\theta') \geq U_H(\theta)$ and $s_H(\theta') \notin S_L(\theta) \cup S_H(\theta)$, ensuring that θ is not H-dominant and hence (Proposition 5) that $\rho(\theta)$ is not a locally stable outcome. □

6 Stable Separating Equilibria

The results in the previous section establish necessary conditions for an outcome to be a locally stable outcome. In this section we search for sufficient conditions under which a separating equilibrium outcome is either a locally stable outcome or the globally stable outcome.

From Remark 3 it is without loss of generality to require, throughout this section:

Assumption 5 *There exists a* ***unique*** *signal s_L^* such that*

$$\bar{U}_L = p_0^*(s_L^*) - c_L(s_L^*). \tag{26}$$

---

[30] Mailath, Okuno-Fujiwara and Postlewaite restrict attention to sequential equilibria in the Cho-Kreps game and require that $S_T(\theta)$ and $S_T(\theta')$ both be singletons for each T (so that they exclude minimal partially separating equilibria from the discussion). They also require more stringent conditions for defeat if the inequality in (25) is weak for one type.

Our first Lemma combines previous results to establish that we can limit our search to outcomes that are Riley outcomes in the sense of the following definition:

Definition 13 *An outcome ρ^* is the* ***Riley outcome*** *of G if there exists a minimal separating equilibrium θ with $\rho(\theta) = \rho^*$ such that*

$$U_L(\theta) = \bar{U}_L$$

and for all quasi-separating equilibria θ' with $\rho(\theta') \neq \rho^*$:

$$U_H(\theta') < U_H(\theta).$$

Note that a game satisfying condition (26) cannot have more than one Riley outcome,[31] so that we can speak of *the* Riley outcome. We let s_T^* denote the signal used by type T in the Riley outcome.

Lemma 5 *Suppose ρ^* is a locally stable outcome of a quasi-separating equilibrium. Then ρ^* is the Riley outcome of G.*

Proof: Let θ be a quasi-separating equilibrium with locally stable outcome ρ^*. Then θ is minimal (Corollary 2) and hence separating, and $U_L(\theta) = \bar{U}_L$ (Lemma 4). Suppose there exists a quasi-separating equilibrium θ' with $\rho(\theta') \neq \rho^*$ satisfying $U_L(\theta') = \bar{U}_L$ and $U_H(\theta') \geq U_H(\theta)$. Since θ is H-dominant (Proposition 5), it cannot be that θ' is separating and satisfies $S_H(\theta') \neq \{s_H(\theta)\}$; since otherwise there would exist a minimal separating equilibrium H-dominating θ (cf. the proof of Proposition 4). Similarly, if θ' is not separating, it must satisfy $S_T(\theta') = \{s_T(\theta), s'\}$ for some $s' \in S \setminus \{s_L(\theta), s_H(\theta)\}$. It is then easy to see that θ' and θ are contained in the same locally stable component, contradicting the fact that θ is a locally stable outcome. Hence ρ^* must be the Riley outcome. □

The only possible locally stable separating equilibrium outcome is thus the Riley outcome. From Proposition 5 and Corollary 4, however, we know that a Riley outcome need not be locally stable. In particular, there may be pooling equilibria H-dominating or defeating the Riley outcome, in which case our market does *not* choose the Riley outcome.

Even an H-dominant Riley outcome may fail to be locally stable. Consider the following example:

	L	P	H	L	P	H
U_L	2	5	8	-8	-3	2
U_H	2	5	8	-4	1	6
		s_0			s_1	

Example 2

---

[31] This outcome may be produced by a large number of different states, where these states differ in conjectures and behavior at unsent signals.

This game has a pooling equilibrium in which all sellers send signal s_0, yielding payoffs $(U_L, U_H) = (5, 5)$; as well as a separating equilibrium in which L sellers send signal s_0 while H sellers send s_1, for payoffs $(2, 6)$. The separating equilibrium yields an H-dominant Riley outcome but is not a locally stable outcome. In particular, given the separating equilibrium, mutations can shift L sellers to signal s_1 without leaving the component containing the separating equilibrium outcome. This can continue until such a mutation takes the system outside this component, at which point buyers are prompted to switch to price P at signal s_1. Learning can now induce all sellers to switch to signal s_0, giving the pooling equilibrium and precluding local stability of the separating equilibrium.

The key to the instability of the Riley outcome in Example 2 is that L type sellers are indifferent between their equilibrium signal and the signal sent by H sellers. If we exclude such indifferences, we achieve local stability:

Proposition 6 *Let $\rho^* = \rho(\theta^*)$ be an H-dominant Riley outcome. If θ^* satisfies*

$$p_0^*(s_L^*) - c_L(s_L^*) \quad > \quad p_1^*(s_H^*) - c_L(s_H^*) \tag{27}$$

$$p_1^*(s_H^*) - c_H(s_H^*) \quad > \quad p_0^*(s_L^*) - c_H(s_L^*), \tag{28}$$

then ρ^* is a locally stable outcome.

To see the intuition behind this result, let state θ^* correspond to the Riley outcome. While a mutation can result in a state that induces sellers to leave their equilibrium signals, subsequent adjustments can never cause buyers to deviate from the competitive prices that they offered at the equilibrium signals. The essential steps in establishing this are that the equilibrium θ^* is separating and (27)–(28) hold with strict inequality. In particular, separation ensures that no matter which subset of sellers leave their equilibrium signal, the competitive price at the signal remains unchanged; and the strict inequalities ensure that all sellers will prefer to return to their equilibrium signals instead of mimicking the equilibrium signal of the other type.[32] Hence, just as it is presumed in most refinement concepts, the original equilibrium payoff always remains "available" for sellers after a deviation from a separating equilibrium has occurred. The availability of the original equilibrium payoff ensures that the learning dynamics cannot converge to an equilibrium in which some type of the seller receives a lower payoff than in the original equilibrium. The remainder of the proof then consists of showing that because the separating equilibrium yielding the Riley outcome is H-dominant, the learning dynamics cannot converge to an equilibrium (or reach a non-singleton absorbing set) that gives some type of seller a higher payoff. Consequently, the learning dynamics must return to a state yielding the Riley

---

[32] Notice that if $p_1^*(s_H^*) = p^*_{|A_H|/(|A_H|+1)}(s_H^*)$, which will hold if the population of sellers is sufficiently large, then (27)-(28) are necessary for the local stability of a separating equilibrium outcome; since otherwise mutations could allow indifferent sellers to drift between equilibrium signals without affecting prices, leading to new equilibria that vitiate the local stability of θ^*.

outcome. Hence, there is a locally stable component in which every state yields outcome $\rho(\theta^*)$.

Proof: It suffices to show that there exists a set $X \in \Theta$ with the following properties:

> **(6.1)** X is closed under the learning dynamic $\Gamma(0)$.
>
> **(6.2)** If θ is a singleton absorbing set of $\Gamma(0)$ with $\rho(\theta) = \rho^*$, then every state that can be reached from θ via a single mutation is contained in X.
>
> **(6.3)** All absorbing sets of $\Gamma(0)$ contained in X are singletons giving outcome ρ^*.

Let X be the set of states θ that satisfy:

$$\sigma_a = s_T^* \Rightarrow a \in A_T, \quad T = L, H \tag{29}$$

$$a \in A_B \Rightarrow \sigma_a(s_L^*) = p_0^*(s_L^*) \tag{30}$$

$$a \in A_B \Rightarrow \sigma_a(s_H^*) = p_1^*(s_H^*) \tag{31}$$

$$a \in A_T \Rightarrow \sigma_{-a}(s_L^*, p_0^*(s_L^*)) = |A_B|, \quad T = L, H \tag{32}$$

$$a \in A_T \Rightarrow \sigma_{-a}(s_H^*, p_1^*(s_H^*)) = |A_B|, \quad T = L, H. \tag{33}$$

(29) states that any sellers that send signal s_L^* (if any do) are L type sellers, while s_H^* is sent (if at all) by H sellers. (30)-(31) state that all buyers offer the competitive price at signal s_L^* (given that only L types send this signal) and s_H^* (given that only H types send this signal). (32)-(33) state that seller conjectures match this behavior.

(6.1) We need to show that X is closed under the learning dynamics, i.e., $\forall \theta \in X$ and $\theta' \in \Theta$, we have $\theta \in B(\theta', 0) \Rightarrow \theta' \in X$. To verify this, fix a state $\theta_t \in X$ and consider the state θ_{t+1} to which the learning dynamics might move the market. Notice first that a buyer who receives the learn draw will change conjectures about the sellers who send signals s_L^* and s_H^* only if some sellers send these signals. Because only type T sellers send signal s_T^* (from (29)), the new conjectures must make prices $p_0^*(s_L^*)$ and $p_1^*(s_H^*)$ optimal, and hence conditions (30) and (31) must hold in θ_{t+1}. Because buyers' offer prices according to (30) and (31) in state θ_t, then sellers' conjectures about the prices at s_L^* and s_H^* cannot change, so that conditions (32) and (33) must be satisfied in θ_{t+1}. Given these conjectures, by (27) s_L^* is a better response for L sellers than s_H^* and by (28) s_H^* a better response for H sellers than s_L^*. Hence, no seller will adopt the equilibrium signal of the other type, implying that (29) also holds in θ_{t+1} and giving the result.

(6.2) Next, suppose θ is an absorbing state of the learning dynamics and $\rho(\theta) = \rho^*$. Then by construction θ is contained in X. In addition, conditions (27) and (28) ensure that a single mutation cannot induce an L seller to adopt signal s_H^* or an H seller to adopt s_L^*. Nor, since these signals are used, can a mutation affect prices at these signals or conjectures about these prices. Hence, every state that can be reached from θ via a single mutation is contained in X.

(6.3) It remains to show that if an absorbing set of the learning dynamics is contained in X, then that absorbing set is a singleton resulting in the Riley outcome.

Absorbing states. First, we show that there cannot be an absorbing *state* $\theta' \in X$ such that $\rho(\theta') \equiv \rho' \neq \rho^*$. If such a state were to exist, by Proposition 1, ρ' is an equilibrium outcome. Because $\theta' \in X$, we have $U_L(\theta') \geq U_L(\theta^*)$ and $U_H(\theta') \geq U_H(\theta^*)$ (because a type T seller could earn payoff $U_T(\theta^*)$ by choosing signal s_T^*) and $s_T^* \notin S_{T'}(\theta')$ for $T \neq T'$. Now suppose $s_H^* \notin S_H(\theta')$. By Proposition 4 there then exists a minimal equilibrium θ'' with $s_H(\theta'') \notin \{s_L^*, s_H^*\}$ and $U_H(\theta'') \geq U_H(\theta') \geq U_H(\theta^*)$, contradicting the assumption that θ is H-dominant. It then follows that $s_H^* \in S_H(\theta')$. We now consider three possibilities.

Case 1: θ' cannot be quasi-separating, because the existence of a quasi-separating equilibrium outcome satisfying $U_L(\theta') \geq U_L(\theta^*)$ and $U_H(\theta') \geq U_H(\theta^*)$ contradicts the assumption that ρ^* is the Riley outcome.

Case 2: θ' cannot be pooling, because every pooling equilibrium is minimal and hence $s_H^* \in S_H(\theta')$ implies $s_L(\theta) = s_H^*$, contradicting $\theta' \in X$. If θ' were quasi-pooling, there would exist a unique signal $s' \notin \{s_L^*, s_H^*\}$ used by both types in θ' and the proof of Proposition 4 then shows that there exists a pooling equilibrium at s' H-dominating the Riley outcome, a contradiction.

Case 3: θ' cannot be partially separating, because $S_H(\theta') \subseteq S_L(\theta')$ contradicts the assumption $\theta' \in X$.

Nonsingleton absorbing sets. Second, there cannot be a nonsingleton absorbing set Y in X. If such a set existed, then Proposition 2 would allow us to conclude that there exists a four-state cycle $F \subseteq Y$. Because $F \subseteq X$, we must have $U_L(F) \geq U_L(\theta^*)$ and hence (by Assumption 5) $s_L(F) = s_L^*$. In addition, $s_H(F) \neq s_H^*$ (because $s_H(F) = s_H^*$ would make it impossible for L sellers to switch to s_H^* (by (27)). Furthermore, we must have, for $T \in \{L, H\}$,

$$\underline{p} - c_T(s_H(F)) < U_T(\theta^*). \tag{34}$$

For type H, this inequality follows from the fact that θ^* is an equilibrium and Assumption 4. For type L, the inequality follows because it is a best response to switch to $s_L(F) = s_L^*$ when $p^*_{\phi^0}(s_H(F)) > \underline{p}$ is offered at $s_H(F)$. But these inequalities contradict the assumption that the four-state cycle F is contained in an absorbing set: Consider the state θ in F, and let a crash reduce the price at $s_H(F)$ to $\underline{p}$. Then let all sellers learn. Since the payoff $U_T(\theta^*)$ is available, all sellers will leave $s_H(F)$. Once they have left, sellers can never return to $s_H(F)$ since (34) implies that choosing s_T^* is always a strictly better response in any subsequent state in X. □

The following example shows that the conditions of Proposition 6 do not suffice to ensure that the Riley outcome is a global outcome. This example also shows that markets exist in which there are two stable outcomes, so that the limit of at most two stable outcomes established in Corollary 3 is tight.

	L	P	H	L	P	H
U_L	2	5	8	-6	-3	0
U_H	2	5	8	-2	1	4
		s_0			s_1	

Example 3

This game has two limit distributions. One gives the pooling equilibrium outcome (in which signal s_0 is sent) as the stable outcome, the other the separating equilibrium (in which L sellers send s_0 and H sellers send s_1), which is the Riley outcome. Two limit distributions exist because mutations cannot lead the system from the pooling to the separating equilibrium or vice versa. Mutations cannot lead the market away from the pooling equilibrium because the pooling equilibrium provides higher payoffs than any payoff available from signal s_1, and hence no conjecture can induce a seller to abandon the pooling payoff at signal s_0 to send signal s_1. Similarly, no mutation can induce either seller to leave his separating equilibrium signal, since doing so requires sending the other signal that, by the construction of a separating equilibrium, must earn a lower payoff.[33]

The difficulty in Example 3 is that the pooling equilibrium (as well as the Riley outcome) is H-dominant, making it impossible for a single mutation to move the market from the pooling to the separating equilibrium. If no H-dominant pooling equilibrium exists, then it is immediate from Corollary 3 that, under the conditions of Proposition 6, the Riley outcome is the unique locally stable outcome. In fact, we can show that there is no other locally stable component in such a case, yielding the much stronger result that the Riley outcome is the *global* outcome:

Proposition 7 *Suppose ρ^* satisfies the assumptions of Proposition 6 and there exists no H-dominant pooling equilibrium. Then ρ^* is the global outcome of G.*

The intuition for this result is that every locally stable component must contain a minimal equilibrium (or four-state cycle) in which type L sellers receive at least $\bar{U}_L$. Consider a component containing such an equilibrium outcome $\rho(\theta)$ (the argument for the case of a four-state cycle is similar) that is not the Riley outcome. Since the Riley outcome is the unique H-dominant outcome, H-type sellers receive a lower utility in $\rho(\theta)$ than in the Riley outcome. H-type sellers thus have an incentive to leave such an θ for their equilibrium signal in the Riley outcome. L-type sellers in turn do not want to mimic such a deviation as long as they conjecture that they are guaranteed the competitive price against L types at any signal. Consequently, a single mutation by an H-type seller suffices to reach the basin of attraction of an equilibrium yielding the Riley outcome, implying that the original component is not locally stable. There is then only one locally stable component, giving the Riley outcome, which is thus the global outcome.

---

[33] Slightly more complicated versions of this example, with more signals, can be constructed in which the separating equilibrium does not use all of the signals.

Proof: We will show that every locally stable component must intersect the set X introduced in the proof of Proposition 6. Our claim follows, since minimality of locally stable components then implies that every locally stable component must be contained in X and hence must yield the Riley outcome ρ^*.

Lemma 4 implies that every locally stable component C contains either an equilibrium θ satisfying $U_L(\theta) \geq \bar{U}_L$ or a four-state cycle F satisfying $U_L(F) \geq \bar{U}_L$. Let us first consider the case in which C contains an equilibrium θ satisfying this condition. By Proposition 4 we may assume that this equilibrium is minimal. Suppose $\rho(\theta) \neq \rho^*$. Then we must have $s_H(\theta) \neq s_H^*$ since there is no equilibrium satisfying $U_L(\theta) \geq \bar{U}_L$ in which type L sellers choose s_H^* (from (27)), implying that if $s_H(\theta) = s_H^*$, then θ must be separating in order to be minimal and hence must yield the Riley outcome. Because there is no H-dominant pooling equilibrium (by assumption) and θ^* is H-dominant, we must have $U_H(\theta) < U_H(\theta^*)$. By Lemma 3 we can obtain a state θ' in C that has the same outcome as θ but has all buyers offering $p_1^*(s_H^*)$ at s_H^*. We can also suppose that in θ' all L sellers conjecture to receive $p_0^*(s)$ at all signals not used in θ. Given such a state $\theta' \in C$ a single mutation by an H seller that causes him to conjecture that $p_1^*(s_H^*)$ is offered at s_H^* will cause this seller to send signal s_H^*. The following learning sequence then reaches a state in X: First, all H sellers receive the learn draw and switch to s_H^*. Next, buyers receive a competitive learn sequence, which can reduce the price at s_L^* no lower than $p_0^*(s_L^*)$, since there are either no sellers sending s_L^* (yielding no price revision) or only L sellers (yielding $p_0^*(s_L^*)$). Then all sellers learn, causing all L sellers to send s_L^*, which must be the unique best response against their conjectures, and all H sellers to remain at s_H^*. Another competitive learn sequence for buyers establishes a state in X.

We now consider the case in which C contains a four-state cycle F satisfying $U_L(F) \geq \bar{U}_L$. By Assumption (5) such a cycle must satisfy $s_L(F) = s_L^*$. Because s_L^* is used by L-types, the price $p_0^*(s_L^*)$ must be offered at s_L^*, and conjectures must match this action. The signal chosen by H types cannot be s_H^*, because L sellers would not leave the competitive price at s_L^* to accept price $p_1^*(s_H^*)$, as required by the four-state cycle. Consider state $\theta \in F$. Suppose a crash yields a state in C at which the price has been reduced to $\underline{p}$ at $s_H(F)$. Now let a single mutation cause an H seller to adapt the conjecture that $\underline{p}$ is offered at all unused signals, except for s_H^* where he conjectures that $p_1^*(s_H^*)$ is offered. Because θ^* is an equilibrium it must be a best reply for this seller to switch to s_H^*. Now let buyers receive a competitive learn sequence. Buyers will then offer $p_0^*(s_L^*)$ at s_L^* and $p_1^*(s_H^*)$ at s_H^*. Furthermore, only L sellers are using s_L^* and only (one) H seller is using s_H^*. A learn draw for all buyers must then yield a state in X, establishing that C intersects X. □

7 Pooling and Partially Separating Equilibria

The previous section has shown that the unique candidate for a locally stable separating equilibrium outcome, namely an H-dominant Riley equilibrium, is indeed locally stable as long as no buyer is indifferent between the two signals used in the equilibrium. If this is the only H-dominant equilibria, then its outcome is a global outcome.

Are similar results true for H-dominant pooling or partially separating equilibrium outcomes? This section answers this question negatively.

The following example shows that an H-dominant pooling equilibrium may not only fail to be locally stable (even though its payoff strictly dominates every other available equilibrium payoff), but may not even appear in a limit set.

	L	P	H	L	P	H	L	P	H
U_L	2	5	8	-6	-3	0	-8	-2	4
U_H	2	5	8	-2	1	4	-6	0	6
		s_0			s_1			s_2	

Example 4

This market contains a H-dominant pooling equilibrium (in which signal s_0 is sent) but yields a unique limit set and global outcome, consisting of the separating equilibrium (with L sending s_0 and H sending s_1). To see that the pooling equilibrium is not contained in the limit set, let the pool be supported by conjectures of price H at signal s_1 and either price P or L at signal s_2. Consider the following sequence that leads away from the pooling equilibrium. First, a mutation causes an H seller to send signal s_2. Buyers then receive a competitive learn sequence, setting price H at signal s_2. H sellers then learn, switching to s_2. Buyers again face a competitive learn sequence, switching to price L at s_0. L sellers now learn, switching to s_2. A competitive learn sequence now switches buyers to price P at s_2. Seller learning then switches L sellers to s_0 and H sellers to s_1, at which point we have the separating equilibrium, ensuring that the pooling equilibrium outcome is not locally stable.

One cannot similarly construct a path from the separating equilibrium to the pooling equilibrium, and the separating equilibrium outcome is a global outcome. The difficulty is that H sellers will never switch to signal s_0 and L sellers will never switch to s_1 as long as price L is conjectured at s_0; while price L will be conjectured at s_0 unless H sellers switch to s_0.

Exploiting the forces that appear in this example, we can show that only under extraordinary circumstances will pooling or partially separating equilibrium outcomes be locally stable. In particular, if there is any possibility for mutations to induce any seller to leave his equilibrium signal and if there exists any alternative minimal sequential equilibrium in which H sellers send a currently unused signal, then an H-dominant pooling or partially separating equilibrium outcome fails to be locally stable:[34]

---

[34] The requirement that the alternative equilibrium be sequential can be substantially relaxed, but at the cost of a more complicated statement of the result and a less intuitive proof.

Proposition 8 *Let θ be a pooling or partially separating equilibrium. Then $\rho(\theta)$ fails to be a locally stable outcome if there exist*

(8.1) a signal $s' \notin S_L(\theta) \cup S_H(\theta)$ and a type T such that

$$\overline{p} - c_T(s') > U_T(\theta)$$

(8.2) and a minimal sequential equilibrium θ' such that

$$s_H(\theta') \notin S_L(\theta) \cup S_H(\theta).$$

Notice that $\rho(\theta)$ is not a locally stable outcome even in the case that θ is H-dominant and hence the alternative equilibrium θ' gives H sellers a lower utility than does θ. $\rho(\theta)$ is not a locally stable outcome because mutations can cause sellers to deviate to signal s' in quest of the payoff $\overline{p} - c_T(s')$. This payoff cannot be supported as part of an equilibrium, causing subsequent strategy adjustments, but these adjustments can lead to θ' rather than θ, ensuring that $\rho(\theta)$ is not locally stable.[35]

Proof: Because they are necessary conditions for local stability, we can assume that θ is minimal (Corollary 2), H-dominant (Proposition 5) and satisfies $U_L(\theta) \geq \bar{U}_L$ (Lemma 4). From Lemma 3 we can assume that at each signal not used in θ, all buyers offer the equilibrium price from θ', except at s' where all buyers offer $\overline{p}$. Since θ is not H-dominated we can assume that all H sellers hold conjectures about price offers at signals unused in θ which match the actual offers from θ'. We then show that a single mutation can induce learning dynamics that lead to an equilibrium with outcome $\rho(\theta')$, precluding the local stability of θ.

Pooling Equilibria. Assume that θ' is a pooling equilibrium. By Lemma 3, we can assume that all L sellers in state θ conjecture the competitive price $p_0^*(s)$ at any unused signal s. Suppose that condition (8.1) holds for type H. Then a single mutation by an H seller, endowing the seller with the conjecture that price $\overline{p}$ is offered at s' will trigger a deviation by this seller to s'. Thereafter let all H sellers learn, who will then all switch to s'. A competitive learn sequence then establishes a state, say θ'', in which $p_0^*(s_H(\theta))$ is offered at $s_H(\theta)$ (and $p_0^*(s_L(\theta))$ at $s_L(\theta)$ if it exists) and $p_1^*(s')$ is offered at s'. First, suppose s' is not a best reply for H sellers in state θ''. Then θ' must be a pooling equilibrium at some message other than s'. Let H sellers receive the learn draw, which must cause them to switch to the equilibrium message in state θ'; and let a competitive learn sequence establish price $p_1^*(s_H(\theta'))$ at $s_H(\theta')$. Then letting L sellers receive the learn draw will cause them to switch to the equilibrium message in state θ' (because θ' is sequential), after which a competitive learn sequence yields an equilibrium with outcome $\rho(\theta')$. Second, suppose that s' is a best reply for H sellers in state θ''. Then a learn draw for L sellers must now induce them to switch to s'. (If this were not the case, then the existing state would be an

---

[35] Proposition 8 clearly implies the weaker result that θ cannot be locally stable if (8.2) holds and some seller can be tempted away from θ by some competitive price, i.e., if (8.1) holds with $\overline{p}$ replaced by $p_1^*(s')$.

equilibrium and would H-dominate θ, a contradiction.) If θ' is a pooling equilibrium at s', then a further competitive learn sequence will establish an equilibrium with outcome $\rho(\theta')$. (In particular, sequentiality of θ' ensures that s' remains a best response for all sellers when the competitive price against the pool is offered at s'.) If θ' is a pooling equilibrium at a different signal, let buyer learning create a crash that reduces the price to $\underline{p}$ at s'. Given their conjectures, it must then be a best response for H sellers to switch to the equilibrium signal in θ'. The sequentiality of θ' then ensures that after a competitive learn sequence, a learn draw for L sellers and then another competitive learn sequence, an equilibrium with outcome $\rho(\theta')$ is established.

Next, suppose condition (8.1) holds only for type L. Following an argument analogous to that of the preceding case (with a mutation and subsequent learn draw by L rather than H sellers, followed by a competitive learn sequence), we can then

reach a state in which all L sellers choose s', all H sellers choose $s_H(\theta)$, price $p_0^*(s')$ is offered at s', and price $p_1^*(s_H(\theta))$ is offered at $s_H(\theta)$ (note that these price offers are the reverse of what they were in the previous case). If L sellers now receive a learn draw they must switch to $s_H(\theta)$. (This holds because L sellers can receive at most $\bar{U}_L$ at signal s’, while $U_L(\theta) \geq \bar{U}_L$ and the price offered at $s_H(\theta)$ has increased from its level in equilibrium θ, which (from the proof of Corollary 3) must be less than p_1^*, to p_1^*.) A crash at $s_H(\theta)$ then establishes a state which is in the basin of attraction of equilibrium with outcome $\rho(\theta')$. In particular, by sequentiality of θ' and Assumption 4 it becomes a best response for H sellers to switch to $s_H(\theta')$, where they still conjecture to receive the equilibrium price from θ' (since they have not yet received a learn draw and $s_H(\theta')$ is unused in the current state). The sequentiality of θ' then ensures that after a competitive learn sequence, a learn draw for L sellers will induce them to switch to $s_L(\theta') = s_H(\theta')$, and a competitive learn sequence then gives an equilibrium with outcome $\rho(\theta')$.

(Partially) separating equilibria. Suppose now that θ' is (partially) separating. This implies $U_L(\theta') = \overline{U}_L \leq U_L(\theta)$ and we can thus suppose that L sellers' conjectures match the actions from θ'. Suppose condition (a) holds for type H. The argument is the same as in the case in which θ' is pooling up to the point in which H sellers have chosen s' and $p_1^*(s')$ is offered at s', which we again call state θ''. First, suppose s' is a best reply for H sellers in state θ''. Let L sellers receive the learn draw, at which point they must switch to s' (otherwise we have an equilibrium that H-dominates θ, a contradiction). Suppose $s_H(\theta') = s'$, which can only occur if θ' is partially separating rather than separating (otherwise θ' would H-dominate θ). In this case, after a competitive learn sequence, it will become a best response for L sellers to switch to $s_L(\theta')$. Assuming that the appropriate proportion of L sellers receives the learn draw, followed by a competitive learn sequence, then allows us to reach an equilibrium state with outcome $\rho(\theta')$. If $s_H(\theta) \neq s'$, then let a crash at s' occur. Type H sellers will then switch to $s_H(\theta')$ and, by sequentiality, stay there after a competitive learn sequence has raised the price at s' back to $p_0^*(s')$. If θ' is separating, then a further learn draw for L sellers and a further competitive learn sequence for buyers establishes a separating equilibrium. If θ' is partially separating, then L sellers will first switch to $s_H(\theta)$ and the argument then continues as in the case $s_H(\theta) = s'$. Second, suppose s' is not a best reply for H sellers in state θ''. Then let H sellers receive the learn draw, which must prompt them to switch to state $s_H(\theta')$. The argument then proceeds as in the case of s' being a best reply for H sellers in state θ''.

If condition (a) holds only for L, the argument from the case in which θ' was a pooling equilibrium can be adjusted in an equivalent way. □

Proposition 8 indicates that H-dominant pooling or partially separating equilibria will be locally stable only if stringent conditions are met.[36] Hence, pooling and partially separating equilibria generally appear in limit sets, if at all, only if other outcomes also appear, with the market then cycling between varying outcomes.

Proposition 8 requires that there is an alternative equilibrium, but the absence of such an equilibrium is *not* sufficient to imply that an H-dominant pooling equilibrium outcome is locally stable:

	L	P	H	L	P	H
U_L	0	2	4	-3	-1	1
U_H	0	2	4	1	3	5
		s_0			s_1	

Example 5

This market has an H-dominant pooling equilibrium, in which signal s_0 is sent. There are no other equilibria, so that condition (8.2) in Proposition 8 fails, though there is a four-state cycle F with $s_L(F) = s_0$ and $s_H(F) = s_1$. The pooling equilibrium is not locally stable. To see this, notice that the component containing the pooling equilibrium contains a state in which buyers offer price H at signal s_1. Then a mutation causing an H seller to send signal s_1, followed by first H sellers receiving the learn draw and switching to s_1 and then a competitive learn sequence reducing the price at signal s_0 to L, yields a state in the four-state cycle. This four-state cycle is contained in an absorbing set of the learning dynamics, precluding the local stability of the pooling equilibrium.

If pressed, we can find sufficient conditions for pooling equilibria to be locally stable, but they are nearly trivial. In particular, if there is no possibility for mutations to induce a seller to switch from his equilibrium signal, then a minimal pooling equilibrium must be locally stable:[37]

---

[36] The proof of Proposition 8 points to a strengthening of Proposition 7. A Riley outcome satisfying the conditions of Proposition 6 is a global outcome even in the existence of an H-dominant pooling outcome as long as the pool occurs at signal s_L^* and satisfies condition (8.1)-(8.2).

[37] Strong as they are, the assumptions of Remark 4 do *not* suffice for the local stability of a partially separating equilibria. In such equilibria there is always the possibility that L types drift from $s_L(\theta)$ to $s_H(\theta)$, causing a crash at $s_H(\theta)$ and opening the possibility to destabilize the equilibrium. The only sufficient conditions for local stability of partially separating equilibria we could find are so strong as to be uninteresting.

Remark 4 *Suppose θ is a minimal pooling equilibrium and for all signals $s' \neq s_H(\theta)$*

$$\bar{p} - c_T(s') < U_T(\theta), \quad T = L, H. \tag{35}$$

Then $\rho(\theta)$ is locally stable.

Even the extremely strong assumptions in Remark 4 are not sufficient to imply that the pooling equilibrium outcome is a *global* outcome. Condition (35) is satisfied in Example 3, but the pooling equilibrium is not a global outcome.

These results lead us to conclude that in the absence of a Riley outcome satisfying the conditions from Proposition 7, the outcome of the market is extremely unlikely to be nicely behaved, in the sense that it is extremely unlikely to have a either a global outcome or even locally stable outcomes.

8 Conclusion

The outcome of our signaling market will be nicely behaved as long as a "good" separating equilibrium exists. In particular, if the Riley outcome provides H-sellers with a higher utility than any other equilibrium (and each type of seller strictly prefers its own signal to that of the other type), then the limiting outcome consists entirely of the Riley outcome. If a Riley outcome satisfying these conditions does not exist, however, as will occur when the Riley outcome is H-dominated, then only under extremely stringent conditions will the limiting outcome be single-valued. The limiting outcome in these cases will generally cycle between various equilibria and, possibly, nonequilibrium outcomes.

Our results allow some interesting contrasts with the refinements literature. First, our argument that an equilibrium is not locally stable begins by identifying conditions under which a seller would prefer to deviate from the equilibrium path. Similar considerations pervade the refinement literature. However, most refinements rest heavily on arguments concerning which types of seller have the greatest incentive to deviate. These considerations are irrelevant for our analysis, being replaced by considerations of whether there is another equilibrium (or absorbing set) upon which the process can settle.

While separating equilibria provide the best hope for stability in our market, Corollary 4 shows that a separating equilibrium cannot be stable if it is defeated. Our analysis thus lends little support to the intuitive criterion of Cho and Kreps [3] (and the refinements of this concept), which show a proclivity to select separating equilibria even when they are defeated by pooling equilibria.

Our results (e.g., Corollary 4) lend some support to the undefeated equilibrium concept. At the same time, our results indicate that the undefeated equilibrium concept is too demanding in the conditions it requires to eliminate an equilibrium from consideration. The undefeated concept rejects an equilibrium if the path leading away from the equilibrium leads *immediately* to a new equilibrium. This certainly suffices to ensure the equilibrium is not a stable outcome. However, there are other, indirect ways to escape an equilibrium. The pooling equilibrium in Example 4 is undefeated, because no path immediately leads to an alternative equilibrium, but this equilibrium is excluded from our limit distribution because of the possibility of indirectly reaching another equilibrium.[38]

Because they generally produce a global outcome only if an H-dominant Riley outcome exists, the dynamics of our model may appear to be somewhat volatile. How are the results affected if we put additional restrictions on either the type of mutations that can occur or the transitions that the learning dynamics can accomplish? Placing further restrictions on the mutation process, by making it more difficult for mutations to destabilize absorbing sets, might create new locally stable outcomes. There is a tradeoff here, however, since at the same time the H-dominant Riley outcome may no longer be a global outcome, with the new locally stable outcomes created by restrictions on mutations disrupting the proof of Proposition 7. Similarly, placing restrictions on the learning dynamics holds out the promise of creating new locally stable outcomes, but also the risk of causing previously locally stable outcomes to no longer be locally stable. In particular, a restricted learning dynamic will yield more absorbing sets, including especially more nonsingleton absorbing sets. Some of these sets may lie a single mutation away from a previously locally stable outcome, causing the latter to no longer be locally stable. There is thus no immediately obvious panacea for the volatility or our model.

9 Appendix

Proof of Lemma 1: A state satisfying the stated properties is clearly competitive: buyers' conjectures are correct by assumption and buyers are playing a best response, as a price higher than the competitive price ensures negative profits while a lower price ensures zero profits. It remains to show the reverse implication, for which it suffices to show that a collection of mutual best replies to any signal used by sellers must call for all buyers to offer the competitive price. Consider any $s \in S_L \cup S_H$. It cannot be a best response for any buyer to offer a price strictly greater than $p^*_{\phi(s)}(s)$ because profits are then negative. Suppose some buyer a offers a price p strictly lower than $p^*_{\phi(s)}(s)$. Let p be the *lowest* such price. If the maximum offer of other buyers exceeds p, then a can improve upon his (zero) payoff by simply offering the competitive price, which ensures a strictly positive payoff (since no other buyer can be offering a higher price). If all other buyers also offer p, then consider a deviation by buyer a to offering $p + \Delta$. The payoff from the original strategy,

---

[38] It may appear paradoxical to claim that undefeated requires overly strong conditions to eliminate equilibria from consideration, given that there are games in which undefeated equilibria do not exist. Notice, however, that a limit distribution always exists in our market. If all equilibria are defeated, then there is no locally stable outcome and the market will cycle between various equilibria in the limit.

conditional on offering p, was

$$\frac{1}{\mid A_B \mid}(v_{\phi(s)}(s) - p)$$

since all other buyers must also be offering p. The corresponding payoff from the new strategy is

$$v_{\phi(s)}(s) - p - \Delta$$

The new strategy is then superior if

$$v_{\phi(s)}(s) > p + \frac{\mid A_B \mid}{\mid A_B \mid -1}\Delta = p + \Delta + \frac{1}{|A_B| - 1}\Delta. \tag{36}$$

Since $p \leq p^*_{\phi(s)}(s) - \Delta$, Assumption 3 ensures that (36) holds. Thus the new strategy is a strict improvement. Hence, in a competitive state all buyers must offer the competitive price at all used signals. □

Proof of Proposition 2: We consider a deterministic dynamic process on Θ that arises as one possible realization of the stochastic learning process. Given a state θ_t, θ_{t+1} is defined as follows. First, let buyers receive a competitive learn sequence, giving the competitive successor $c(\theta_t)$. Next, let all sellers receive the learn draw. If seller $a \in A_T$ finds that he is not playing a best reply to his current conjecture, then seller a switches to a signal contained in $S_T(c(\theta_t))$ if $S_T(c(\theta_t))$ contains a best reply to a's conjecture, and chooses a signal from the complement of $S_T(c(\theta_t))$ otherwise. Let $k = 1, \cdots, |S|$ be an indexing of the set S. Seller a chooses the lowest-indexed best reply in $S_T(c(\theta_t))$ (if a best reply exists in $S_T(c(\theta_t))$ and otherwise chooses the lowest-indexed best reply in the complement of $S_T(c(\theta_t))$. This "lowest-index rule" determines a unique successor to $c(\theta_t)$ and this state is θ_{t+1}. The process then begins again with a buyer competitive learn sequence.

We refer to this sequence of learn draws and best-reply choices as the "alternating-best-response" dynamics and denote it by Γ^*. Because a non-singleton absorbing set X is closed under the learning dynamics, X is also closed under Γ^*. Because X is finite, it must then contain a cycle Λ of Γ^*. This cycle cannot be a singleton, because every steady

state of Γ^* is an equilibrium state and thus cannot be contained in X.

Let the set $\Lambda \subseteq X$ denote a non-singleton cycle of the alternating best response dynamics and $S_T(\Lambda) = \cup_{\theta\in\Lambda} S_T(\theta)$. The following arguments will make heavy use of the fact that for every state in Λ it must be the case that all sellers hold the same conjectures about price offers at signals in $S_L(\Lambda) \cup S_H(\Lambda)$ (because when any seller learns in the alternating-best-reply sequence, all sellers learn). In particular, in any state in Λ all sellers of a given type agree on the set of best responses contained in $S_T(\Lambda)$.

The three following steps show that every nonsingleton cycle of Γ^* has exactly two states - corresponding to states θ and θ' in the statement of the proposition:

(2.1): Each type of seller sends only one signal in each state of the cycle, i.e., for $T \in \{L, H\}$ and for all $\theta \in \Lambda$, $S_T(\theta)$ is a singleton.

(2.2): H sellers send the same signal in every state of the cycle, i.e., there exists $\hat{s}_H$ such that for all $\theta \in \Lambda : s_H(\theta) = \hat{s}_H$.

(2.3): L sellers alternate between two signals over the cycle, and the states in which these signals are sent correspond to states θ and θ' in the statement of the proposition.

(2.1) For $T \in \{L, H\}$ and for all $\theta \in \Lambda$, $S_T(\theta)$ is a singleton.

To establish this, first note that if there exists any state θ in Λ such that $S_T(\theta)$ is a singleton, then for all θ in Λ, $S_T(\theta)$ must be a singleton. In particular, if all type T sellers currently use the same signal and one type T seller switches to a signal s then the lowest index rule implies that all type T sellers will switch to s.

Second, the lowest index rule also implies that if there are states $\theta, \theta' \in \Lambda : \exists s \in S_T(\theta')$ such that $s \notin S_T(\theta)$, then for all cycle states S_T must be a singleton. In particular, there must be a first state appearing after θ in the cycle at which some type T seller chooses signal $s \notin S_T(\theta)$. But then *all* type T sellers must switch to s at this state, implying that for the subsequent state, and thus by the previous observation for all states, S_T is a singleton.

Hence, if there exists a state $\theta \in \Lambda$ such that $S_T(\theta)$ is not a singleton then we must have

$$\forall \theta \in \Lambda : S_T(\theta) = S_T(\Lambda). \tag{37}$$

Now notice that if $S_T(\theta)$ is not a singleton for some T and θ, then *all* signals in $S_T(\Lambda)$ are best responses for all type T sellers in *all* states in Λ. This in turn implies that no type T seller ever switches signals during the cycle. This observation implies that there can be at most one type T for whom there is a state $\theta \in \Lambda$ with $S_T(\theta)$ a nonsingleton, since otherwise no sellers would ever switch strategies and Λ would consist of a single state.

We now have two cases to consider. Suppose first that there is a state $\theta \in \Lambda$ at which $S_H(\theta)$ is not a singleton, so that condition (37) holds for type H, but $S_L(\theta)$ is a singleton for all cycle states θ. We use $s_L(\theta)$ to denote the signal used by all type L sellers in state θ. There must exist states $\theta, \theta' \in \Lambda$ such that $s_L(\theta) \neq s_L(\theta')$ (otherwise L sellers never switch strategies and Λ consists of a single state). It cannot be the case that $s_L(\theta) \notin S_H(\Lambda)$ and $s_L(\theta') \notin S_H(\Lambda)$. In particular, if

$$p_0^*(s_L(\theta')) - c_L(s_L(\theta')) > p_0^*(s_L(\theta)) - c_L(s_L(\theta)), \tag{38}$$

then in all cycle states L sellers must conjecture that a strictly higher utility can be obtained from choosing $s_L(\theta')$ instead of $s_L(\theta)$ (since neither of these signals is ever used by a type H seller, conjectures about prices at $s_L(\theta)$ and $s_L(\theta')$ must always be given by the $p_0^*(s_L(\theta))$ and $p_0^*(s_L(\theta'))$). Hence $s_L(\theta) \notin S_L(\Lambda)$, a contradiction. An analogous argument rules out the possibility that the inequality in (38) is reversed. Finally, if equality holds in (38), in the cycle learning can never cause L sellers to switch from $s_L(\theta)$ to $s_L(\theta')$ or vice versa. Hence, either $s_L(\theta) \in S_H(\Lambda)$ or $s_L(\theta') \in S_H(\Lambda)$.

Let us suppose, without loss of generality, that $s_L(\theta) \in S_H(\Lambda)$. There must be a state in the cycle at which L sellers are induced to switch to $s_L(\theta)$. The utility they expect from choosing $s_L(\theta)$ must be given by

$$p_1^*(s_L(\theta)) - c_L(s_L(\theta))$$

since only H type sellers are currently sending $s_L(\theta)$. After L sellers have switched and buyers have received the competitive learn sequence, the utility type L sellers will receive from choosing $s_L(\theta)$ is given by

$$p_{\phi'}^*(s_L(\theta)) - c_L(s_L(\theta)),$$

where $\phi' \leq \phi^0$ is the probability that results from a mixture of all L sellers and the fixed number of H sellers choosing $s_L(\theta)$. It then follows from Assumption 2 that

$$p_1^*(s_L(\theta)) > p_{\phi'}^*(s_L(\theta)) \tag{39}$$

and it follows that H type sellers' utility from choosing $s_L(\theta)$ decreases after L sellers have made the switch to $s_L(\theta)$. Since at all other signals in $S_H(\Lambda)$ prices cannot have been decreasing, (39) implies that $s_L(\theta)$ must cease to be a best response for type H sellers. This contradicts (37). Hence, there cannot be a state $\theta \in \Lambda$ at which $s_H(\theta)$ is not a singleton.

Next, suppose that the roles of the sellers are reversed, so there is a state $\theta \in \Lambda$ at which $S_L(\theta)$ is not a singleton, so that condition (37) holds for type L, but $S_H(\theta)$ is a singleton for all cycle states θ. Then an argument analogous to the preceding argument, with roles reversed, shows that at some cycle state type H must switch into a signal that is currently only used by L sellers. However, if type H sellers were ever to switch to a signal in $S_L(\Lambda)$, then the subsequent competitive learn sequence on part of the buyers ensures that the price at this signal increases and hence H can never subsequently leave that signal; implying that H sellers, and hence all sellers, never change signals, and hence that Λ has only a single state, yielding a contradiction.

Having established that both $S_L(\theta)$ and $S_H(\theta)$ must be singletons in all cycle states, we may use $s_T(\theta)$ to denote the signal used by all type T sellers in cycle state θ.

(2.2) There exists $\hat{s}_H$ such that for all $\theta \in \Lambda : s_H(\theta) = \hat{s}_H$.

Since all type T sellers are using the same signal in every cycle state, then for all $s \in S_L(\Lambda) \cup S_H(\Lambda)$ and all cycle states, the price that is offered at s and conjectures about that price must be one of $p_0^*(s)$, $p_1^*(s)$, or $p_{\phi^0}^*(s)$. Refer to these as the low, high, and pooling competitive price.

First, there can be no cycle state in which type H sellers switch out of a signal s at which the high competitive price is offered: When H sellers switched to s they cannot have conjectured a better price from choosing s than the high competitive price. Since prices at signals other than s cannot increase while all H sellers are choosing s, this implies that s must continue to be a best response for H sellers as long as the high competitive price is offered. It follows that if H sellers ever leave a signal s during the cycle, they must do so while receiving the pooling competitive price (as long as H sellers use s the price offered at s cannot fall below the pooling competitive price). Second, whenever H sellers switch to a signal s, they must conjecture the low or pooling competitive price at this signal (they cannot conjecture the high competitive price, since there must have been a previous state in the cycle at which they left s, fixing their conjectures at the pooling competitive price). But then H sellers cannot switch out of s again: As long as H sellers use s they must receive at least the pooling competitive price at s. Since only L sellers use any other signal, H's conjectures about prices at other signals cannot increase, implying that s must continue to be a best response for H throughout the cycle. Hence, no subsequent switch away from s can occur. Hence, there must exist $\hat{s}_H$ with $s_H(\theta) = \hat{s}_H$ for all $\theta \in \Lambda$.

(2.3) Given that $s_H(\theta)$ is constant, $s_L(\theta)$ cannot be constant (otherwise we would have an equilibrium and hence Λ contains only a single state). Hence, L sellers must be alternating between $\hat{s}_H$ and a single alternative signal $\hat{s}_L$. In particular, from Step 1, there cannot be multiple signals other than $\hat{s}_H$ sent by L sellers, because the L seller must conjecture the low competitive price at each of these signals, and the only one to which the L seller will switch is then the one with the lowest index among the set that maximizes the seller's payoff. Given the specification of the alternating best response dynamics Γ^*, we then have a four-state cycle with the desired properties. □

Proof of Proposition 4:

Quasi-separating equilibria. Suppose that θ is a quasi-separating equilibrium. By Assumption 1, there is a single signal, say s', sent by both types of seller. We can also assume that there is a single signal s_L sent by only L sellers and a single signal s_H sent by only H sellers. For example, if there were two (or more) signals sent by only L sellers, then a sequence of mutations, switching L sellers to only one of these signals would yield a new quasi-separating equilibrium contained in the component C and satisfying (20)-(21). Now consider a mutation that switches an H seller from signal s' to s_H, followed by a competitive learn sequence. If the latter reduces the market price at signal s', then seller learning can switch all L sellers to s_L and H sellers to s_H, yielding a separating equilibrium in component C satisfying (20)-(21). If the price at s' does not fall, then we have a new quasi-separating equilibrium in component C satisfying (20)-(21) in which one less seller sends signal s'. We can then repeat the argument, considering another mutation switching an H seller (if there exists an H seller sending signal s', and otherwise an L seller) away from signal s', and can continue until the desired separating equilibrium is reached.

Quasi-pooling equilibria. Let θ be a quasi-pooling equilibrium. Then there is a single (by Assumption 1) signal, say $s_H(\theta)$, with $s_H(\theta) \in S_L(\theta) \cap S_H(\theta)$ and $S_L(\theta) = \{s_H(\theta)\}$. If θ is not minimal, then there are other signals that are sent by some H sellers (and only by H sellers). We can assume that there is only one such signal, say s', since a sequence of mutations switching the H sellers from messages other than s_H to s' yields an equilibrium in C that preserves (20)-(21) and in which only signals s' and s_H are sent. Now consider a mutation that causes an H seller to switch from signal s' so s_H. Let this be followed by a competitive learn sequence. If this sequence increases the market price at signal s_H, then a buyer learning will switch all buyers to signal s_H, and a competitive learn sequence then yields a minimal pooling equilibrium at signal s_H contained in C and satisfying (20)-(21). If the market price at signal s_H does not change, then we have a new equilibrium in the component C, satisfying (20)-(21), with one less H seller sending signal s'. We can then repeat the argument, considering another mutation switching an H seller form s' so s_H, and can continue until the desired minimal pooling equilibrium is reached.

Partially separating equilibria. Let θ be a partially separating equilibrium. Then there is a single (by Assumption 1) signal, say s_H, sent by both L and H sellers, as well as a collection $\{s_1, \ldots, s_n\}$ of additional signals sent only by L sellers. Then any mutation that switches an L seller from a signal in $\{s_1, \ldots, s_{n-1}\}$ to signal s_n yields a new equilibrium in the component C satisfying (20)–(21), and a sequence of such mutations then yields a minimal partially separating equilibrium in which only the signals s_n and s_H are sent, with this equilibrium contained in the component C and satisfying (20)-(21). □

Proof of Lemma 4: Let

$$\bar{S}_L = \{s \in S : p^*_0(s) - c_L(s) = \bar{U}_L\}.$$

By Propositions 4 and 2 every locally stable component must contain either a minimal equilibrium or a four-state cycle. To prove the result it thus suffices to show that from every minimal equilibrium (or four-state cycle) that violates condition (22) (or condition (23)), a single mutation suffices to reach the basin of attraction of either an equilibrium or a non-singleton absorbing set containing a four-state cycle satisfying the required inequalities.

Let θ be either a minimal equilibrium violating (22), or a state in a four- state cycle violating (23), with $c(\theta)$ in the latter case being the state where sellers are both choosing $s_H(F)$ and buyers are offering price $p^*_{\phi^0}(s_H(F))$. If θ is an equilibrium, then $s_L(\theta) \notin \bar{S}_L$. A single mutation by a type L seller can then endow the seller with the conjecture that price $p^*_0(s^*_L)$ is offered at signal $s^*_L \in \bar{S}_L$ (leaving all other conjectures unchanged). This mutation will cause the L seller to switch to s^*_L because, by assumption, $p^*_0(s^*_L) - c_L(s^*_L) = \bar{U}_L > U_L(\theta)$. Similarly, if F is a four state cycle, then $s_L(F) \notin \bar{S}_L$. When the cycle is in state $c(\theta)$, then a single mutation can again cause L sellers to switch to s^*_L, since the four-state cycle calls for this seller to switch to signal $s_L(F)$ in order to receive utility $p^*_0(s_L(F)) - c_L(s_L(F)) < \bar{U}_L$.

Once this mutation has occurred, consider the alternating-best-response-dynamics, beginning with buyers. Buyers will first switch to the competitive price $p^*_0(s^*_L)$ at signal s^*_L. The argument used in the proof of Proposition 2 then allows us to conclude that this learning sequence leads either to an equilibrium or to a four-state cycle F'. Because the competitive learn sequences embedded in the alternating-best-response dynamics cannot drive prices at s^*_L below $p^*_0(s^*_L)$, we must have $U_L(\theta') \geq \bar{U}_L$ in the case the dynamics reach an equilibrium state θ'. For the same reason we must have $U_L(F') \geq \bar{U}_L$. No further argument would be required if F' is contained in an absorbing set. Unfortunately, we cannot be sure of this. The unrestricted learning dynamics allows the possibility that a crash at $s_H(F')$ causes the market price at $s_H(F')$ to fall below $p^*_{\phi^0}(s_H(F'))$, inducing sellers to leave $s_H(F')$ and leading to an absorbing set, possibly even an equilibrium, that contains F' in its basin of attraction. However, it suffices to show that every four-state cycle satisfying $s_L(F') \in \bar{S}_L$ lies in the basin of attraction of an absorbing set satisfying (22)-(23). We consider two cases.[39]

[Case 1] Let there exist a signal s such that $\underline{p} - c_H(s) > p^*_0(s_L(F')) - c_H(s_L(F'))$. Then, the learning dynamics can never induce an H seller to send signal $s_L(F')$. Hence, learning can affect neither beliefs nor buyers' actions at $s_L(F')$. Every absorbing set that can be reached from the cycle must then contain either an equilibrium satisfying (22) or a four-state cycle satisfying (23).

[Case 2] Suppose for all signals in $S \setminus \{s_L(F')\}$, we have $\underline{p} - c_H(s) < p^*_0(s_L(F')) - c_H(s_L(F'))$. Consider the following learning sequence, beginning in the four-state cycle with the state $\theta \in F'$ in which L and H sellers send the signal $s_H(F')$, with buyers offering price $p^*_1(s_H(F'))$. Let a crash occur, i.e. let the buyers all receive the learn draw, discover that their profits are negative at signal $s_H(F')$, and switch to price $\underline{p}$, followed by a learn draw for all sellers (which induces all of them to leave $s_H(F')$) and then the alternating-best-response dynamics. If the alternating best response dynamics reaches an equilibrium θ', then the equilibrium state must satisfy $U_L(\theta') \geq \bar{U}_L$ since the alternating-best-response dynamics can never endow sellers with a conjecture that prices less than $p^*_0(s_L(F'))$ are offered at $s_L(F')$. If the alternating best response dynamics reaches a four-state cycle F'', then no seller can use the signal $s_H(F')$, because whenever sellers learn during the alternating best response dynamics $s_L(F')$ is a strictly better response than $s_H(F')$, where all sellers conjecture $\underline{p}$ to be offered. (For L sellers this claim follows from $s_H(F') \notin \bar{S}_L$, as otherwise L sellers would never switch to $s_L(F')$ in the original four-state cycle.)

If F'' is contained in an absorbing set or satisfies the condition from Case 1 (which could occur because L sellers might be choosing a different signal contained in $\bar{S}_L$), we are done. If the new cycle satisfies the condition from Case 2, a crash occurring at the signal $s_H(F'')$, followed by the alternating-best-reply-dynamics, must reach either an equilibrium satisfying (22) or a four-state cycle in which neither $s_H(F')$ nor $s_H(F'')$ are used by any seller. This process can be repeated (if it does not first reach an absorbing set) until only one signal remains that has not been the object of a crash. This signal must be contained in $\bar{S}_L$ and must be chosen by all sellers, ensuring that we have reached an equilibrium satisfying (22).

Proof of Proposition 5: Suppose θ is not H-dominant but $\rho(\theta)$ is locally stable. Lemma 4 and the local stability of $\rho(\theta)$ then imply $U_L(\theta) \geq \bar{U}_L$.

We first argue that θ is H-dominated by an equilibrium θ' satisfying $U_L(\theta') \geq \bar{U}_L$. We know that there exists a minimal equilibrium θ' that H-dominates θ, so that $s_H(\theta') \notin$

---

[39] Assumption 4 ensures that these cases are exhaustive.

$S_L(\theta) \cup S_H(\theta)$ and $U_H(\theta') \geq U_H(\theta)$. Suppose θ' does not satisfy $U_L(\theta') \geq \bar{U}_L$. Let $s_L^* \in \bar{S}_L \equiv \{s \in S : p_0^* s) - c_L(s) = \bar{U}_L\}$. Consider a state that differs from θ' only in that L sellers choose s_L^* (which must differ from both $s_L(\theta')$ and $s_H(\theta')$ for θ' to be an equilibrium with $U_L(\theta') < \bar{U}_L$) and receive the competitive price $p_0^*(s_L^*)$ at this signal, with matching conjectures. Then this state is either a separating equilibrium state, which we denote θ'', or

$$p_0^*(s_L^*) - c_H(s_L^*) > U_H(\theta') \geq U_H(\theta), \tag{40}$$

which implies that there exists a pooling equilibrium at s_L^*, which we denote θ'''. In the first case, θ'' is a separating equilibrium that H-dominates θ because $U_L(\theta'') \geq \bar{U}_L$, $U_H(\theta'') = U_H(\theta')$, and $s_H(\theta'') = s_H(\theta')$. In the second case, it follows immediately from (40) that s_L^* must have been unused in θ. Furthermore, we have $U_H(\theta'') > U_H(\theta')$, so that θ''' H-dominates θ.

Hence, θ is H-dominated by an equilibrium, which we now call θ', that satisfies $U_L(\theta') \geq \bar{U}_L$. From Lemma 3 we can suppose that θ satisfies the condition that all L sellers conjecture to receive price $p_0^*(s)$ at all signals that are not used in θ. Furthermore, we can also suppose that in θ all sellers offer price $p_1^*(s_H(\theta'))$ at $s_H(\theta')$. We then distinguish two cases.

Case 1: Suppose θ' is separating. Beginning with state θ, let a single mutation occur, causing an H seller to switch to signal $s_H(\theta')$, followed by a competitive learn sequence. If either $p_1^*(s_H(\theta')) - c_H(s_H(\theta')) > U_H(\theta)$ or the competitive learn sequence causes the price at $s_H(\theta)$ to decline, let all H sellers receive the learn draw, which will cause all H sellers to switch to $s_H(\theta')$. Another competitive learn sequence and learn draw for all sellers will cause H sellers to remain at $s_H(\theta')$ while L sellers choose some signal in $\bar{S}_L$, giving a separating equilibrium outcome that is not $\rho(\theta)$ and hence precluding the local stability of $\rho(\theta)$. If $p_1^*(s_H(\theta')) - c_H(s_H(\theta')) = U_H(\theta)$ and the competitive learn sequence does not affect the price at $s_H(\theta)$, then we have immediately reached a new equilibrium outcome that is not $\rho(\theta)$, again precluding local stability of the latter.

Case 2: Suppose θ' is a pooling or partially separating equilibrium. Then by Assumption 2, we have $p_1^*(s_H(\theta')) - c_H(\theta') > p_{\phi^0}^*(s_H(\theta')) - c_H(\theta') \geq U_H(\theta)$. If θ' is a pooling equilibrium this implies $p_1^*(s_H(\theta')) - c_H(\theta') > U_H(\theta)$. If θ' is a partially separating equilibrium, then we can also assume that $p_1^*(s_H(\theta')) - c_H(\theta') > U_H(\theta)$.[40] Beginning with state θ, let a mutation cause an H seller to switch to signal $s_H(\theta')$. Now let all H sellers receive the learn drawn, which will cause them to switch to $s_H(\theta')$. A competitive learn sequence for buyers then results in a state in which the unique best response for all sellers is to choose $s_H(\theta')$. If θ' is a pooling equilibrium, seller learning and a competitive learn sequence then gives a pooling equilibrium at $s_H(\theta')$ with outcome $\rho(\theta')$. If θ' is partially separating, let only the equilibrium proportion of L sellers learn and switch to $s_H(\theta')$.

---

[40] It could be that θ' is a partially separating equilibrium with $p_1^*(s_H(\theta')) - c_H(\theta') = U_H(\theta)$, but then it must be that $p_1^*(s_H(\theta')) = p^*_{\phi(s_H(\theta'),\theta')}(s_H(\theta'))$. Switching all L sellers to $s_L(\theta')$ then gives us a minimal separating equilibrium that H-dominates θ and satisfies $U_L(\theta') \geq \bar{U}_L$, in which case we can apply the argument of Case 1.

After a further competitive learn sequence it becomes a best response for the remaining L sellers to switch to $s_L(\theta')$ (since $U_L(\theta') \geq \bar{U}_L$ implies that $s_L(\theta') \in \bar{S}_L$ and L sellers hold the appropriate conjectures). A further learn draw for all agents then yields an equilibrium with outcome to θ'. In either case the ability of the learning process to reach an equilibrium with an outcome different from $\rho(\theta)$ precludes local stability of $\rho(\theta)$. $\square$

References

[1] J. Banks and J. Sobel. Equilibrium selection in signaling games. *Econometrica*, 55:647–662, 1987.

[2] Andreas Blume, Yong-Gwan Kim, and Joel Sobel. Evolutionary stability in games of communication. *Games and Economic Behavior*, 5:547–575, 1993.

[3] In-Koo Cho and David M. Kreps. Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102:179–221, 1987.

[4] In-Koo Cho and Joel Sobel. Strategic stabililty and uniqueness in signalling games. *Journal of Economic Theory*, 50:381–413, 1990.

[5] Drew Fudenberg and David M. Kreps. A theory of learning, experimentation, and equilibrium in games. Mimeo, Stanford University and Massachusetts Institute of Technology, 1988.

[6] Drew Fudenberg and David K. Levine. Self-confirming equilibrium. *Econometrica*, 61:523–446, 1993.

[7] Drew Fudenberg and David K. Levine. Steady state learning and Nash equilibrium. *Econometrica*, 61:547–574, 1993.

[8] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, Cambridge, Massachusetts, 1991.

[9] Douglas Gale. A Walrasian theory of markets with adverse selection. *Review of Economics Studies*, 59:229–256, 1992.

[10] Sanford J. Grossman and Motty Perry. Perfect sequential equilibrium. *Journal of Economic Theory*, 39:97–119, 1986.

[11] Ehud Kalai and Ehud Lehrer. Private-beliefs equilibrium. Center for Mathematical Studies in Economics and Management Science Discussion Paper 926, Northwestern University, 1991.

[12] Michihiro Kandori, George J. Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56, 1993.

[13] George J. Mailath, Masahiro Okuno-Fujiwara, and Andrew Postlewaite. Belief-based refinements in signalling games. *Journal of Economic Theory*, 60:241–276, 1993.

[14] John H. Nachbar. Perturbed best–response dynamics. Mimeo, Washington University, St. Louis, 1993.

[15] Georg Nöldeke and Larry Samuelson. The evolutionary foundations of backward and forward induction. SFB discussion paper B-216, University of Bonn, 1992.

[16] Georg Nöldeke and Larry Samuelson. An evolutionary analysis of backward and forward induction. *Games and Economic Behavior*, 5:425–454, 1993.

[17] John Riley. Informational equilibrium. *Econometrica*, 47:331–359, 1979.

[18] M. Rothschild and J. Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics*, 80:629–649, 1976.

[19] Larry Samuelson. Stochastic stability in games with alternative best replies. *Journal of Economic Theory*, 1993. Forthcoming.

[20] Joel Sobel. Evolutionary stabiliy in communication games. *Economics Letters*, 1993. Forthcoming.

[21] A. M. Spence. Job market signaling. *Quarterly Journal of Economics*, 90:225–243, 1973.

[22] A. M. Spence. *Market Signaling, Information Transfer in Hiring and Related Processes*. Harvard University Press, Cambridge, 1974.

[23] J. Stiglitz and A. Weiss. Sorting out the differences between screening and signaling models. In M. Bachrach, editor, *Oxford Essakys in Mathematical Economics*. Oxford University Press, Oxford, 1993.

[24] Eric van Damme. Refinements of Nash equilibrium. Discussion Paper 9107, CentER for Economics Research, 1991.

[25] Charles Wilson. A model of insurance markets with incomplete information. *Journal of Economic Theory*, 16:167–207, 1977.

[26] Peyton Young. The evolution of conventions. *Econometrica*, 61:57–84, 1993.