IHS Economics Series
Working Paper 27
February 1996

A Dynamic Model of Equilibrium Selection in Signaling Markets

Georg Nöldeke
Larry Samuelson

Impressum

Author(s):
Georg Nöldeke, Larry Samuelson

Title:
A Dynamic Model of Equilibrium Selection in Signaling Markets

ISSN: Unspecified

1996 Institut für Höhere Studien - Institute for Advanced Studies (IHS)
Josefstädter Straße 39, A-1080 Wien
E-Mail: office@ihs.ac.at
Web: www.ihs.ac.at

Institut für Höhere Studien (IHS), Wien
Institute for Advanced Studies, Vienna

Reihe Ökonomie / Economics Series No. 27

A Dynamic Model of Equilibrium Selection in Signaling Markets

Georg Nöldeke, Larry Samuelson

A Dynamic Model of Equilibrium Selection in Signaling Markets

Georg Nöldeke, Larry Samuelson

Reihe Ökonomie / Economics Series No. 27

February 1996

Georg Nöldeke
Department of Economics
University of Bonn
Adenauerallee 24-26
D-53113 Bonn
Germany
Phone: +49-228-73-7996
e-mail: georg@glider.econ3.uni-bonn.de

Larry Samuelson
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, Wisconsin 53706
U.S.A.
Phone: +1-608-263-7791
e-mail: larrysam@cournot.econ.wisc.edu

Institut für Höhere Studien (IHS), Wien
Institute for Advanced Studies, Vienna

Abstract

In his work on market signaling, Spence proposed a dynamic model of a signaling market in which a buyer revises prices in light of experience and sellers choose utility-maximizing signals given these prices. Spence also suggested that subjecting the dynamic process to rare perturbations might allow one to choose between multiple equilibria. This paper examines the effect of introducing such perturbations into Spence's dynamic model. We find that refinement results arise naturally from the dynamic analysis. In a broad class of markets, our model selects a separating equilibrium outcome if and only if the equilibrium outcome satisfies a version of the undefeated equilibrium concept, whereas a pooling equilibrium outcome is selected if and only if the equilibrium outcome is both undefeated and satisfies $D1$.

Keywords
Equilibrium selection, evolution, signaling

JEL-Classifications
C70, C72, D82, D83

Comments
Previous versions of this paper were circulated under the title "Learning to Signal in Markets". First draft January 13, 1993. Financial support from National Science Foundation grant SBR 9320678 and the Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 303 at the University of Bonn, is gratefully acknowledged. We thank Drew Fudenberg for helpful comments. We are also thankful for the helpful comments of three referees and an associate editor.

1 Introduction

Signaling models often have many equilibrium outcomes, some of which seem more "plausible" than others.[1] The equilibrium refinements literature attempts to formalize this plausibility criterion by placing restrictions on out-of-equilibrium-beliefs which are in turn motivated by considering the incentives for the various types of the informed player to deviate from equilibrium.

Spence [11, 12] suggested an alternative approach to equilibria in signaling models. He proposed a dynamic model in which agents adjust their beliefs and actions in response to past market outcomes.[2] This dynamic model has many equilibrium (stationary) outcomes and may also reach a cycle. Spence focussed his attention on the characterization of equilibria, but suggested that the multiplicity of equilibrium outcomes and the possibility of cycles might be addressed by considering a perturbed version of the dynamic process. In particular, Spence [12, Appendix H] suggested eliminating equilibrium outcomes or cycles that are not robust against arbitrarily small perturbations.

Kandori, Mailath and Rob [4] and Young [14], have recently shown that a dynamic model incorporating both learning from past outcomes and the possibility of perturbations or mistakes can yield strong equilibrium selection results. Kandori, Mailath and Rob and Young examine (primarily) two-player 2×2 normal-form games with two strict Nash equilibria. In this paper, we use related techniques to investigate whether models of perturbed dynamics can produce equilibrium selection results for signaling models.

The basic dynamic we investigate, presented in Section 2, is Spence's. Section 3 adds the perturbations and introduces the notion of a recurrent set, the tool we use to describe the outcomes which emerge when considering arbitrarily small perturbations. Our perturbations differ from those suggested by Spence, with these differences and their implications discussed in Section 5.

Section 4 contains our main results. We find that there can exist at most two recurrent sets of our dynamic process. One of these must contain all states corresponding to Riley equilibria (the separating equilibria giving the sellers the highest possible payoffs among separating equilibria). The other recurrent set must contain either the pooling equilibria in which sellers use the signal that maximizes the high-quality sellers' utility, given that a pooling price is received, or must contain a cycle using this signal if such an equilibrium does not exist. If such equilibria exist, we call them Hellwig equilibria.[3] In a broad class of markets, there is a unique recurrent set of the dynamic process. For example, this is the case if the Riley and Hellwig equilibria use different signals. It is also the case in the

[1] See Kreps and Sobel [5] for a recent survey of signaling models.

[2] Stiglitz and Weiss [13] also advocate a dynamic approach to signaling and screening models.

[3] Hellwig [3] presents a game-theoretic analysis focusing on these equilibria.

commonly-studied variant of Spence's model in which signals do not affect productivity and a single-crossing property holds.

Our approach is based on following a dynamic model to its conclusions. However, we find many refinement ideas reappearing in our results. A necessary and sufficient condition for there to exist a recurrent set consisting entirely of separating equilibria yielding the same equilibrium outcome is that they satisfy a version of Mailath, Okuno-Fujiwara and Postlewaite's [6] undefeated equilibrium concept, suitably adapted to our framework. A sufficient condition for the existence of a recurrent set consisting entirely of pooling equilibria yielding the same outcome is that the latter be undefeated and satisfy Cho and Sobel's [2] $D1$ concept, and these conditions are also necessary in a broad class of markets.

For markets in which the undefeated equilibrium concept selects a separating equilibrium outcome (which is then the outcome generated by the Riley equilibria and commonly also satisfies $D1$), our analysis thus supports the selection of this particular outcome. The undefeated equilibrium and $D1$ concepts often conflict, generally because $D1$ selects a Riley equilibrium whereas the undefeated equilibrium concept selects a Hellwig equilibrium. Neither the set of Hellwig equilibria nor the set of Riley equilibria will constitute a recurrent set of our model in such a case. Our approach accordingly does not select a unique equilibrium outcome, though we can show that a unique recurrent set exists, which must contain the set of Hellwig equilibria (as well as other equilibria or cycles).

The refinements literature has concentrated on choosing a single refinement to be applied to all equilibria. In contrast, we find that different refinement criteria are relevant for assessing separating and pooling equilibria, with the undefeated concept alone sufficing to evaluate separating equilibria while both undefeated and $D1$ are relevant for pooling equilibria. The applicability of these criteria arises out of the differing dynamics that surround separating and pooling equilibria. The informal intuition offered to motivate equilibrium refinements is often dynamic, though the refinements are simply criteria for evaluating outcomes of a static model. One contribution of our work is to replace such intuition with a model in which refinement ideas can appear directly as a result of the underlying adjustment process.

Section 5 examines extensions of our model. Section 6 concludes. Proofs are gathered in the Appendix.

2 The Spence Model

Despite the common identification of Spence with signaling, his is a screening model.[4] A single buyer first posts a price for each of a number of signals, where this price equals the buyer's expected value of purchasing from a seller who sends such a signal. A number of sellers, some of low quality and some of high quality, then enter the market, observe these prices, and choose the signals which maximize their payoffs. The buyer purchases from each seller at the price posted for the appropriate signal. The buyer then revises her

[4]Spence used the label "signaling" to describe equilibria of his screening model in which different types of seller send different signals.

expected values to match the quality of the sellers at each signal that was sent by some seller. The process is repeated in the next period with a new collection of sellers who come in the same proportions of low-quality and high-quality as in the previous period.

Spence assumes that there is a single buyer who sets prices equal to expected values. Why does the buyer set such "competitive" prices, and why is it reasonable to assume that the buyer purchases from every seller? For Spence, these issues are unimportant, as his interest lies in the question of how, *given* competitive prices, sellers may be able to signal their types. We view Spence's buyer as a convenient representation of a market process in which there may be many buyers subject to forces that induce them to price competitively. Modelling the process by which buyers are led to price competitively is an important issue, but it is not the issue we want to pursue here. We accordingly follow Spence in simply assuming that there is a single buyer who sets prices equal to expected values. We then concentrate on the sellers' choices of signals, leaving competitive behavior for future work.

2.1 The Market

We consider a market in which trading occurs in each of the time periods $t \in \{1, 2, \ldots\}$. There is a single buyer. At the beginning of a period, the buyer posts a price schedule $p : X \to \Re$ that attaches a price $p(x)$ to each of the signals in a finite set $X = \{x_1, \cdots, x_n\}$. A finite number of sellers then enter the market, H of whom are endowed with one high-quality unit of a commodity and L of whom are endowed with one low-quality unit of a commodity. We use the subscripts ℓ and h to identify low-quality sellers and high-quality sellers and let $\phi^0 = H/(H + L)$ denote the fraction of high quality-sellers. Sellers know the quality of the commodity with which they are endowed. Each seller observes the price schedule, chooses a signal x, and sells his commodity to the buyer at price $p(x)$. The sellers then exit the market, with H new high-quality and L new low-quality sellers entering next period.

The buyer purchases from all of the sellers. The payoff to the buyer from each unit purchased at price p from a seller who is of type $q \in \{\ell, h\}$ and who has sent signal $x \in X$ is given by:

$$v_q(x) - p,$$

where $v_q(x)$ denotes the value of a unit purchased from a q-quality seller who has sent signal x. The expected value of a unit which is high-quality with probability ϕ is then $v(x, \phi) = \phi v_h(x) + (1 - \phi) v_\ell(x)$. The payoff to a seller from choosing signal x and receiving price p is given by

$$p - c_q(x),$$

where $c_q : X \to \Re$ is a function identifying the cost of sending signal x for a seller of type $q \in \{\ell, h\}$.

The buyer begins each period with a belief $\Phi(x)$ that associates, with every signal x, the proportion of sellers sending signal x who are high quality (or the probability that a single such seller is high quality). Given her belief, the buyer prices competitively, meaning that the buyer sets the price $p(x)$ equal to her expected value of buying from a seller who sends signal x, or:

$$p(x) = v(x, \Phi(x)). \tag{1}$$

In order to work with a finite Markov process, we follow Spence [12, Appendix H] in assuming that the buyer's beliefs are contained in Υ, a finite subset of $[0,1]$ that includes $\{0, \phi^0, 1\}$.[5]

Let $u_q : X \times \Upsilon \to \Re$ be defined by

$$u_q(x, \phi) = v(x, \phi) - c_q(x).$$

Then $u_q(x, \Phi(x))$ gives the utility to a seller of type q sending signal x given that the buyer has belief $\Phi(x)$ and hence sets $p(x) = v(x, \Phi(x))$. We assume:

Assumption 1

(1.1) For all $x \in X$,

$$v_h(x) > v_\ell(x).$$

(1.2) For all $(x, \phi) \neq (x', \phi')$ and $q \in \{\ell, h\}$,

$$u_q(x, \phi) \neq u_q(x', \phi').$$

Assumption 1.1 indicates that a high-quality seller is always more valuable to the buyer than a low-quality seller. This in turn implies that $u_q(x, \phi)$ is strictly increasing in ϕ. Assumption 1.2 is a genericity assumption, ensuring that for any belief $\Phi(x)$, each type of seller has a uniquely determined utility-maximizing signal, given by

$$x_q = \arg\max_x u_q(x, \Phi(x)),\ q \in \ell, h. \tag{2}$$

All low-quality sellers choose the signal x_ℓ and all high-quality sellers choose the signal x_h.

A tuple $\theta(t) = (\Phi(t), x_\ell(t), x_h(t))$, where $x_q(t)$ satisfies (2) for $q \in \{\ell, h\}$, specifies a state of the market at time t. We use Θ to denote the set of possible states. Because (2) uniquely fixes sellers' actions given Φ we will often speak of $(\Phi(t), x_\ell(t), x_h(t))$ as the **state induced by** $\Phi(t)$.

2.2 Dynamics

After transactions at time t are completed, the buyer's belief $\Phi(x)$ is adjusted to match the buyer's experience. For any state $\theta = (\Phi, x_\ell, x_h)$, we let $\nu(x_\ell)$ and $\nu(x_h)$ (where we suppress the dependence of $\nu(\cdot)$ on θ) be given by:

$$\begin{aligned} \nu(x_\ell) = 0, \qquad & \nu(x_h) = 1 \qquad \textit{if } x_\ell \neq x_h \\ \nu(x_\ell) \ = \ \phi^0 \ & = \ \nu(x_h) \qquad \textit{if } x_\ell = x_h. \end{aligned}$$

[5] Υ may be very large. The important points are that Υ is finite and Υ contains those beliefs that can be generated by the buyer's experience. Throughout the main part of our analysis these beliefs are $\{0, \phi^0, 1\}$. Other beliefs can appear in some of the extensions of the model considered in Section 5, and will then be assumed to be in Υ.

Then define state $s(\theta) = (\Phi', x'_\ell, x'_h)$, the successor of state θ, as

$$\Phi'(x) = \Phi(x) \quad if \ x \notin \{x_\ell, x_h\} \tag{3}$$
$$\Phi'(x) = \nu(x) \quad if \ x \in \{x_\ell, x_h\}. \tag{4}$$
$$x'_q = \arg\max_x u_q(x, \Phi'(x)), \quad q \in \{\ell, h\}. \tag{5}$$

Hence, $s(\theta)$ is the state obtained from θ if the buyer's belief remains unchanged at those signals that are not used in θ, and is updated to reflect the buyer's experience at those signals, namely x_q, that are sent in θ.

A deterministic dynamic process is then obtained by specifying that each state is followed by its successor. More formally, letting $\theta(t)$ be the state at time t, the dynamics are specified by an initial state $\theta(0)$ and the adjustment rule that for all $t > 0$, $\theta(t) = s(\theta(t-1))$. We refer to this deterministic process as the **Spencian dynamic.**

For any state $\theta \in \Theta$ let

$$S(\theta) = \{\theta' | \exists T \geq 0 : \theta' = s^T(\theta)\},$$

where $s^0(\theta) = \theta$ and $s^T(\theta) = s(s^{T-1}(\theta))$. Then $S(\theta)$ is the set of states which arise if the Spencian dynamic is started from the initial state θ. A set $C \subset \Theta$ is closed under the Spencian dynamic if $\forall \theta \in C : S(\theta) \subset C$. A non-singleton set C is a cycle if $\forall \theta \in C$: $S(\theta) = C$. A state $\theta \in \Theta$ is a stationary state of the Spencian dynamic if $S(\theta) = s(\theta) = \theta$. Since the state space of the Spencian dynamic is finite, the dynamic must lead, in finite time, either to a stationary state or a finite cycle.

It follows immediately from the definition of the dynamics that:

Lemma 1 *A state $(\Phi^*(x), x^*_\ell, x^*_h)$ is a stationary state of the Spencian dynamic if and only if*

$$\Phi(x^*_q) = \nu(x^*_q) \quad q \in \{\ell, h\}.$$

A stationary state of the Spencian dynamic is thus a state induced by a belief for the buyer that correctly identifies the proportion of high-quality sellers at each signal that is sent by any seller, given that the buyer sets prices equal to expected values (from 1)) and given that sellers choose utility-maximizing signals given prices (from (2)). These conditions correspond precisely to the ones Spence used to define an equilibrium and are also the conditions which characterize pure strategy sequential equilibria in the signaling game Cho and Kreps [1] proposed as a model of market signaling. We will accordingly often refer to stationary states as **equilibria.** If $x^*_\ell = x^*_h$, then we have a **pooling equilibrium** while $x^*_\ell \neq x^*_h$ gives a **separating equilibrium.**

Let

$$\underline{x} = \arg\max_x u_\ell(x, 0). \tag{6}$$

Then $\underline{x}$ is the signal low-quality sellers would prefer to send if, at whatever signal they send, they receive the price that the buyer sets when she expects sellers to be all low-quality. In every separating equilibrium, all low-quality sellers use $\underline{x}$.

Spence notes that his dynamic process will lead either to an equilibrium or to a cycle, but he does not characterize cycles. We can obtain a precise and useful characterization of such cycles. The proof of the following (and the proofs of all subsequent Lemmas and Propositions, unless otherwise noted) is contained in the Appendix:

Lemma 2

(2.1) Every cycle of the Spencian dynamic contains two states. For every such cycle there exists a signal x^c such that one of these states, say $\theta = (\Phi, x_\ell, x_h)$, satisfies $x_\ell = x_h = x^c$ and $\Phi(x^c) = 1$ whereas the other state, $\theta' = (\Phi', x'_\ell, x'_h)$, satisfies $x'_h = x^c$, $x'_\ell = \underline{x}$, and $\Phi'(x^c) = \phi^0$.

(2.2) A two-cycle of the Spencian dynamic in which $x_h = x'_h = x^c$ exists if and only if

$$u_h(x^c, \phi^0) > \max_x \{u_h(x, 0)\} \quad (7)$$

$$u_\ell(x^c, 1) > u_\ell(\underline{x}, 0) > u_\ell(x^c, \phi^0). \quad (8)$$

Lemma 2.1 indicates that high-quality sellers always send the same signal in every state of a cycle. Letting x^c denote this signal, we then have that low-quality sellers alternate between sending x^c and sending $\underline{x}$. The buyer always expects sellers at $\underline{x}$ to be low-quality, while beliefs at x^c alternate between high-quality and the pooling belief ϕ^0. In particular, we can think of low-quality sellers first sending $\underline{x}$ and high-quality sellers sending x^c. Upon revising her beliefs accordingly, the buyer then offers price $v(x^c, 1)$ at signal x^c. Low-quality sellers then find it optimal to switch to x^c. This causes the buyer to reduce the price at x^c to $v(x^c, \phi^0)$, inducing the low-quality seller to return to signal $\underline{x}$, and beginning the cycle anew. The existence conditions in (7)–(8) state that it must always be optimal for the high-quality seller to send x^c (condition (7)) while the low-quality seller prefers the high-quality price $v(x^c, 1)$ at x^c to the low-quality price $v(\underline{x}, 0)$ at $\underline{x}$, which he prefers to the pooling price $v(x^c, \phi^0)$ at x^c.

If a cycle exists in which high-quality sellers are always using signal x^c, we call it a **cycle at** x^c. We call θ the pooling state and θ' the separating state of the cycle.

It is immediate that a two-cycle involving signals $\underline{x}$ and x^c exists only if there exists a mixed-strategy equilibrium of the corresponding Cho-Kreps [1] signaling game in which high-quality sellers send x^c and low-quality sellers mix between $\underline{x}$ and x^c. Two-cycles thus correspond to mixed equilibria of the Cho-Kreps signaling game in the sense that the set of signals sent in the former by each type of seller is the *support* of the distribution of signals sent by that seller in the latter.[6]

Section 5 describes how the model could be altered to eliminate cycles. We do not do so because the resulting analysis is more tedious and because we encounter no difficulties in analyzing the model with cycles. However, we find it easiest to interpret the results if one thinks of the cycles as corresponding to mixed equilibria.

3 Perturbations

The Spencian dynamic reaches either an equilibrium or a two-cycle. But there may be many equilibria and two-cycles. Which ones are worthy of our attention?

[6]The average frequencies with which the signals are used across the states in the cycle need not match the probabilities in the mixed equilibrium. The Cho-Kreps signaling game may have other mixed strategy equilibria which do not correspond to any of the limits of the Spencian dynamic.

To answer this question, we associate with each state $\theta = (\phi, x_\ell, x_h)$ a non-empty set of perturbed states $P(\theta)$ with the following property:

Assumption 2 *For all states $\theta \in \Theta$:*

(2.1) $s(\theta) \in P(\theta)$

(2.2) $P(\theta) \subset \{\theta' | \phi'(x_q) = \nu(x_q), q \in \{\ell, h\}\}$

(2.3) $\{\theta' | \exists! x \notin \{x_\ell, x_h\} : \Phi'(x) \neq \Phi(x),\ \phi'(x_q) = \nu(x_q), q \in \{\ell, h\}\} \subset P(\theta)$.

Assumption 2.1 states that perturbations may have no effect on the state, leaving us simply with the successor $s(\theta)$ produced by the Spencian dynamic in the absence of a perturbation. This ensures that $P(\theta)$ is always non-empty. Assumption 2.2 places an upper bound on the states that can arise through perturbations. In particular, perturbed states that differ from $s(\theta)$ must satisfy (4)–(5), though they need not satisfy (3). Hence, a perturbation can be interpreted as arising from an alteration in the buyer's beliefs, where this alteration affects only beliefs at signals that were not sent in the previous period. Finally, Assumption 2.3 requires that all those states in which the seller's belief at one (and only one) unused signal is changed in an arbitrary way are feasible perturbations.

We then define the perturbed Spencian dynamic as follows: Let θ be the state of the dynamic at time t. With probability $1 - \lambda$, the period $t + 1$ state is unperturbed and hence $\theta(t + 1)$ is given by $s(\theta)$. With probability λ, however, a perturbation occurs, in which case $\theta(t + 1)$ is drawn from a probability measure that depends only on θ and has full support on the set of possible perturbations $P(\theta)$.

For every perturbation rate $\lambda \in (0, 1)$ (fixing all other probabilities), the perturbed dynamics constitute a Markov process, which we refer to as $\Gamma(\lambda)$. Let $A_1, \ldots, A_n$ be the absorbing sets of the process $\Gamma(\lambda)$, where an absorbing set is a minimal set of states that is closed under $\Gamma(\lambda)$. Notice that there may be multiple absorbing sets, and that the identity of these sets is independent of λ. For each absorbing set A_j, there is a unique stationary distribution with support A_j. We are interested in the behavior of our perturbed dynamics as the probability of a perturbation becomes small. We accordingly focus on limit distributions:

Definition 1 *A* **limit distribution** *is given by $\zeta_j^* = \lim_{\lambda \to 0} \zeta_j(\lambda)$, where $\zeta_j(\lambda)$ is the unique stationary distribution of $\Gamma(\lambda)$ with support on absorbing set A_j.*

Our basic tool for characterizing limit distributions is the concept of a recurrent set:

Definition 2 *A nonempty set of states R is* **recurrent** *if R is a minimal set of states such that*

(2.1) R is closed under the Spencian dynamic;

(2.2) If $\theta \in R$ and $\theta' \in P(\theta)$, then $S(\theta') \cap R \neq \emptyset$, that is, the (unperturbed) Spencian dynamic reaches a state in R starting from any perturbed state θ' in $P(\theta)$.

Lemma 3 [7] *A recurrent set exists. Every recurrent set R is the union of equilibria and cycles of the Spencian dynamic and hence $\cup_{\theta \in R} S(\theta) = R$. Recurrent sets are disjoint.*

A recurrent set is stable against single perturbations in the following sense. From any state in a recurrent set that can be reached via a single perturbation, the Spencian dynamic must ultimately reach an equilibrium or a cycle of the Spencian dynamic contained in the recurrent set. Hence, once the unperturbed dynamic has reached a state in a recurrent set, single perturbations cannot cause the dynamic to converge to an equilibrium state or two-cycle not contained in the recurrent set.

The minimality requirement in the definition of a recurrent set ensures that a recurrent set does not contain a subset of states from which it is impossible to reach the other states in the set with a single perturbation. This in turn implies that, through a sequence of perturbations followed by adjustment to an equilibrium or two-cycle, it is possible to reach any state in a recurrent set from any other state in the recurrent set.

Rabin and Sobel [8] also examine recurrent sets, though they restrict attention to recurrent sets consisting of equilibrium states and do not investigate the possibility or recurrent sets consisting of cycles. For Rabin and Sobel, the concept of a recurrent set is their basic solution concept, and [8] contains an intuitive justification for using such a concept.[8] We are interested in recurrent sets because:

Lemma 4

(4.1) If there is a unique recurrent set R then there is a unique limit distribution ζ^* and the support of ζ^* is R.

(4.2) If there are multiple recurrent sets, then the support of any limit distribution is a union of recurrent sets.

In all but one of the cases we encounter, we shall find that for every specification of perturbations satisfying Assumption 2, a unique recurrent set exists. From Lemma 4.1, there is then a unique limit distribution whose support coincides with the unique recurrent set. In addition, we will establish conditions under which all of the states in the unique recurrent set give the same outcome, obviating the need to characterize the limit distribution beyond its support.

4 Recurrent Sets and Equilibrium Selection

For every signal x, define $E(x)$ to be the set of states that are equilibria in which the high-quality seller sends signal x or are members of two-cycles in which the high-quality seller sends x. We refer to the sets $E(x)$ as **components.** Note that if a component $E(x)$

[7] Lemmas 3 and 4 holds for any specification of perturbations $P(\theta)$.

[8] Rabin and Sobel [8] differs from our work in two other important respects. Their work with a different dynamic based on iterating the best response correspondence of a Cho-Kreps signaling game. They also start with a given equilibrium refinement, use this refinement to derive a set of feasible perturbations, and then ask whether the recurrent sets of equilibria which emerge from such perturbations coincide with the predictions made by the given refinement concept. If given a different refinement, they will derive a different set of perturbations and a different collection of recurrent sets to be checked.

contains a pooling equilibrium, then it contains only pooling equilibria. The same is true of separating equilibria or two cycles.

We begin our study of recurrent sets with the observation that recurrent sets select components rather than states. This follows from the fact that states which yield the same outcome and hence differ only in the prices offered at unused signals can be connected through a sequence of perturbations:

Lemma 5 *Let R be a recurrent set. For all x, either $E(x) \subset R$ or $E(x) \cap R = \emptyset$.*

Lemma 5 tells us that recurrent sets are unions of components, but does not tell us *which* components will appear in a recurrent set. To address this question we need to investigate whether perturbations allow the system to move from one component to another. Towards this end, define

$$r(x) = \{\tilde{x} | \exists \theta \in E(x), \tilde{\theta} \in E(\tilde{x}), \theta' \in P(\theta) : \tilde{\theta} \in S(\theta')\}. \tag{9}$$

We say that the component $E(\tilde{x})$ can be **reached** from the component $E(x)$ if $\tilde{x} \in r(x)$. In such a case, there exists a state θ in $E(x)$ from which a perturbation yields a state θ' from which the Spencian dynamic in turn leads to a state $\tilde{\theta}$ in $E(\tilde{x})$. Note that from Assumption 2.1, we have $x \in r(x)$ for all non-empty sets $E(x)$. The following result is then an immediate consequence of the definition of a recurrent set and Lemma 5 (and hence its proof is omitted):

Lemma 6

(6.1) Let $x' \in r(x)$. If $E(x)$ is contained in a recurrent set then $E(x')$ is contained in the same recurrent set.

(6.2) Let $E(x)$ be non-empty. Then $E(x)$ is a recurrent set if and only if $r(x) = \{x\}$.

4.1 Uniqueness of Recurrent Sets

We now establish conditions under which there is a unique recurrent set. We first show that there exist signals x^* and $\overline{x}$ such that the component $E(x^*)$ can be reached from all components of pooling equilibria or two-cycles, whereas the component $E(\overline{x})$ can be reached from all components of separating equilibria. It then follows from Lemma 6.1 that there can be at most two recurrent sets. The first signal is

$$x^* = \arg\max_{x} u_h(x, \phi^0). \tag{10}$$

The signal x^* is the signal high-quality sellers would choose if the buyer were to believe that signals convey no information. In particular, if a pooling equilibrium at x^* exists, then such an equilibrium gives the high-quality sellers the highest utility achievable in a pooling equilibrium. Note that if a pooling equilibrium fails to exist at x^*, then there must exist either a separating equilibrium in which $x_h = x^*$ or a two cycle with $x^c = x^*$, so we always have $E(x^*) \neq \emptyset$.[9]

[9]To verify this, let $\Phi(x) = 0$ for all $x \neq x^*$ and $\Phi(x^*) = \phi^0$. Then the definition of x^* ensures that high-quality sellers send signal x^*. If low-quality sellers also send x^*, we have a pooling equilibrium. If not, then low-quality sellers send $\underline{x}$. If $u_\ell(\underline{x}, 0) < u_\ell(x^*, 1)$, then conditions (7)–(8) hold and we have a two-cycle with $x^c = x^*$. If $u_\ell(\underline{x}, 0) > u_\ell(x^*, 1)$, then setting $\Phi(x^*) = 1$ yields a separating equilibrium with $x_h = x^*$.

For our next signal, let

$$\overline{X} = \{x | u_\ell(x,1) < u_\ell(\underline{x},0),\ u_h(x,1) > \max_{x'} u_h(x',0)\}. \tag{11}$$

The set $\overline{X}$ is the set of signals x such that there exists a separating equilibrium in which type h chooses x. Then let

$$\overline{x} = \begin{cases} \arg\max_{x \in \overline{X}} u_h(x,1) & \text{if } \overline{X} \neq \emptyset \\ x^* & \text{if } \overline{X} = \emptyset \end{cases}$$

Hence, if the set $\overline{X}$ is nonempty, meaning that at least one separating equilibrium exists, then $\overline{x}$ is the signal sent by the high-quality seller in a separating equilibrium that gives high-quality sellers the highest utility level achievable in a separating equilibrium. It may happen, however, that there is no separating equilibrium. In this case, we find it convenient to set $\overline{x} = x^*$ in order to ensure that the signal $\overline{x}$ is still well-defined and the set $E(\overline{x})$ is non-empty.

If $E(x^*)$ is a set of pooling equilibria and $E(\overline{x})$ is a set of separating equilibria, then these sets are familiar from the refinement literature:

Definition 3 *A separating equilibrium θ with $x_h = \overline{x}$ is a* **Riley equilibrium**. *A pooling equilibrium θ with $x_\ell = x_h = x^*$ is a* **Hellwig equilibrium**. *We call the set of all Riley equilibria the* **Riley set** *and the set of all Hellwig equilibria the* **Hellwig set**.

These terms are motivated by Riley [9], who offers Riley equilibria as the preferred solution in a screening model, and Hellwig [3], who uses the concept of a stable equilibrium to select a Hellwig equilibrium in a variant of a screening model with many buyers who simultaneously set price schedules and can refuse to trade with some sellers after sellers have chosen their signals.

Proposition 1 *Let R be a recurrent set. If R contains a separating equilibrium then it contains $E(\overline{x})$. If R contains a pooling equilibrium or a cycle then it contains $E(x^*)$. Hence, every recurrent set contains either $E(\overline{x})$ or $E(x^*)$ and there are at most two recurrent sets.*

The intuition for Proposition 1 is straightforward. Consider, for example, a separating equilibrium which is not in the Riley set. Given such an equilibrium a single perturbation (resulting in the price offer $v(\overline{x},1)$ at signal $\overline{x}$) will induce high-quality sellers (but not low-quality sellers) to switch from their equilibrium signal to $\overline{x}$. The resulting state is a Riley equilibrium and it then follows from Lemma 6.1 that any recurrent set containing the original separating equilibrium must also contain all Riley equilibria. The argument for pooling equilibria or two-cycles is similar, the idea again being that a perturbation which yields the price $v(x^*, \phi^0)$ at signal x^* will suffice to induce high-quality sellers to switch to x^*.

We can give a simple example of a market with two recurrent sets.

Example 1 Consider the market shown in Figure 1. There are two signals in this

	$v(\underline{x},0)$	$v(\underline{x},\phi^0)$	$v(\underline{x},1)$	$v(\overline{x},0)$	$v(\overline{x},\phi^0)$	$v(\overline{x},1)$
u_ℓ	2	5	8	−6	−3	0
u_h	2	5	8	−2	1	4
		$\underline{x}=x^*$			$\overline{x}$	

Figure 1: A market with two recurrent sets

market, $\underline{x}$ (which is also x^*) and $\overline{x}$. The set Υ contains three beliefs, 0, ϕ^0 or 1. Hence, for each signal, there are three possible prices, given by $v(\cdot,0)$, $v(\cdot,\phi^0)$, and $v(\cdot,1)$. The first row gives the payoffs of an L seller from each signal and price combination. The second row gives the payoffs of an H seller from each signal and price combination.

The Hellwig set $E(x^*)$ consists of pooling equilibria in which both sellers send signal x^*. The Riley set $E(\overline{x})$ consists of the unique separating equilibrium in which $x_\ell = \underline{x}$ and $x_h = \overline{x}$. Both the Hellwig and Riley sets are recurrent. Because $u_q(x^*,\phi^0) = 5 > u_q(\overline{x},\phi)$ for all $\phi \in \{0,\phi^0,1\}$ and $q \in \{\ell,h\}$, the Hellwig set $E(x^*)$ contains all states θ satisfying $x_\ell = x_h = x^* = \underline{x}$ and $\Phi(x^*) = \phi^0$. Since perturbations only affect beliefs at unused signals, it follows that every perturbation of any state in $E(x^*)$ results in another state in $E(x^*)$. Hence, the Hellwig set is recurrent. Since there are no unused signals in the Riley equilibrium, the set of feasible perturbations consists of the Riley equilibrium itself and the Riley set is thus recurrent.[10] □

The market in Figure 1 is rather special in that the Hellwig equilibrium uses the same signal as low-quality sellers use in the Riley equilibrium ($\underline{x} = x^*$). The two recurrent sets also *coincide* with the Riley and Hellwig sets (rather than simply containing the latter, as is ensured by Proposition 1). We can show that this is not an artifact of this example. All markets with two recurrent sets have these properties:

Proposition 2 *If $x^* \neq \underline{x}$ then there exists a unique recurrent set. If $x^* = \underline{x}$ then either there exists a unique recurrent set or both the Hellwig set and the Riley set are recurrent sets.*

In some cases the existence of a unique recurrent set is immediate from Proposition 1. If the market has no separating equilibrium, then every recurrent set must contain $E(x^*)$ and uniqueness then follows from Lemma 3. At the other extreme, if $E(x^*)$ is a set of separating equilibria then every recurrent set containing $E(x^*)$ must also contain the Riley set and there is again a unique recurrent set.[11] Hence, suppose that the Riley set is non-empty and $E(x^*)$ is either a component of pooling equilibria or two-cycles. If $x^* \neq \underline{x}$, then the proof of Proposition 2 proceeds by showing that either the Hellwig set can be reached from the Riley set, so that $x^* \in r(\overline{x})$, or the Riley set can be reached

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

[11] It follows from Step 1 in the proof of Proposition 2 that in this case the unique recurrent set is exactly the Riley set.

from the Hellwig set, giving $\overline{x} \in r(x^*)$. In either case uniqueness of the recurrent set is assured from Lemma 6.1. If $\underline{x} = x^*$ the uniqueness argument is more complicated. It is shown that unless both the Riley set and the Hellwig set are recurrent, there must exist a component $E(\tilde{x})$ such that either it is the case that $\tilde{x} \in r(\overline{x})$ and $x^* \in r(\tilde{x})$ or it is the case that $\tilde{x} \in r(x^*)$ and $\overline{x} \in r(\tilde{x})$, so that it is either possible to reach the Riley set from the Hellwig set with an intermediate step through $E(\tilde{x})$ or the other way round.

Note that $x^* = \underline{x}$ is only a necessary and not a sufficient condition for the existence of two recurrent sets. The following sections show that the additional condition requiring both the Riley set and the Hellwig set to be recurrent is very stringent. There will thus often be a unique recurrent set, and hence a unique limit distribution with this support, even though $x^* = \underline{x}$.[12]

4.2 Equilibrium Selection

Proposition 1 and Lemma 6.2 tell us that the Riley set and the Hellwig set are our only candidates for recurrent sets whose members all produce a single equilibrium outcome. This section establishes conditions under which these sets are recurrent.

Identifying such conditions is an exercise analogous to that of the equilibrium refinements literature, which attempts to identify those equilibrium outcomes which are impervious against further deviations once they are reached. It is generally taken for granted when studying a favorite refinement that only equilibria satisfying the refinement deserve further study, presumably because other equilibria are susceptible to deviations. This ignores the issue of whether there are larger sets of equilibria which deserve attention because the set as a whole is stable against deviations, even if its individual members are not. As Rabin and Sobel [8] note, the members of such a set would be stable not because it is impossible for deviations to lead away from them but because it is easy for deviations from other equilibria to lead to them. In our framework, Proposition 2 implies there can be no such set if either the Riley set or the Hellwig set meets our "refinement condition" of being recurrent. Hence, if just one of the Riley or Hellwig sets is recurrent, then our model selects that component.

We consider Riley equilibria first. The Riley set is recurrent if and only if it is impossible to lure high-quality sellers away from the Riley equilibria with the promise of a pooling price at another, unused signal:

Proposition 3 *Suppose the Riley set is non-empty. Then the Riley set is a recurrent set if and only if*

$$\forall x \neq \underline{x}, \qquad u_h(x, \phi^0) < u_h(\overline{x}, 1). \tag{12}$$

To see the intuition behind this result, consider a Riley equilibrium. While a perturbation can produce a state that induces sellers to leave their equilibrium signals, subsequent adjustments can never cause the buyer to alter the price at these equilibrium signals. Each type of seller will prefer to return to his own equilibrium signal instead of sending

[12] If two recurrent sets exist, then Proposition 2 and Lemma 4.2 imply that the support of any limit distribution contains only Riley equilibria or Hellwig equilibria.

the equilibrium signal of the other type, so the buyer never has cause to revise beliefs at the original equilibrium signals. Hence, just as 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 Spencian dynamic then cannot converge to an equilibrium (or cycle) in which some type of the seller receives a lower payoff than in the original equilibrium. If the Riley set fails to be recurrent, it must accordingly be that a perturbation allows the Spencian dynamic to reach an equilibrium or cycle in which sellers earn higher payoffs than in the Riley equilibrium. From the definition of the Riley equilibrium, no other separating equilibrium provides such payoffs. A sufficient condition for recurrence is then that no pooling equilibrium or cycle provides such a payoff, which is just what (12) ensures. The remainder of the proof consists of showing that condition (12) is also necessary, meaning that if there is a signal x with the property that high-quality sellers prefer signal x and price $v(x, \phi^0)$ to the Riley equilibrium, then we can can find a perturbation such that the Spencian dynamic converges to a pooling equilibrium or cycle in which the high-quality sellers receive a higher payoff than in a Riley equilibrium and low-quality sellers receive at least as high a payoff.

While it is easy to find sufficient conditions for the recurrence of a Hellwig set, necessary conditions are a bit more subtle:[13]

Proposition 4 *Suppose the Hellwig set is non-empty. Then:*
(4.1) The Hellwig set is recurrent if for all $x \neq x^*$,

$$\forall \phi : \quad u_h(x, \phi) > u_h(x^*, \phi^0) \Rightarrow u_\ell(x, \phi) > u_\ell(x^*, \phi^0). \tag{13}$$

(4.2) If there exists a signal $x \neq x^*$ violating (13) and either (a) $E(x) \neq \emptyset$, or (b) the Riley set is non-empty, or (c) $x^* \neq \arg\max_x u_h(x, 0)$, then the Hellwig set is not recurrent.

If condition (13) holds for all signals, then it is impossible to find a perturbation such that only high-quality sellers, but not low-quality sellers, are induced to leave x^*. This implies that the buyer will never revise her belief downward at x^*. As a result, starting with a perturbation of a Hellwig equilibrium, throughout the ensuing Spencian dynamic both types can ensure at least their equilibrium payoff from the Hellwig equilibrium by using x^* (just as in the case of the Riley equilibrium). In addition, since there is no equilibrium or pooling state of a cycle giving both types a higher utility level than the one from a Hellwig equilibrium, it follows that the Spencian dynamic must converge back to a Hellwig equilibrium, ensuring that the Hellwig set is recurrent.

[13] Because we think of cycles as approximations to mixed strategy equilibria, it is interesting to note that necessary and sufficient conditions for $E(x^*)$ to be a recurrent set when $E(x^*)$ is a set of two-cycles are quite similar to the ones obtained in Proposition 4 for Hellwig sets. (It is obvious from Proposition 1 that no other set of two-cycles can be recurrent.) The main difference is that one needs to consider two cases, namely perturbations of pooling states and perturbations of separating states. First, a necessary and sufficient condition which ensures the Spencian dynamic must return to a cycle in $E(x^*)$ starting from any perturbation of a pooling state in this set is that there does not exist a Riley equilibrium satisfying $u_h(\overline{x}, 1) > u_h(x^*, \phi^0)$. Replacing condition (13) by the condition that $\forall \phi : u_h(x, \phi) > u_h(x^*, 1) \Rightarrow u_\ell(x, \phi) > u_\ell(\underline{x}, 0)$ throughout the statement of Proposition 4 (and requiring this condition for all $x \notin \{x^*, \underline{x}\}$) gives necessary and sufficient conditions for stability against perturbations which affect the separating state of a cycle in $E(x^*)$.

	$v(x^*,0)$	$v(x^*,\phi^0)$	$v(x^*,1)$	$v(\underline{x},0)$	$v(\underline{x},\phi^0)$	$v(\underline{x},1)$	$v(\tilde{x},0)$	$v(\tilde{x},\phi^0)$	$v(\tilde{x},1)$
u_ℓ	0	10	20	2	15	25	−40	−30	5
u_h	0	10	20	−8	5	15	−20	−10	25
		x^*			$\underline{x}$			$\tilde{x}$	

Figure 2: A recurrent Hellwig set that fails condition (13)

Turning to the necessary conditions, note that if condition (13) fails for some signal, then there exists a perturbation causing only high-quality sellers to leave signal x^* while low-quality sellers remain. The price offered at x^* will then drop to $v(x^*, 0)$, implying that the payoff from the original Hellwig equilibrium is no longer available. The additional conditions appearing in Proposition 4.2 ensure that, starting from such a perturbation, the Spencian dynamic will converge to an equilibrium or cycle not contained in the Hellwig set. Without such additional conditions it is possible that the Spencian dynamic must converge back to an Hellwig equilibrium. For example, it may be that condition (13) fails but *all* limits of the Spencian dynamic are Hellwig equilibria.

The following example demonstrates that the conditions in Proposition 4.2 cannot be weakened to simply require the existence of some component different from the Hellwig set. The Hellwig set fails (13) in this market, but any perturbation of a Hellwig equilibrium yields a state from which the the Spencian dynamic converges to a Hellwig equilibrium.

Example 2 Consider the market shown in Figure 2. By simply considering all of the possibilities, one can verify that this market has no separating equilibria or two-cycles. As a result, there is only one recurrent set. The only equilibria are pooling equilibria at x^* (the Hellwig set) and $\underline{x}$. The Hellwig set $E(x^*)$ fails condition (13) in Proposition 4 (the signal-belief pair $(\tilde{x}, 1)$ violates (13)). Nevertheless the Hellwig set is recurrent. To see this, note that every state in $E(\underline{x})$ must satisfy $\Phi(x^*) = 0$ and $\Phi(\tilde{x}) \leq \phi^0$. Since both types are using x^* in a Hellwig equilibrium and perturbations only affect unused signals, making the transition from a Hellwig equilibrium to an equilibrium in $E(\underline{x})$ thus requires the Spencian dynamic to reach a state in which low-quality sellers are using x^* (to achieve $\Phi(x^*) = 0$) and high-quality sellers are using $\underline{x}$ (to ensure $\Phi(\tilde{x}) \leq \phi^0$) . But there are no such states: whenever x^* is a best response for low-quality sellers, either x^* or $\tilde{x}$ is the best response for high-quality sellers. Hence, perturbations of states in the Hellwig set lead to dynamics that cannot converge to $E(\underline{x})$ and accordingly must converge to $E(x^*)$, making the Hellwig set recurrent and hence making it the unique recurrent set of this market. □

4.3 Interpretation

The results of Propositions 3 and 4 resemble those that have emerged from the equilibrium refinements literature. First, consider the Riley equilibrium. Mailath, Okuno-Fujiwara and Postlewaite [6] have recently introduced the notion of undefeated equilibrium. In our model, their definition becomes[14]

Definition 4 *Let θ be an equilibrium. Then an equilibrium θ'* **defeats** *θ if there exists a signal x' with $x' \notin \{x_\ell, x_h\}$ and $x' \in \{x'_\ell, x'_h\}$, such that, for $q \in \{\ell, h\}$,*

$$x' = x'_q \Rightarrow u_q(x', \Phi'(x')) > u_q(x, \Phi(x)). \tag{14}$$

A two-cycle defeats θ if its separating state θ' satisfies (14). An equilibrium θ is **undefeated** *is there is no equilibrium or cycle defeating it.*

It follows immediately from (14) and (12) that the Riley set is recurrent if and only if the Riley equilibria are undefeated. Since all separating equilibria which are not in the Riley set are defeated by a Riley equilibrium, Proposition 3 is thus equivalent to:

Corollary 1 *A component of separating equilibria is a recurrent set if and only if the equilibria are undefeated.*

Hence, our results do not match the recommendations of the Intuitive Criterion of Cho and Kreps [1] (and the refinements of this concept), which shows a proclivity to select separating equilibria even when they are defeated by pooling equilibria.

Now consider pooling equilibria. The criterion (13) in Proposition 4.1 is the statement that Hellwig equilibria satisfy the counterpart of the $D1$ criterion in our model.[15] An undefeated pooling equilibrium must be the Hellwig equilibrium, since any other pooling equilibrium is defeated by the members of $E(x^*)$, whether these members are Hellwig equilibria or cycles. We thus have that being undefeated and satisfying $D1$ suffices for a component of pooling equilibria to be recurrent. Proposition 4 provides conditions under which undefeated and $D1$ are necessary, giving:[16]

Corollary 2 *Let $E(x) \neq \emptyset$ for all x or let (b) or (c) of Proposition 4.2 hold. A component of pooling equilibria is recurrent if and only if the equilibria are undefeated and satisfy $D1$.*

One interesting aspect of these results is that we obtain different refinement criteria for assessing separating and pooling equilibria. For an equilibrium to be recurrent in

[14]We make two modifications to Mailath, Okuno-Fujiwara and Postlewaite's definition. First, Assumption 1.2 allows us simplify their definition by avoiding issues connected with payoff ties. Second, we allow cycles as well as equilibria to defeat an equilibrium. We view this as the counterpart of allowing mixed strategy equilibria to defeat other equilibria in Mailath, Okuno-Fujiwara and Postlewaite's model. If we allowed only equilibria to defeat other equilibria in our model, then being undefeated would be necessary but not sufficient in the following results.

[15]In particular, the $D1$ criterion, as defined by Cho and Sobel [2], formalizes the requirement that it be possible to support an equilibrium outcome with beliefs whose support contains only sellers with the strongest incentive to deviate. Under our assumptions, this requirement implies only that if (x, ϕ) violates the condition that for all ϕ, $u_h(x, \phi) > u_h(x_h, \Phi(x_h)) \Rightarrow u_\ell(x, \phi) > u_\ell(x_\ell, \Phi(x_\ell))$ with $x \notin \{x_\ell, x_h\}$, then $\Phi(x) = 1$. But then x_h is not a best reply for high-quality sellers and the equilibrium cannot be supported by beliefs with $\Phi(x) = 1$. A Hellwig equilibrium accordingly satisfies $D1$ in our model if and only if (13) holds for all $x \neq x^*$.

[16]To see that being undefeated is necessary, notice that if a Hellwig equilibrium is defeated, then it is defeated by a separating equilibrium, in which case $\overline{x}$ violates (13).

our market, it must be impossible for a perturbation to prompt learning dynamics that lead to an alternative component. If we replace "perturbation" with "deviation" and replace "learning dynamics" with "inferences drawn from the deviation", this appears very similar to the heuristic justification underlying the undefeated equilibrium concept. In particular, we share with the undefeated concept the stress on the existence of an alternative equilibrium. It is then no surprise that our results come close to that of the undefeated equilibrium concept. Differences arise because the undefeated concept shares with the rest of the equilibrium refinements literature the assumption that the original equilibrium always remains available while agents contemplate the effects of a deviation. In contrast, our model allows the possibility that in the course of the adjustments prompted by a perturbation, the original equilibrium prices are affected and sellers are accordingly induced to choose signals they would not have chosen had the original equilibrium prices remained available. This possibility does not arise in the case of a separating equilibria and being undefeated is accordingly necessary and sufficient for a separating equilibrium to be recurrent. However, it is to preclude this possibility that the conditions for the recurrence of a component of pooling equilibria include $D1$.

4.4 Recurrent Sets with Multiple Outcomes

Propositions 3 and 4 provide conditions under which the Riley and Hellwig sets are recurrent. If one of these sets is recurrent, then we know from Propositions 1 and 2 that either it is the unique recurrent set or both the Riley and Hellwig sets are recurrent. What if neither the Riley nor Hellwig sets is recurrent?

The proof of Proposition 2 shows that if the Riley set fails to be recurrent but is nonempty, then there exists a component of pooling equilibria or two-cycles which can be reached from the Riley set. In conjunction with Propositions 1 and 2, this immediately gives the following (which is accordingly stated without proof):

Proposition 5 *If neither the Riley nor Hellwig sets is recurrent, then there is a unique recurrent set containing $E(x^*)$.*

We thus have a unique recurrent set containing either the Hellwig set or a set of cycles at x^*.[17] Because the Hellwig set is not recurrent, however, the unique recurrent set must contain other outcomes in addition to the Hellwig equilibria whenever the latter is nonempty.[18]

We have an algorithm for calculating this unique recurrent set: Begin with $E(x^*)$ and iteratively apply the operation r. We know that $E(x^*)$ is contained in the recurrent set (from Proposition 5), and hence so are all components $E(x)$ with $x \in r(x^*)$ (from Lemma 6.2), and hence so are all components which can be reached from some $E(x)$ which can be reached from $E(x^*)$, and so on until a fixed point is obtained. This calculation is often quite tedious, however, preventing the achievement of general results. The following

[17] Examples are easily constructed in which a Riley equilibrium exists but is not contained in this unique recurrent set.

[18] Footnote 13 provides conditions under which the recurrent set consists precisely of the set of cycles $E(x^*)$.

example illustrates the calculation for the commonly-studied version of the Spence model in which values do not depend on signals and the single-crossing property holds.

Example 3 We assume there exists a function $v(\phi) : \Upsilon \rightarrow \Re$ such that for all $x \in X$,

$$v(x, \phi) = v(\phi), \tag{15}$$

and that for all $x_i \in \{x_1, \ldots, x_n\}$ with $i < n$, there exists a $\phi \in \Upsilon$ such that

$$\begin{aligned} 0 &< c_h(x_{i+1}) - c_h(x_i) \\ &< v(\phi) - v(\phi^0) \\ &< c_\ell(x_{i+1}) - c_\ell(x_i) \\ &< v(1) - v(\phi^0). \end{aligned} \tag{16}$$

Condition (15) states that signals do not affect values. To interpret (16), think of x_{i+1} as being a "higher" signal than x_i. Then (16) implies that higher signals are more costly for sellers to send. In addition, the single-crossing property holds, in that the marginal cost of sending a higher signal is higher for low-quality sellers than it is for high-quality sellers. Finally, signals are sufficiently close together that the single crossing property has some effect. In particular, given any signal x_i with $i < n$, the difference between receiving some price $v(\phi)$ and price $v(\phi^0)$ is sufficient to compensate high-quality but not low-quality sellers for the extra cost of sending signal x_{i+1} and the difference between $v(1)$ and $v(\phi^0)$ is sufficient to compensate both sellers for the extra cost of sending x_{i+1}.[19]

Conditions (15)–(16) immediately give $x^* = x_1 = \underline{x}$. Hence, the market has a Hellwig set at the signal with the lowest cost, x_1, which is also the signal send by low-quality sellers in any separating equilibrium. The Hellwig set is recurrent only if it is the unique component of equilibrium outcomes or cycles in the market. To see this, first note that from (16) the signal x_2 will violate (13). If $E(x_2)$ is non-empty or a separating equilibrium exists, then it follows from Proposition 4.2 that the Hellwig set is not recurrent. The Hellwig set can accordingly be recurrent only if $E(x_2)$ is empty and no separating equilibrium exists, in which case it is easy to check from condition (16) that the Hellwig set is the unique component of equilibria or cycles.

We thus have a unique recurrent set. In particular, if an undefeated Riley equilibrium exists, then the Riley set is recurrent (Proposition 3) and the Hellwig set is not, so the Riley set is the unique recurrent set (Proposition 2). If an undefeated Riley equilibrium does not exist, so that the Riley set is not recurrent, then there is again a unique recurrent set (Proposition 5). In this last case the unique recurrent set contains not only $E(x^*)$ but also every pooling equilibrium of the market. This follows from the observations that (a) if $E(x_k)$ is a component of pooling equilibria then $E(x_i)$ is a component of pooling equilibria for all $i < k$ and (b) if $E(x_i)$ and $E(x_{i+1})$ are components of pooling equilibria, then $x_{i+1} \in r(x_i)$.[20]

[19] The market in Figure 1 satisfies (15)–(16) (with $x_1 = \underline{x}$ and $x_2 = x_n = \overline{x}$) except for these final two conditions. In particular, no price can induce either seller to prefer sending x_2 instead of receiving price $v(\phi^0)$ at signal x_1.

[20] To verify this, let θ be the pooling equilibrium at x_i induced by $\Phi(x_i) = \phi^0$ and $\Phi(x) = 0$ for $x \neq x_i$.

Suppose now that there is a component of pooling equilibria at the highest-cost signal x_n (which will occur if $v_\ell(x_n, \phi^0) > v_\ell(x_1, 0)$), in which case no separating equilibria or cycles exist and there must be a pooling equilibrium at every signal. The unique recurrent set will then contain *all* equilibria. In contrast, the only undefeated equilibria are the Hellwig equilibria and the only equilibria satisfying $D1$ are the pooling equilibrium at x_n.

This last case makes two points. First, we again see that neither satisfying $D1$ nor being undefeated suffices by itself for a component of pooling equilibria to be recurrent. Instead, both conditions must hold. Second, being undefeated and satisfying the Intuitive Criterion of Cho and Kreps [1] does not suffice for a pooling equilibrium to be recurrent. In particular, suppose the previous market also satisfies $u_\ell(x_n, 1) > u_\ell(x^*, 0)$. Then the Hellwig equilibrium passes the Intuitive Criterion but continues to fail $D1$ and hence is not recurrent.[21] □

5 Extensions

We have worked with a very special model. In this section, we examine the role played by various features of the model in driving the results.

Genericity. Assumption 1.2 imposes a strong genericity assumption. Inspection of our proofs, however, indicates that a weaker assumption would suffice. In particular, what is required is that sellers have unique best replies in the Riley and Hellwig equilibria. If this condition is met, then we can reformulate the dynamic to cope with nongenericities elsewhere by assuming that sellers randomly choose a best reply when indifferent among best replies, at which point our propositions continue to hold.

Inertia. The Spencian dynamic assumes that all agents adjust immediately to the observation of the current market conditions. What if the updating of the buyer's beliefs and sellers' signal choices were subject to inertia? We model this inertia by assuming either that the same set of sellers trade in every period or each new seller is matched with a (unique) seller of the same quality from the previous generation whose signal can be observed and mimicked. Then suppose that in each period, each seller chooses a signal optimally with probability μ and simply repeats his previous signal with probability $1-\mu$. The buyer similarly updates beliefs with probability μ and retains her previous beliefs with probability $1-\mu$. This now makes the learning dynamics a stochastic rather than deterministic process.

From (16), $P(\theta)$ then contains the state θ^1 induced by $\Phi^1(x) = \Phi(x)$ for $x \neq x_{i+1}$ and $\Phi^1(x_{i+1}) = \phi$ such that $u_h(x_{i+1}, \phi) > u_h(x_i, \phi^0)$ and $u_\ell(x_{i+1}, \phi) < u_\ell(x_i, \phi^0)$. Then $x_h^1 = x_{i+1}$ and $x_\ell^1 = x_i$. Hence, $\theta^2 = s(\theta^1)$ is induced by $\Phi^2(x_{i+1}) = 1$ and $\Phi(x) = 0$ for $x \neq x_{i+1}$, giving $x_\ell^2 = x_h^2 = x_{i+1}$ and ensuring that $s(\theta^2)$ is a pooling equilibrium at x_{i+1}. This gives $x_{i+1} \in r(x_i)$.

[21] To verify the Intuitive Criterion, we note that any signal x must satisfy $u_\ell(x, 1) > u_\ell(x^*, 0)$ (because $u_\ell(x_n, 1) > u_\ell(x^*, 0)$). The Intuitive Criterion then allows the buyer to believe that x is sent by the pool of low-quality and high-quality buyers and to set price $v(\phi^0)$ at signal x. But this suffices to make x^* a best response for all sellers, given price $v(\phi^0)$ at x^*, which in turn suffices to preserve the Hellwig equilibrium outcome.

The addition of inertia does not change the equilibrium states of the learning dynamics. However, cycles will now have many more than two states, with the additional states reached in instances when some but not all agents learn.[22] It remains a possibility that every agent learns in a given period, however, and we can use this to show that every cycle contains two states that are a two-cycle under the Spencian dynamic without inertia. Hence, every cycle still contains the support of a mixed-strategy equilibrium.

Proposition 1 and 2 and the necessary conditions for recurrence established in Propositions 3 and 4 continue to hold. These proofs relied on the ability to construct paths from one equilibrium or cycle to another that involved a single perturbation followed by learning dynamics. It suffices for the results to hold that the stochastic dynamics with inertia might follow these paths.

The proofs of sufficiency in Propositions 3 and 4 rely on the fact that once a perturbation has arisen, subsequent adjustments cannot affect the prices offered at the equilibrium signals. In the case of the Riley equilibrium, this continues to hold in the presence of inertia because high-quality (low-quality) sellers will never choose $\underline{x}$ ($\overline{x}$) as long as the prices $v(\overline{x}, 1)$ and $v(\underline{x}, 0)$ are offered at $\overline{x}$ and $\underline{x}$, and the buyer's experience will never lead her to alter these prices given that signals have this property.

In the case of the Hellwig equilibrium, inertia introduces the possibility that (13) may be satisfied but a perturbation may cause some but not all sellers to switch away from x^*, an impossibility without inertia, with the change in the proportion of high-quality sellers at x^* then prompting a change in the price at x^*. The conditions for recurrence of the Hellwig equilibrium thus become more stringent. In particular, it must be impossible to induce high-quality sellers to leave x^*, since only then can we ensure that the price at x^* will not fall following a perturbation. We can accordingly obtain a sufficient condition for the recurrence of Hellwig equilibria in the face of inertia by replacing (13) with the stronger condition that there is no signal x such that $u_h(x^*, \phi^0) < u_h(x, 1)$. The effect of introducing inertia into the model is thus to make the conditions under which the Hellwig equilibrium is recurrent more demanding.

Mixed Equilibria. Three things would be required in order to eliminate cycles from our model and allow convergence to a state that exactly matches a mixed-strategy equilibrium of the Cho-Kreps model. First, we would have to either abandon our assumption that beliefs are contained in the finite set Υ or assume that Υ contains the belief needed to support the sellers' mixture, so that sellers of a given type can optimally choose different actions. Either option sacrifices our genericity assumption. This would have to be coupled with an assumption concerning how agents who are indifferent choose between strategies, so that the correct mixture over signals could be retained once it has appeared.[23] Second,

[22] Notice that by "cycle" we now mean an absorbing set of the stochastic learning dynamics that is not a singleton. The path of the learning dynamics through this set will be random, so that any given state in the set will recur at irregular intervals.

[23] For example, we might assume that sellers persist in the market or can observe their predecessors' choices, as in our discussion of inertia, and that agents randomly choose among best replies over which they are indifferent when they (or their predecessor) are not currently playing a best reply, and simply retain their strategy when playing a best reply (even if alternative best replies exist).

we would require inertia in the process by which sellers choose their actions to avoid cases in which the dynamic gets stuck in a two-cycle. Third, we would need low-quality and high-quality sellers to occur in our population in just the right proportions, so that the required mixture can be exactly achieved as a population polymorphism.

Signaling vs. Screening. Like Spence's, our model is a screening model. To construct a signaling version of this model, let sellers first choose their signals in each period. Let the buyer respond by offering the price $v(x, \Phi(x))$ at every signal that is sent by some seller. The new feature here is that sellers must choose their signals before seeing the buyer's prices. We can model this choice by assuming that sellers have beliefs about the prices offered by the buyer at the various signals and choose optimally given their beliefs. A dynamic on this enlarged state space can then be defined by assuming that sellers (just as the buyer) update their beliefs after every period to match observed behavior.[24] A result analogous to Lemma 2 no longer holds with such a dynamic, raising the possibility that new types of cycles exist. For example, we can construct a two-cycle in which the low-quality seller always sends a signal x_ℓ and the buyer always offers price $v(x_h, 1)$ at some signal x_h, and with states θ^1 and θ^2 characterized by:

$$x_\ell^1 = x_\ell \quad x_h^1 = x_h, \quad \Phi^1(x_\ell) = \phi^0, \qquad x_\ell^2 = x_h^2 = x_\ell, \quad \Phi^2(x_\ell) = 0.$$

Hence, the agents alternate between state θ^1, in which high-quality sellers choose x_h thinking the wage at x_ℓ will be $v(x_\ell, 0)$ while this wage is actually $v(x_\ell, \phi^0)$, and state θ^2, in which high-quality sellers choose x_ℓ expecting wage $v(x_\ell, \phi^0)$ but receiving $v(x_\ell, 0)$. In contrast to the two-cycles of the Spencian dynamic, these cycles—in which seller's are continually mistaken about the prices offered by the buyer—disappear once one introduces inertia into the learning process. We accordingly think that an investigation of signaling is most fruitfully coupled with inertia in learning.

Once we have a signaling model with inertia in learning, attention turns to perturbations. If we make no other changes, then perturbations of equilibrium states that affect the buyer's beliefs about the types of sellers who send unused signals will have no effect because sellers will not observe the prices induced by these beliefs. As a result, stationary points of the dynamic process will be self-confirming equilibria but need not be Nash equilibria, and every pure-strategy self-confirming equilibrium outcome corresponds to a recurrent set.

Consequently, obtaining refinement results in a signaling model requires a different specification of perturbations. For example, one could add perturbations in the beliefs of sellers concerning the prices that buyers would offer at currently unused signals. Such perturbations will potentially induce sellers to send different signals and hence observe perturbations in the buyer's price offers at unused signals. This creates possibilities for reaching one component of (self-confirming) equilibria from another component, where these pos-

[24] This again requires that sellers are either long-lived or inherit their beliefs. An alternative approach would be to follow Young [14] and assume that there is a historical record of prices offered at the various signals which sellers can consult before making their choices.

sibilities include all of the possibilities generated in a screening model with inertia.[25] As a result, Proposition 1 and the necessary conditions in Propositions 3 and 4, all of which hinge upon the ability to reach one component from another one, will continue to hold in such a model. The sufficient condition in Proposition 3 also continues to hold. As was the case with inertia in a screening model, however, the sufficient conditions for the Hellwig set to be recurrent become more stringent, requiring that there be no signal and price pair offering high-quality sellers a higher utility than earned in the Hellwig equilibrium.

Perturbations. The perturbations in our model can make any change that affects the buyer's beliefs about any single unused signal. What if we allowed perturbations more scope for changing the state?

If we allow perturbations to affect not only the buyer's actions but also the signals sent by sellers, and impose no further restrictions, then our selection results no longer hold (whereas Propositions 1 and 2 are unaffected). Consider, for example, an undefeated Riley equilibrium. A perturbation causing a low-quality seller to send signal $\overline{x}$ will induce the buyer to reduce the price at $\overline{x}$, causing it to no longer be the case that the equilibrium prices at $\underline{x}$ and $\overline{x}$ remain available after a perturbation and hence disrupting our proof that an undefeated Riley equilibrium is recurrent. A sufficient condition for our selection results to hold in the presence of perturbations in sellers' choices is that the deviation of a single seller from his equilibrium choice in a Hellwig or Riley equilibrium does not cause a price adjustment at equilibrium signals which either upsets any of the strict inequalities appearing in Proposition 3 and Proposition 4 or induces the remaining sellers to abandon their equilibrium choices.[26]

Spence [12] suggested modeling perturbations as mistakes made by sellers when choosing their signals, thus allowing perturbations to affect the actions of sellers but not the actions of buyers. Coupled with the assumptions described in the preceding paragraph, the result is much like that of a model in which perturbations affect the buyer's beliefs at unused signals but can change these beliefs to only zero or one, since these are the beliefs that could be created by a single seller switching to the unused signal. Lemma 5 will no longer hold if the set of feasible perturbations is restricted this way. It will then no longer be the case that recurrent sets must consist of unions of components of equilibria or two-cycles, invalidating Propositions 1 and 2. However, every recurrent set containing separating equilibria must still contain some element of $E(\overline{x})$, whereas every recurrent set containing pooling equilibria or cycles must contain some element of $E(x^*)$, so that every recurrent set must contain either a Riley or a Hellwig equilibrium if such equilibria exist. Condition (12) in Proposition 3 will be necessary and sufficient for the existence of a recurrent set which contains only Riley equilibria. Necessary and sufficient conditions

[25] Perturbations can never cause sellers to choose signals which are not best responses against some belief about prices at unused signals, and hence can never lead to the revelation of the buyer's choices at these signals, but such signals play no role in allowing transitions from one component to another.

[26] The precise conditions are as follows. Let x^* satisfy $u_q(x^*, \frac{H-1}{L+H}) > u_q(x, 0)$, $\forall x, q \in \{\ell, h\}$. Then Proposition 4 continues to hold if ϕ^0 is replaced by $\frac{H-1}{L+H}$ in the statement of the proposition. Let $\overline{x}$ satisfy $u_h(\overline{x}, \frac{H}{H+1}) > u_h(\underline{x}, 0)$ and $u_h(\underline{x}, \frac{1}{L+1}) < u_h(\overline{x}, 1)$. Then Proposition 3 continues to hold if $u_h(\overline{x}, 1)$ is replaced by $u_h(\overline{x}, \frac{H}{H+1})$ in the statement of the proposition.

for the existence of a recurrent set which contains only Hellwig equilibria are less stringent than in our model. Condition (13) only needs to hold for $\phi = 1$, instead of for all ϕ, since $\phi = 1$ is the only belief for the buyer which could arise through a perturbation and could also cause some type of seller to leave x^*.

6 Conclusion

We have pursued an approach to equilibrium selection in signaling models based on examining the limiting outcome of a dynamic process subjected to arbitrarily rare perturbations. On the one hand, we view this is as simply following through Spence's suggestion that dynamic models can be used to choose between equilibria. On the other hand, the motivation offered for many equilibrium refinements is implicitly dynamic, even though the equilibria and the models are static, and we view a dynamic model as the natural setting for investigating such refinements. Despite the differences between the adaptive agents of our model and the rational agents of most refinements, we find refinement ideas emerging from the dynamic analysis.

In a previous version of this paper ([7]) we worked with a model in which sellers sent signals before observing buyers' prices and which included an explicit model of price-setting behavior by multiple buyers, inertia in learning, and a much weaker genericity assumption. The most important difference is that the more complicated model in [7] comes at the cost of fewer results. As we have noted, the presence of inertia in [7] also leads to stronger necessary conditions for the recurrence of pooling equilibria than those in the current paper.

Our analysis leaves a great deal of work to be done. Our learning model is quite crude and it is important to ascertain how its results will be affected by moving to more sophisticated models. Perhaps the most important innovation here would be to relax the memoryless nature of the learning process. A single observation of the market outcome currently suffices for buyers to abandon their previous beliefs, whereas belief rules that make greater use of past information may be more realistic. We also allow perturbations to make large changes in beliefs, and hence prices, at unused signals, just as conventional equilibrium refinements allow deviations to produce drastic belief revisions. A more realistic model might allow beliefs at unused signals to drift slowly, so that only small changes in belief could occur in a single period. Finally, there remains the issue of how to model competitive behavior on the part of buyers. We have chosen not to attempt a model of competition here in order to focus attention on signaling issues, but modeling competition remains an important topic for future work.

7 Appendix: Proofs

Proof of Lemma 2. Let C be a cycle.

STEP 1: This step shows that

$$\exists x^c \in X \ s.t. \ \forall \theta \in C: \quad x_h = x^c. \tag{17}$$

To verify (17), we first show that for any state $\hat{\theta} = (\hat{\Phi}, \hat{x}_\ell, \hat{x}_h)$,

$$\hat{\Phi}(\hat{x}_h) \leq \phi^0 \Rightarrow x_h = \hat{x}_h, \quad \forall \theta \in S(\hat{\theta}). \tag{18}$$

To verify (18), fix $\hat{\theta}$ with $\hat{\Phi}(\hat{x}_h) \leq \phi^0$ and consider any state θ satisfying

$$\Phi(\hat{x}_h) \geq \hat{\Phi}(\hat{x}_h) \tag{19}$$

$$\Phi(x) \leq \hat{\Phi}(x) \quad x \neq \hat{x}_h \tag{20}$$

Then $x_h = \hat{x}_h$ (because $\hat{x}_h = \arg\max u_h(x, \hat{\Phi}(x)) \Rightarrow \hat{x}_h = \arg\max u_h(x, \Phi(x))$) and hence (from (3) and (4)), for $\theta' = s(\theta)$,

$$\begin{aligned} \Phi'(\hat{x}_h) &\geq \phi^0 \geq \hat{\Phi}(\hat{x}_h) \\ \Phi'(x) &\leq \Phi(x) \leq \hat{\Phi}(x) \quad x \neq \hat{x}_h \end{aligned}$$

and $x'_h = \hat{x}_h$. Hence, if a state θ satisfies (19)–(20), then $x_h = \hat{x}_h$ and the successor of θ also satisfies (19)–(20). Since $\hat{\theta}$ satisfies (19)–(20), we have established (18).

Next, let $\theta' = s(\theta)$. Then we show

$$x_h \neq x_\ell \Rightarrow x'_h = x_h. \tag{21}$$

To verify (21), notice that if $x_h \neq x_\ell$, then $\Phi'(x_h) \geq \Phi(x_h)$ and $\Phi'(x_\ell) \leq \Phi(x_\ell)$. Hence $u_h(x_h, \Phi'(x_h)) \geq u_h(x_h, \Phi(x_h)) > u_h(x, \Phi(x)) \geq u_h(x, \Phi'(x))$, where the last two inequalities hold for all $x \neq x_h$, implying (21).

Now let C be a cycle containing states $\hat{\theta}$ and $\tilde{\theta}$ with $\hat{x}_h \neq \tilde{x}_h$. Without loss of generality, we can take $\tilde{\theta} = s(\hat{\theta})$. From (21), we then have $\hat{x}_\ell = \hat{x}_h$. From (18), we have $\hat{\Phi}(\hat{x}_h) = 1$. Because $\tilde{\theta} = s(\hat{\theta})$, we have $\tilde{\Phi}(\hat{x}_h) = \phi^0$. Consider $S(\tilde{\theta})$. Either $S(\tilde{\theta})$ does not contain a state θ' such that $x'_h = \hat{x}_h$, a contradiction because then $\hat{\theta} \notin S(\tilde{\theta})$, or there exists $\theta' \in S(\tilde{\theta})$ such that $\phi'(\hat{x}_h) \leq \phi^0$ and $x'_h = \hat{x}_h$ (where θ' is the first state following $\tilde{\theta}$ in which high-quality sellers send $\hat{x}_h$). But then $\tilde{\theta} \notin S(\theta')$ (from (18)), a contradiction. Hence, C cannot contain states $\hat{\theta}$ and $\tilde{\theta}$ with $\hat{x}_h \neq \tilde{x}_h$, establishing (17).

STEP 2: We show that

$$\forall \theta \in C, \quad x_\ell \in \{\underline{x}, x^c\}. \tag{22}$$

Let $\hat{\theta} \in C$ satisfy $\hat{x}_\ell \neq x^c$. If we can show that

$$\hat{\Phi}(\hat{x}_\ell) = 0, \tag{23}$$

then it follows that $\hat{x}_\ell = \underline{x}$ (because $u_\ell(\hat{x}_\ell, 0) > u_\ell(x, \hat{\Phi}(x))$ for $x \neq x_\ell$ implies $x_\ell = \arg\max u_\ell(x, 0) = \underline{x}$), yielding (22). To show (23) note that, from (4), $\theta' = s(\hat{\theta})$ satisfies

$\Phi'(\hat{x}_\ell) = 0$ (since $\hat{x}_\ell \neq x^c$) and that (from (3)–(4)) for all $\theta \in S(\theta') = C$, we have $\Phi(\hat{x}_\ell) = 0$. Since $\hat{\theta} \in C$, (23) follows.

STEP 3: Let $\theta^1, \theta^2 \in C$ with $\theta^2 = s(\theta^1)$. Then $x_\ell^2 \neq x_\ell^1$ since otherwise both $x_\ell^2 = x_\ell^1$ and $x_h^2 = x_h^1 = x^c$, which implies that θ^2 is an equilibrium state, a contradiction. We can let $x_\ell^1 = x^c$ and $x_\ell^2 = \underline{x}$. Then

$$\Phi^2(x) = \begin{cases} \Phi^1(x) & \text{if } x \notin \{\underline{x}, x^c\} \\ \phi^0 & \text{if } x = x^c \\ 0 & \text{if } x = \underline{x} \end{cases}$$

Let $\theta^3 = s(\theta^2)$. Then

$$\Phi^3(x) = \begin{cases} \Phi^2(x) = \Phi^1(x) & \text{if } x \notin \{\underline{x}, x^c\} \\ 1 & \text{if } x = x^c \\ 0 & \text{if } x = \underline{x} \end{cases}$$

Then $x_\ell^3 = x^c$ (otherwise θ^3 is an equilibrium, a contradiction). Hence $\theta^4 = s(\theta^3)$ satisfies $\Phi^4 = \Phi^2$ and hence $\Phi^3 = \Phi^1$, yielding a two state cycle in which high-quality sellers send x^c in every period and low-quality sellers alternate between x^c and $\underline{x}$.

STEP 4: Conditions (7)–(8) are clearly necessary for the optimality of the sellers' actions in the two-cycle. To see sufficiency, notice that if (7)–(8) hold, then states $\theta = s(\theta')$ and $\theta' = s(\theta)$ are a two cycle, where $x_h = x_h' = x^c$, $x_\ell = x^c$, $x_\ell' = \underline{x}$, $\Phi(x) = \Phi'(x) = 0$ for $x \neq x^c$, $\Phi(x^c) = 1$, and $\Phi'(x^c) = \phi^0$. □

Proof of Lemma 5. Let $\theta, \hat{\theta}$ both be elements of the same component. Then $x_h = \hat{x}_h$ and we can also suppose without loss of generality that $x_l = \hat{x}_l$ (If this equality does not hold then the component is a component of two-cycles and $\hat{\theta}$ can be replaced with $s(\hat{\theta})$ in the following arguments without affecting the conclusion). Recall that the signals in X are ordered $\{x_1, \ldots, x_n\}$. Let $\theta^0 = \theta$ and let θ^i, $i = 1, \ldots n$, be induced by

$$\Phi^i(x) = \begin{cases} \Phi^{i-1}(x) & x \neq x_i \\ \hat{\Phi}(x) & \text{if } x = x_i \end{cases}$$

In the case in which both θ and $\hat{\theta}$ are equilibria, then for all i, θ^i is also an equilibrium state with $x_q^i = x_q$. Furthermore, from Assumptions 2.1 and 2.3, $\theta^i \in P(\theta^{i-1})$ (because $x_q^i = x_q$ for all $i \in \{1, \ldots, n\}$ and Φ^i differs from Φ^{i-1} only if $x_i \notin \{x_\ell, x_h\}$). Similarly, in the case of two-state cycles, either $\theta^i = \theta^{i-1}$ or $\theta^i \in P(s(\theta^{i-1}))$. In either case, if R is a recurrent set, $\theta^{i-1} \in R$ thus implies $\theta^i \in R$. Hence, $\theta^0 = \theta \in R$ implies $\theta^n = \hat{\theta} \in R$. □

Proof of Proposition 1.

Since recurrent sets are disjoint (cf. Lemma 3) the result follows from Lemma 6.1 if we can show $\overline{x} \in r(\hat{x})$ for every component of separating equilibria $E(\hat{x})$ and can show $x^* \in r(\hat{x})$ for every component of pooling equilibria or two-cycles $E(\hat{x})$.

CASE 1: Let $E(\hat{x})$ be a component of separating equilibria and let $\theta^1 = (\Phi^1, \underline{x}, \hat{x})$ be any state in $E(\hat{x})$. Then consider the state $\theta^2 \in P(\theta^1)$ satisfying

$$\Phi^2(x) = \begin{cases} \Phi^1(x) & \text{if } x \neq \overline{x} \\ 1 & \text{if } x = \overline{x} \end{cases}$$

Then $x_\ell^2 = \underline{x}$ because $u_\ell(\underline{x}, 0) > u_\ell(x, \Phi^1(x))$ for $x \neq \overline{x}$ (because θ^1 is an equilibrium) and $u_\ell(\underline{x}, 0) > u_\ell(\overline{x}, 1)$ (by definition of $\overline{x}$). Similarly, $x_h^2 = \overline{x}$ because $u_h(\overline{x}, 1) > u_h(x, \Phi^1(x))$ for $x \neq \overline{x}$ (from the definition of $\overline{x}$ and the fact that θ^1 is an equilibrium). Hence $\theta^2 \in E(\overline{x})$ and thus $\overline{x} \in r(\hat{x})$.

CASE 2: Let $E(\hat{x})$ be a component of pooling equilibria and let $\theta^1 \in E(\hat{x})$ be the state induced by

$$\Phi^1(x) = \begin{cases} 0 & \text{if } x \neq \hat{x} \\ \phi^0 & \text{if } x = \hat{x} \end{cases}$$

Then consider state $\theta^2 \in P(\theta^1)$ with

$$\Phi^2(x) = \begin{cases} \Phi^1(x) & \text{if } x \neq x^* \\ \phi^0 & \text{if } x = x^* \end{cases}$$

Then $x_h^2 = x^*$ because $u_h(x^*, \phi^0) > u_h(x, \phi^0) > u_h(x, 0), \forall x \neq x^*$ (by definition of x^*). If $x_\ell^2 = x^*$, then θ^2 is a pooling equilibrium in $E(x^*)$ and thus $x^* \in r(\hat{x})$.

If $x_\ell^2 \neq x^*$, then $x_\ell^2 = \hat{x}$ (because $x_\ell = \hat{x}$ in the pooling equilibrium θ_1). Then $\theta^3 = s(\theta^2)$ satisfies

$$\Phi^3(x) = \begin{cases} 0 & \text{if } x \neq x^* \\ 1 & \text{if } x = x^* \end{cases}$$

Then either $x_\ell^3 = \underline{x}$ (in which case θ^3 is a separating equilibrium in $E(x^*)$) or $x_\ell^3 = x^*$. In the latter case θ^3 is either the pooling state of a two cycle in $E(x^*)$ or $s(\theta^3)$ is a pooling equilibrium in $E(x^*)$. In either case $x^* \in r(\hat{x})$.

CASE 3: Let $E(\hat{x})$ be a component of two-cycles and let $\theta^1 \in E(\hat{x})$ be the state induced by

$$\Phi^1(x) = \begin{cases} 0 & \text{if } x \neq \hat{x} \\ 1 & \text{if } x = \hat{x} \end{cases}$$

Let $\theta^2 \in P(\theta^1)$ be the state induced by

$$\Phi^2(x) = \begin{cases} \phi^0 & \text{if } x \in \{\hat{x}, x^*\} \\ \Phi^1(x) & \text{if } x \notin \{\hat{x}, x^*\} \end{cases}$$

Then $x_h^2 = x^*$ (from the definition of x^*) and either $x_\ell^2 = x^*$ or $x_\ell^2 = \underline{x}$. In the first case θ^2 is a pooling equilibrium contained in $E(x^*)$. In the second case $\theta^3 = s(\theta^2)$ is either a separating equilibrium contained in $E(x^*)$ or the pooling state of a two-cycle contained in $E(x^*)$. In each case $x^* \in r(\hat{x})$. □

Proof of Proposition 2.

If no separating equilibria exist, Proposition 1 implies that there is a unique recurrent set containing $E(x^*)$. If $E(x^*)$ is a component of separating equilibria, then Proposition 1 implies that every recurrent set containing $E(x^*)$ also contains $E(\overline{x})$ and there is thus a unique recurrent set containing $E(\overline{x})$. It thus remains to consider the case in which $E(x^*)$ is a component of pooling equilibria or two-cycles and $E(\overline{x})$ is the Riley set. Our first three steps establish three preliminary results.

STEP 1: Let $\theta^1 \in E(\overline{x})$ and $\theta^2 \in P(\theta^1)$. Then $\forall \theta \in S(\theta^2)$:

$$\phi(\underline{x}) = 0, \quad \phi(\overline{x}) = 1. \tag{24}$$

In particular, if $E(\hat{x})$ is a component of separating equilibria with $\hat{x} \neq \overline{x}$ then

$$\hat{x} \notin r(\overline{x}), \tag{25}$$

since (from the Definition of $\overline{x}$) $E(\hat{x})$ contains no state θ satisfying $\Phi(\overline{x}) = 1$.

To verify (24), note that θ^1 satisfies (24) because it is a Riley equilibrium. θ^2 then satisfies (24) from Assumption 2.2. It then suffices to show that if (24) holds for some state, then (24) also holds for $\hat{\theta} = s(\theta)$. But if (24) holds for some state θ, then $x_h \neq \underline{x}$ and $x_\ell \neq \overline{x}$ (because $u_h(\overline{x}, 1) > u_h(\underline{x}, 0)$ and $u_\ell(\underline{x}, 0) > u_\ell(\overline{x}, 1)$), implying $\hat{\Phi}(\underline{x}) = \Phi(\underline{x}) = 0$ and $\hat{\Phi}(\overline{x}) = \Phi(\overline{x}) = 1$.

STEP 2: Let $\hat{x} \neq \underline{x}$ satisfy

$$u_h(\hat{x}, \phi^0) > u_h(\overline{x}, 1). \tag{26}$$

Then

$$E(\hat{x}) \text{ is a component of pooling equilibria or two-cycles satisfying } \hat{x} \in r(\overline{x}). \tag{27}$$

To show (27), let $\theta^1 \in E(\overline{x})$ be the state induced by

$$\Phi^1(x) = \begin{cases} 0 & \text{if } x \neq \overline{x} \\ 1 & \text{if } x = \overline{x} \end{cases}$$

Consider $\theta^2 \in P(\theta^1)$ induced by

$$\Phi^2(x) = \begin{cases} \Phi^1(x) & \text{if } x \neq \hat{x} \\ \phi^0 & \text{if } x = \hat{x} \end{cases}$$

Then $x_h^2 = \hat{x}$ (because, from (26), $u_h(\hat{x}, \phi^0) > u_h(\overline{x}, 1) > u_h(x, \Phi^1(x))$ for $x \neq \overline{x}$). If $x_\ell^2 = \hat{x}$, then θ^2 is a pooling equilibrium in $E(\hat{x})$, implying $\hat{x} \in r(\overline{x})$. Otherwise, $x_\ell^2 = \underline{x}$ (because θ^1 is an equilibrium). Then $\theta^3 = s(\theta^2)$ satisfies:

$$\Phi^3(x) = \begin{cases} 0 & \text{if } x \notin \{\hat{x}, \overline{x}\} \\ 1 & \text{if } x \in \{\hat{x}, \overline{x}\} \end{cases}$$

Then $x_h^3 = \hat{x}$. In addition, $x_\ell^3 = \hat{x}$ (because otherwise θ^3 would be a separating equilibrium with $u_h(x_h^3, 1) > u_h(x_h^3, \phi^0) > u_h(\overline{x}, 1)$, contradicting the definition of $\overline{x}$). But then $\{\theta^2, \theta^3\}$ is a two-cycle contained in $E(\hat{x})$, implying $\hat{x} \in r(\overline{x})$.

STEP 3: Suppose $E(\hat{x})$ is a component of pooling equilibria or two-cycles and let $\hat{x}$ satisfy

$$u_h(\hat{x}, \phi^0) < u_h(\overline{x}, 1). \tag{28}$$

Then

$$\overline{x} \in r(\hat{x}). \tag{29}$$

To show (29), suppose first that $E(\hat{x})$ is a set of pooling equilibria. Let $\theta^1 \in E(\hat{x})$ be induced by

$$\Phi^1(x) = \begin{cases} 0 & \text{if } x \neq \hat{x} \\ \phi^0 & \text{if } x = \hat{x} \end{cases}$$

If $E(\hat{x})$ is a set of two-cycles, let θ^1 be induced by

$$\Phi^1(x) = \begin{cases} 0 & \text{if } x \neq \hat{x} \\ 1 & \text{if } x = \hat{x} \end{cases}$$

In either case θ^2, induced by

$$\Phi^2(x) = \begin{cases} 0 & \text{if } x \notin \{\hat{x}, \overline{x}\} \\ \phi^0 & \text{if } x = \hat{x} \\ 1 & \text{if } x = \overline{x} \end{cases}$$

is in $P(\theta^1)$ and, from (28), satisfies $x_h^2 = \overline{x}$. Furthermore, from the definition of $\underline{x}$ and the fact that $E(\overline{x})$ is the Riley set, $x_\ell^2 \in \{\underline{x}, \hat{x}\}$. If $x_\ell^2 = \underline{x}$ then $\theta^2 \in E(\overline{x})$, implying (29). If $x_\ell^2 = \hat{x}$ then $\theta^3 = s(\theta^2)$ is induced by

$$\Phi^3(x) = \begin{cases} 0 & \text{if } x \neq \overline{x} \\ 1 & \text{if } x = \overline{x} \end{cases}$$

and thus $\theta^3 \in E(\overline{x})$, implying (29).

Now consider the two cases in the statement of Proposition 2.

CASE 1: Suppose $x^* \neq \underline{x}$. Either

$$u_h(\overline{x}, 1) > u_h(x^*, \phi^0) \tag{30}$$

or

$$u_h(x^*, \phi^0) > u_h(\overline{x}, 1) \tag{31}$$

must be satisfied. If (30) holds, then $\overline{x} \in r(x^*)$ (from Step 3) and it then follows from Lemma 6.1 that every recurrent set containing $E(x^*)$ must also contain the Riley set. Hence, from Proposition 1, there is a unique recurrent set containing the Riley set. If (31) holds, then $x^* \in r(\overline{x})$ (from Step 2) and it then follows from Lemma 6.1 that every recurrent set containing the Riley set must also contain $E(x^*)$. Hence, from Proposition 1, there is a unique recurrent set containing $E(x^*)$.

CASE 2: Suppose $x^* = \underline{x}$. If the Riley set is not recurrent, it follows from (25) in Step 1 and Lemma 6.2 that there must exist a component of pooling equilibria or two-cycles $E(\hat{x})$ such that $\hat{x} \in r(\overline{x})$. It then follows from Lemma 6.1 that every recurrent set containing the Riley set must contain $E(\hat{x})$. Hence (from Proposition 1), there is a unique recurrent set containing $E(x^*)$. Suppose then that the Riley set is recurrent. If $E(x^*)$ is not recurrent, it follows from Lemma 6.2 that there exists $\hat{x} \neq \underline{x} \in r(x^*)$. If $E(\hat{x})$ is a component of separating equilibria it follows from Lemma 6.2 and Proposition 1 that every recurrent set containing $E(x^*)$ must also contain the Riley set. Since the Riley set is recurrent it is thus the unique recurrent set. If $E(\hat{x})$ is a component of pooling equilibria or two-cycles, condition (28) must be satisfied (otherwise (26) holds, contradicting the recurrence of the Riley set). Hence, $\overline{x} \in r(\hat{x})$ and it follows from Lemma 6.1 and Proposition 1 that the Riley set is the unique recurrent set.

Since $x^* = \underline{x}$ implies that $E(x^*)$ is a component of pooling equilibria it thus follows that two recurrent sets exist if and only if both the Riley set and the Hellwig set are recurrent. □

Proof of Proposition 3.

IF: Let θ^1 be a Riley equilibrium and let $\theta^2 \in P(\theta^1)$. Then, from Step 1 in the proof of Proposition 2, every state in $S(\theta^2)$ satisfies (24). From condition (12) every pooling equilibrium or two-cycle state $\theta = (\Phi(x), x_\ell, x_h)$ must satisfy $\Phi(\overline{x}) < 1$ (otherwise $x_h \notin \arg\max u_h(x, \Phi(x))$, contradicting the definition of an equilibrium state) and $S(\theta^2)$ thus does not contain any pooling equilibria or two-cycle states. It then follows from (25) in Step 1 of the proof of Proposition 2 that the only equilibrium states contained $S(\theta^2)$ are Riley equilibria. Hence $r(\overline{x}) = \{\overline{x}\}$. From Lemma 6.2 the Riley set thus is a recurrent set.

ONLY IF: If condition (12) fails, (27) in Step 2 of the proof of Proposition 2 implies that $r(\overline{x}) \neq \{\overline{x}\}$. Hence, the Riley set is not recurrent. □

Proof of Proposition 4.

IF: Let (13) hold for all $x \neq x^*$. Let $\theta^1 \in E(x^*)$ and $\theta^2 \in P(\theta^1)$. Then it suffices to show that

$$\forall \theta \in S(\theta^2), \quad \Phi(x^*) \geq \phi^0. \tag{32}$$

In particular, $S(\theta^2)$ then cannot contain a pooling equilibrium (other than a Hellwig equilibrium) or a two-cycle (because $u_h(x^*, \phi^0) > u_h(x, \phi^0)$ for all $x \neq x^*$) and cannot contain a separating equilibrium (because $u_\ell(\underline{x}, 0) < u_\ell(x^*, \phi^0)$). Hence, $r(x^*) = \{x^*\}$, ensuring that the Hellwig set is recurrent.

It remains to verify (32). First, $\Phi^2(x^*) = \phi^0$. Next, if θ satisfies $\Phi(x^*) = \phi^0$ or satisfies $\Phi(x^*) = 1$ and $x_h = x^*$, then $\theta' = s(\theta)$ also satisfies one of these conditions. To show this, we consider six cases. (1) $\Phi(x^*) = \phi^0$ and $x_\ell = x_h = x^*$. Then $\Phi'(x^*) = \phi^0$. (2) $\Phi(x^*) = \phi^0$ and $x^* \notin \{x_\ell, x_h\}$. Again, $\Phi'(x^*) = \phi^0$. (3) $\Phi(x^*) = \phi^0$ and $x_\ell = x^*$ and $x_h \neq x^*$. This contradicts the condition that there is no (x, ϕ) satisfying (13). (4) $\Phi(x^*) = \phi^0$ and $x_h = x^*$ and $x_\ell \neq x^*$. Then $\Phi'(x^*) = 1$ and $x_h' = x^*$. (5) $\Phi(x^*) = 1$ and $x_h = x_l = x^*$. Then $\Phi'(x^*) = \phi^0$. (6) $\Phi(x^*) = 1$, $x_h = x^*$ and $x_\ell \neq x^*$. Then $\Phi'(x^*) = 1$ and $x_h' = x^*$.

ONLY IF: Assume (13) fails for some signal-belief pair $(\hat{x}, \hat{\phi})$.

CASE 1 ($E(\hat{x}) \neq \emptyset$): Let θ^1 be the state in $E(x^*)$ induced by

$$\Phi^1(x) = \begin{cases} \phi^0 & \text{if } x = x^* \\ 0 & \text{if } x \neq x^* \end{cases}$$

Consider the state $\theta^2 \in P(\theta^1)$ induced by

$$\Phi^2(x) = \begin{cases} \Phi^1(x) & \text{if } x \neq \hat{x} \\ \hat{\phi} & \text{if } x = \hat{x} \end{cases}$$

Then $x_\ell^2 = x^*$ and $x_h^2 = \hat{x}$. Consequently, $\theta^3 = s(\theta^2)$ satisfies

$$\Phi^3(x) = \begin{cases} 0 & \text{if } x \neq \hat{x} \\ 1 & \text{if } x = \hat{x} \end{cases}$$

If $E(\hat{x})$ is a set of separating equilibria then $\theta^3 \in E(\hat{x})$. If $E(\hat{x})$ is a set of two-cycles then θ^3 is the pooling state of a two-cycle contained in $E(\hat{x})$. Finally, if $E(\hat{x})$ is a set of pooling equilibria then $x_\ell^3 = \hat{x}$ and $s(\theta^3)$ is a pooling equilibrium contained in $E(\hat{x})$. In any case $\hat{x} \in r(x^*)$ and the Hellwig set thus fails to be recurrent.

CASE 2: (The Riley set is non-empty): Suppose first that $u_h(\overline{x}, 1) > u_h(x^*, \phi^0)$. Since $u_\ell(x^*, \phi^0) > u_\ell(\underline{x}, 0) > u_\ell(\overline{x}, 1)$, it follows that $(x, \phi) = (\overline{x}, 1)$ satisfies the two inequalities in (13) and the result then follows from the previous case. Hence, suppose that $u_h(\overline{x}, 1) < u_h(x^*, \phi^0)$. Then there exists $(\tilde{x}, \tilde{\phi})$ satisfying (13) with $\tilde{x} \neq \overline{x}$. The state θ^1 induced by

$$\Phi^1(x) = \begin{cases} \phi^0 & \text{if } x = x^* \\ 1 & \text{if } x = \overline{x} \\ 0 & \text{if } x \notin \{\overline{x}, x^*\} \end{cases}$$

is an element of $E(x^*)$. Consider $\theta^2 \in P(\theta^1)$ induced by

$$\Phi^2(x) = \begin{cases} \Phi^1(x) & \text{if } x \neq \tilde{x} \\ \tilde{\phi} & \text{if } x = \tilde{x} \end{cases}$$

Then (from (13)) $x_\ell^2 = x^*$ and $x_h^2 = \tilde{x}$. Hence $\theta^3 = s(\theta^2)$ satisfies

$$\Phi^3(x) = \begin{cases} 0 & \text{if } x \notin \{\tilde{x}, \overline{x}\} \\ 1 & \text{if } x \in \{\tilde{x}, \overline{x}\} \end{cases}$$

Then for all $\theta \in S(\theta^3)$,

$$\Phi(\overline{x}) = 1 \qquad \Phi(x^*) = 0. \tag{33}$$

To verify (33), notice that θ^3 satisfies (33). But if $\theta \in S(\theta^3)$ satisfies (33), then so does $s(\theta)$ (because $u_h(\overline{x}, 1) > u_h(x^*, 0)$ and hence $x_h \neq x^*$, and $u_\ell(\underline{x}, 0) > u_\ell(\overline{x}, 1)$ and hence $x_\ell \neq \overline{x}$), yielding (33) for all states in $S(\theta^3)$. Since all states in $S(\theta^3)$ satisfy $\Phi(x^*) = 0$ it follows that the set of Hellwig equilibria is not recurrent.

CASE 3 ($x^* \neq \arg\max_x u_h(x, 0)$): Let θ be any state satisfying $\Phi(x^*) = 0$. Then $\tilde{\theta} = s(\theta)$ satisfies $\tilde{\Phi}(x^*) = 0$ since $u_h(x^*, 0) < \max u_h(x, 0) \leq \max u_h(x, \Phi(x))$ and thus $x_h \neq x^*$. Now consider states θ^1 and $\theta^2 \in P(\theta^1)$ defined as in Case 1. Then $\theta^3 = s(\theta^2)$ satisfies $\Phi^3(x^*) = 0$. Consequently:

$$\theta \in S(\theta^3) \Rightarrow \Phi^3(x^*) = 0. \tag{34}$$

and the set of Hellwig equilibria is thus not recurrent. □

References

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

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

[3] Martin Hellwig. Some recent developments in the theory of competition in markets with adverse selection. *European Economic Review*, 1/2:319–325, 1987.

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

[5] David M. Kreps and Joel Sobel. Signalling. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory, Volume 2*, pages 849–867. North Holland, New York, 1994.

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

[7] Georg Nöldeke and Larry Samuelson. Learning to signal in markets. SSRI working paper 9409, University of Wisconsin, 1994.

[8] Matthew Rabin and Joel Sobel. Deviations, dynamics, and equilibrium refinements. Mimeo, University of California-Berkeley and University of California- San Diego, 1994.

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

[10] E. Seneta. *Non-Negative Matrices and Markov Chains.* Springer Verlag, New York, 1981.

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

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

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

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

8 Appendix II: Potentially Omitted Proofs

Proof of Lemma 3 The existence of a recurrent set follows from the facts that the entire state space satisfies (2.1)–(2.2) and there is a finite number of subsets of the state space, so there must be a minimal set satisfying (2.1)–(2.2). Next, if R is any set satisfying (2.1)–(2.2), then the set obtained by removing from R all states that are not part of equilibria or cycles also satisfies (2.1)–(2.2), and hence minimality ensures that R contains only equilibria and members of cycles. Finally, if two sets satisfying (2.1)–(2.2) intersect, then their intersection also satisfies (2.1)–(2.2), and minimality then ensures that no recurrent set can intersect another recurrent set other than itself. □

Proof of Lemma 4 It suffices to show that if θ is contained in the support of a limit distribution ζ^*, then (a) θ is contained in a recurrent set R, and (b) all $\theta' \in R$ are contained in the support of the limit distribution ζ^*.

STEP 1: We first collect the necessary tools. Let $A_1, \ldots, A_n$ be the absorbing sets of $\Gamma(\lambda)$, $\lambda \in (0,1)$. Fix A_j. We consider the irreducible Markov process given by restricting $\Gamma(\lambda)$ to the state space A_j. This process has a unique stationary distribution with support contained in A_j (Seneta [10, Theorem 4.1]). We let $\zeta(\theta, \lambda)$ denote this distribution (suppressing the subscript "j").

For state $\theta \in A_j$, a θ-tree is a directed graph with the properties that (1) there is one node for each state in A_j, (2) each state in $A_j \setminus \{\theta\}$ is the origin of a single edge, (3) θ is not the origin of an edge, (4) the graph contains no cycles. Hence, the graph contains a path from θ' to θ for all $\theta' \in A_j \setminus \{\theta\}$. We let $(\theta' \rightarrow \theta'')$ denote an edge (from θ' to θ'') in a θ-tree. We let t_θ denote a θ-tree and let T_θ denote the set of θ-trees.

For each edge $(\theta' \rightarrow \theta'')$ in a θ-tree, let $\psi(\theta', \theta'', \lambda)$ be the probability of a transition from θ' to θ'' in $\Gamma(\lambda)$ and let $\psi(\theta, \theta', 0)$ be the probability of this transition in the unperturbed dynamic. Let

$$\Psi(t_\theta, \lambda) = \prod_{(\theta' \rightarrow \theta'') \in t_\theta} \psi(\theta', \theta'', \lambda)$$

be the product of these transition probabilities in a θ-tree t_θ. Then Freidlin and Wentzell[27] (Lemma 3.1, P. 177) show that

$$\zeta(\theta, \lambda) = \frac{\sum_{t_\theta \in T_\theta} \Psi(t_\theta, \lambda)}{\sum_{\theta' \in A_j} \sum_{t_{\theta'} \in T_{\theta'}} \Psi(t_{\theta'}, \lambda)}. \tag{35}$$

We now investigate the limit distribution $\zeta^*(\theta) = \lim_{\lambda \rightarrow 0} \zeta(\theta, \lambda)$. If $\Psi(t_\theta, \lambda)$ is nonzero, let $N(t_\theta)$ be number of transitions $(\theta' \rightarrow \theta'')$ in t_θ for which $\psi(\theta', \theta'', 0) = 0$ and let $N^*(\theta) = \min_{t_\theta \in T_\theta} N(t_\theta)$. Then it follows from (35) that

$$\zeta^*(\theta) > 0 \Longleftrightarrow N^*(\theta) = \min_{\theta' \in A_j} N^*(\theta'). \tag{36}$$

In particular, $\Psi(t_\theta, \lambda)$ is of the form $\lambda^{N(t_\theta)}(1-\lambda)^{(|A_j|-1-N(t_\theta))}(k(t_\theta) + O(\lambda))$, where $k(t_\theta)$ does not depend on λ. (36) follows from (35) and the fact that $\zeta^*(\theta) = \lim_{\lambda \rightarrow 0} \zeta(\theta, \lambda)$.

[27] M. I. Freidlin and A. D. Wentzell, *Random Perturbations of Dynamical Systems*, Springer-Verlag, New York, 1984.

STEP 2: We show that if $\theta^1 = s(\theta)$, then $N^*(\theta^1) \le N^*(\theta)$. This implies that if one state of a two-cycle of the Spencian dynamic minimizes $N^*(\theta)$, then so does the other state of the cycle. Hence, neither or both states of a two-cycle of the Spencian dynamic are in the support of the limit distribution. To verify this, let t^*_θ be a θ-tree with $N(t^*_\theta) = N^*(\theta)$. Note that t^*_θ contains a transition $(\theta^1 \to \theta^2)$. Now construct a θ^1-tree t_{θ^1} by deleting $(\theta^1 \to \theta^2)$ from t^*_θ and adding $(\theta \to \theta^1)$. Since $\theta^1 = s(\theta)$, $N(t_{\theta^1}) \le N(t^*_\theta) = N^*(\theta)$, giving the result.

STEP 3: We now show that if each of Q^1 and Q^n is a set containing a single equilibrium or two-cycle, and if both are contained in a single recurrent set R, then the elements of Q^1 appear in the limit distribution if and only if the elements of Q^n also appear. From Lemma 3, this implies that if one element in R appears in the limit distribution, then so do all elements in R, yielding (b).

Let Q^1 and Q^n be as stated. Then there exist states $\theta^1 \in Q^1$, $\theta^n \in Q^n$, and states $\theta^2, \ldots, \theta^{n-1}$ with $\theta^2 \in P(\theta^1)$ and $\theta^{i+1} = s(\theta^i)$, $i = 2, \ldots, n-1$. Let $t^*_{\theta^1}$ be a θ^1-tree with $N(t^*_{\theta^1}) = N^*(\theta^1)$. Then $t^*_{\theta^1}$ must contain a transition $(\theta' \to \theta'')$ for some $\theta' \in Q^n$ and $\theta'' \notin Q^n$. Then construct a θ^n-tree by deleting from $t^*_{\theta^1}$ the transitions beginning with $\theta^2, \ldots, \theta^n$ and θ' while adding the transitions $\theta^i \to \theta^{i+1}$, $i = 1, \ldots n-1$ and also adding the transition $\theta' \to \theta^n$ if $\theta' \neq \theta^n$. The only one of the added transitions for which $\phi(\cdot,\cdot,0) = 0$ is $(\theta^1 \to \theta^2)$. Since $\phi(\theta', \theta'', 0) = 0$, we then have $N(t_{\theta^n}) \le N(t^*_{\theta^1}) = N^*(\theta^1)$. This establishes that θ^1 appears in the limit distribution only if θ^n does, which suffices for the result.

STEP 4: We now prove (a). For this, it suffices to show that if θ^1 is not contained in a recurrent set and θ^n is either an equilibrium or an element of a two-cycle with $\theta^n \in S(P(\theta^1)) \equiv \{\theta | \exists \theta' \in P(\theta^1) : \theta \in S(\theta')\}$, then $N^*(\theta^n) \le N^*(\theta^1)$ and $N^*(\theta^n) < N^*(\theta^1)$ if θ^n is contained in a recurrent set. The first inequality follows from Step 3. Hence, we consider the second. Because $\theta^n \in S(P(\theta^1))$, there exists $\theta^2, \ldots, \theta^{n-1}$ with $\theta^2 \in P(\theta^1)$ and $\theta^{i+1} = s(\theta^i)$, $i = 2, \ldots, n-1$. Let Q^n be the set consisting of θ^n if θ^n is an equilibrium and otherwise let Q^n be the two-cycle containing θ^n. Let $t^*_{\theta^1}$ be a θ^1-tree with $N(t^*_{\theta^1}) = N^*(\theta^1)$. Then $t^*(\theta^1)$ contains transitions $(\theta^{n+i}, \theta^{n+i+1})$ for $i = 0, \ldots, k-1$ and (θ^{n+k}, θ^1) with the properties that there exist $i', i'' \in \{0, \ldots k\}$ such that $\theta^{n+i'} \in Q^n$ but $\theta^{n+i'+1} \notin Q^n$ (because $\theta^1 \notin Q^n$) and $\theta^{n+i''} \in S(P(Q^n)) \equiv \cup_{\theta \in Q^n} S(P(\theta))$ but $\theta^{n+i''+1} \notin S(P(Q^n))$ (because Q^n is contained in a recurrent set that does not include θ^1). We now construct a θ^n-tree by deleting from $t^*_{\theta^1}$ every transition beginning with $\theta^2, \ldots, \theta^n$ or beginning with any element of $S(\theta^{n+i'}) \cup S(\theta^{n+i''})$. Then add the transitions:

$$
\begin{array}{lcll}
\theta^i & \to & \theta^{i+1} & i = 1, \ldots, n-1 \\
\theta & \to & s(\theta) & \forall \theta \in \{S(\theta^{n+i'}) \cup S(\theta^{n+i'+1})\} \setminus \{\theta^n\}.
\end{array}
$$

The result is a θ^n-tree. In addition, of the transitions added to $t^*_{\theta^1}$, only $(\theta^1 \to \theta^2)$ satisfies $\phi(\cdot,\cdot,0) = 0$, while of the deleted transitions $\phi(\theta^{n+i'}, \theta^{n+i'+1}, 0) = 0$ (because $\theta^{n+i'} \in Q^n$ but $\theta^{n+i'+1} \notin Q^n$) and $\phi(\theta^{n+i''}, \theta^{n+i''+1}, 0) = 0$ (because $\theta^{n+i''} \in S(P(Q^n))$ but $\theta^{n+i''+1} \notin S(P(Q^n))$). Then $N(t_{\theta^n}) < N(t_{\theta^1}) = N^*(\theta^1)$, giving the result. □

Institut für Höhere Studien
Institute for Advanced Studies
Stumpergasse 56
A-1060 Vienna
Austria

Phone: +43-1-599 91-145
Fax: +43-1-599 91-163