



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

The Stability of Price Dispersion under Seller and Consumer Learning

Citation for published version:

Hopkins, E & Seymour, RM 2004 'The Stability of Price Dispersion under Seller and Consumer Learning' ESE Discussion Papers, Edinburgh School of Economics, University of Edinburgh.
<<http://ideas.repec.org/p/edn/esedps/52.html>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



The Stability of Price Dispersion under Seller and Consumer Learning

Ed Hopkins*
Department of Economics
University of Edinburgh
Edinburgh EH8 9JY, UK

Robert M. Seymour†
Department of Mathematics
University College
London WC1E 6BT, UK

December, 2000

Abstract

In many markets it is possible to find rival sellers charging different prices for the same good. Earlier research has attempted to explain this phenomenon by demonstrating the existence of dispersed price equilibria when consumers must make use of costly search to discover prices. We ask whether such equilibria can be learnt when sellers adjust prices adaptively in response to current market conditions. With consumer behaviour fixed, convergence to a dispersed price equilibrium is possible in some cases. However, once consumer learning is introduced, the monopoly outcome first found by Diamond (1971) is the only stable equilibrium.

Journal of Economic Literature classification numbers: C72, D83.

Keywords: Learning, Evolution, Search, Price Dispersion.

*This paper has benefitted from comments from Dan Friedman, Carmen Matutes and Karl Schlag, two referees and the Editor, George Mailath. Errors remain our own. E.Hopkins@ed.ac.uk, <http://www.ed.ac.uk/~ehk>

†rms@math.ucl.ac.uk

1 Introduction

It is a common experience to find that prices vary between different sellers, giving consumers an incentive to search for low prices. Rothschild (1973) set economists a challenge. Sellers, presumably, would only charge prices different from those of their competitors if they could make a profit by doing so. To explain price dispersion, economists must show that such price-setting behaviour was a rational response by traders to the search behaviour of consumers, and vice-versa. Subsequently, several models have been developed in which price dispersion is indeed a Nash equilibrium. But in fact, Diamond (1971) had already gone one stage further by investigating what form of equilibrium pricing was dynamically stable under a plausible adaptive process.

Here we ask whether the models of price dispersion that followed from Rothschild's challenge can also pass this additional test. We examine a number of existing models which exhibit equilibrium price dispersion and show how they can be placed in simple, common framework. We extend current models of learning to deal with the large strategy sets present in these price setting games. We then investigate whether the equilibrium price distributions are stable, first, under learning by sellers alone, and then, under simultaneous learning by buyers and sellers.

Diamond pioneered this adaptive approach to price setting but his main result is usually viewed as a paradox. He was able to show that for any positive search costs, in equilibrium, no consumer would search, and all firms would charge prices at monopoly levels. This is clearly a Nash equilibrium: when prices are identical, there is no incentive to search; when there is no search, there is no incentive to cut prices to increase sales. Note that the converse state where all consumers are fully-informed and all firms charge a competitive price cannot be a Nash equilibrium. For positive search costs and with all prices identical, active search is not optimal. Since consumers are not fully informed, firms can raise prices without losing all customers. While those economists raised on the "Law of One Price" might have expected price dispersion to be fragile it was surprising that the collapse was in this direction.

Faced with this challenge, subsequent authors, (a partial list includes Salop and Stiglitz, 1977; Varian, 1980; Burdett and Judd, 1983; Rob, 1985; Wilde, 1992; Benabou, 1993), produced models with dispersed price equilibria. However, none of these results have passed the test that Diamond imposed on his model. The striking difference about the model of Diamond (1971) is that it is "A Model of Price Adjustment" not of equilibrium. The advantage to such a disequilibrium approach is that it can answer the question of how one equilibrium is chosen over another. This question is particularly relevant because in this type of model there are often multiple equilibria. In the models we examine, Varian (1980), Burdett and Judd (1983), two dispersed price equilibria coexist with the Diamond monopoly outcome. As Burdett and Judd themselves remarked, we need some way of reducing this multiplicity of equilibria.

Until this question is resolved, the existence of dispersed price equilibria is not a fully satisfactory explanation for the existence of price dispersion.

Furthermore, a disequilibrium approach allows for additional possibilities. The dispersed price equilibria we model are mixed strategy equilibria. Sellers are indifferent between charging the different prices in the support of the equilibrium distribution. This is troubling in that theorists from Shapley (1964) onwards have shown non-convergence of learning models in games with mixed strategy equilibria. The example often chosen to illustrate this is a “Rock-Scissors-Paper” (RSP) game, a game of three choices with a cyclical best-response structure, rock beats scissors which beats paper which in turn beats rock. There is a unique mixed strategy equilibrium. This equilibrium will be stable under a wide range of learning rules (Hofbauer, 2000; Hopkins 1999a, 1999b) if payoffs satisfy a negative definiteness condition, linked to evolutionary stability. However, it is just as easy to find examples where the unique equilibrium is unstable, and there is no convergence to Nash equilibrium at all.

Oligopolistic pricing games when not all consumers are fully informed about prices are like RSP games in the following way. The best response to a rival charging a given price p is often, like in Bertrand-style competition, to charge a price fractionally below p . In this way a seller can steal away relatively well-informed consumers who respond to prices. However, unlike in Bertrand-style competition, there will be some lowest price \underline{p} , greater than marginal cost, charging below which is not a best response. This is because of the presence of badly-informed consumers. When prices are sufficiently low, a seller may have an incentive to give up competing for informed buyers and rather raise her price and sell only to the uninformed. That is, just as in RSP games there is a cycle of best responses, a high price is “beaten” by a medium price which is beaten by a low price which is beaten by a high price. The major difference is that, in fact, there are not just three possible prices or strategies but an infinite number.

The contribution of this present paper is twofold. First, to examine the problem that sellers face we develop learning dynamics for use when the strategy set is a continuum and payoffs are nonlinear. Economists are used to treating price as a continuum. Retailers, however, are forced to price in whole pennies. Similarly, most often when dealing with learning with a continuum of strategies, researchers divide the strategy space into a discrete grid. However, mixed strategy equilibria which take the form of continuous distributions over prices are not easily represented using this method. Hence, we take a hybrid approach. We use a discrete approximation to the strategy set in order to examine dynamics in finite dimensions. But we determine the stability of this learning process by analysis of the sellers’ profit function treated as an operator on the function space L_2 .

Second, this approach enables us to obtain clear results on the stability of dispersed price equilibria. The condition for dynamic stability in these price-setting games is that payoffs should satisfy a negative definiteness condition similar to that for basic RSP games. Our results are to a certain extent positive. We find that agents

following simple learning rules may learn their way to equilibrium, even though equilibrium strategies require randomisation over an infinite number of prices. However, convergence to equilibrium is only possible for particular distributions of information amongst consumers. This tightens considerably the existing predictions concerning the nature of dispersed price equilibrium. The natural next step is to make consumer behaviour endogenous. That is, what happens when consumers learn too?

The condition for stability under seller learning is what we call “sufficient ignorance”, that is, the number of informed consumers must be sufficiently low. Therefore any equilibrium stable under seller learning must be close to the no-search Diamond outcome. This in turn means that when consumer learning is introduced, even a small perturbation from the dispersed price equilibrium will move the system into the basin of attraction of the Diamond outcome. Thus we can show that this is the only stable equilibrium under the joint dynamics. We discuss the implications of this result in our final section.

2 Pricing Games with the Rock-Scissors-Paper Property

In this section we will argue that a number of oligopoly games, where sellers compete to attract consumers who have imperfect information about prices, have a distinct similarity with the simple game of Rock-Scissors-Paper (RSP). Such a game is usually of the form,

	R	S	P
Rock	0, 0	$a, -b$	$-b, a$
Scissors	$-b, a$	0, 0	$a, -b$
Paper	$a, -b$	$-b, a$	0, 0

$a, b > 0.$ (1)

There is a cycle of best responses, Rock beats Scissors which beats Paper which beats Rock, and the only equilibrium is in mixed strategies (in the above version at $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$). We claim that making this analogy has two advantages. First, it brings home the actual problem that price-setters face in such markets, they must randomise over prices in order to be unpredictable. Second, it enables a simple, unified treatment of several different models. Third, it clarifies what happens when we introduce learning into these models.

One result from the recent literature is that the mixed equilibrium of the simple game above is only stable under adaptive learning schemes, such as fictitious play, if the benefit to winning is bigger than the cost to losing, that is if $a > b$. The intuition is that when the penalty to losing, $-b$, is particularly high, the prospect of a draw along the diagonal is relatively attractive. Agents are tempted into playing the same strategy as their opponent. Translating this into the price-setting games

we will consider, when price competition is particularly fierce, price dispersion is not sustainable. There is a similar pressure toward price conformity. As we will see, dispersed price equilibria will only be stable under learning, if on average consumers are not well informed and hence price competition is muted.

We will show that a number of models of price competition have what we will call the Rock-Scissors-Paper property because they have the same basic best reply structure as the simple game above. We are concerned with a market for a homogeneous good. For example, the same book or computer from a particular producer is often sold by many different outlets, often at different prices. The sellers we can think of as a continuum of identical small shops, which buy the good from a wholesaler for a constant unit cost, which here we assume to be zero. We assume (an average of) μ customers per seller. Sellers' utility is identical with their expected profits. Consumers seek to buy exactly one unit providing it is offered to them at less than or equal to their reservation price, $p^* > 0$. An action for a seller is then simply a choice of a price. A continuum of consumers are uninformed about which firms charge which prices. They must engage in costly non-systematic search in order to obtain price quotations.

We assume the current choice of all sellers can be summarised by a cumulative distribution function $F(p)$. The profit for any one seller is given by

$$\Pi(p) = pD(F(p)) \quad (2)$$

Clearly there are a number of differing assumptions that we can make about what information consumers possess and what choices they will make in response to that. This will feed through to different functional forms for the demand function $D(F(p))$. For example, if all consumers were completely informed about prices and always purchased from the cheapest supplier, we would have Bertrand-style competition and demand is discontinuous at the lowest price. However, if some consumers are not fully informed and therefore not sensitive to prices, demand may be continuous in $F(p)$. Second, if sellers cut prices to compete for informed customers, at a certain point margins become so low that it becomes more attractive to raise prices and sell only to the uninformed customers. This we formalise as:

Definition 1: We say that a price-setting game has the Rock-Scissors-Paper (RSP) property if :

1. Demand is continuously differentiable and strictly decreasing in $F(p)$.
2. If the highest price charged is $p \leq p^*$, then demand at that price is strictly positive, that is $D(1) > 0$, so that there exists some $\underline{p} \in (0, p^*)$ such that $\underline{p}D(0) = p^*D(1)$.

Condition 1 implies that if all sellers charge a price p , where $p^* \geq p > \underline{p}$, then any seller, just as under Bertrand style competition, can gain a discrete jump in sales

by the smallest possible price cut. However, Condition 2 says that even the highest priced seller has positive demand. This is what differentiates this case from strict Bertrand competition. It follows that there must be a lowest price beyond which price-cutting makes no sense. Even if a seller is the lowest priced ($F(p) = 0$) and hence she sells the most, if her price is less than \underline{p} she could do better by charging p^* . A game with the RSP property will have an equilibrium in mixed strategies which of course is a dispersed price equilibrium. We can characterise any such equilibrium a bit further.¹

Lemma 1 *Price setting games with the Rock-Scissors-Paper property possess a mixed equilibrium with a continuous density on the interval $[\underline{p}, p^*]$ where $\underline{p} > 0$ and this is the only symmetric equilibrium.*

Proof: For a symmetric equilibrium, we have to find, out of possible price distributions $F(p)$, an equilibrium price distribution denoted $\Phi(p)$ such that $pD(\Phi(p))$ is constant for all prices in its support. Suppose p^* is such a price. Then, by Condition 2 of the RSP property, we have the following relations,

$$p^*D(1) = \underline{p}D(0) = pD(\Phi(p)). \quad (3)$$

We must set $\Phi(p)$ such that $D(\Phi(p)) = p^*D(1)/p$ or $\Phi(p) = D^{-1}(p^*D(1)/p)$. Because demand is continuously differentiable and strictly decreasing in $F(p)$, we know that such an inverse function exists and is unique. Furthermore, the equilibrium distribution function $\Phi(p)$ is continuously differentiable and strictly increasing on $[\underline{p}, p^*]$. The only other possible equilibrium would be one without p^* in its support, but then a profitable deviation to p^* exists. \square

There now follow some examples of models that have the RSP property.

Model I. Informed/Uninformed consumers. This is based on the model of Varian (1980) where consumers have a stark choice between being either very well informed or completely uninformed. It is assumed that information is not freely available, but at a cost c consumers can obtain (almost) complete information about prices. One can imagine that this cost is to purchase some authoritative newspaper or to use an internet search engine. The essential point is that the exact amount of information obtained is outside the control of the consumer but is determined by an information provider, a situation that corresponds quite closely to many real markets. Here we assume that, for the cost c , one receives N price quotations, where N is large.² The consumer then purchases from the cheapest of the N sellers. If this information is not

¹Similar results for specific models can be found in Varian (1980), Burdett and Judd (1983) and Rob (1985). The unified framework is new.

²In Varian's original model, the number of sellers is finite and equal to N . So possession of N quotations means complete information. If N is large, the difference between the specifications is small.

purchased then consumers must choose a seller at random. Let q_1 be the proportion of these uninformed consumers, and $1 - q_1$ the proportion of consumers who are informed.

Thus, in any period, all sellers get μq_1 uninformed customers, as they are randomly distributed without regard to price. The expected number of informed customers for a seller charging a price p is $\mu(1 - q_1)N(1 - F(p))^{N-1}$. Demand for an individual seller at a price p is therefore

$$D(F(p)) = \mu[q_1 + N(1 - q_1)(1 - F(p))^{N-1}]. \quad (4)$$

Model IIa. Fixed sample size search. Again consumers are not fully informed about prices. Instead, the consumer must decide how many quotations to obtain at a constant cost c per quotation (the convention is that the first quotation is free). Only once all price quotations have arrived can the consumer purchase from the firm that offers the lowest price. Such nonsequential search can be optimal, and fits the case where a consumer must write away for quotations, or where a number of quotations can be obtained by buying a magazine or newspaper. This case has been analysed by Burdett and Judd (1983).

In particular, we use the notation that a proportion q_1 of consumers choose to obtain one price, q_2 have two price quotations and so on. The measure of consumers for whom a given price p is the lowest that they find with two quotations is $2(1 - F(p))$, after three $3(1 - F(p))^2$. Hence, demand for each seller is given by

$$D(F(p)) = \mu \sum_{k=1}^{\infty} q_k k (1 - F(p))^{k-1}. \quad (5)$$

Model IIb. Noisy sequential search. In the standard model of sequential search, a consumer looks at one price and then must decide whether to buy or to sample another price at cost c . It is well known that it is optimal in this case to adopt the strategy of buying if and only if offered a price at or below a reservation price p^* . If consumers share a common search cost c , and hence share a single reservation price p^* , no seller will have positive sales at any price $p > p^*$. However, suppose search is “noisy” in that at each search, a consumer has a possibility of seeing more than one price. In particular, the (exogenous) probability that k prices are observed is q_k . Then, as Burdett and Judd (1983) point out this case becomes almost identical to the previous one. Here the difference is that the distribution of the q_k is exogenous, and does not arise out of consumer choice. Nonetheless, the sales for each seller charging a price which is acceptable to consumers, that is, below p^* , are given by (5), just as in the non-sequential case.

If the proportion of uninformed buyers denoted q_1 , is neither one nor zero, it is easy to verify that all three models, I, IIa, IIb, will have the RSP property.³

³In giving the demand functions (4), (5), we have in effect assumed no mass points in the

Other Models. There are three other types of model that we do not analyse here but we note do have the RSP property. The first class of model is where consumers search sequentially with a non-degenerate distribution of search costs, and hence, reservation prices, as considered by, for example, Rob (1985). Consequently, sellers face a similar choice between targeting the more or less price sensitive consumers. The second class consists of models with switching costs. Suppose that in contrast to the above models, consumers are perfectly informed about prices. However, imagine that there is some difficulty or cost to switching between suppliers, either because of geographical distance, Shilony (1977), or because consumers become attached to a particular supplier in the first period of a multiperiod situation, Padilla (1995). Then, as before, some consumers will only go to the seller with the lowest price, while others are effectively not sensitive to prices because they are locked in. The third class, Bertrand-Edgeworth competition, is where individual sellers do not have the capacity to serve the whole market. Hence, even the highest priced seller may have positive sales (Condition 2 above) as lower priced sellers may run out of stock or capacity.

In both model I and II, it is possible to derive exactly the equilibrium distribution. In equilibrium, the profit from charging \underline{p} is equal to the profit from charging p^* which is equal to the profit from charging some arbitrary price p , $\underline{p} < p < p^*$. Given the relations (3), we can solve for both \underline{p} and $F(p)$ (see Section 6 below). We highlight here two important characteristics. First, for $0 < q_1 < 1$ the equilibrium is given by a density function that we denote ϕ , which is positive and continuous on (\underline{p}, p^*) . Second, as the level of captive customers q_1 approaches one, then \underline{p} approaches p^* and we have the Diamond monopoly outcome.

3 Positive Definite Adaptive Market Dynamics

Having described some possible equilibria, we now deal with disequilibrium. We imagine the above one-shot game is repeated many times. That is, at each point in time firms must choose prices, buyers must choose a level of search. As is common in the literature on learning, agents do not play some complex intertemporal equilibrium. Instead they adjust their play of the stage game. In this context, firms change prices in the direction of increasing profits. Positive Definite Adaptive (PDA) dynamics are a simple way to model the aggregate effects of this individual learning behaviour. We show (Proposition 2) that a negative definiteness condition related to the concept of an “evolutionarily stable strategy” (ESS) is a sufficient condition for stability under all PDA dynamics. We also establish a sufficient condition for instability (Proposition 3).

distribution of prices. If there were any mass points, we leave to the reader to confirm that, provided $0 < q_1 < 1$, first, demand would still be polynomial and hence continuous in $F(p)$, and, second, that $D(1) > 0$.

By modelling the sellers as a single infinite population we are doing as much as possible to make dispersed price outcomes possible. First, given there is a continuum of sellers, there is some freedom about whether agents actually randomise. Some agents could play mixed strategies, some could play pure. This could still support a continuous density of prices in aggregate. Second, this is a situation where firms are small relative to the size of the market and have little strategic power. Therefore the likelihood of a collusive outcome is relatively low.

If individual sellers do indeed behave in a way which is consistent with some model of learning, then we want to characterise the aggregation of that learning behaviour. Does the market as a whole approach a dispersed price equilibrium? Fortunately, we know that the aggregation of a number of different learning models described below all belong to the class of PDA dynamics. These were first introduced by Hofbauer and Sigmund (1990) and linked to human learning by Hopkins (1999a,b). PDA dynamics represent a class of different models and so it is difficult to give a simple characterisation. However, as we will see, they do embody in a general way the simple idea that better performing strategies become more popular. The reason we use PDA dynamics is that they arise quite naturally from the aggregation of several better known dynamics, which we now detail. Thus the results we obtain also hold for a wide variety of specific learning models.

Learning Model A. Fictitious Play. This is the learning model which has perhaps attracted the most interest in recent years (see for example, Fudenberg and Kreps, 1993; Young, 1993). Agents keep track of the play of their opponents and play a best response to past play. In this context, this assumption would mean that sellers would have to observe the prices of all other sellers. Fudenberg and Levine (1998) discuss extensions to the case where the actions of other players are not observed. In this case, fictitious play becomes similar to reinforcement learning, our next model.

Learning Model B. Reinforcement Learning. Here we take this to mean either what Erev and Roth (1998) call their “basic model”, which seems to fit quite well the behaviour of experimental subjects, or the model of Cross examined in Börgers and Sarin (1997). In both cases, the demands on information available to agents and on their rationality are much less than for fictitious play. The essence of these models is simply that an agent has a probability distribution over possible actions and if an action is taken and the resulting payoff is satisfactory the probability distribution is updated to place a higher probability on that action. Consequently, the only information needed is access to one’s own payoffs.

Learning Model C. Learning by Imitation. As an example of this approach, we can take the model of Schlag (1998). Here, it is imagined that each agent observes the action and payoff of one other randomly-drawn agent. Schlag finds that in this context the optimal behavioural rule is a proportional imitation rule which when aggregated over a population of agents yields the evolutionary replicator dynamics.

We first describe the general form of the PDA learning dynamics for finite games, before extending it to the more complicated continuum cases we shall need. Consider a population of agents who play a symmetric game with n possible actions. Each agent plays a (possibly mixed) strategy $y \in S_n$, where S_n is the standard n -simplex, $S_n = \{y = (y_1, \dots, y_n) \in \mathbb{R}^n : \sum_i y_i = 1, \text{ and } y_i \geq 0 \text{ for } 1 \leq i \leq n\}$. We suppose that the strategy mix in the population is described by a distribution function G on S_n . Then $x = \int y dG$ is the vector which determines the expected mixed strategy played by a randomly chosen agent. In particular, if all agents play pure strategies, then x is simply the vector whose components are the proportions of the population playing each strategy. We suppose that the expected return to each strategy, given an aggregate population state x , is determined by a (possibly nonlinear) profit function, $\Pi(x) = (\Pi_1(x), \dots, \Pi_n(x))$. That is, an agent playing a strategy $y \in S_n$ receives payoff $y \cdot \Pi(x)$. The possible non-linearity of Π will be important.

We have to describe how the population state x changes in time as a result of some learning process such as the learning models A, B and C discussed above.⁴ The PDA dynamics we consider have the general form

$$\dot{x} = Q(x)\Pi(x), \quad (6)$$

where $Q(x)$ is an $n \times n$ matrix which operates on the vector $\Pi(x)$. For example, the best-known dynamics of the form (6) are the replicator dynamics, with $Q(x)$ being in this case the matrix with $x_i(1 - x_i)$ in each diagonal position and $-x_i x_j$ on the off-diagonal.

In general, $Q(x)$ must satisfy some restrictions in order that (6) should represent a plausible learning dynamic. Recall that, for $x \in S_n$, the *support* of x is the set $T(x) = \{i : x_i > 0\}$, and for any subset $T \subseteq \{1, 2, \dots, n\}$, we define the following subspaces of \mathbb{R}^n , $\mathbb{R}_T^n = \{y \in \mathbb{R}^n : y_i = 0 \text{ for } i \notin T\}$, $\mathbb{R}_{T^c}^n = \{y \in \mathbb{R}^n : \sum_i y_i = 0\}$ and $\mathbb{R}_{T_1}^n = \{y \in \mathbb{R}^n : y_i = \text{constant for } i \in T\}$. When $T = \{1, 2, \dots, n\}$, we write \mathbb{R}_0^n for $\mathbb{R}_{T^c}^n$ and similarly for \mathbb{R}_1^n . Note that \mathbb{R}^n may be decomposed as an orthogonal direct sum, $\mathbb{R}^n = \mathbb{R}_{T^c}^n \oplus \mathbb{R}_{T_1}^n$. This allows us to specify the following restrictions on $Q(x)$.

Definition 2: A (differentiable) *PDA operator* on \mathbb{R}^n is an $n \times n$ - matrix operator $Q(y)$, defined and continuously differentiable for $y \in \mathbb{R}^n$, and satisfying the following properties for $x \in S_n$ with support $T = T(x)$:

1. $Q(x)$ maps \mathbb{R}^n into $\mathbb{R}_{T^c}^n$.
2. $Q(x)y = 0$ for $y \in \mathbb{R}_{T_1}^n$.

⁴These learning processes in fact unfold in discrete time, PDA dynamics in continuous time. However, results in the theory of stochastic approximation show the two are linked. See Fudenberg and Levine (1998, Ch4) for a survey of results on (stochastic) fictitious play and Hopkins (1999b) for the connection in this context with PDA dynamics.

3. $Q(x)$ is positive definite on $\mathbb{R}_{T_0}^n$; i.e. $y \cdot Q(x)y > 0$ for all $y \in \mathbb{R}_{T_0}^n$ with $y \neq 0$.
4. $Q(x)$ is symmetric.

The dynamics (6) is called a (differentiable) *PDA dynamics with payoff function* Π , if $\Pi(x)$ is defined and continuously differentiable for $x \in \mathbb{R}^n$, and $Q(x)$ is a differentiable PDA operator on \mathbb{R}^n .

The long list of conditions should not hide the generality of the dynamics specified. The substantive conditions placed on Q are positive definiteness and symmetry. Positive definiteness ensures that the vector of changes in strategy frequencies \dot{x} is at less than a 90° angle to the vector of payoffs $\Pi(x)$. This is thus a very weak formulation of the assumption that strategies with a high payoff grow at the expense of those with a lower return. Symmetry is naturally present in the learning models considered here (see Hopkins, 1999a). It simply implies that the different strategies receive equal treatment. More specifically, if an increase of some amount Δ_1 in the payoff to a strategy i leads to a reduction of the growth rate of strategy j of Δ_2 , then symmetry requires that an increase of Δ_1 in the payoff to j would lead to a decrease of Δ_2 in the growth rate of i .⁵

Property 1 is principally there so that x will continue to sum to one. However, it also implies that no strategy present in the initial distribution will disappear in finite time, nor will any new strategy be created. Thus we will want to look at cases where all prices are present in the initial distribution. This may seem somewhat restrictive, but it should be remembered that any distribution, including those where all firms charge the same price, can be approximated arbitrarily closely by a distribution with full support. Second, this formulation does not prevent the limit of the dynamic process being a state like the no-search outcome, where all sellers charge the same price. Property 2 means that a mixed strategy equilibrium, that is where all strategies have the same return, is an equilibrium for the dynamic.

However, the fact that a Nash equilibrium is an equilibrium for some learning dynamic does not mean that it will be asymptotically stable. As a consequence we introduce the idea of *evolutionarily stable strategy* (ESS), a refinement on Nash equilibrium. This concept originates in biology where it is defined as a strategy profile that is “uninvadable”. Agents playing some alternative strategy would not be able to supplant agents playing the established equilibrium strategy. However, despite its biological origins, it turns out to be a sufficient condition for stability under many learning processes. It differs from Nash equilibrium in that it considers deviations from equilibrium not just by individual agents but also by small but positive measures of the population.

⁵Positive definiteness without symmetry would generate dynamics which Friedman (1991) calls “weak compatible”. Friedman shows that ESS’s may be unstable under such dynamics and therefore Proposition 2 below does not hold without symmetry.

The conditions for a state ϕ to be an ESS are first that ϕ should be a Nash equilibrium, and second that if there is an alternative best reply to ϕ , then ϕ should do better against such an alternative than that alternative does against itself. In terms of our (possibly nonlinear) payoff function Π , these conditions may be expressed formally as:

$$\phi \cdot \Pi(\phi) \geq x \cdot \Pi(\phi) \quad \text{for all } x \in S_n, \quad (7)$$

and for all $x \in S_n$ for which equality holds in (7)

$$\phi \cdot \Pi(x) > x \cdot \Pi(x) \quad (8)$$

What this last condition implies is a kind of concavity of the payoff function. For example, if Π is linear, such as the payoff matrix for a normal form game, then (8) implies that $(x - \phi) \cdot \Pi(x - \phi) < 0$ for any alternative best response x (because $(x - \phi) \cdot \Pi(\phi) = 0$). Now, as x and ϕ are both vectors whose components sum to unity, $z = x - \phi$ is an element of \mathbb{R}_0^n . It therefore follows from these observations that Π must be negative definite on $\mathbb{R}_{T_0}^n$ where T is the support of ϕ .

When Π is not linear, we shall, instead of (8), require the related condition that the linear approximation at the equilibrium point, given by the derivative $\Pi'_\phi = \frac{d\Pi}{dx}(\phi)$, be negative definite on $\mathbb{R}_{T_0}^n$. That is,

$$z \cdot \Pi'_\phi z < 0 \quad \text{for all } z \in \mathbb{R}_{T_0}^n \text{ with } z \neq 0. \quad (9)$$

As we have seen, when Π is linear, (8) implies (9), but only a restricted converse holds: (9) implies (8) only for alternative best replies having support subordinate to T . In general, for nonlinear payoff functions, neither of these implications is valid. However, it can be shown that the second implication is valid *locally*. That is, condition (9) implies that there is a neighbourhood of ϕ in S_n , such that (8) holds for any alternative best reply in this neighbourhood with support subordinate to T . In any case, (9) is a weaker condition than (8). We therefore introduce, by way of compensation, the following refinement of Nash equilibrium, specifically with non-linear payoff functions in mind.

Definition 3: A point $\phi \in S_n$ is a *regular ESS* with respect to the differentiable payoff function Π , if conditions (7, 9) are satisfied, and also

$$\Pi_i(\phi) < \pi^* \quad \text{for } i \notin T, \quad (10)$$

where T is the support of ϕ , and $\pi^* = \phi \cdot \Pi(\phi)$ is the equilibrium payoff.

Condition (10) means that any strategy not in the support of ϕ receives a strictly lower return than those which are. Thus, every alternative best reply to ϕ has support subordinate to T , which, as observed above, means that (9) and (10) imply (8), at least locally. In the special case in which ϕ is a pure strategy, T contains only one point and $\mathbb{R}_{T_0}^n = \{0\}$. Thus, condition (9) does not apply, and a pure-strategy regular ESS is just a strict Nash equilibrium.

Clearly a regular ESS is a Nash equilibrium, and thus is a stationary point of any PDA dynamics. Furthermore, the following fundamental result holds. The proof of this result (and all subsequent results) is in the Appendix but the intuition is that PDA dynamics shift strategies in the direction of better responses. Thus, strategies not in the support of ϕ will tend to decrease in frequency because of (10), and will eventually disappear. On the other hand, strategies in the support of ϕ constitute an (local) ESS, and so ϕ is resistant to small perturbations from equilibrium.

Proposition 2 *A regular ESS for a (possibly nonlinear) differentiable payoff function Π is a locally asymptotically stable stationary point for any differentiable PDA dynamics (6) with payoff function Π .*

We shall also be interested in conditions under which a Nash equilibrium ϕ is *unstable* with respect to a general PDA dynamics. If, for example, the derivative Π'_ϕ is positive definite on $\mathbb{R}_{T_0}^n$, then the mixed equilibrium will be unstable for all PDA dynamics. But it is also possible to show that if Π'_ϕ is neither positive nor negative definite, then the equilibrium will be unstable for some PDA dynamic.

Proposition 3 *Let $\phi \in S_n$ be a Nash equilibrium with support T for the payoff function Π . If Π'_ϕ is positive definite on $\mathbb{R}_{T_0}^n$, then ϕ is unstable under all PDA dynamics. Suppose Π'_ϕ is neither positive nor negative definite then there is a non-empty class of PDA operators, such that ϕ is unstable under the PDA dynamics (6) defined by any Q in this class.*

In fact, the class of PDA operators referred to in Proposition 3 is very large, even when there is only a 1-dimensional subspace of $\mathbb{R}_{T_0}^n$ on which Π'_ϕ is positive definite. Of course, if T contains only two elements, then $\mathbb{R}_{T_0}^n$ is 1-dimensional, so the hypothesis of Proposition 3 implies that ϕ is unstable for any PDA dynamics. In the even simpler case in which ϕ is a pure strategy, then $\mathbb{R}_{T_0}^n = \{0\}$, and so Proposition 3 cannot apply. However, our main interest is in dispersed price equilibria, which are mixed strategy equilibria. These we now investigate.

4 Deriving a Stability Condition

As Burdett and Judd themselves suggested

“Examples of further possible work include stability analysis which may give further information concerning the durability of equilibrium price dispersion and reduce the multiplicity of equilibria in the nonsequential model.” (1983, p967)

In this section, we do indeed carry out a stability analysis, for the fixed sample size model of Burdett and Judd, their model of noisy sequential search, Varian's model of sales, and indeed any market model which has the RSP property. For now we consider only the behaviour of sellers, treating consumer behaviour as fixed. More specifically, we assume that for both Model I and II, the distribution of consumer behaviour $\{q_k\}$ is constant and $0 < q_1 < 1$, so that a dispersed price equilibrium exists for sellers. We go on to treat the problem of consumer dynamics in Section 6.

In the previous section, we introduced the idea of modelling learning dynamics as a positive definite transformation of payoffs. However, to deal with Models I and II within the conceptual framework of the learning theory introduced in the previous section, we must confront the fact that these models are formulated for a continuum of prices, while our learning models were described only for finite strategy sets. We could proceed either by approximating the market game models of Section 2 by finite-dimensional games over a finite set of prices, or by introducing an infinite-dimensional version of the learning theory of Section 3.⁶ We take a hybrid approach which requires some words of justification.

The most common approach to learning with a continuum of strategies has been to take a discrete approximation to the strategy space. This is undoubtedly more simple, and is to be recommended if one is concerned with equilibria that are pure. A pure equilibrium in a game of continuum of strategies will typically remain an equilibrium in a game which is a discrete approximation. However, the models of price dispersion we are concerned with have mixed strategy equilibria which consist of continuous distributions. Of course, if only a finite number of prices were allowed, it would be possible to construct a discrete distribution which approximated the original equilibrium. But it is not guaranteed that this approximation would be an equilibrium of the discrete game with a finite number of strategies.⁷ Yet, the other possible approach, truly infinite dimensional dynamics, offers a host of complications. We would argue that our framework allows us the "best of both worlds".

Our dynamics are finite dimensional but the use of some aspects of the Hilbert space L_2 allows us to determine stability of equilibrium by the evaluation of a single integral rather than the daunting prospect of determining the eigenvalues of an arbitrarily large matrix. We have shown in Proposition 2 that with a finite number of strategies a mixed strategy equilibrium will be locally asymptotically stable if the linearisation of the profit function is negative definite. That is, the sign of the quadratic form $z \cdot \Pi'_\phi z = \sum z_i (\Pi'_\phi z)_i$ determines stability. Similarly here, if we construct an alternative price distribution $f(p)$ close to the equilibrium distribution $\phi(p)$ then over time the distribution of prices will approach $\phi(p)$ if, letting $z = f - \phi$, the

⁶Examples of evolutionary dynamics with continuous strategy spaces are to be found in Friedman and Yellin, 1997; Oechssler and Reidel, 2001; and Seymour, 2000.

⁷For a fine enough grid and for an $\epsilon > 0$, no price in the support of the discrete equilibrium would earn a profit more than ϵ greater than any other. But a mixed equilibrium demands equality.

integral $\int z(p)\Pi'_\phi z(p) dp$ is negative. Finally, this in turn means that there is a relatively simple condition on the demand function, given in Proposition 4 below, which determines whether this integral will be positive or negative, and hence whether the mixed equilibrium will be unstable or stable. The interpretation of this condition is discussed in Section 5.

In effect, what we assume is that buyers, faced with a price distribution with continuous support, round up any prices in fractions of pennies mentally to the nearest penny. That is, in so far as customers can ‘see’ such small price differences, they can only respond rather crudely in their buying behaviour, in terms of prices set in finite units. Because of this, sellers have no incentive to make price changes below this unit threshold. Thus, assuming a large, but finite number of equally spaced prices, and from here on that p^* is normalised to 1, $0 < p_1 < p_2 < \dots < p_n = p^* = 1$, we define the *finite-dimensional approximation to the profit operator* Π by:

$$\Pi^n(f)(p) = \Pi_i^n(f) = n \int_{p_{i-1}}^{p_i} rD(F(r)) dr, \quad \text{for } p_{i-1} < p \leq p_i \quad (1 \leq i \leq n). \quad (11)$$

That is, the profit to a seller charging p in a market characterised by the density f and associated distribution function F , is the *average* over the price interval $p_{i-1} < p \leq p_i$ of the profits obtained from all possible prices in the interval (given the differentiable demand function D).⁸ This means that when we consider the stability of a dispersed price equilibrium we only have to deal with deviations from that equilibrium that involve changes above the unit threshold. Hence the dynamics will be finite dimensional.

Nonetheless, as previously mentioned, we do make some use of the infinite dimensional function space $E = L_2[0, 1]$. This is the linear space of those real-valued functions $f(p)$ which are square integrable on $[0, 1]$; *i.e.* $\int_0^1 f^2(p)dp < \infty$. Then E is a Hilbert space with inner-product $\langle f, g \rangle = \int_0^1 f(p)g(p)dp$. The associated norm is $\|f\| = \langle f, f \rangle^{\frac{1}{2}}$. Of course, we will be particularly interested in the subset of probability density functions, $S_E = \{f \in E : f \geq 0, \text{ and } \int_0^1 f(p)dp = 1\}$. Thus, functions $f \in S_E$ define the allowable price distributions. Equation (2) then defines a profit function, $\Pi(f) \in E$, given by $\Pi(f)(p) = pD(F(p))$. As outlined in Section 2, a dispersed price equilibrium in both Model I and Model II takes the form of an equilibrium probability density function ϕ which is non zero on an interval $T = [\underline{p}, 1]$ called the support of ϕ as defined in Lemma 1. Finally, let $\Phi(p) = \int_0^p \phi(r)dr$.

We now consider a PDA dynamics for the profit function (11), defined for any $f \in E$, but with a fixed number n of prices:

$$\dot{\hat{f}} = Q(\hat{f})\Pi^n(f), \quad f \in E. \quad (12)$$

where $\hat{f} = (\hat{f}_1, \dots, \hat{f}_n)$ is the associated finite probability distribution over prices, $\hat{f}_i = \int_{p_{i-1}}^{p_i} f(r)dr$, and $F_i = \sum_{j \leq i} \hat{f}_j = F(p_i)$. Thus, $Q(\hat{f})$ is a finite-dimensional,

⁸Note that, because prices are equally spaced, $p_i - p_{i-1} = \frac{1}{n}$, so that $1/(p_i - p_{i-1}) = n$, which explains the factor n in front of the integral in (11).

differentiable PDA operator, as in Definition 2. In particular, any $x \in \mathbb{R}^n$ defines an element of E by: $x(0) = 0$, and $x(p) = nx_i$ for $p_{i-1} < p \leq p_i$. Then the associated cumulative function is $X(p) = \int_0^p x(r)dr = X(p_{i-1}) + x_in(p - p_{i-1})$ for $p_{i-1} < p \leq p_i$. Thus, $X(p_i) = X_i = \sum_{j \leq i} x_j$ gives the cumulative distribution associated to the vector x , and $\hat{x}_i = \int_{p_{i-1}}^{p_i} x(p)dp = x_i$ recovers the vector x from its functional form. It follows that (12) defines a PDA dynamics on \mathbb{R}^n ,

$$\dot{x} = Q(x)\Pi^n(x), \quad x \in \mathbb{R}^n, \quad (13)$$

with profit function,

$$\Pi_i^n(x) = n \int_{p_{i-1}}^{p_i} pD(X(p)) dp. \quad (14)$$

Nevertheless, it will be convenient to retain the infinite dimensional formulation (12), principally because it is technically easier to work with the continuous density $\phi \in S_E$, rather than just the associated equilibrium distribution $\hat{\phi} \in S_n$. This technicality centers around the fact that the cumulative distribution $\Phi(p) = \int_0^p \phi(r)dr$ is not equal to the cumulative function $\hat{\Phi}(p)$, derived from the finite probability distribution $\hat{\phi} = (\hat{\phi}_1, \dots, \hat{\phi}_n)$, with $\hat{\phi}_i = \int_{p_{i-1}}^{p_i} \phi(r)dr$. Of course, the equilibrium condition (3), that all prices in the support of the equilibrium earn the same profit, holds for the continuum of prices $p \in T$. But it does not hold with $\Phi(p)$ replaced by $\hat{\Phi}(p)$. Because of this, profits generated by the discrete approximation $\Pi^n(\hat{\phi})$ will not be equal to $\Pi^n(\phi)$. Consequently, the discrete distribution $\hat{\phi}$ need not be even a stationary point of the dynamics (13). For this reason, our strategy will be to start with the infinite dimensional space E and project to a finite dimensional \mathbb{R}^n , as in (11), (12), rather than to work with dynamics that are explicitly finite dimensional, as in (13), (14).

It is the stability of the dispersed-price equilibrium $\phi \in S_E$, regarded as a stationary point of the dynamics (12), which we investigate in the remainder of this section. More precisely, for $z \in \mathbb{R}_0^n = \{y \in \mathbb{R}^n : \sum_i y_i = 0\}$, we may define $f = \phi + z \in E$, where, as above, the vector z is identified with the function, $z(p) = nz_i$ for $p_{i-1} < p \leq p_i$. Then $\int_0^1 f(p)dp = 1$, and hence there is a subset $S_n(\phi) \subset \mathbb{R}_0^n$ consisting of those z for which f is a probability density function; *i.e.* for which $\phi + z \geq 0$. We now use the dynamics (12) to define a PDA dynamics on \mathbb{R}_0^n , relativized at ϕ , by

$$\dot{z} = \gamma^n(z) = Q(\hat{\phi} + z)\Pi^n(\phi + z), \quad z \in \mathbb{R}_0^n, \quad (15)$$

under which $S_n(\phi)$ is invariant, and which has a stationary point at $z = 0$. We shall find conditions under which $z = 0$ is locally asymptotically stable, and conditions under which it is unstable, with respect to the dynamics (15) restricted to $S_n(\phi)$.

Our main result of this section can now be stated. Recall that $T = [\underline{p}, 1]$ is the support of ϕ , with $\underline{p} = p_{\underline{n}}$. We also use T to denote the associated finite set $\{\underline{n}, \underline{n} + 1, \dots, n\}$.

Proposition 4 *Suppose the demand function $D(F)$ is twice differentiable (almost everywhere) for $F \in [0, 1]$, and define*

$$\Theta(F) = D''(F)D(F) - D'(F)^2. \quad (16)$$

Then:

(a) *If $\Theta(F) > 0$ for (almost) all $F \in [0, 1]$, the derivative of the profit function (11), $\Pi_\phi^{n'}$, is negative definite on $\mathbb{R}_{T_0}^n$, and the dispersed price equilibrium at $z = 0$ is a regular ESS, and hence is a locally asymptotically stable equilibrium of the dynamics (15).*

(b) *If $\Theta(F) < 0$ for (almost) all $F \in [0, 1]$, the derivative of the profit function (11), $\Pi_\phi^{n'}$, is positive definite on $\mathbb{R}_{T_0}^n$, and hence the dispersed price equilibrium at $z = 0$ is unstable under the dynamics (15).*

5 Stability under Sufficient Ignorance

In this section, we use the condition we have just derived to examine whether there are any dispersed price equilibria of the models we consider which are in fact asymptotically stable. We find that the set of stable equilibria, while not empty, is rather smaller than the set of unstable equilibria. Furthermore, it is characterised by a condition which we could call “sufficient ignorance”. That is, one has stability only when the average amount of information held by consumers is small.⁹

To see this, remember the condition for stability derived in Proposition 4 is that the function $\Theta(F) = D''(F)D(F) - D'(F)^2$ must be positive. Second, remember from our discussion in Section 2, that results from RSP games suggested that stability would be inversely related to the severity of price competition. Obviously $D'(F)$, the slope of the demand curve, represents one possible measure of price competition. Furthermore, one can see that the larger it is in absolute value the less likely it is the stability condition will be met. We will see below that in both Models I and II, the absolute value of $D'(F)$ is increasing in the proportion of informed consumers. Ignorance certainly promotes stability, if not bliss.

There is another effect we can identify. The condition for stability (16) will more likely be satisfied if $D''(F)$ is large and positive. This would mean that the demand curve in terms of F would be relatively steep around \underline{p} and relatively flat near p^* . This initial steepness implies that the premium in sales for the very lowest price over those near the lowest is particularly large. This is roughly equivalent to the payoff from winning being large in RSP games. The later flatness could be compared to

⁹Though clearly at least some consumers need to be informed or else we would have the Diamond outcome.

the penalty for failing to be the cheapest, “losing”, not being large. Both effects, as noted in Section 2, would tend to lead to stability. We now examine when exactly our models of price dispersion satisfy the criterion (16).

Model I From the definition of Model I in Section 2 and in particular the demand function (4), we obtain

$$\begin{aligned} D(F) &= q_1 + N(1 - q_1)(1 - F)^{N-1} \\ D'(F) &= -(1 - q_1)N(N - 1)(1 - F)^{N-2} \\ D''(F) &= (1 - q_1)N(N - 1)(N - 2)(1 - F)^{N-3}. \end{aligned}$$

Immediately, we can see that $D''(F) = 0$ for the case $N = 2$, so that Π' must be positive definite. However, we are principally interested in the case when N is large. When $N > 2$, the condition for the linearisation of the profit function to be negative definite boils down to

$$(1 - q_1)N(N - 1)(1 - F)^{N-3}[q_1(N - 2) - (1 - q_1)N(1 - F)^{N-1}] > 0$$

for all $F \in [0, 1]$. However, this holds if and only if it holds for $F = 0$. Then we have,

$$q_1 > \frac{N}{2(N - 1)}$$

That is, for the dispersed price equilibrium to be stable the proportion of uniformed has to be sufficiently large. As N goes to infinity, the minimum proportion falls to one half. One can also calculate that if the above inequality is not met then Π' is neither positive nor negative definite.

Model II

We classify distributions, $q = \{q_k\}$, for which $\Theta(F)$ is positive or negative. To keep things as simple as possible, we shall only consider the case in which $q_k = 0$ for $k > 3$. Then given the original demand function (5)

$$\begin{aligned} D(F) &= q_1 + 2q_2(1 - F) + 3q_3(1 - F)^2, \\ D'(F) &= -2q_2 - 6q_3(1 - F), \\ D''(F) &= 6q_3, \end{aligned}$$

and the condition for stability that $\Theta(F) > 0$ reduces to

$$2q_2^2 - 3q_1q_3 + 6q_2q_3(1 - F) + 9q_3^2(1 - F)^2 < 0 \quad \text{for all } F \in [0, 1].$$

Again this holds if and only if it holds for $F = 0$; *ie* if and only if

$$2q_2^2 - 3q_1q_3 + 6q_2q_3 + 9q_3^2 < 0.$$

Substituting $q_3 = 1 - q_1 - q_2$, this condition reduces to

$$H(q_1, q_2) = 12q_1^2 + 15q_1q_2 + 5q_2^2 - 21q_1 - 12q_2 + 9 < 0. \quad (17)$$

We now note that $H(q_1, q_2) = 0$ is the equation of an ellipse in the (q_1, q_2) -plane (with centre at $(2, -\frac{9}{5})$), and that (17) holds for points in the interior of this ellipse. The points of interest are therefore in the intersection of the ellipse with the projection of the (q_1, q_2, q_3) -simplex onto the (q_1, q_2) -plane. It is this intersection which is labelled “Stable” in Fig. 1. Clearly it is non-empty. Also observe that the values of q_1 in the stable region all satisfy $q_1 > \frac{2}{5}(5 - \sqrt{10}) \approx 0.74$. That is, q_1 must be fairly large. However, stability also requires the presence of some well informed consumers, that is, q_3 must be greater than zero.

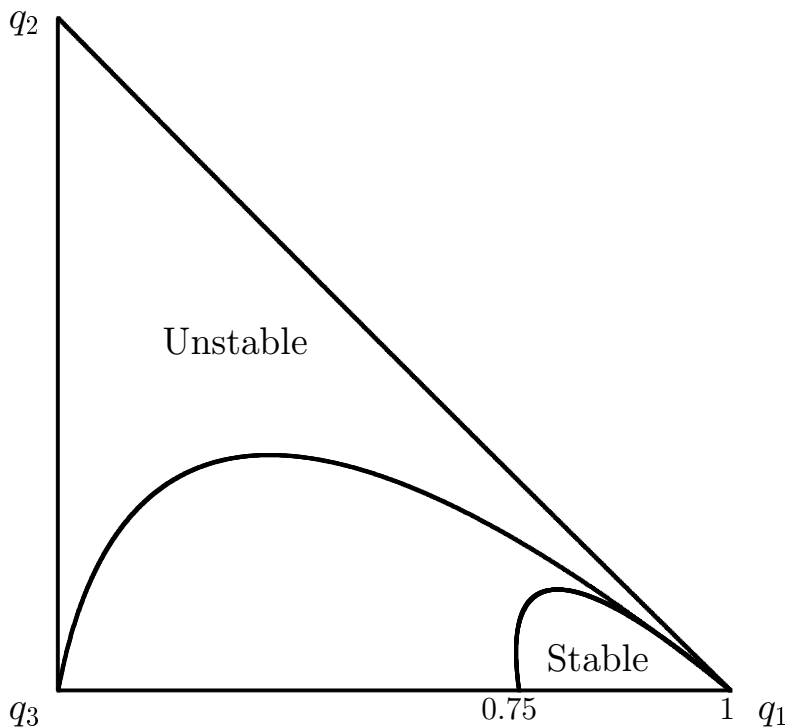


Figure 1: Distributions for which equilibria are negative definite.

The points in the interior of the (q_1, q_2) -simplex which lie outside the ellipse $H(q_1, q_2) = 0$, given in (17), all give rise to unstable dispersed price equilibria. Those equilibria which are strictly positive definite satisfy (17) with the inequality reversed. A similar argument to the one used above shows that such points lie outside the ellipse in the (q_1, q_2) -plane, $3q_1^2 + 3q_1q_2 + 2q_2^2 - 3q_1 = 0$. This ellipse is the curve which separates the purely unstable region (see Fig. 1) from the rest.¹⁰

If one looks at the results on Models I and II, one thing is apparent. The conditions for stability effectively demand a sufficient level of ignorance. In Model I

¹⁰In the intermediate area, the linearisation of the profit function Π' would be neither negative nor positive definite, and any equilibrium would be unstable under some PDA dynamic by Proposition 3, a result which is discussed further at the end of the next section.

the proportion of the uninformed has to be relatively high, in Model II, the average number of prices known has to be close to 1. This in turn means, as \underline{p} is endogenous, that for stable equilibria prices will be relatively concentrated near \underline{p}^* .

6 Consumer Dynamics

We now examine what happens if we remove the assumption that consumer behaviour is fixed and assume instead that it follows a learning process similar to that we have analysed for sellers. One might wonder how the results of the previous two sections might change. Will consumer learning stabilise equilibria which are unstable under seller dynamics alone? Is there a possibility of joint stability? As we shall see, the answers to both questions are negative.¹¹ We first characterise the dynamics for consumers. We then show that the Diamond no-search outcome is locally stable under the joint dynamics. We go on to characterise the possible dispersed price equilibria and show that they cannot be robustly stable under the joint buyer-seller dynamic.

Let $f(p)$ be the price distribution for sellers, and $q = (q_1, q_2, \dots, q_N)$ be the proportion of consumers having 1, 2, ..., N price quotations respectively. In Model I, all elements of q are zero except q_1 , the proportion uninformed, and q_N , the proportion informed. In Model IIa, N is an integer large enough so that $q_k = 0$ for $k > N$. Let $c > 0$ be the fixed cost per quotation. We do not consider Model IIb in this section because in that model the distribution of consumer information $\{q_k\}$ is exogenous.

We will analyse general consumer dynamics of the form,

$$\dot{q} = R(q)K(f), \quad (f, q) \in E \times \mathbb{R}^N, \quad (18)$$

where R is the PDA operator for consumers. Here, $R(q)$ is an $N \times N$ matrix, and $K(f) \in \mathbb{R}^N$ gives the consumers' payoffs when sellers are using the price distribution f . More precisely, the expected payoff to a consumer in possession of k quotations (the convention is that the first quotation is free) is $-(k-1)c - \int_0^1 f(p)kp[1-F(p)]^{k-1} dp \equiv K_k(f)$. These payoffs and hence the dynamics (18) are defined in terms of a continuous price distribution f . However, as explained in Section 4, we can restrict these dynamics to a finite number of dimensions in the neighbourhood of an equilibrium.

As discussed in Section 2, the Diamond equilibrium for sellers is the mass point distribution at which all firms charge the maximum price $p = 1$. As it is a strict Nash equilibrium, one would expect it to be locally stable under learning. However, to show this certain technical difficulties have to be overcome. Whereas up to now we have been working with continuous density functions, the Diamond equilibrium is the Dirac delta-function at $p = 1$, defined to be the continuous linear operator on (bounded,

¹¹This does not follow from the standard results on the instability of mixed strategy equilibria in asymmetric games, which apply only to games with linear payoffs.

measurable) functions given by $\delta_1 f = f(1)$. The full Diamond equilibrium for sellers and consumers is the pair (δ_1, e_1) , where $e_1 = (1, 0, \dots, 0) \in \mathbb{R}^N$ is the sample distribution for consumers. That is, all consumers sample only once (are uninformed in Model I). We can show that this equilibrium is indeed locally asymptotically stable.

Proposition 5 *The subspace $f = \delta_1$ is invariant and locally attracting under the joint seller-consumer dynamics. Within this subspace, the Diamond equilibrium $q = e_1$, is globally asymptotically attracting for all initial conditions other than a finite set of unstable isolated equilibria of the consumer dynamics.*

Having considered the Diamond equilibrium, we now look at the possible stability of dispersed price equilibria under the joint seller-consumer dynamics. We can show that there exist a distribution of prices and a value of q_1 such that for both Model I and Model IIa buyers and sellers are simultaneously in equilibrium. In Model I buyers have to choose whether to be informed at a cost c before they purchase, and in Model IIa they must choose how many prices to sample each at a cost c . Thus, while in Model I buyers have effectively two strategies to choose between, in Model IIa, they potentially have many. However, Burdett and Judd (1983) show that in Model IIa a joint dispersed price Nash equilibrium is only possible when consumers sample at most two prices.

This follows from the fact that an equilibrium q^* for consumers must give an equal return for the different levels of search in the support of q^* : the benefit of each additional search must equal its cost c . However, there are decreasing returns to search. That is, the reduction in expected price from one further search is decreasing in the number of searches, or in other words, the expected price paid is a convex decreasing function of k . Hence, if the benefit of a second search over the first is c , the benefit of the third is less than c , and consumers cannot be indifferent between one, two and three searches. Lastly, a dispersed price equilibrium for sellers is only possible if $0 < q_1 < 1$, so a joint equilibrium is only possible where $q_1 + q_2 = 1$.

In view of these remarks, we restrict attention to the case in which $q = (q_1, 1 - q_1)$, with $0 \leq q_1 \leq 1$. In this case, the properties of PDA operators imply that the consumer dynamics (18) are equivalent to a 1-dimensional dynamics of the form

$$\dot{q}_1 = \rho(q_1) (c - V(F)), \quad (19)$$

where $c > 0$ is a constant, giving the cost of being informed in Model I and the cost of moving from one to two quotations in Model IIa. The function $\rho(q_1)$ is a scalar-valued PDA operator, positive for $0 < q_1 < 1$, with $\rho(0) = \rho(1) = 0$. The function $V(F)$ is defined as

$$\begin{aligned} V(F) &= \int_0^1 p f(p) dp - \int_0^1 p f(p) N [1 - F(p)]^{N-1} dp \\ &= \int_0^1 [1 - F(p)] dp - \int_0^1 [1 - F(p)]^N dp, \end{aligned} \quad (20)$$




Figure 2: Equilibria for consumers

(the second equality follows on integration by parts). The function $V(F)$ gives the expected difference in price paid in Model I between an uninformed and an informed buyer. In Model IIa for $N = 2$, it gives the expected difference in price paid by a buyer searching once and a buyer searching twice. A mixed equilibrium for buyers, therefore, is a distribution of prices F such that $V(F) = c$. That is, given sellers' behaviour, buyers are indifferent between being informed and uninformed.

The task is to find a joint equilibrium. From (3), for fixed $0 < q_1 < 1$, there is a unique equilibrium for sellers given by:

$$\Phi(p) = \begin{cases} 1 - \left(\frac{q_1(1-p)}{N(1-q_1)p} \right)^{\frac{1}{N-1}} & \text{for } \underline{p} \leq p \leq 1, \\ 0 & \text{for } 0 \leq p < \underline{p}, \end{cases} \quad (21)$$

where

$$\underline{p} = \frac{q_1}{q_1 + N(1 - q_1)}. \quad (22)$$

Substituting this into (20), we obtain $V(\Phi)$. Note that if we can find a value of q_1 such that $V(\Phi) = c$, we have a joint equilibrium between buyers and sellers. In the following proposition we show $V(\Phi)$ to be concave. Consequently, generically, for a

given equilibrium level of prices, and provided that c is not too high, two such values of q_1 will exist. These we label $(\underline{q}_1, \bar{q}_1)$. This result is illustrated in Figure 2. There is another equilibrium, the Diamond outcome with $q_1 = 1$ and all sellers charging $p^* = 1$.

Proposition 6 *In Models I and IIa, for $c > 0$ and fixed $N \geq 2$, there is at least 1 and at most 3 values of $q_1 > 0$ such that buyers and sellers are simultaneously at equilibrium. One of these is the Diamond equilibrium $q_1 = 1$.*

We now address the stability of these equilibria. Given the results of the last section, we know that for Model IIa the linearisation of the sellers' profits at both equilibria will be positive definite (because for these equilibria to exist $q_k = 0, k > 2$). We confirm that this dooms stability under the joint dynamic. However, for Model I the prospects look more hopeful. The *a priori* candidate for stability would be the mixed equilibrium with the higher proportion of uninformed, that is (ϕ, \bar{q}_1) , which potentially would satisfy our condition of sufficient ignorance. However, we can show that this equilibrium too is unstable. The essence of the argument can be seen in Figure 2. To the right of \bar{q}_1 , that is, when there are slightly more uninformed consumers than in equilibrium, $c > V(q_1)$. Or in other words the return to search is less than its cost. So the number of uninformed consumers will tend to increase in the direction of the no-search Diamond outcome.

Proposition 7 *Consider Models I and IIa. For $c > 0, N > 2$ and n sufficiently large, a dispersed price equilibrium is unstable for the joint seller-consumer dynamics defined by some non-empty class of pairs of PDA operators (Q, R) . For $N = 2$, a dispersed price equilibrium is unstable for all PDA seller-consumer dynamics.*

For $N = 2$, the result is conclusive: any dispersed price equilibrium, including for example, the equilibrium constructed by Burdett and Judd (1983) (here Model IIa), is unstable under any PDA learning dynamic. For $N > 2$, there is instability of dispersed price equilibria under some PDA dynamics but not necessarily all. However, we would argue that it is very difficult to be confident about exactly how people learn and so stability results that only hold for one type of learning must be suspect. This is because stability and instability are asymmetric. In these games with large numbers of players, it would only need a positive measure of agents to adopt a learning rule for which the equilibrium was unstable to destabilise the equilibrium as a whole. In contrast, for an equilibrium which is not robustly stable to be an attractor it would require all agents to adopt an appropriate rule.

7 Conclusion

We have shown that the market pricing models that we consider are in many ways like Rock-Scissors-Paper games, both in terms of their equilibria and the learning behaviour they induce out of equilibrium. We have developed techniques to determine the stability of dispersed price equilibria under learning. We have shown that most such equilibria are in fact unstable. The ones that survive are where average consumer information is low and where prices are close to the monopoly level. Finally, once consumer learning is introduced even these are unstable.

Once again we have a problem. We continue to see price dispersion on an everyday level. Indeed, recent research (Brynjolfsson and Smith, 1999) shows this is just as true for internet-based sellers as it is for conventional retailers. However, it remains enormously difficult to construct dispersed price equilibria which are not extremely fragile. The possible exception is Model IIb, where search is noisy, and the flow of information to consumers is exogenous and so perhaps higher than they would choose themselves. Even there stability of equilibrium is only possible in a relatively small area of the parameter space.

What this paper leaves open is price dispersion as a disequilibrium phenomenon. In particular, one can think of a constant flow of new consumers who search more frequently than they would with experience.¹² Or one can imagine price cycles that never converge to any equilibrium. What evidence is there for such behaviour? One difficulty is that although the collection of evidence of dispersed prices is common, explicit comparison with the predictions of particular models is rare. Cason and Friedman (2000) test the noisy search model of Burdett and Judd (1983) (our Model IIb) experimentally. Prices chosen by their experimental subjects were not generated by a mixed strategy in that there was strong serial correlation. Rather, prices seem to follow cycles which match the pattern of best responses present in the Rock-Scissors-Paper structure. That is, prices fall in a manner reminiscent of Bertrand undercutting, but then when prices reach the minimum price we call \underline{p} , prices jump back up again. It is difficult to say in the time scale of the experiments whether the cycles are convergent or divergent.

This suggests two paths for future research. First, further empirical or experimental investigation might reveal whether existing equilibrium models of price dispersion explain the data well, or whether non-equilibrium approaches fare better. Second, on a theoretical level, the models we investigate here rely heavily on homogeneity assumptions. One can speculate whether learning in an environment such as that proposed by Benabou (1993) with heterogeneity amongst both buyers and sellers would fare better. But what we have shown here is that, ironically, learning by consumers can lead to the Diamond equilibrium and monopoly prices.

¹²One recent paper very much in this spirit is Haruvy and Erev (2000).

Appendix

Proof of Proposition 2: Before proving Proposition 2, we need a Lemma, which shows that a PDA operator satisfies an apparently stronger positivity condition than Property 3.

Notation: For $x \in S_n$ and $y \in \mathbb{R}^n$, let $\text{Var}_x(y) = x \cdot y^2 - (x \cdot y)^2$ be the variance of y with respect to the probability distribution x . For $T \subseteq \{1, 2, \dots, n\}$, let $S_T = \{x \in S_n : \text{support}(x) = T\}$, and take $e_i = (0, \dots, 1, \dots, 0) \in \mathbb{R}^n$ to be the vector having 1 in the i th coordinate and zero elsewhere. Let $e_T = \sum_{i \in T} e_i$, and write $e = \sum_{i=1}^n e_i$.

Lemma A: Let Q be a PDA operator on \mathbb{R}^n . Then there is a constant $m > 0$ (depending only on Q), such that, for all $x \in S_n, y \in \mathbb{R}^n$,

$$y \cdot Q(x)y \geq m \text{Var}_x(y).$$

Proof: Let \mathcal{T} be the set of triples (T, x, u) where $T \subseteq \{1, 2, \dots, n\}, x \in S_T$ and $u \in \mathbb{R}_T^n$ satisfies $x \cdot u = 0$ and $|u| = 1$. Then \mathcal{T} is a compact subset of $\mathcal{P}\{1, 2, \dots, n\} \times S_n \times \mathbb{R}^n$, where $\mathcal{P}X$ denotes the power set of X . If Q is continuously differentiable, the function $u \cdot Q(x)u$ is continuous on \mathcal{T} . By Properties 1, 2 and 3 of PDA operators, $u \cdot Q(x)u \geq 0$, with equality if and only if u is constant on $T = T(x)$. However, if $u_i = c$ for $i \in T$, then $0 = x \cdot u = cx \cdot e = c$, since $x \cdot e = 1$, so u is zero on T , and hence $u = 0$ because $u \in \mathbb{R}_T^n$. This contradicts $|u| = 1$. It follows that $u \cdot Q(x)u > 0$ on \mathcal{T} . Let $m = \inf_{\mathcal{T}} \{u \cdot Q(x)u\}$. Then $m > 0$ by the compactness of \mathcal{T} . It now follows that $y \cdot Q(x)y \geq m|y|^2$, for any $y \in \mathbb{R}_T^n$ with $x \cdot y = 0$.

Next, observe that $\text{Var}_x(y) = x \cdot y^2 = \sum_i x_i y_i^2 \leq \sum_i y_i^2 = |y|^2$. Hence, $y \cdot Q(x)y \geq m \text{Var}_x(y)$ for all $y \in \mathbb{R}_T^n$ with $x \cdot y = 0$. For general $y \in \mathbb{R}^n$, we have $y' = y - (x \cdot y)e_T$ satisfies $y' \in \mathbb{R}_T^n$ and $x \cdot y' = 0$ (because $x \cdot e_T = 1$). Thus, by Properties 1 and 2 of PDA operators, $y \cdot Q(x)y = y' \cdot Q(x)y' \geq m \text{Var}_x(y') = m \text{Var}_x(y)$.

Finally, consider an arbitrary $y \in \mathbb{R}^n$. Write $y = y_T + z$, where $y_T \in \mathbb{R}_T^n$, and z is zero on T . Then $x \cdot y = x \cdot y_T$ and $\text{Var}_x(y) = \text{Var}_x(y_T)$. Also, $z \in \mathbb{R}_{T^c}^n$, and hence $Q(x)z = 0$ by Property 2 of PDA operators, which also, by Property 4, implies that $z \cdot Q(x) = 0$. Hence, $y \cdot Q(x)y = y_T \cdot Q(x)y_T$, and the general result follows. \square

We now proceed with the proof of Proposition 2. Recall that the derivative of the vector-valued profit function, $\Pi(x)$, at a point $x \in \mathbb{R}^n$, is the matrix Π'_x , defined by

$$\Pi'_x y = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \{\Pi(x + \varepsilon y) - \Pi(x)\}, \quad y \in \mathbb{R}^n. \quad (\text{A1})$$

Similarly, the derivative of the matrix-valued PDA operator, $Q(x)$, at a point $x \in \mathbb{R}^n$, is the *bilinear* function, Q'_x , defined by

$$Q'_x(y, z) = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \{Q(x + \varepsilon y)z - Q(x)z\}, \quad y, z \in \mathbb{R}^n. \quad (\text{A2})$$

By the product rule, we therefore obtain the Jacobian derivative, J_x , for the PDA dynamics (6):

$$J_x y = [Q\Pi]'_x y = Q(x)\Pi'_x y + Q'_x(y, \Pi(x)). \quad (A3)$$

Let T be the support of ϕ . Note first that the restriction of J_ϕ to $\mathbb{R}_{T_0}^n$ will, by (A2) and Property 2 of PDA operators, be equal to $Q(\phi)\Pi'_\phi$. Since Π'_ϕ is negative definite, and $Q(\phi)$ is symmetric and positive definite on $\mathbb{R}_{T_0}^n$, it follows from a theorem of Hines (1980, pp348-9) that this restriction of J_ϕ to $\mathbb{R}_{T_0}^n$ has only eigenvalues with negative real part. Hence, ϕ is locally asymptotically stable under the restriction of the dynamics to S_T . If T contains $m \geq 1$ elements, this implies that the stable manifold of ϕ has dimension $\geq m - 1$. If $m = n$ (*i.e.* ϕ has full support), we are home. If $m < n$, we proceed by induction by successively adding faces of S_n to the stable manifold. To do this, proceed as follows.

Choose a sequence $T = T_0 \subset T_1 \subset \dots \subset T_{n-m} = \{1, 2, \dots, n\}$, with T_r containing $m + r$ elements. Suppose, inductively, that ϕ is locally asymptotically stable under the restriction of the dynamics (6) to $S_r = S_{T_r}$. We must show that ϕ is locally asymptotically stable under the restriction of the dynamics to S_{r+1} .

Let $\alpha = \pi^* e - \Pi(\phi)$. Then from Definition 3, $\alpha_i = 0$ for $i \in T$ and $\alpha_i > 0$ for $i \notin T$. In particular, $\alpha \in \mathbb{R}_{T_1}^n$. It follows that, for $x \in S_n$, $\alpha \cdot x = 0$ if and only if $x \in S_T$. From (A1) and Properties 3 and 4 of PDA operators, we have, for $z \in \mathbb{R}_0^n$,

$$\alpha \cdot J_\phi z = \alpha \cdot Q'_\phi(z, \Pi(\phi)) = -\alpha \cdot Q'_\phi(z, \alpha). \quad (A4)$$

To estimate the right hand term, consider $S_\phi = \{x - \phi : x \in S_n\}$. Then, $\alpha \cdot Q(\phi + z)\alpha \geq 0$ for each $z \in S_\phi$, by Property 3 of PDA operators, with equality if and only if α is constant on the support of $\phi + z$, by Property 2. Note that $z_i \geq 0$ for $i \notin T$. By Lemma A, there is a constant $m > 0$, such that $\alpha \cdot Q(\phi + z)\alpha \geq m \text{Var}_{\phi+z}(\alpha)$. Hence,

$$\frac{1}{\varepsilon} \alpha \cdot Q(\phi + \varepsilon z)\alpha \geq m \frac{1}{\varepsilon} \text{Var}_{\phi+\varepsilon z}(\alpha) = m(\alpha^2 \cdot z) + \mathcal{O}[\varepsilon],$$

the latter equality since $\alpha \cdot \phi = \alpha^2 \cdot \phi = 0$ and $z \in \mathbb{R}_0^n$. In the limit $\varepsilon \rightarrow 0$, we therefore obtain

$$\alpha \cdot Q'_\phi(z, \alpha) \geq \ell(\alpha \cdot z), \quad (A5)$$

where $\ell = m\alpha_* > 0$, with $\alpha_* = \min\{\alpha_i : i \notin T\}$.

Let U_r be a neighbourhood of ϕ in S_n such that $S_r \cap U_r$ is contained in the stable manifold of ϕ . Suppose that the unstable or centre manifold of ϕ has a non-empty intersection with $(S_{r+1} - S_r) \cap U_r$. Then, since S_r has dimension $m + r - 1$, and S_{r+1} has dimension $m + r$, the unstable or centre manifold in $S_{r+1} \cap U_r$ can be at most 1-dimensional. Since S_{r+1} and S_r are invariant under the dynamics, this implies that J_ϕ has a real eigenvalue $\lambda \geq 0$, and a real eigenvector v with $\phi + v \in S_{r+1}$. By hypothesis $v \notin S_r$, and hence $\alpha \cdot v > 0$. Thus $\alpha \cdot J_\phi v = \lambda(\alpha \cdot v) \geq 0$. But, by (A4) and

(A5), $\alpha \cdot J_\phi v \leq -\ell(\alpha \cdot v) < 0$. This is a contradiction. We conclude that the unstable or centre manifolds do not intersect S_{r+1} , and hence that there is a neighbourhood $U_{r+1} \subseteq U_r$ of ϕ in S_n such that $S_{r+1} \cap U_{r+1}$ is contained in the stable manifold. The result now follows by induction. \square

Proof of Proposition 3: If Π'_ϕ is positive definite on $\mathbb{R}_{T_0}^n$, then the restriction of $Q(\phi)\Pi'_\phi$ to $\mathbb{R}_{T_0}^n$ has only eigenvalues with positive real part, by the theorem of Hines (1980, pp348-9). On the other hand, if Π'_ϕ is not negative definite on $\mathbb{R}_{T_0}^n$, then by the same theorem, there is a positive definite symmetric operator A on $\mathbb{R}_{T_0}^n$ such that $AP\Pi'_\phi$ has an eigenvalue with positive real part on $\mathbb{R}_{T_0}^n$, where $P : \mathbb{R}^n \rightarrow \mathbb{R}_{T_0}^n \subset \mathbb{R}^n$ is the standard orthogonal projection. We shall extend any such A to a PDA operator, $Q(x)$, on \mathbb{R}^n so that $Q(\phi) = AP$. We do this by starting with any PDA operator, $\tilde{Q}(x)$, on \mathbb{R}^n and modifying it locally in a neighbourhood of ϕ .

For $\delta > 0$, let $B_\delta(\phi) \subset \mathbb{R}^n$ be the open ball of radius δ , centred at ϕ . Choose δ sufficiently small so that $T \subseteq T(x)$ for all $x \in S_n \cap B_\delta(\phi)$. Let $\mu : \mathbb{R}^n \rightarrow [0, 1]$ be a bump function at ϕ ; *i.e.* μ is continuously differentiable and satisfies:

(a): $\mu(y) = 0$ for $y \notin B_\delta(\phi)$; (b): $\mu(\phi) = 1$; (c): $0 < \mu(y) < 1$ for $y \in B_\delta(\phi) - \{\phi\}$.

Now define a modified operator Q by:

$$Q(y) = (1 - \mu(y))\tilde{Q}(y) + \mu(y)Q(\phi), \quad (\text{A6})$$

where $Q(\phi)$ is as constructed previously. It remains to check that Q defines a PDA operator. Clearly, Q is continuously differentiable on \mathbb{R}^n , and $Q(x)$ is symmetric. For $x \in S_n$ and $|x - \phi| \geq \delta$, we have $Q(x) = \tilde{Q}(x)$ by property (a) of μ , and hence Properties 1-4 of PDA operators are satisfied for such x . By construction of $Q(\phi)$ and property (b) of μ , these properties are satisfied for $x = \phi$. For $x \in S_n \cap B_\delta(\phi)$, $x \neq \phi$, we have $T \subseteq T(x)$, so that $\mathbb{R}_{T_0}^n \subseteq \mathbb{R}_{T(x)0}^n$ and $\mathbb{R}_{T(x)1}^n \subseteq \mathbb{R}_{T_1}^n$. Since $\tilde{Q}(x)$ maps \mathbb{R}^n to $\mathbb{R}_{T(x)0}^n$, and $Q(\phi)$ maps \mathbb{R}^n to $\mathbb{R}_{T_0}^n$, by Property 1 of PDA operators, it follows that $Q(x)$ maps \mathbb{R}^n to $\mathbb{R}_{T(x)0}^n$. This proves Property 1 for $Q(x)$. If $y \in \mathbb{R}_{T(x)1}^n$, then $\tilde{Q}(x)y = 0$ by Property 2 for $\tilde{Q}(x)$, and $Q(\phi)y = 0$ since $y \in \mathbb{R}_{T_1}^n$. This proves Property 2 for $Q(x)$. Finally, any non-zero $y \in \mathbb{R}_{T(x)0}^n$ may be orthogonally decomposed in the form $y = y_0 + y_1$, with $y_0 \in \mathbb{R}_{T_0}^n$ and $y_1 \in \mathbb{R}_{T_1}^n \cap \mathbb{R}_{T(x)0}^n$. Then, from (A6),

$$y \cdot Q(x)y = (1 - \mu(x))y \cdot \tilde{Q}(x)y + \mu(x)y_0 \cdot Q(\phi)y_0.$$

By construction, $y_0 \cdot Q(\phi)y_0 \geq 0$, with equality if and only if $y_0 = 0$. Also, $y \neq 0$ implies $y \cdot \tilde{Q}(x)y > 0$ by Property 3 for $\tilde{Q}(x)$. Since $x \neq \phi$, we have $(1 - \mu(x)) > 0$ by property (c) of μ . It now follows that $y \cdot Q(x)y > 0$, which proves Property 3 for $Q(x)$. \square

Proof of Proposition 4:

We first show

$$\langle f, \Pi^n(g) \rangle = \hat{f} \cdot \Pi^n(g), \quad \text{for all } f, g \in E. \quad (\text{A7})$$

To see this:

$$\begin{aligned}
\langle f, \Pi^n(g) \rangle &= \int_0^1 f(p) \Pi^n(g)(p) dp \\
&= \sum_{i=1}^n \int_{p_{i-1}}^{p_i} f(p) \Pi^n(g)(p) dp \\
&= \sum_{i=1}^n \Pi_i^n(g) \int_{p_{i-1}}^{p_i} f(p) dp \\
&= \sum_{i=1}^n \hat{f}_i \Pi_i^n(g) \\
&= \hat{f} \cdot \Pi^n(g).
\end{aligned}$$

Suppose $i \notin T$. Then, since $\phi(p)$, and hence $\Phi(p)$, is zero for $p < \underline{p}$,

$$\Pi_i^n(\phi) = n \int_{p_{i-1}}^{p_i} p D(\Phi(p)) dp = 0.$$

On the other hand, because of (3), the equilibrium payoff is $\pi^* = D(1) > 0$. This shows that condition (10) of the definition of regular ESS is satisfied. For condition (7), we have, using (A7), for $f \in S_E$,

$$\langle f, \Pi^n(\phi) \rangle = \hat{f} \cdot \Pi^n(\phi) = \pi^* \sum_{i \in T} \hat{f}_i \leq \pi^* = \langle \phi, \Pi^n(\phi) \rangle = \hat{\phi} \cdot \Pi^n(\phi),$$

since $\hat{f} \in S_n$, and hence $\sum_{i \in T} \hat{f}_i \leq \sum_{i=1}^n \hat{f}_i = 1$, with equality if and only if the support of \hat{f} is subordinate to the support of $\hat{\phi}$. This shows that all alternative best replies to ϕ have support subordinate to T .

(a): It remains to show that the negative definiteness condition (9) for a regular ESS holds in case (a). The result then follows by Proposition 2.

The derivative matrix $\Pi_\phi^{n'}$ may be calculated from (A1), which gives

$$\Pi_\phi^{n'} z(p) = n \int_{p_{i-1}}^{p_i} r D'(\Phi(r)) Z(r) dr, \quad \text{for all } p \in (p_{i-1}, p_i].$$

Using (3), this may be written

$$\Pi_\phi^{n'} z(p) = n\pi^* \int_{p_{i-1}}^{p_i} \left[\frac{D'(\Phi(r))}{D(\Phi(r))} \right] Z(r) dr, \quad \text{for all } p \in (p_{i-1}, p_i]. \quad (\text{A8})$$

Let $z \in \mathbb{R}_{T_0}^n$, so that $z(p) = Z(p) = 0$ for $p < \underline{p}$, and $Z(1) = 0$. Because $\Pi_\phi^{n'} z$ is constant on each price interval $(p_{i-1}, p_i]$, the formula (A7) holds with $\Pi^n(g)$ replaced by $\Pi_\phi^{n'} z$. Thus,

$$\langle z, \Pi_\phi^{n'} z \rangle = \hat{z} \cdot \Pi_\phi^{n'} z = \int_{\underline{p}}^1 \Pi_\phi^{n'} z(p) z(p) dp = \frac{1}{2} n\pi^* \int_{\underline{p}}^1 \left[\frac{D'(\Phi(p))}{D(\Phi(p))} \right] dZ^2(p).$$

Write $F = \Phi(p) \in [0, 1]$, and note that F increases monotonically from 0 to 1 as p increases from \underline{p} to p . Now integrate by parts to obtain,

$$\langle z, \Pi_\phi^{n'} z \rangle = -\frac{1}{2} n\pi^* \int_{\underline{p}}^1 \frac{d}{dp} \left[\frac{D'(F)}{D(F)} \right] Z^2(p) dp = -\frac{1}{2} n\pi^* \int_0^1 \frac{d}{dF} \left[\frac{D'(F)}{D(F)} \right] Z(\Phi^{-1}(F))^2 dF. \quad (\text{A9})$$

By hypothesis,

$$\frac{d}{dF} \left[\frac{D'(F)}{D(F)} \right] = \frac{\Theta(F)}{D(F)^2} > 0, \quad (A10)$$

for (almost) all $F \in [0, 1]$. Then (A9) gives $\langle z, \Pi_\phi^n z \rangle < 0$ for all $z \neq 0$ in $\mathbb{R}_{T_0}^n$, which proves (a).

(b) The reversed inequality (A10), together with (A9), show that Π_ϕ^n is positive definite on $\mathbb{R}_{T_0}^n$, and so the equilibrium $z = 0$ is unstable by Proposition 3, which proves (b). \square

Proof of Proposition 5: Even though δ_1 is not an element of $E = L_2[0, 1]$, we may still form its projection $\hat{\delta}_1 \in \mathbb{R}^n$. Thus,

$$\hat{\delta}_{1i} = \int_{p_{i-1}}^{p_i} \delta_1(p) dp = \begin{cases} 0 & \text{for } i < n \\ 1 & \text{for } i = n \end{cases}. \quad (A11)$$

That is, $\hat{\delta}_1 = e_n$. Similarly, the associated profit function $\Pi^n(\delta_1)$ is defined by equation (11), where the cumulative probability distribution Δ_1 associated with δ_1 , is the function on $[0, 1]$ defined by $\Delta_1(p) = 0$ if $p < 1$ and $\Delta_1(1) = 1$. Thus,

$$\Pi_i^n(\delta_1) = n \int_{p_{i-1}}^{p_i} p D(\Delta_1(p)) dp = \frac{1}{2}(p_{i-1} + p_i) D(0). \quad (A12)$$

Note that $\Pi^n(\delta_1)$ depends on the distribution of consumer samples, q , through the term $D(0)$.

Consider a perturbation from the Diamond equilibrium of the form $(f, q) = (\delta_1, e_1) + (z, r)$, with $z \in \mathbb{R}_0^n, r \in \mathbb{R}_0^N$, such that f is a probability density, and q a probability distribution. It is possible to show that such a perturbation must decay to zero. This proof is in two parts.

Lemma B The dynamics (15) on \mathbb{R}_0^n with $\phi = \delta_1$, has a locally asymptotically stable equilibrium at $z = 0$, independently of any consumer dynamics.

Proof: First note from (A11) that the support of $\hat{\delta}_1$, is $T = \{n\}$. Thus, \mathbb{R}_T^n is the 1-dimensional subspace of \mathbb{R}^n generated by e_n , from which it follows that $\mathbb{R}_{T_0}^n = \{0\}$ (see (??)). We therefore have $\mathbb{R}^n = \mathbb{R}_{T_1}^n$ and hence, by Property 1 of PDA operators, that $Q(\hat{\delta}_1) = 0$ for any PDA operator Q . This certainly implies that $z = 0$ is a stationary point of the dynamic (15) with $\phi = \delta_1$.

Let $z \in \mathbb{R}_0^n$ be such that $\hat{\delta}_1 + z \in S_n$. Then $z \geq 0$ on $T^c = \{1, \dots, n-1\}$. From the condition $Z_n = \sum_i z_i = 0$, we have $z_n = -Z_{n-1}$. Thus, $z = z^1 - Z_{n-1} e_n$, where $z^1 = \sum_{i < n} z_i e_i$. From (A11), $\hat{\delta}_1 + z = z^1 + (1 - Z_{n-1}) e_n$. Thus, z is completely determined by z^1 , subject only to the requirement that $Z_{n-1} \leq 1$.

Let $\gamma^n(z)$ be defined as in (15) with $\phi = \delta_1$, and $z \in \mathbb{R}_0^n$. Then, from (A3), the Jacobian derivative of γ^n at $z = 0$ is the operator on \mathbb{R}_0^n

$$J_0 z = Q(\hat{\delta}_1) \Pi_{\hat{\delta}_1}^n z + Q'_{\hat{\delta}_1}(z, \Pi^n \delta_1) = Q'_n(z, \Pi^n \delta_1), \quad (A13)$$

where we have written Q'_n for $Q'_{\hat{\delta}_1}$. Now consider the element $\alpha \in \mathbb{R}^n$ defined by

$$\alpha_i = \hat{p}_n - \hat{p}_i, \quad (\text{A14})$$

where $\hat{p}_i = \frac{1}{2}(p_{i-1} + p_i)$. Then $\alpha_n = 0$ and $\alpha_i \geq \alpha_{n-1} = \frac{1}{n} > 0$. By Property 1 of PDA operators, together with (A12) and (A13), it follows that

$$J_0 z = -D(0)Q'_n(z, \alpha). \quad (\text{A15})$$

Note that $z \cdot \alpha = z^1 \cdot \alpha = \sum_{k=1}^{n-1} z_k \alpha_k \geq \alpha_{n-1} Z_{n-1} = \frac{1}{n} Z_{n-1}$. Thus, if $z \in \mathbb{R}_0^n$ with $z_i \geq 0$ for $i < n$, then, as observed earlier, z is completely determined by $z^1 = \sum_{i < n} z_i e_i$, and hence $z \cdot \alpha \geq 0$, with equality if and only if $z = 0$. We denote by \mathbb{R}_{0+}^n the cone consisting of those $z \in \mathbb{R}_0^n$ for which $z_i \geq 0$ for $i < n$.

We can now make the key calculation as follows. Let $z \in \mathbb{R}_{0+}^n$. Then, for $\varepsilon > 0$, Lemma A gives

$$\frac{1}{\varepsilon} \alpha \cdot Q(\hat{\delta}_1 + \varepsilon z) \alpha \geq \frac{1}{\varepsilon} m \text{Var}_{\hat{\delta}_1 + \varepsilon z}(\alpha) = m(z \cdot \alpha^2) + \mathcal{O}[\varepsilon],$$

since $\hat{\delta}_1 \cdot \alpha = e_n \cdot \alpha = \alpha_n = 0$, and similarly $\hat{\delta}_1 \cdot \alpha^2 = 0$. Thus, taking the limit $\varepsilon \rightarrow 0$, and using (A15), we obtain

$$\alpha \cdot J_0 z \leq -mD(0)(z \cdot \alpha^2) \leq -m\mu(z \cdot \alpha^2), \quad (\text{A16})$$

where $\mu > 0$ is a constant, independent of any consumer dynamics, satisfying $D(0) \geq \mu$.

Now define a pseudo-inner product $\langle \cdot, \cdot \rangle_\alpha$ on \mathbb{R}_0^n , by $\langle z, w \rangle_\alpha = (z \cdot \alpha)(w \cdot \alpha)$. This is clearly bilinear and symmetric, with associated pseudo norm $\|z\|_\alpha = \sqrt{\langle z, z \rangle_\alpha} = |(z \cdot \alpha)|$. Thus, $\|z\|_\alpha \geq 0$ for all $z \in \mathbb{R}_0^n$. However, equality does not necessarily imply that $z = 0$, unless, as we have shown above, $z \in \mathbb{R}_{0+}^n$.

Consider the function $\ell : \mathbb{R}_{0+}^n \times [0, \infty) \rightarrow \mathbb{R}$ defined by $\ell(z, \varepsilon) = \alpha \cdot \gamma^n(\varepsilon z)$. Then ℓ is differentiable with respect to ε , and $\frac{\partial \ell}{\partial \varepsilon}(z, \varepsilon) = \alpha \cdot D\gamma_{\varepsilon z}^n z$. Hence, by the Mean Value Theorem, there exists $\theta(z) \in (0, 1)$ such that

$$\ell(z, \varepsilon) = \varepsilon \alpha \cdot D\gamma_{\varepsilon \theta(z)}^n z. \quad (\text{A17})$$

Now let $S_\alpha^n = \{z \in \mathbb{R}_{0+}^n : \|z\|_\alpha = 1\}$ be the closed unit $\|\cdot\|_\alpha$ -sphere segment in \mathbb{R}_{0+}^n . Then the restricted function, $d\ell : S_\alpha^n \times [0, 1] \rightarrow \mathbb{R}$, given by $d\ell(z, \varepsilon) = \alpha \cdot D\gamma_{\varepsilon z}^n z$, is uniformly continuous on its compact domain. Furthermore, from (A16)

$$d\ell(z, 0) = \alpha \cdot D\gamma_0^n z \leq -m\mu(z \cdot \alpha^2) = -m\mu\|z\|_\alpha \leq -m\mu\alpha_{n-1} = -\frac{m\mu}{n}. \quad (\text{A18})$$

Set $m_1 = m\mu/n$. Then, since $d\ell$ is continuous, there is an $\varepsilon_1 > 0$ such that $d\ell(z, \varepsilon) \leq -\frac{1}{2}m_1 < 0$ for all $(z, \varepsilon) \in S_\alpha^n \times [0, \varepsilon_1]$. We therefore conclude from (A17) that $\alpha \cdot \gamma^n(\varepsilon z) = \ell(z, \varepsilon) \leq -\frac{1}{2}\varepsilon m_1$ for all $(z, \varepsilon) \in S_\alpha^n \times [0, \varepsilon_1]$.

For general $z \in \mathbb{R}_{0+}^n$, we have $z = \varepsilon u$, with $u \in S_\alpha^n$, provided $\|z\|_\alpha = \varepsilon$. It now follows that

$$\frac{d}{dt}\|z\|_\alpha = \alpha \cdot \gamma^n(z) \leq -\frac{1}{2}m_1\|z\|_\alpha \quad \text{whenever } z \in \mathbb{R}_{0+}^n \text{ and } \|z\|_\alpha \leq \varepsilon_1.$$

Thus, $0 \leq \|z(t)\|_\alpha \leq \|z(0)\|_\alpha e^{-\frac{1}{2}m_1 t}$ for all $t \geq 0$ and $\|z(0)\|_\alpha \leq \varepsilon_1$. Hence, $\|z(t)\|_\alpha \rightarrow 0$ as $t \rightarrow \infty$. Since $z(t) \in \mathbb{R}_{0+}^n$, it follows that $z(t) \rightarrow 0$ as $t \rightarrow \infty$. \square

For the second part of the proof of Proposition 5, we must show that $z = 0$ is globally stable on the subspace $f = \delta_1$. First we require a preliminary lemma.

Define a function $\sigma(f) : [0, 1] \rightarrow \mathbb{R}^N$ by: $\sigma(f) = (\sigma_1(f), \dots, \sigma_N(f))$, where

$$\sigma_k(f)(p) = kp[1 - F(p)]^{k-1}, \quad 1 \leq k \leq N.$$

Lemma C. $\langle \sigma(\delta_1), \delta_1 \rangle = e$, where $e = (1, 1, \dots, 1) \in \mathbb{R}^N$.

Proof: For $\varepsilon > 0$, let $\delta_\varepsilon(p) = 0$ for $p < 1 - \varepsilon$, and $\delta_\varepsilon = \frac{1}{\varepsilon}$ for $1 - \varepsilon \leq p \leq 1$. Thus, $\delta_\varepsilon \in L_2[0, 1]$, and $\delta_1 = \lim_{\varepsilon \rightarrow 0} \delta_\varepsilon$, uniformly in p . Then

$$\Delta_\varepsilon(p) = \int_0^p \delta_\varepsilon(p') dp' = \begin{cases} 0 & \text{for } p < 1 - \varepsilon, \\ \frac{1}{\varepsilon}(p - 1 + \varepsilon) & \text{for } p \geq 1 - \varepsilon. \end{cases}$$

Hence,

$$\langle \sigma_k(\delta_\varepsilon), \delta_\varepsilon \rangle = \int_0^1 p \frac{k}{\varepsilon} [1 - \frac{1}{\varepsilon}(p - 1 + \varepsilon)]^{k-1} dp = 1 - \left(\frac{k}{k+1}\right)\varepsilon.$$

Thus, $\langle \sigma_k(\delta_1), \delta_1 \rangle = \lim_{\varepsilon \rightarrow 0} \langle \sigma_k(\delta_\varepsilon), \delta_\varepsilon \rangle = 1$, as required. \square

Lemma D The subspace $f = \delta_1$ is invariant and locally attracting under the joint seller-consumer dynamics. The Diamond equilibrium $q = e_1$, is globally attracting for all initial conditions inside this subspace, other than a finite set of isolated equilibria of the consumer dynamics.

Proof: Let $\kappa = (0, 1, 2, \dots, N-1) \in \mathbb{R}^N$. First note that if $q = e_1 + r$ is a probability distribution then $r \cdot e = 0$, $r_k \geq 0$ for $k > 1$, and $r_1 = -\sum_{k>1} r_k$. Thus,

$$\kappa \cdot r = \sum_{k>1} (k-1)r_k \geq 0. \tag{A19}$$

It follows that $\kappa \cdot r = 0$ if and only if $r = 0$. From (18) and Lemma C, it therefore suffices to show that $\kappa \cdot r(t) \rightarrow 0$ as $t \rightarrow \infty$ under the consumer dynamics

$$\dot{r} = -cR(e_1 + r)\kappa. \tag{A20}$$

To prove this, we have, by Lemma A for the PDA operator $R(q)$,

$$\kappa \cdot R(e_1 + r)\kappa \geq m \text{Var}_{e_1+r}(\kappa) = m\{\kappa^2 \cdot r - (\kappa \cdot r)^2\}, \tag{A21}$$

where $m > 0$ is a fixed constant. If $r \neq 0$, then $r_k > 0$ for at least one $k > 1$, and so $|r_1| = -r_1 = \sum_{k>1} r_k > 0$. Let $\hat{r} = (\hat{r}_2, \dots, \hat{r}_N) \in \mathbb{R}^{N-1}$, be the probability distribution defined by $\hat{r}_k = r_k/|r_1|$. By Jensen's inequality, we have

$$\kappa^2 \cdot r = |r_1| \kappa^2 \cdot \hat{r} \geq |r_1| (\kappa \cdot \hat{r})^2, \quad (\text{A22})$$

with equality if and only if $k - 1 = \text{constant}$ for all $k > 1$ for which $r_k > 0$. That is, equality holds if and only if $r_k = 0$ for all but one $k > 1$. In this case, $r = \alpha(e_k - e_1)$ for some $\alpha \in (0, 1]$ and some $k > 1$.

We now have,

$$\kappa^2 \cdot r - (\kappa \cdot r)^2 \geq |r_1|(1 - |r_1|)(\kappa \cdot \hat{r})^2 = \frac{1 - |r_1|}{|r_1|} (\kappa \cdot r)^2.$$

Substituting in (A20) and (A21), we obtain

$$\frac{d}{dt}(\kappa \cdot r) \leq -mc \left(\frac{1 - |r_1|}{|r_1|} \right) (\kappa \cdot r)^2.$$

whenever $r \neq 0$. Suppose that $1 - |r_1| \geq \eta > 0$, where η is arbitrarily small (but fixed). Then $\frac{d}{dt}(\kappa \cdot r) \leq -mc\eta(\kappa \cdot r)^2$. Integrating this, we obtain, for an initial condition $r(0) \neq 0$ satisfying $1 - |r_1(0)| \geq \eta$,

$$0 \leq \kappa \cdot r(t) \leq \frac{\kappa \cdot r(0)}{1 + mc\eta(\kappa \cdot r(0))t}, \quad (\text{A23})$$

Since $\kappa \cdot r(t) \geq |r_1(t)|$, we conclude that $|r_1(t)| \leq |r_1(0)| \leq 1 - \eta$, for all $t \geq 0$, and hence that (A23) remains valid for all $t \geq 0$. Thus, $r(t) \rightarrow 0$ as $t \rightarrow \infty$. Since we can take η arbitrarily small, this shows that $r = 0$ is globally asymptotically attracting outside the set of measure zero, $|r_1| = 1$.

Finally, consider what happens when $|r_1| = 1$. Then (A22) reduces to $\kappa^2 \cdot r \geq (\kappa \cdot r)^2$. However, this inequality is strict unless $r = \alpha(e_k - e_1)$ for some $\alpha \in (0, 1]$. On the other hand, in this case $|r_1| = 1$ if and only if $\alpha = 1$, and so equality holds if and only if $q = e_1 + r$ belongs to the set of isolated stationary points, $\mathcal{S} = \{e_k \mid k \geq 2\}$. From (A20) and (A21), we therefore obtain that $\frac{d}{dt}(\kappa \cdot r) < 0$, except on \mathcal{S} , where $\frac{d}{dt}(\kappa \cdot r) = 0$. Since $\kappa \cdot r$ is instantaneously decreasing outside \mathcal{S} , it follows that $|r_1|$ is instantaneously decreasing, and hence that a trajectory with initial condition $r \notin \mathcal{S}$, but with $|r_1| = 1$, moves instantaneously into the region $|r_1| < 1$, and is therefore asymptotically attracted to $r = 0$. This shows that $r = 0$ ($q = e_1$) is globally attracting for all trajectories with initial condition not in the finite set of unstable consumer equilibria \mathcal{S} . \square

Lemmas B and D complete the proof of Proposition 5. \square

Proof of Proposition 6: To simplify notation, write $\alpha = 1/(N - 1)$, so $\alpha \in (0, 1]$. Now use (22) to express q_1 in terms of \underline{p} , and substitute in (20) and (21) to obtain

$$V(\Phi) = \underline{p}^\alpha (1 - \underline{p})^{-(1+\alpha)} \int_{\underline{p}}^1 (p - \underline{p})(1 - p)^\alpha p^{-(1+\alpha)} dp.$$

Now set $p = \underline{p} + (1 - \underline{p})x$, and $\beta = (1 - \underline{p})/\underline{p} \in (0, \infty)$. Then

$$V(\Phi) = \beta \int_0^1 x(1 - x)^\alpha (1 + \beta x)^{-(1+\alpha)} dx.$$

For $\alpha < 1$ ($N < 2$), this integral may be expressed in terms of a hypergeometric function, ${}_2F_1(a, b; c, u)$,

$$V(\Phi) = \frac{\beta}{(1 + \alpha)(2 + \alpha)} {}_2F_1(2, 1 + \alpha; 3 + \alpha, -\beta) = \beta W(\alpha, \beta),$$

(see Whittaker and Watson, 1978, Chapter XIV), and for $\alpha = 1$ ($N = 2$),

$$V(\Phi) = \beta^{-1} \left\{ -2 + (1 + 2\beta^{-1}) \ln(1 + \beta) \right\}.$$

Thus, $V(\Phi)$ is a known function of two variables (α, \underline{p}) (equivalently, (N, q_1)), defined over the bounded domain $0 < \alpha \leq 1, 0 < \underline{p} \leq 1$. A typical example of a cross section of this surface (with α fixed) is shown in Figure 2. The important points to note are that $V(\Phi)(\underline{p}) \geq 0$, $V(\Phi)(0) = V(\Phi)(1) = 0$, with $V(\Phi)'(0) = \infty$ and $V(\Phi)'(1) = -1$, and $V(\Phi)(\underline{p})$ is concave with a unique maximum at $\hat{\underline{p}} = \hat{\underline{p}}(\alpha) \in (0, 1)$.

It now follows that, if $c > V(\Phi)(\hat{\underline{p}})$, then there are no mixed equilibria, and if $0 < c < V(\Phi)(\hat{\underline{p}})$, then there are two such equilibria. Since the Diamond equilibrium, $\underline{p} = 1$ is always present, the Proposition is proved. \square

Proof of Proposition 7: First linearise the joint dynamics (15, 19) about the given equilibrium (ϕ, q_1^*) to obtain the Jacobian operator which, regarded as a function on $\mathbb{R}_{T_0}^n \times \mathbb{R} \subset E \times \mathbb{R}$, has the form

$$J = \Lambda \Omega' = \begin{pmatrix} Q(\hat{\phi}) & 0 \\ 0 & \rho(q_1^*) \end{pmatrix} \begin{pmatrix} \Pi_\phi^{n'} & \Sigma(\hat{\phi}) \\ -V_\phi' & 0 \end{pmatrix},$$

where $\Sigma(f) = \sigma_1(f) - \sigma_N(f)$ with $\sigma_k(f)(p) = kp[1 - F(p)]^{k-1}$, $1 \leq k \leq N$. That is, using (23, 24) and writing $\alpha = 1/(N - 1)$, $\hat{\underline{p}} = (1 + \alpha)\underline{p}/(\alpha + \underline{p})$,

$$\Sigma(\phi)(p) = \left[\frac{(\alpha + \underline{p})}{\alpha(1 - \underline{p})} \right] (p - \hat{\underline{p}}).$$

Also, from (20), V_ϕ' is the \mathbb{R} -valued linear operator on E given by

$$V_\phi' z = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \{V(\Phi + \varepsilon Z) - V(\Phi)\} = - \int_{\underline{p}}^1 p^{-1} \Sigma(\phi)(p) Z(p) dp.$$

Note that $\hat{p} - \underline{p} = \underline{p}(1 - \underline{p})/(\alpha + \underline{p})$ and $1 - \hat{p} = \alpha(1 - \underline{p})/(\alpha + \underline{p})$. Thus, $\underline{p} < \hat{p} < 1$ for fixed $\underline{p} \in (0, 1)$ (equivalently, fixed $q_1 \in (0, 1)$).

First suppose $N > 2$, and take $\tilde{z} = (\varepsilon z, 1)$, with $z \in \mathbb{R}_{T_0}^n$. Then, from the definition of Ω' above,

$$(\varepsilon z, 0) \cdot \Omega' \tilde{z} = \varepsilon \left\{ \langle z, \Sigma(\phi) \rangle + \varepsilon z \cdot \Pi_\phi^{n'} z \right\}, \quad \text{and} \quad (0, 1) \cdot \Omega' \tilde{z} = -\varepsilon V'_\phi z.$$

Hence if there exists z such that $\langle z, \Sigma(\phi) \rangle > 0$ and $V'_\phi z < 0$, then we can choose $\varepsilon > 0$ sufficiently small so that $(\varepsilon z, 0) \cdot \Omega' \tilde{z} > 0$ and $(0, 1) \cdot \Omega' \tilde{z} > 0$. The result then follows from Proposition 3.¹³ It therefore suffices to construct such a z .

Integration by parts using $Z(\underline{p}) = Z(1) = 0$, gives

$$\langle z, \Sigma(\phi) \rangle = - \left[\frac{\alpha + \underline{p}}{\alpha(1 - \underline{p})} \right] \int_{\underline{p}}^1 Z(p) dp.$$

Now, recall that $\underline{p} < \hat{p} < 1$, and that $\underline{p} = p_{\underline{n}}$. Hence, for fixed \underline{p} , we may choose n large enough so that $\underline{p} < p_{\underline{n}+2} \leq \hat{p}$. Choose $z_i = 0$ for $i \neq \underline{n} + 1, \underline{n} + 2$, and $z_{\underline{n}+1} = -1, z_{\underline{n}+2} = 1$. Then $Z(p) < 0$ for $\underline{p} < p < p_{\underline{n}+2}$, and $Z(p) = 0$ for $p \geq p_{\underline{n}+2}$. Thus, $\langle z, \Sigma(\phi) \rangle > 0$, and

$$V'_\phi z = - \int_{\underline{p}}^{p_{\underline{n}+2}} p^{-1} \Sigma(\phi)(p) Z(p) dp < 0,$$

since $\Sigma(\phi)(p) < 0$ for $\underline{p} < p < \hat{p}$.

Finally, suppose $N = 2$. By the argument of Section 5, we know that $\Pi_\phi^{n'}$ is positive definite on $\mathbb{R}_{T_0}^n$, and hence that $Q(\hat{\phi})\Pi_\phi^{n'}$ has only eigenvalues with positive real parts. But the trace of J is equal to the trace of $Q(\hat{\phi})\Pi_\phi^{n'}$, which is positive. Hence, J must have at least one eigenvalue with positive real part. This completes the proof. \square

References

- Benabou, R. (1993). "Search market equilibrium, bilateral heterogeneity and repeat purchases," *J. Econ. Theory*, **60**, 140-158.
- Börgers, T., Sarin, R. (1997) "Learning through reinforcement and replicator dynamics," *J. Econ. Theory*, **77**, 1-14.

¹³The situation we are considering here is an asymmetric game, unlike in Proposition 3, which was formulated for symmetric games. The *two* positivity conditions, $(v, 0) \cdot \Omega' \tilde{z} > 0$ and $(0, 1) \cdot \Omega' \tilde{z} > 0$, are required to hold separately in this case, in order to preserve the product structure, $\Lambda(\phi, q_1^*) = Q(\hat{\phi}) \times \rho(q_1^*)$, of the destabilizing PDA dynamics. The proof of the appropriate form of Proposition 3 for asymmetric games is a straightforward elaboration of the proof for symmetric games.

- Brynjolfsson, E., Smith, M. (1999). "Frictionless commerce? A Comparison of Internet and Conventional Retailers," MIT (Sloan) working paper.
- Burdett, K., Judd, K. (1983). "Equilibrium price dispersion," *Econometrica*, **51**, 955-969.
- Cason, T., Friedman, D. (2000). "Buyer search and price dispersion: a laboratory study," working paper, University of California, Santa Cruz.
- Diamond, P. (1971). "A model of price adjustment," *J. Econ. Theory*, **3**, 156-168.
- Erev, I., Roth, A.E. (1998) "Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria," *Amer. Econ. Rev.*, **88**, 848-881.
- Friedman, D. (1991). "Evolutionary games in economics," *Econometrica*, **59**, 637-666.
- Friedman, D., Yellin, J. (1997). "Evolving landscapes for population games," working paper, University of California, Santa Cruz.
- Fudenberg, D., Kreps D. (1993). "Learning mixed equilibria," *Games Econ. Behav.*, **5**, 320-367.
- Fudenberg, D., Levine D. (1998). *The Theory of Learning in Games*. Cambridge, MA: MIT Press.
- Haruy E., Erev I. (2000). "Variable pricing: a customer learning perspective," working paper, Technion, Haifa.
- Hines, W.G.S. (1980). "Three characterizations of population strategy stability," *J. Appl. Prob.*, **17**, 333-340.
- Hofbauer, J. (2000). "From Nash and Brown to Maynard Smith: equilibria dynamics and ESS," forthcoming *Selection*.
- Hofbauer, J., Sigmund, K. (1990). "Adaptive dynamics and evolutionary stability," *App. Math. Lett.*, **3**, 75-79.
- Hofbauer, J., Sigmund K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge: Cambridge University Press.
- Hopkins, E. (1999a). "Learning, matching and aggregation," *Games Econ. Behav.*, **26**, 79-110.
- Hopkins, E. (1999b). "A note on best response dynamics," *Games Econ. Behav.*, **29**, 138-150.

- Oechssler, J., Riedel, F. (2001). "Evolutionary dynamics on continuous strategy spaces," *Econ. Theory*, **17**, 141-162
- Padilla, A.J. (1995). "Revisiting dynamic duopoly with consumer switching costs," *J. Econ. Theory*, **67**, 520-530.
- Rob, R. (1985). "Equilibrium price distributions," *Rev. Econ. Stud.*, **52**, 487-504.
- Rothschild, M. (1973). "Models of market organization with imperfect information: a survey," *J. Polit. Econ.*, **81**, 1283-1308.
- Salop, S., Stiglitz, J. (1977). "Bargains and rip-offs: a model of monopolistically competitive price dispersion," *Rev. Econ. Stud.*, **44**, 493-510.
- Schlag, K. (1998). "Why imitate and if so how?" *J. Econ. Theory*, **78**, 130-156.
- Seymour, R.M. (2000). "Dynamics for infinite dimensional games," working paper, University College, London.
- Shapley, L. (1964). "Some topics in two person games," in eds. M. Dresher et al., *Advances in Game Theory*. Princeton: Princeton University Press.
- Shilony, Y. (1977). "Mixed pricing in oligopoly," *J. Econ. Theory*, **14**, 373-88.
- Varian, H.R. (1980). "A model of sales," *Amer. Econ. Rev.*, **70**, 651-659.
- Whittaker, E. T., Watson, G.N. (1978). *A Course of Modern Analysis*. 4th Edition. Cambridge: CUP.
- Wilde, L.L. (1992). "Comparison shopping as a simultaneous move game," *Econ. J.*, **102**, 562-569.
- Young, H.P. (1993) "The evolution of conventions," *Econometrica*, **61**, 57-84.