Discussion Paper No. 1013

The Diamond Paradox: A Dynamic Resolution

by

Kyle Bagwell and Gary Ramey*

November 1992

Abstract

We consider the role of repeat business in resolving the paradox of Diamond [1971]. In each period, consumers engage in sequential price search at a positive search cost. Consumers enforce pricing discipline via a simple *loyalty-boycott search rule* that directs future-period searches away from firms that raise prices in the current period. In consumers' best equilibria, the equilibrium price decreases continuously with the level of search costs, and the competitive outcome obtains as search costs approach zero. We show further that Rotemberg and Saloner's [1986] finding of countercyclical markups does not arise in the presence of positive search costs.

---

*Northwestern University/Hoover Institution and University of California, San Diego. We thank Ian Domowitz, Michael Peters, Valerie Ramey, Joel Sobel and Bob Staiger for helpful discussions.

I. INTRODUCTION

Any economist studying search and pricing ultimately encounters the paradox of Diamond [1971]: in a static market with a homogeneous good and price–setting firms, if (i) firms have identical costs of production, (ii) consumers search sequentially for price information, and (iii) the search cost is positive for each consumer, then all equilibrium transactions must occur at the (common) monopoly price. Paradoxically, this is true no matter how small is the level of search costs, but when search costs reach zero, the equilibrium price jumps discontinuously to the competitive level. Diamond's key idea is that the presence of positive search costs allows firms to shade prices upward without losing any customers; such price shading will continue until each firm charges its most preferred price.

Previous research has attempted to resolve the Diamond paradox by relaxing one of the conditions (i), (ii) and (iii). Thus equilibrium prices below monopoly levels may arise because firms have different costs, and thus different monopoly prices; this permits consumers to pose a credible threat to search out the low–cost firms. Consumers may employ "nonsequential" search technologies by which multiple prices are observed in a single search, or sequential search may be undertaken at zero cost by some subset of consumers.[1]

---

[1] Reinganum [1979] studies search equilibria in the presence of exogeneous cost differences. Nonsequential search is considered by Salop and Stiglitz [1977] and Burdett and Judd [1983], among others. Varian [1980], Rob [1985], Stahl [1989] and others study models in which a subset of consumers may search at zero cost. Bagwell and Ramey [1992a] explore another route to resolving the Diamond paradox, which preserves (i)–(iii) but introduces "money–burning" advertising that firms use to coordinate consumer purchases. In that model, search threats are credible because advertising rivalry generates *endogenous* differences in firms' production costs and monopoly prices. Finally, Wolinsky [1986] proposes a resolution of the paradox that involves a slight relaxation of the assumption of

In this paper we provide a resolution of the Diamond paradox that maintains assumptions (i), (ii) and (iii), and instead reconsiders the standard search problem in an explicitly dynamic environment, in which market trading occurs period after period. Our central idea is that firms' incentives to shade price upward may be affected by the implications of price–shading for repeat business: a price–shading firm may not lose any business in the current period, due to positive search costs, but consumers can credibly boycott the firm in future periods by directing their future searches to rival firms. Conversely, consumers can show loyalty to a firm that eschews price–shading by returning to the firm in the following period. This *loyalty–boycott search rule* represents a simple and empirically plausible consumer search strategy, especially for low–cost, frequently–purchased items (groceries, restaurants, hardware, etc.).

We find that when consumers use the loyalty–boycott search rule, dynamic equilibria exist in which prices are below monopoly levels; moreover, reductions in search costs allow equilibria with lower prices to be sustained, and equilibrium prices approach the competitive limit in a continuous fashion as search costs are reduced to zero. Thus in the presence of repeat purchases, the natural monotonic relationship between price and search costs can arise, even when firms are identical, consumers search sequentially within each period, and search is costly for all buyers.

Our model confronts firms with a fundamental tradeoff between the incentive to "cheat" consumers in the current period with unexpected price hikes, and the long–run loss of equilibrium profits caused by this price–gouging. Consumers use their ability to search to alter this tradeoff. Intuitively, when search costs are high, firms can raise prices a great deal without losing current–period sales, and thus the gains from cheating are high; in this case, we show that high equilibrium prices are needed to dissuade firms from cheating, both because they generate high equilibrium profits and because they dampen

---

product homogeneity.

the current–period gain from cheating. For low search costs, however, only small price increases are feasible in the current period, and the gains from cheating are small; it is then possible to sustain equilibria with low equilibrium prices and profits. Gains from cheating are completely eliminated as search costs approach zero, and so firms are willing to accept equilibrium prices arbitrarily close to competitive levels.

We extend our model by introducing random period–to–period fluctuations in consumer demand. This allows us to reconsider the results of Rotemberg and Saloner [1986], who study the implications of demand fluctuations for equilibrium pricing in a dynamic price–setting model with zero search costs. Rotemberg and Saloner find that in *firm–optimal equilibria*, i.e. equilibria that give firms the highest profits over the set of equilibria, equilibrium markups are lower in periods of greater demand. We show, in contrast, that when positive search costs are introduced, equilibrium markups in firm–optimal equilibria are *independent* of the level of demand; this arises from the fact that with positive search costs, monopoly pricing can always be sustained via simple repetition of the static Nash equilibrium. Moreover, in *consumer–optimal equilibria*, periods of greater demand are characterized by *higher* equilibrium markups. Thus Rotemberg and Saloner's finding of countercyclical markups does not extend to the case of positive search costs.[2]

Our work is related to past literature on reputation and product quality, including Farrell [1980], Klein and Leffler [1981], Shapiro [1983] and Allen [1984]. In these papers, the prospect of repeat–business punishments gives firms an incentive to invest in *experience attributes* of the product, i.e. attributes that consumers can observe only after

---

[2] The empirical literature on the cyclicality of price–cost margins is not conclusive, but our prediction of procyclical margins does find some support. In particular, using panel data for 284 industries over 1958–1981, Domowitz, Hubbard and Petersen [1986a,1986b,1988] find that price–cost margins are procyclical in concentrated nondurable goods industries. Bils [1989], Rotemberg and Saloner, and Rotemberg and Woodford [1991] report evidence of countercyclical markups.

purchasing the product. We build on this work by considering how repeat business influences *search attributes*, which can be observed prior to purchase, and in particular we show that the effectiveness of repeat–business punishments is critically linked to the level of search costs. The idea that repeat–business punishments can discipline pricing has been considered by Okun [1980], who informally discusses a variety of links between pricing and repeat business, and Bagwell and Peters [1988], who consider the role of cost conditions in determining the credibility of repeat–business punishments in a two–period model. Our ideas are also related to some recently–developed dynamic price–setting models that incorporate search costs, but not repeat–business punishments; these include Benabou [1988], Fershtman and Fishman [forthcoming], Fishman and Rob [1991], and McMillan and Morgan [1988].

The plan of the paper is as follows. Section II lays out the basic model, and Section III considers a class of equilibria in which firms choose a common price along the equilibrium path. Section IV establishes a relationship between the equilibrium set and search costs that gives a resolution to the Diamond paradox. Section V extends the basic model to incorporate random demand, and considers the relationship between demand levels and equilibrium prices. Section VI concludes.

II. BASIC MODEL

Our model is comprised of a finite number of firms and a continuum of consumers, who trade a single homogeneous good in each of a countable infinity of periods numbered $t = 1,2,\ldots$. Firms are indexed by $i = 1,\ldots,N$, while consumers are taken to be uniformly distributed on $[0,1]$ with unit mass. The utility that each consumer derives in a given period from purchasing q units of the good at price p is $U(q) - pq$, where $U' > 0$, $U'' < 0$. Let $D(p)$ denote the demand function, which gives the utility–maximizing q for given p; we have $D'(p) < 0$ for $p < U'(0)$. The maximized level of utility is denoted by $V(p)$.

Each firm has a constant unit cost of c, which is assumed to satisfy $0 < c < U'(0)$. Fixed costs are taken to be negligible.[3] The profit per consumer of a firm charging price p is written:

$$\Pi(p) \equiv (p - c)D(p)$$

A firm that captures mass m of consumers earns $m\Pi(p)$. We assume that Π is strictly concave in p, and thus it has a unique maximizer $p^* \in (c,U'(0))$. Note that the monopoly price, which is p^*, is independent of m under our assumption of constant unit costs.

Consumers are assumed to be initially uninformed about the prices that firms charge in a given period, but consumers can acquire price information through search. We model this situation by means of the following *one–shot trading game*, which governs the trading that occurs within each period. The one–shot trading game consists of three stages:

Stage 1. Firms simultaneously choose prices p_i, $i = 1,...,N$. At the same time, each consumer chooses which firm he will visit initially. Consumers incur no cost for this initial search.

Stage 2. Each consumer observes the Stage 1 price choice of the firm that he visited in Stage 1. Consumers then simultaneously decide whether or not to visit a second firm. A search cost of $k > 0$ is deducted from utility if a second search is undertaken.

---

[3] Our results are altered somewhat by the addition of nonnegligible avoidable fixed costs; see note 10 below.

Stage 3. Each consumer that undertook a second search in Stage 2 observes the Stage 1 price choice of the second firm visited. Consumers then purchase as many units as desired from firms whose prices they have observed.

This one–shot trading game is a simple version of the Diamond [1971] model. Specifying a zero search cost for the initial search is a standard method for ruling out cases in which the sole equilibrium outcome has consumers staying out of the market. The limitation to two searches within a period creates no loss of generality, as our results can be easily extended to the usual case of sequential search. In either case, the key assumption that drives our results is that a consumer can observe only one additional price each time he incurs the search cost.

It is well known that in every sequential equilibrium of the one–shot trading game, all firms choose $p_i \geq p^*$, and all firms with positive equilibrium market share choose $p_i = p^*$. This remains true no matter how small k is. If $k = 0$ is specified, however, then there appear equilibria in which all active firms charge $p_i = c$, and all equilibria in which consumers search twice have this form. This severe discontinuity between competitive and monopolistic equilibrium pricing at $k = 0$ is known as the *Diamond paradox.*

Now consider the infinite repetition of the game, in which the one–shot trading game is repeated period after period. The structure of information with respect to current–period choices is given above, but now we need to say what information agents possess with respect to choices that were made in preceding periods. Let h_{it} denote the information that Firm i possesses in period t concerning the pricing, search and purchase decisions that have been made in periods 1,2,...,t–1, which is called Firm i's *history at t.* Correspondingly, $h_{\omega t}$ denotes the history at t that is possessed by Consumer $\omega \in [0,1]$.

There are many possible ways to specify the relationship between agents' histories and the actual choices made in preceding periods, and it turns out that our results hold for a broad class of such specifications. In particular, we impose only two restrictions: (i) $h_{\omega t}$ contains at least the names of the firms that Consumer ω had himself visited in the preceding period, together with the prices observed during those visits; and (ii) h_{it} does not identify the search or purchase decisions of any individual Consumer ω, nor does $h_{\omega t}$ identify the search or purchase decisions of any Consumer $\omega' \neq \omega$, although histories may contain the distributions of these decisions across consumers.[4] Beyond this, histories for firms and consumers are unrestricted, and they may contain anything from the full record of past decisions and observations to no additional information.

A *strategy* for Firm i is given by a sequence of mappings σ_{it}: $H_{it} \to \mathbb{R}_+$, $t = 1,2,...$, where H_{it} is the set of possible h_{it}, and σ_{it} maps histories to period t price choices. A strategy for Consumer ω is taken to be a pair of sequences of mappings $f_{\omega t}$: $H_{\omega t} \to \{1,...,N\}$, $s_{\omega t}$: $H_{\omega t} \times \mathbb{R}_+ \to \{1,...,N\}$, $t = 1,2,...$, where $H_{\omega t}$ gives the possible $h_{\omega t}$, $f_{\omega t}$ maps from histories to initial search decisions in period t, and $s_{\omega t}$ maps from histories and price observations to second search decisions in period t. By convention, $s_{\omega t} = f_{\omega t}$ denotes a decision not to make a second search in period t. A consumer's Stage 3 purchase decision is specified simply as the quantity D(p) purchased from the firm having the lowest price observed in the period.[5]

Payoffs for the repeated game are given by discounted payoff streams derived from the sequence of one–shot games, where $r > 0$ gives the common rate of discount. We consider a form of *sequential equilibrium* for this game, which is given by sequences

---

[4] Note that condition (i) allows a consumer to be "forgetful" with respect to all but the most recent price observations. The "consumer anonymity" assumption embodied in condition (ii) is needed to rule out equilibria in which consumers are deterred from making search or purchase decisions that maximize current–period utility by the threat of triggering a low–utility outcome in future periods. In such equilibria the consumers' threat to make a second search can be substantially neutralized. We feel that equilibria of this sort place an unrealistically large strategic weight on the decisions of individual consumers.

[5] This specification of Stage 3 decisions causes no loss of generality under the consumer anonymity assumption.

$\hat{\sigma}_{it}, \hat{f}_{\omega t}, \hat{s}_{\omega t}$, t = 1,2,..., for all i and ω, that satisfy the following conditions.

(A) For each Firm i, for each t and each $h_{it} \in H_{it}$, $\hat{\sigma}_{is}$ for s = t,t+1,... maximizes Firm i's discounted profit stream, conditional on h_{it}, from period t forward, where the path of profits is computed relative to some conjecture by Firm i as to the strategies of other firms and consumers. The path of decisions for periods 1,2,...,t–1 implied by Firm i's period t strategy conjecture must be consistent with h_{it}. Further, Firm i must conjecture that other firms and consumers will follow their equilibrium strategies for periods t,t+1,... .

(B) For each Consumer ω, for each t and each $h_{\omega t} \in H_{\omega t}$ and $p \in \mathbb{R}_+$, $\hat{f}_{\omega t}(h_{\omega t})$ and $\hat{s}_{\omega t}(h_{\omega t},p)$ maximize period t utility, where prior to visiting any firm, Consumer ω conjectures that the firm's period t price is chosen according to its equilibrium strategy.

Equilibrium condition (A) reflects the "perfection" requirement that each firm's decisions are payoff–maximizing at every decision point that might be reached, conditional on some conjecture by the firm as to how other agents play the game in past, present and future periods. Further, the firm's strategy conjecture for past periods must agree with the observed history of decisions, and the firm must conjecture that other agents play according to their equilibrium strategies in the present and all future periods. Condition (B) gives the corresponding perfection requirement for consumers' visitation decisions; consumers maximize current–period utility because, under our restriction on allowable histories, a consumers' own decisions cannot affect the future path of play. Prior to making visits, consumers conjecture that prices are determined according to firms' equilibrium strategies. Importantly, if a consumer observes a price on his first visit that is inconsistent with the visited firm's equilibrium strategy, the consumer still conjectures that other firms follow their equilibrium strategies.[6]

III. SINGLE–PRICE EQUILIBRIA

We focus attention on a simple class of sequential equilibria in which along the equilibrium path, all firms choose a common price in every period. A sequential equilibrium is a *single–price equilibrium* if there is a price $\hat{p}$ such that the following is true: for every Firm i and every t, if h_{it} contains no evidence of a deviation by any firm or any consumer, then $\hat{\sigma}_{it}(h_{it}) = \hat{p}$. The range of prices $\hat{p}$ that support single–price equilibria is particularly easy to characterize, as we demonstrate in the following proposition.

Proposition 1. The price $\hat{p}$ supports a single–price equilibrium if and only if $\hat{p} > c$ and:

(1) $$\Pi(p_d(\hat{p},k)) \leq \frac{1 + r}{r} \Pi(\hat{p})$$

where:

(2) $$p_d(\hat{p},k) \equiv \min\{ \hat{p} + d(\hat{p},k),\ p^* \}$$

---

[6] Conditions (A) and (B) give a form of the *sequential rationality* criterion of Kreps and Wilson [1982] that incorporates *structural consistency* in the formation of the firms' beliefs. The requirement that agents must conjecture that other agents follow their equilibrium strategies in the current and future periods, even when other agents have been observed to have deviated from their equilibrium strategies in the past, reflects the principle of *backward–induction rationality*, according to which agents posit that subsequent decisions are rational conditional on any decision point that might be reached. Our results would be completely unaltered if agents were instead free to form any conjecture about another agent who has been observed to have deviated from his equilibrium strategy. It is important, however, that agents should regard other agents' deviations from the equilibrium strategies as being uncorrelated; see Kreps and Ramey [1987] for further discussion of the latter point.

and $d(\hat{p},k) > 0$ is defined by:

(3) $$V(\hat{p} + d(\hat{p},k)) = V(\hat{p}) - k$$

Proof. Given in the Appendix.

Inequality (1) gives a necessary and sufficient condition for firms to be willing to choose their strategy in the single–price equilibrium. On the left–hand side of (1) is the maximum profit that a firm can obtain in a given period by deviating from the equilibrium strategy: $\hat{p} + d(\hat{p},k)$ gives the maximum price that a firm can charge in a period without inducing its initial visitors to make a second search, as indicated by (3); a firm's best deviation sets $p_d(\hat{p},k) = \hat{p} + d(\hat{p},k)$ if this price lies below the monopoly price, and otherwise it sets $p_d(\hat{p},k) = p^*$. On the right–hand side of (1) is the discounted profit stream that a firm obtains from following its equilibrium strategy. If a deviation leads the firm to lose all sales in future periods, then the equilibrium payoff is preferred to the deviation if and only if (1) holds.

Thus the key to the characterization is that a firm loses all future sales when it deviates. Consumers enforce this threat by means of a simple *loyalty–boycott search rule*, by which they continue to return to a firm if it charges $\hat{p}$, but switch to another firm in the following period if the firm deviates from $\hat{p}$. It is important to note that in supporting single–price equilibria with $\hat{p} < p^*$, it suffices for consumers to boycott only firms that *raise* their price above $\hat{p}$, while remaining loyal to firms that charge $\hat{p}$ or less. The boycott threat is credible because under the equilibrium conditions, consumers may conjecture that the deviating firm chooses a price no lower than $\hat{p}$ in the periods following the deviation, while other firms are conjectured to charge $\hat{p}$. This search rule, which places minimal informational and computational requirements on consumers, appeals to intuition as a realistic hypothesis concerning consumer behavior. Nonetheless it serves to deliver the largest credible punishment to deviating firms, and thereby supports the largest possible range of equilibrium prices.[7]

IV. A RESOLUTION OF THE DIAMOND PARADOX

In this section we show that the Diamond paradox is resolved via selection of single–price equilibria that have the lowest $\hat{p}$ among all single–price equilibria; we refer to these as the *consumer–optimal equilibria*. Let the consumer–optimal equilibrium price be written p_c.[8] Denote the static or one–time incentive to cheat as:

$$\Omega(\hat{p},k) \equiv \Pi(p_d(\hat{p},k)) - \Pi(\hat{p})$$

With this, (1) may be rewritten in the alternative form:

$$\Omega(\hat{p},k) \leq \frac{1}{r}\,\Pi(\hat{p})$$

This expression indicates clearly the tradeoff between the static incentive to cheat,

---

[7] In the language of Abreu [1988], the loyalty–boycott search rule serves as an *optimal penal code* since it gives a sequential equilibrium that drives a firm to its minmax value. The simplicity of the strategies needed to support the optimal penal code in this case contrasts strongly with the complexity of the decision rules that are often needed in other contexts, e.g. oligopolistic collusion with quantity–setting (Abreu [1986]). Note further that outcomes induced by a firm's deviation from the single–price equilibrium are on the Pareto frontier of outcomes for the game, and thus the punishments are renegotiation–proof (see Abreu and Pearce [1991]).

[8] Consumer–optimal equilibrium outcomes differ only in the distribution of consumers across firms.

represented by the left–hand side, and the future equilibrium profits forgone due to cheating, given on the right–hand side.

Observe that future equilibrium profits are zero when the equilibrium price is c, but they increase as $\hat{p}$ rises toward p^*, as shown in Figure 1. The incentive to cheat is positive at $\hat{p} = c$, and is zero at $\hat{p} = p^*$; thus it is necessary that the curves $\Omega(\hat{p},k)$ and $\Pi(\hat{p})/r$ intersect for some $\hat{p}$. Further, Figure 1 depicts the incentive to cheat as a decreasing function of $\hat{p}$, which implies a unique intersection between c and p^* of the curves $\Omega(\hat{p},k)$ and $\Pi(\hat{p})/r$. This intersection gives the consumer–optimal equilibrium price p_c, since (1) holds with equality at $\hat{p} = p_c$, and is violated for all $\hat{p} < p_c$.

Is it plausible to assume that the incentive to cheat decreases as the price rises? For $\hat{p}$ close to p^* we have $p_d(\hat{p},k) = p^*$, and thus $\Omega(\hat{p},k) = \Pi(p^*) - \Pi(\hat{p})$, which decreases in $\hat{p}$. But the same may not hold for smaller levels of $\hat{p}$, where $p_d(\hat{p},k) = p + d(p,k)$. Two effects emerge in this case. On the one hand, the concavity of the profit function ensures that the benefit of a price hike of a fixed size diminishes as the equilibrium price rises toward p^*. This effect suggests that the incentive to cheat is lower at higher prices. On the other hand, however, as the equilibrium price rises, the value to consumers of a second search is reduced, and they will therefore tolerate a larger price hike before seeking out a second firm. This effect suggests that the incentive to cheat increases with the equilibrium price.

To see this second effect, we may use (3) and $V'(p) = -D(p)$ to obtain:

$$\frac{\partial p_d(\hat{p},k)}{\partial \hat{p}} = \frac{D(\hat{p})}{D(p_d(\hat{p},k))} > 1$$

Thus an increase in $\hat{p}$ raises the best deviant price by even more. This effect tends to raise the incentive to cheat as $\hat{p}$ rises, and for $\Omega(\hat{p},k)$ to be decreasing in $\hat{p}$, the profit function itself must be sufficiently concave in price to offset the effect.

The needed concavity condition may be derived as follows. Differentiating $\Omega(\hat{p},k)$ for the case $p_d(\hat{p},k) = \hat{p} + d(\hat{p},k)$ and rearranging gives:

$$\frac{\partial\Omega(\hat{p},k)}{\partial\hat{p}} = D(\hat{p}) \int_{\hat{p}}^{p_d(\hat{p},k)} \psi'(p)dp$$

where:

$$\psi(p) \equiv \frac{(p - c)D'(p)}{D(p)}$$

It follows that $\psi'(p) < 0$ is a sufficient condition for the incentive to cheat to decrease in price. Further, $\psi'(p) < 0$ holds if and only if:

(4) $$\frac{\partial}{\partial p}\frac{-pD'(p)}{D(p)} > c\left[\frac{D'(p)^2 - D''(p)}{D(p)^2}\right]$$

On the left–hand side of (4) we have the slope of the demand elasticity, which customarily is taken to be positive to assure existence of a profit–maximizing price. (4) indicates that a more stringent condition on the slope of elasticity may be needed, in order to introduce a sufficient amount of concavity into the profit function to make $\Omega(\hat{p},k)$ decrease in $\hat{p}$. It is straightforward to show that $D''(p) \leq 0$ implies (4) when $p \geq c$.

Henceforth we assume that (4) holds, and this leads us to the following proposition.

Proposition 2. (a) p_c exists and satisfies $p_c \in (c,p^*)$.

(b) p_c is continuous and increasing in k and r, and satisfies:

$$\lim_{k\to 0} p_c = \lim_{r\to 0} p_c = c$$

The proof of part (b) is given in Figure 1. When the search cost declines, $d(\hat{p},k)$ is reduced. This lowers the gain from cheating for prices on the lower range of (c,p^*), i.e. for $\hat{p} \le p'$, where p' is defined by $p' + d(p',k) = p^*$. This shifts the curve $\Omega(\hat{p},k)$ downward in a continuous way, which reduces p_c continuously. Since $d(\hat{p},k)$ converges to zero as k approaches zero, the curve $\Omega(\hat{p},k)$ approximates the horizontal axis for sufficiently small k, and p_c is driven to zero. The effect of r is demonstrated in a similar way, except that a decline in r shifts the curve $\Pi(\hat{p})/r$ upward.[9]

It follows that the consumer–optimal equilibria exhibit a natural relationship between market power and search costs: markups are reduced continuously as search costs fall, and competitive pricing emerges as search costs approach zero. What drives this result is that consumers' threat to make a second search constrains firms' gains from deviating in a given period: if the cost of a second search is small, then a firm's best deviation from the single–price equilibrium involves only a small increase in price, which yields a small increase in profits. Thus lower search costs result in a smaller incentive to cheat. This in turn implies that lower equilibrium prices may be supported when search costs decline: the lower equilibrium prices act to restore the balance between the static incentive to cheat and the corresponding loss of future equilibrium profits, both by reducing the level of future equilibrium profits and by increasing the incentive to cheat.

---

[9] (4) is needed to ensure that p_c is continuous in k and r, but the remaining parts of Proposition 2 continue to hold when (4) is relaxed.

As search costs approach zero, the gains from cheating are everywhere driven to zero, and outcomes with arbitrarily small equilibrium profits become sustainable.[10]

It is interesting to note as k approaches zero, the set of equilibrium outcomes expands to the entire set of outcomes that generate positive profit, i.e. the outcomes that dominate firms' minmax payoff. Further, in contrast to the usual treatments of repeated games, this "folk theorem"–type result holds for a fixed positive level of the discount rate; here the gains from deviating are driven to the minmax level, in contrast to the standard technique in which the punishment for deviating from a positive–profit equilibrium is made arbitrarily large by reducing r to zero. Proposition 2 also considers the effect of r on the consumer–optimal equilibria: p_c becomes smaller as r is reduced, and p_c appoaches c as r approaches zero, in conformity with the conventional folk theorem.[11]

Proposition 2 shows that the natural relationship between search costs and equilibrium prices is *possible* in the presence of repeat business, but *necessity* of the relationship does not follow due to the existence of single–price equilibria with $\hat{p} > p_c$. Equilibria of the latter sort might reasonably be excluded, however, if firms are able actively to coordinate the selection of a single–price equilibrium, and there are at least three firms. To see this, note that any two firms can credibly propose to serve the market at a price $\hat{p}$ that satisfies (1), as such a proposal gives a single–price equilibrium for the

[10] The result is altered if there is an avoidable fixed cost of $F > 0$ per period that firms must incur if they post a price in Stage 1. A "shutdown" strategy that allows firms to leave the market must be added to firms' available Stage 1 choices, and it is assumed that consumers search only among the firms that have chosen not to shut down. In this case, the avoidable fixed cost bounds the equilibrium prices above c: in order to support a single–price equilibrium with $\hat{p} \neq p^*$, at least two firms must choose not to shut down in each period, and this is possible if and only if $(\hat{p} - c)D(\hat{p}) \geq 2F$. As $k \to 0$, the latter constraint becomes binding in the consumer–optimal equilibrium. Note that the constraint is weakened as F falls, and the results of the text emerge as $F \to 0$.

[11] Note however that the limiting outcomes differ as k and r become large: we have $\lim_{k\to\infty} p_c < p^*$, as can be seen from Figure 1, while clearly $\lim_{r\to\infty} p_c = p^*$.

subcoalition consisting of consumers and the two firms. Further, for any single–price equilibrium with $\hat{p} > p_c$, there is at least one pair of firms that can improve their payoff by dividing the consumers among themselves at a price slightly below $\hat{p}$; this means that for any such equilibrium, some subcoalition consisting of consumers and the two firms can collectively benefit by agreeing to another equilibrium.

Consumer–optimal equilibria cannot be overturned in this manner, however, since no group of firms can credibly propose a price below p_c. Thus when there are three or more firms, the consumer–optimal equilibria are uniquely robust to this type of coordination.[12] Note that this selection procedure can be regarded as a form of price competition, in which groups of firms make price proposals from the set of credible proposals; here firms are able to communicate information concerning their intended pricing policies, even though they are unable to communicate a commitment to an actual price level.[13]

V. DEMAND FLUCTUATIONS AND CYCLICALITY OF MARKUPS

In an important recent paper, Rotemberg and Saloner [1986] develop a model of repeated trading in which (i) consumers can costlessly observe all current–period prices; and (ii) the demand function shifts randomly from period to period according to an i.i.d.

---

[12] The selection procedure closely mimics the logic of *coalition–proof Nash equilibrium*, introduced by Bernheim, Peleg and Whinston [1987]. The argument of the text establishes that if coalitions are restricted to selecting single–price equilibria, then any coalition–proof equilibrium must be a consumer–optimal equilibrium. Further, a consumer–optimal equilibrium satisfies coalition–proofness if and only if it cannot be overturned by a coalition consisting only of firms; the latter can be assured in many cases through appropriate specification of consumers' equilibrium visitation strategies. It is argued in the Conclusion that the restriction to single–price equilibria is actually quite plausible.

[13] This kind of competition in price proposals is most sensible for firms having simple pricing structures. Firms with extensive product lines or complex pricing rules may find communication of pricing policy to be prohibitively difficult, which suggests that equilibria with prices above the consumer–optimal level might be more easily sustained in such cases.

random process. In Rotemberg and Saloner's context, the unique Nash equilibrium outcome of the one–shot game is the competitive outcome, and firms' threat to revert to this outcome allows dynamic equilibria with price above marginal cost to be sustained. Rotemberg and Saloner consider *firm–optimal equilibria*, i.e. equilibria that maximize firms' profits over the set of equilibria, and they find that firm–optimal equilibrium prices must be lower in periods with greater realized demand. This result is important in that it gives a reason why markups might move countercyclically with fluctuations in aggregate demand.

In this section we consider the sensitivity of Rotemberg and Saloner's findings to the introduction of imperfect price information. This can be done in a straightforward way by incorporating random demand fluctuations into the model of the preceding sections. We find that the presence of only a small search cost has a major effect on Rotemberg and Saloner's result: as long as $k > 0$, equilibrium markups are *independent* of the demand realization in the firm–optimal equilibrium. Further, in the consumer–optimal equilibrium, markups must be *higher* in periods with greater realized demand.

Assume now that consumers are uniformly distributed on $[0,M^z]$ with total mass M^z, where $z \in \{H,L\}$ and $M^H > M^L > 0$; state $z = H$ is the "high–demand state," and state $z = L$ is the "low–demand state." The realization of z is determined independently in each period, with $\gamma \in (0,1)$ giving the probability that $z = H$.[14] Following Rotemberg and Saloner, it is assumed that the realization z_t of the demand variable in period t is determined before any agents' actions are chosen, and further that the realization is

[14] This demand structure can be interpreted most directly in terms of a fixed population of consumers who may be active or inactive in the market in any given period. In particular, let there be a fixed population $[0,M^H]$ of consumers, each of whom is active with probability one in the high–demand state, and with probability M^L/M^H in the low–demand state.

public information; thus firms' and consumers' period t strategies may be conditioned on z_t. We require also that consumers' period t histories contain the realization z_{t-1}.

We focus on *two-price equilibria*, which are sequential equilibria with the following property: for every Firm i and every h_{it}, if h_{it} contains no evidence of a deviation by any agent, then $\hat{\sigma}_{it}(h_{it},z_t) = \hat{p}^H$ for $z_t = H$, and $\hat{\sigma}_{it}(h_{it},z_t) = \hat{p}^L$ for $z_t = L$. Thus two–price equilibria are indexed by a pair of prices $\hat{p}^H$ and $\hat{p}^L$ that give the equilibrium–path choices in the high–demand and low–demand states, respectively. The following proposition gives a characterization of two–price equilibria that is similar to the earlier characterization of single–price equilibria.

Proposition 3. The prices $\hat{p}^H$ and $\hat{p}^L$ support a two–price equilibrium if and only if:

(5a) $$\Pi(p_d(\hat{p}^H,k))M^H \leq \Pi(\hat{p}^H)M^H + \frac{1}{r}[\gamma\Pi(\hat{p}^H)M^H + (1-\gamma)\Pi(\hat{p}^L)M^L]$$

(5b) $$\Pi(p_d(\hat{p}^L,k))M^L \leq \Pi(\hat{p}^L)M^L + \frac{1}{r}[\gamma\Pi(\hat{p}^H)M^H + (1-\gamma)\Pi(\hat{p}^L)M^L]$$

where:

(6) $$p_d(\hat{p},k) \equiv \max\{ \min\{ \hat{p} + d(\hat{p},k), p^* \}, c \}$$

and $d(\hat{p},k)$ is defined in (3).

The proof of this proposition is a straightforward adaptation of the proof of Proposition 1 and is therefore ommitted. Inequalities (5a) and (5b) give the conditions that rule out deviations by firms in the high–demand and low–demand states, respectively. Unlike the case of single–price equilibria, we may have $\hat{p}^z \leq c$ in one of the states, supported by the prospect of a high price in the other state. In such a case, it is possible that $\hat{p}^z + d(\hat{p}^z,k) < c$, so that if the firm deviated in state z it would do best by driving its customers away; in (6) the definition of p_d is modified accordingly. As with Proposition 1, this proposition relies on the use of a loyalty–boycott search rule whereby a firm that follows the equilibrium strategy retains its customers in the following period, whereas a firm that deviates loses its customers. Consumers need more information to implement the search rule in the present context, however, since they must know the state of demand in the current and preceding periods; this is not an unreasonable assumption if the demand state is associated with broad and widely–publicized aggregate economic conditions.

It is straightforward to verify that (5a–b) are always satisfied by $\hat{p}^H = \hat{p}^L = p^*$, i.e. it is always possible to select an equilibrium in which firms choose the monopoly price in each period. This follows simply from the fact that the unique sequential equilibrium outcome of the one–shot game is the monopoly outcome when k is strictly positive. Since these prices clearly give a firm–optimal outcome, it follows that markups in the firm–optimal equilibria bear *no relation at all* to the demand realization once positive search costs are introduced. Of course, this result relies on our particular specification of demand fluctuations and production costs, in which the monopoly price is unaffected by demand fluctuations; more generally, it may sometimes be plausible to suppose that high–demand states would be linked to *higher* monopoly prices, which would lead to *procyclical* markups in firm–optimal equilibria.[15]

We now consider the consumer–optimal equilibria of the modified model. This allows us to check the robustness of our earlier results to the possibility of demand

[15] Bagwell and Ramey [1992a,1992b], however, consider a *better deal property* in which monopoly prices are decreasing with the volume of demand; in this case countercyclical markups may be the plausible prediction for firm–optimal equilibria.

fluctuations; further, as argued above, consumer–optimal equilibria may be the most reasonable outcomes when firms have the power to coordinate equilibrium selection, and thus it is important to consider the relationship between markups and demand in these equilibria.

When looking for consumer–optimal equilibria of the modified model, there are two key forces present, and it is useful to identify them at the outset. First, when the demand state is high, the incentive to deviate is particularly acute, because an added price hike is received on a larger volume of sales. Thus if equilibrium is to be maintained, the high–demand equilibrium price must be adjusted to dampen the already–large incentive to defect. Second, as in the preceding section, this dampening is accomplished via the selection of a higher equilibrium price in the high–demand state, and this gives rise to procyclical markups in the consumer–optimal equilibria.

The consumer–optimal equilibria will be derived in two steps: first, we describe the set of $\hat{p}^H,\hat{p}^L$ that satisfy (5a–b); and second, the consumer–optimal equilibrium prices are located by maximizing consumer utility on this set.[16] To begin, let $P^a(\hat{p}^L)$ give the minimum level of $\hat{p}^H$ that satisfies the high–demand equilibrium condition (5a) for a given low–demand price $\hat{p}^L \leq p^*$, if (5a) can possibly be satisfied for the given $\hat{p}^L$; similarly, let $P^b(\hat{p}^L)$ give the minimum $\hat{p}^H$ that satisfies the low–demand equilibrium condition (5b) for given $\hat{p}^L \leq p^*$. The following lemma derives properties of the curves P^a and P^b.

---

16 Rotemberg and Saloner derive firm–optimal equilibria using a simple procedure that exploits the fact that the objective function to be maximized, which is the firms' equilibrium profit stream, also gives the punishment for deviation. This procedure cannot be used to derive consumer–optimal equilibria, however, since the objective and the punishment are distinct. In approaching this more complex problem, we have opted to restrict attention to the case of two possible demand states.

Lemma 1. (a) P^a is a continuous and strictly decreasing function of $\hat{p}^L$ on $[\underline{p}^a, p^*]$, where $\underline{p}^a > 0$ and $P^a(\underline{p}^a) = p^*$. Similarly, P^b is a continuous and strictly decreasing function of $\hat{p}^L$ on $[\underline{p}^b, p^*]$, where $\underline{p}^b > 0$ and $P^b(\underline{p}^b) = p^*$;

(b) $\underline{p}^a < \underline{p}^b$;

(c) $P^b(\hat{p}^L) < P^a(\hat{p}^L)$ at the level of $\hat{p}^L$ such that $P^a(\hat{p}^L) = \hat{p}^L$;

(d) P^b is strictly steeper than P^a at any point of intersection.

Proof. Given in the Appendix.

Figure 2 illustrates the properties of P^a and P^b that are established in Lemma 1. In particular, part (b) of the lemma ensures that P^a intersects the line $\hat{p}^H = p^*$ at a strictly lower level of $\hat{p}^L$ than does P^b, while (c) implies that P^a intersects the 45° line at a strictly higher level of $\hat{p}^L$; thus P^a and P^b must intersect each other at some point above the 45° line, labeled X in Figure 2. Further, accorting to (d) this is the unique point of intersection: P^b must be steeper than P^a at X because $\hat{p}^L$ exerts a relatively stronger effect on P^b through its effect on the profits from deviating in the low–demand state, while the influence of $\hat{p}^L$ on P^a is relatively weaker since only $\hat{p}^H$ affects the profits from deviating in the high–demand state.

The fact that the intersection point X must lie above the 45° line may be understood intuitively in terms of the link between the current–period demand state and the gains from deviating. When $\hat{p}^H = p^*$, the very low price $\hat{p}^L = \underline{p}^a$ may preserve equilibrium in the high–demand state, but it gives a strong incentive to deviate when the low–demand state is realized; thus P^b must lie above P^a for small $\hat{p}^L$. On the 45° line, in contrast, prices are independent of state, and so at a given price the incentive to deviate is greater in the high–demand state; due to (4), the per–customer returns from a price hike are lower when the base price is increased, so that prices must be *higher* in the high–demand state to preserve equilibrium. This ensures that P^a lies above P^b where the curves cross the 45° line. Rotemberg and Saloner's model differs in that *lower* prices are required to preserve equilibrium in the high–demand state, since their assumption of zero search costs implies that firms deviate by reducing, rather than increasing, their prices.

The shaded area of Figure 2 gives the set of $\hat{p}^H$ and $\hat{p}^L$ that satisfy both (5a) and (5b), and which can thus support two–price equilibria. The consumer–optimal equilibrium selects consumers' preferred equilibrium prices from the constraint set given by the shaded area. The consumers' objective function, denoted $W(\hat{p}^H,\hat{p}^L)$, is taken to be the expected value of total consumer utility, where expectation is taken with respect to the total mass of consumers in the market:[17]

$$W(\hat{p}^H,\hat{p}^L) \equiv \gamma M^H V(\hat{p}^H) + (1-\gamma)M^L V(\hat{p}^L)$$

The following lemma relates the indifference curves of W to the curves P^a and P^b.

Lemma 2. (a) For $\hat{p}^H \leq \hat{p}^L$, indifference curves of W are strictly steeper than P^a at every point of intersection;

(b) For $\hat{p}^H \geq \hat{p}^L$, indifference curves of W are strictly flatter than P^b at every point of intersection.

Proof. Given in the Appendix.

The relationship between indifference curves and the constraint set is depicted in Figure 3. Below the 45° line, consumers strictly prefer to move upward along the

[17] The function W/M^H gives the expected utility obtained by each Consumer $\omega \in [0,M^H]$ under the interpretation of note 14.

boundary of the constraint set, which is P^a, toward the 45° line; further movement along P^a is preferred once the 45° line is reached. It follows that $\hat{p}^H > \hat{p}^L$ is necessary in any consumer–optimal equilibrium, i.e. consumer optimality requires that markups are *higher* in the high–demand state. Thus the consumer–optimal equilibrium implies a relationship between markups and demand that is the *opposite* of the relationship found in Rotemberg and Saloner's firm–optimal equilibrium.

The following intuition underlies this result. As shown in the proof of Lemma 2, consumers' desire to substitute between $\hat{p}^H$ and $\hat{p}^L$ is governed by the likelihood ratio of consumer demand in the two states. Similarly, the effect of movements in $\hat{p}^H$ and $\hat{p}^L$ on firms' equilibrium profits is related to this likelihood ratio, but the ratio is weighted by the effect on firms' market power in the two states. Given (4), firms find price increases to be more desirable when the base price is lower, and in particular price increases in the high–demand state are relatively more valuable when $\hat{p}^H < \hat{p}^L$. Thus consumers will strictly prefer an increase in $\hat{p}^H$ and a reduction in $\hat{p}^L$ that leaves firms' equilibrium profits unchanged. The curve P^a combines firms' desire to substitute with the added effect of $\hat{p}^H$ on firms' gains from deviating in the high–demand state; since higher $\hat{p}^H$ reduces the gains from deviating, this gives consumers even more reason to prefer upward movements along P^a.

Similar reasoning underlies consumers' preference for downward movements along P^b at points above the 45° line. It follows that the consumer–optimal equilibrium, denoted by p_c^H and p_c^L, lies either at the intersection point X, as shown in Figure 3, or at a point of tangency between P^a and a consumer indifference curve, lying below X but above the 45° line. Note finally that as $k \to 0$, both P^a and P^b must pass arbitrarily close to the point $\hat{p}^H = \hat{p}^L = c$, and it follows that X converges to this point as $k \to 0$; by the above reasoning the consumer–optimal equilibrium will also converge to this point. The same argument holds for the case of $r \to 0$. These findings are summarized in the following proposition.

Proposition 4. (a) p_c^H and p_c^L exist and satisfy $p_c^L < p_c^H < p^*$;

(b)

$$\lim_{k \to 0} p_c^z = \lim_{r \to 0} p_c^z = c \quad \text{for } z = H,L$$

Observe that for small levels of search costs, the consumer–optimal equilibria approximate the competitive outcome, in which markups are zero in either demand state.

Note finally that our concavity condition (4), which is needed for Lemmas 1 and 2, may be strained when it is possible for two–price equilibria to have prices below c in one of the demand states; this happens when k and r are small. This is because for such prices, the assumption $D'' \leq 0$ is no longer sufficient for (4). It is straightforward to verify that (4) holds for every $p \geq 0$ under the assumption of linear demand. More generally, however, it is possible that the incentive to cheat increases in p for a range of p $< c$, and for low k and r this might give rise to counterychical consumer–optimal equilibria, in which prices are below unit costs in the high–demand state.

VI. CONCLUSION

This paper has shown that the Diamond paradox can be resolved without recourse to cost asymmetries, nonsequential search techologies, or zero–search–cost consumers, if trade occurs repeatedly over time. In such cases, consumers are able to discipline pricing by means of repeat business punishments, and moreover these punishments can be implemented using very simple and plausible search rules. The natural monotonic relationship between search costs and equilibrium prices is present in consumer–optimal equilibria, and further the competitive outcome emerges as search costs approach zero.

Our framework has proven useful for examining the robustness of Rotemberg and Saloner's finding of countercyclical markups to the introduction of positive search costs. With positive search costs, the countercyclicality result does not hold in firm–optimal equilibria, and is actually reversed in consumer–optimal equilibria. These results raise new difficulties for the idea that oligopolistic rivalry is a significant source of macroeconomic "real rigidity;" further, the contrast between our results and Rotemberg and Saloner's suggests that it may prove quite difficult to forge any robust link between real rigidity and imperfect competition.[18]

We have restricted attention to dynamic equilibria that are stationary, in the sense that equilibrium strategies are independent of the particular period in which actions are chosen. This restriction is nontrivial in the present context, since nonstationary equilibria with upward–sloping price paths may yield greater consumer payoffs than the stationary consumer–optimal equilibria considered above. Allowing such nonstationary equilibria would not change our basic conclusions, however: reductions in search costs would still reduce the gains from deviating in any given period, and this would make possible a downward shift in the entire future path of equilibrium prices.

More importantly, nonstationary equilibria strain our interpretation of consumers' equilibrium strategies in terms of simple search rules. Whether actual consumers are able to implement strategies carefully calibrated to particular time periods is certainly open to question, and further there may be difficulty establishing the focal "initial time period" that plays an important role in the theoretical analysis. Thus the restriction to single–price equilibria can be viewed as a reflection of consumers' lack of knowledge of the strategically relevant time period, and/or their inability to link search behavior to particular time periods.

---

[18] Haltiwanger and Harrington [1991] and Staiger and Wolak [1992] discuss other factors that can lead to procyclical markups in a repeated oligopoly.

Finally, we have made use of the artifice that the interval of firms' price adjustments correspond exactly to the time between customers' repeat purchases. It is straightforward, however, to extend the model to allow for a price adjustment interval that is shorter than the repeat purchase interval. Suppose consumers' purchases are distributed uniformly through continuous time and are repeated every Δ^r periods. Let prices be readjusted every $\Delta^p < \Delta^r$ periods. When consumers use the loyalty–boycott search rule, a firm is deterred from raising its price at time t because it will lose its flow of customers for the intervals $[t + \kappa\Delta^r, t + \Delta^p + \kappa\Delta^r)$, $\kappa = 1,2,\ldots$. Equations (1) and (5a–b), with slight adjustments in the discounting, continue to give the equilibrium conditions for single–price and two–price equilibria, and our results can be derived just as above.

APPENDIX

Proof of Proposition 1. Suppose $\hat{p} \in (c, U'(0))$ supports a single–price equilibrium but fails (1). Let $\hat{m}_{it}$ give Firm i's equilibrium market share in period t, and put:

$$\overline{m}_i = r\, \Sigma_{t=2}^{\infty} \left(\frac{1}{1+r}\right)^{t-1} \hat{m}_{it}$$

$\overline{m}_i$ gives the constant market share that would give Firm i the same discounted stream of market shares for periods $t = 2,3,\ldots$ as does $\hat{m}_{it}$. It follows that Firm i's equilibrium payoff is given by:

$$\Pi(\hat{p})(\hat{m}_{i1} + \frac{\overline{m}_i}{r})$$

Since $\Sigma_{i=1}^{N} \hat{m}_{i1} = \Sigma_{i=1}^{N} \overline{m}_i = 1$, there is some j with either $\hat{m}_{j1} \geq \overline{m}_j > 0$ or $\hat{m}_{j1} > \overline{m}_j = 0$. Consider a deviation by Firm j in period 1 to $p_d(\hat{p},k) - \epsilon$ for some small $\epsilon > 0$. According to equilibrium condition (B), all consumers who visit Firm j on their first search in period 1 conjecture that all other firms choose $\hat{p}$ in period 1, in accordance with the equilibrium strategies. Condition (B) thus implies that the consumers' equilibrium strategies specify no second search following this deviation, since by (3) consumers anticipate higher period 1 utility by staying with Firm j, while the restriction on firms' histories ensures that each individual consumer regards the future path of prices, and thus his own future payoff stream, as being independent of his own period 1 search behavior. By this reasoning, consumers who initially visit firms other than Firm j in period 1 also do not make a second search. Thus Firm j's payoff following the deviation is:

$$\Pi(p_d(\hat{p},k) - \epsilon)\hat{m}_{j1} + W_j > \frac{1 + r}{r}\,\Pi(\hat{p})\hat{m}_{j1} \geq \Pi(\hat{p})(\hat{m}_{j1} + \frac{\overline{m}_j}{r})$$

where W_j gives Firm j's profits in the continuation game following the deviation, discounted to period 1, and the strict inequality follows from $W_j \geq 0$ and failure of (1) for ϵ sufficiently small. Thus $\hat{\sigma}_{j1} = \hat{p}$ does not actually give a payoff–maximizing strategy for Firm j. It is straightforward to rule out prices $\hat{p} \notin (c,U'(0))$.

Now suppose $\hat{p} > c$ and satisfies (1). Let firms' strategies be given by $\hat{\sigma}_{it}(h_{it}) = \hat{p}$ for all i and t and *all* $h_{it} \in H_{it}$, and let this also give consumers' conjectures of firms' period t,t+1,... strategies for every $h_{\omega t}$, and firms' conjectures of other firms' period t,t+1,... strategies for every h_{it}. Clearly these conjectures satisfy the restrictions in (A) and (B), and according to these conjectures any specification of $\hat{f}_{\omega t}$ satisfies (B) since consumers conjecture that all firms choose the same price in every period. Thus we may define $\hat{f}_{\omega t}$ as follows: Let $\hat{f}_{\omega 1}$ be an arbitrary measurable function from [0,1] to {1,...,N}, and let $\hat{f}_{\omega t}$ be defined to satisfy:

$$\hat{f}_{\omega t}(h_{\omega t}) \begin{cases} = \hat{f}_{\omega(t-1)}(h_{\omega(t-1)}), & \text{if Firm } \hat{f}_{\omega(t-1)}(h_{\omega(t-1)}) \text{ chose } \hat{p} \text{ in period t–1} \\ \neq \hat{f}_{\omega(t-1)}(h_{\omega(t-1)}), & \text{otherwise} \end{cases}$$

where $h_{\omega t}$ is a continuation of $h_{\omega(t-1)}$. Thus Consumer ω continues to make his initial visit to the firm that he initially visited in the preceding period if that firm did not deviate from the equilibrium strategy in the preceding period. If the firm was observed to deviate, then the consumer visits some other firm.

For the given specification of $\hat{f}_{\omega t}$, (A) and (B) imply that $\hat{s}_{\omega t}$ must satisfy $\hat{s}_{\omega t}(h_{\omega t},p) \neq \hat{f}_{\omega t}(h_{\omega t})$ whenever $p > \hat{p} + d(\hat{p},k)$, since by (3) we have $V(p) < V(\hat{p}) - k$ in this case, and $V(\hat{p}) - k$ gives the utility that consumers conjecture they will obtain from any firm that would be visited on a second search. Similarly, we may specify $\hat{s}_{\omega t}(h_{\omega t},p) = \hat{f}_{\omega t}(h_{\omega t})$ for $p \leq \hat{p} + d(\hat{p},k)$. Let all firms and consumers conjecture that consumers will follow these strategies in periods t,t+1,... no matter what h_{it} or $h_{\omega t}$ is observed.

Finally we show that firms' strategies satisfy (A) and (B). Consider the price choice of Firm i in period 1. For any h_{it}, Firm i's conjecture of other agents' strategies implies some mass m_{it} of consumers that are expected to make their initial visit to Firm i in period t. Since Firm i conjectures that other agents will follow the equilibrium strategies in the current and future periods, it expects that consumers who initially visit other firms in period t will continue to visit those firms in all future periods. Further, it expects that choosing $\sigma_{it}(h_{it}) \neq \hat{p}$ in period t will lead its own period t customers to switch firms in the following period, never to return to Firm i. Thus Firm i anticipates zero sales in all future periods if it deviates in period t. Given its strategy conjecture, Firm i maximizes period t profit by choosing $\sigma_{it}(h_{it}) = \min \{ \hat{p} + d(\hat{p},k), p^* \}$, and thus the highest payoff it can gain by deviating in period t is given by the left–hand side of (1) times m_{it}. By following the equilibrium strategy, however, Firm i expects to obtain the larger payoff given by the right–hand side of (1) times m_{it}. *Q.E.D.*

Proof of Lemma 1. (a) (5a) may be rewritten:

$$\text{(A1)} \qquad \gamma M^H [\frac{1}{\gamma} \Pi(p_d(\hat{p}^H,k)) - (\frac{1}{\gamma} + \frac{1}{r})\Pi(\hat{p}^H)] \leq (1-\gamma)M^L \frac{1}{r} \Pi(\hat{p}^L)$$

Using (4) it can be verified that the term in brackets is strictly decreasing in $\hat{p}^H$ for $\hat{p}^H < p^*$, and is minimized by $\hat{p}^H = p^*$. Our assumption $U'(q) > 0$ for all q implies that $D(p) \to \infty$ as $p \to 0$, and so the right–hand side of (A1) may be made arbitrarily small by taking $\hat{p}^L$ sufficiently close to zero. Thus (A1) cannot be satisfied for small enough $\hat{p}^L$. Since (5a) is satisfied for $\hat{p}^H = \hat{p}^L = p^*$, there is a critical value $\hat{p}_L = \underline{p}^a < p^*$ such that (A1) just holds for $\hat{p}^H = p^*$; we have $P^a(\underline{p}^a) = p^*$, and $P^a(\hat{p}^L) < p^*$ for $\underline{p}^a < \hat{p}^L \leq p^*$. The above–noted monotonicity of the left– and right–hand sides of (A1) assures that P^a is continuous and strictly decreasing as long as $P^a > 0$, and latter follows on the whole domain of P^a since the left–hand side of (A1) becomes arbitrarily small as $\hat{p}^H \to 0$. The properties of P^b may be established similarly.

(b) At $\hat{p}^L = \underline{p}^a$, $\hat{p}^H = p^*$, we have:

$$\text{(A2)} \qquad (\Pi(p_d(\underline{p}^a,k)) - \Pi(\underline{p}^a))M^L > 0 = (\Pi(p_d(p^*,k)) - \Pi(p^*))M^H$$

$$= \frac{1}{r}[\gamma\Pi(p^*)M^H + (1-\gamma)\Pi(\underline{p}^a)M^L]$$

where the first inequality follows from $\underline{p}^a < p^*$ and the second equality follows from the fact that (5a) holds with equality at $\hat{p}^L = \underline{p}^a$, $\hat{p}^H = p^*$. From (A2) it follows that (5b) cannot be satisfied for any $\hat{p}^H$ when $\hat{p}^L = \underline{p}^a$, and so $\underline{p}^b > \underline{p}^a$.

(c) Note first that $\underline{p}^a < p^*$, $P^a(p^*) < p^*$, and P^a continuous and strictly decreasing assure that there is a single value of $\hat{p}^L$ such that $P^a(\hat{p}^L) = \hat{p}^L$, and further this value satisfies $\hat{p}^L < p^*$. At this value we have:

$$\text{(A3)} \qquad 0 < (\Pi(p_d(\hat{p}^L,k)) - \Pi(\hat{p}^L))M^L < (\Pi(p_d(\hat{p}^L,k)) - \Pi(\hat{p}^L))M^H$$

$$= \frac{1}{r}\Pi(\hat{p}^L)(\gamma M^H + (1-\gamma)M^L)$$

where the first inequality follows from $\hat{p}^L < p^*$, the second inequality follows from $M^L <$

M^H, and the third equality follows from the fact that (5a) holds with equality at $\hat{p}^H = P^a(\hat{p}^L)$, and also $\hat{p}^H = \hat{p}^L$ at the selected value of $\hat{p}^L$. (A3) implies that (5b) holds with strict inequality at $\hat{p}^H = P^a(\hat{p}^L)$, and so $P^b(\hat{p}^L) < P^a(\hat{p}^L)$.

(d) Treating (5a) as an equality and differentiating gives:

$$\text{(A4)} \qquad \frac{\partial P^a}{\partial \hat{p}^L} = \frac{-(1 - \gamma)M^L}{\gamma M^H} \frac{\Pi'(\hat{p}^L)/r}{\Pi'(P^a)/r - \Omega'(P^a)/\gamma}$$

Similarly, treating (5b) as an equality and differentiating gives:

$$\text{(A5)} \qquad \frac{\partial P^b}{\partial \hat{p}^L} = \frac{-(1 - \gamma)M^L}{\gamma M^H} \frac{\Pi'(\hat{p}^L)/r - \Omega'(\hat{p}^L)/(1 - \gamma)}{\Pi'(P^b)/r}$$

The result follows from comparison of (A4) and (A5) at any $\hat{p}^L$ and $\hat{p}^H = P^a = P^b$ (the argument is easily amended should Ω fail to be differentiable at p'). Q.E.D.

Proof of Lemma 2. Let p' satisfy $p' + d(p',k) = p^*$. Treating (5a) as an equality, it can be shown that, for $P^a < p'$:

$$\text{(A6)} \qquad \frac{\partial P^a}{\partial \hat{p}^L} = \frac{-(1 - \gamma) M^L D(\hat{p}^L)}{\gamma M^H D(P^a)}$$

$$\times \frac{(1 + \psi(\hat{p}^L))/r}{(1 + \psi(P^a))/r + [\psi(P^a) - \psi(P^a + d(P^a,k))]/\gamma}$$

while for $P^a > p'$:

$$\text{(A7)} \qquad \frac{\partial P^a}{\partial \hat{p}^L} = \frac{-(1-\gamma)M^L D(\hat{p}^L)}{\gamma M^H D(P^a)} \; \frac{(1+\psi(\hat{p}^L))/r}{(1+\psi(P^a))(1/r+1/\gamma)}$$

Treating (5b) as an inequality, it can be shown that, for $P^b < p'$:

$$\text{(A8)} \qquad \frac{\partial P^b}{\partial \hat{p}^L} = \frac{-(1-\gamma)M^L D(\hat{p}^L)}{\gamma M^H D(P^b)}$$

$$\times \frac{(1+\psi(\hat{p}^L))/r + [\psi(\hat{p}^L) - \psi(\hat{p}^L + d(\hat{p}^L,k))]/(1-\gamma)}{(1+\psi(P^b))/r}$$

while for $P^b > p'$:

$$\text{(A9)} \qquad \frac{\partial P^b}{\partial \hat{p}^L} = \frac{-(1-\gamma)M^L D(\hat{p}^L)}{\gamma M^H D(P^b)} \; \frac{(1+\psi(\hat{p}^L))(1/r+1/(1-\gamma))}{(1+\psi(P^b))/r}$$

Next, using $V'(p) = -D(p)$, we have:

$$\text{(A10)} \qquad \left.\frac{d\hat{p}^H}{d\hat{p}^L}\right|_{W=\text{constant}} = \frac{-(1-\gamma)M^L D(\hat{p}^L)}{\gamma M^H D(\hat{p}^H)}$$

Parts (a) and (b) of the lemma follow directly from comparison of (A6)–(A9) with (A10), using (4). *Q.E.D.*

REFERENCES

Abreu, Dilip, "Extremal Equilibria of Oligopolistic Supergames," *Journal of Economic Theory* XXXIX [1986], 191–225.

———, "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica* LXVI [1988], 383–396.

———, and David Pearce, "A Perspective on Renegotiation in Repeated Games," in Reinhard Selten, ed., *Game Equilibrium Models II* (Berlin: Springer–Verlag), 1991.

Allen, Franklin, "Reputation and Product Quality," *Rand Journal of Economics* XV [1984], 311–327.

Bagwell, Kyle, and Michael Peters, "Dynamic Monopoly Power when Search is Costly," CMSEMS D.P. No. 772, Northwestern University, April 1988.

———, and Garey Ramey, "Coordination Economies, Advertising and Search Behavior in Retail Markets," Discussion Paper 92–05, Department of Economics, University of California, San Diego, January 1992a.

———, and ———, "Advertising and Coordination," Discussion Paper 92–07, Department of Economics, University of California, San Diego, January 1992b.

Benabou, Roland, "Search, Price Setting and Inflation," *Review of Economic Studies* LV [1988], 353–376.

Bernheim, B. Douglas, Bezalel Peleg, and Michael D. Whinston, "Coalition–Proof Nash Equilibria I. Concepts," *Journal of Economic Theory* XLII [1987], 1–12.

Bils, Mark, "Pricing in a Customer Market," *Quarterly Journal of Economics* CIV [1989], 699–718.

Burdett, Kenneth, and Kenneth L. Judd, "Equilibrium Price Dispersion," *Econometrica* VI [1983], 955–969.

Diamond, Peter A., "A Model of Price Adjustment," *Journal of Economic Theory* III [1971], 156–168.

Domowitz, Ian, Glenn Hubbard and Bruce Petersen, "Business Cycles and the Relationship between Concentration and Price–Cost Margins," *RAND Journal of Economics* XVII [1986a], 1–17.

———, ——— and ———, "The Intertemporal Stability of the Concentration–Margins Relationship," *Journal of Industrial Economics* XXXV [1986b], 13–34.

———, ——— and ———, "Market Structure and Cyclical Fluctuations in U.S. Manufacturing," *Review of Economics and Statistics* LXX [1988], 55–66.

Farrell, Joseph, *Prices as Signals of Quality*, Ph.D. Dissertation, Brasenose College, 1980.

Fershtman, Chaim, and Arthur Fishman, "Price Cycles and Booms: Dynamic Search Equilibrium," *American Economic Review* (forthcoming).

Fishman, Arthur, and Rafael Rob, "The Durability of Information, Market Efficiency and The Size of Firms: Search with Repeated Transactions," University of Pennsylvania D.P. #91–39, January 1991.

Haltiwanger, John, and Joseph E. Harrington, Jr., "The Impact of Cyclical Demand Movements on Collusive Behavior," *RAND Journal of Economics* XXII [1991], 86–106.

Klein, Benjamin, and Keith B. Leffler, "The Role of Market Forces in Assuring Contractual Performance," *Journal of Political Economy* LXXXIX [1981], 615–641.

Kreps, David M., and Garey Ramey, "Structural Consistency, Consistency and Sequential Rationality," *Econometrica* LV [1987], 1331–1389.

———, and Robert Wilson, "Sequential Equilibrium," *Econometrica* V [1982], 863–894.

McMillan, John, and Peter B. Morgan, "Price Dispersion, Price Flexibility, and Repeat Purchasing," *Canadian Journal of Economics* XXI [1988], 884–902.

Okun, Arthur, *Prices and Quantities: A Macroeconomic Analysis* (Washington, D.C.: Brookings Institution), 1980.

Reinganum, Jennifer F., "A Simple Model of Equilibrium Price Dispersion," *Journal of Political Economy* LXXXVII [1979], 851–858.

Rob, Rafael, "Equilibrium Price Distributions," *Review of Economic Studies* LII [1985], 487–504.

Rotemberg, Julio J., and Garth Saloner, "A Supergame–Theoretic Model of Price Wars during Booms," *American Economic Review* LXXVI [1986], 390–407.

———, and Michael Woodford, "Markups and the Business Cycle," in Olivier Jean Blanchard and Stanley Fischer, eds., *NBER Macroeconomics Annual 1991* (Cambridge, MA: MIT Press), 1991.

Shapiro, Carl, "Premiums for High Quality Products as Returns to Reputation," *Quarterly Journal of Economics* XCVIII [1983], 659–679.

Stahl, Dale O., "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review* LXXIX [1989], 700–712.

Staiger, Robert W., and Frank A. Wolak, "Collusive Pricing with Capacity Constraints in the Presence of Demand Uncertainty," *RAND Journal of Economics* XXIII [1992], 203–220.

Varian, Hal R., "A Model of Sales," *American Economic Review* LXX [1980], 651–659.

Wolinsky, Asher, "True Monopolistic Competition as a Result of Imperfect Information," *Quarterly Journal of Economics* CI [1986], 493–511.

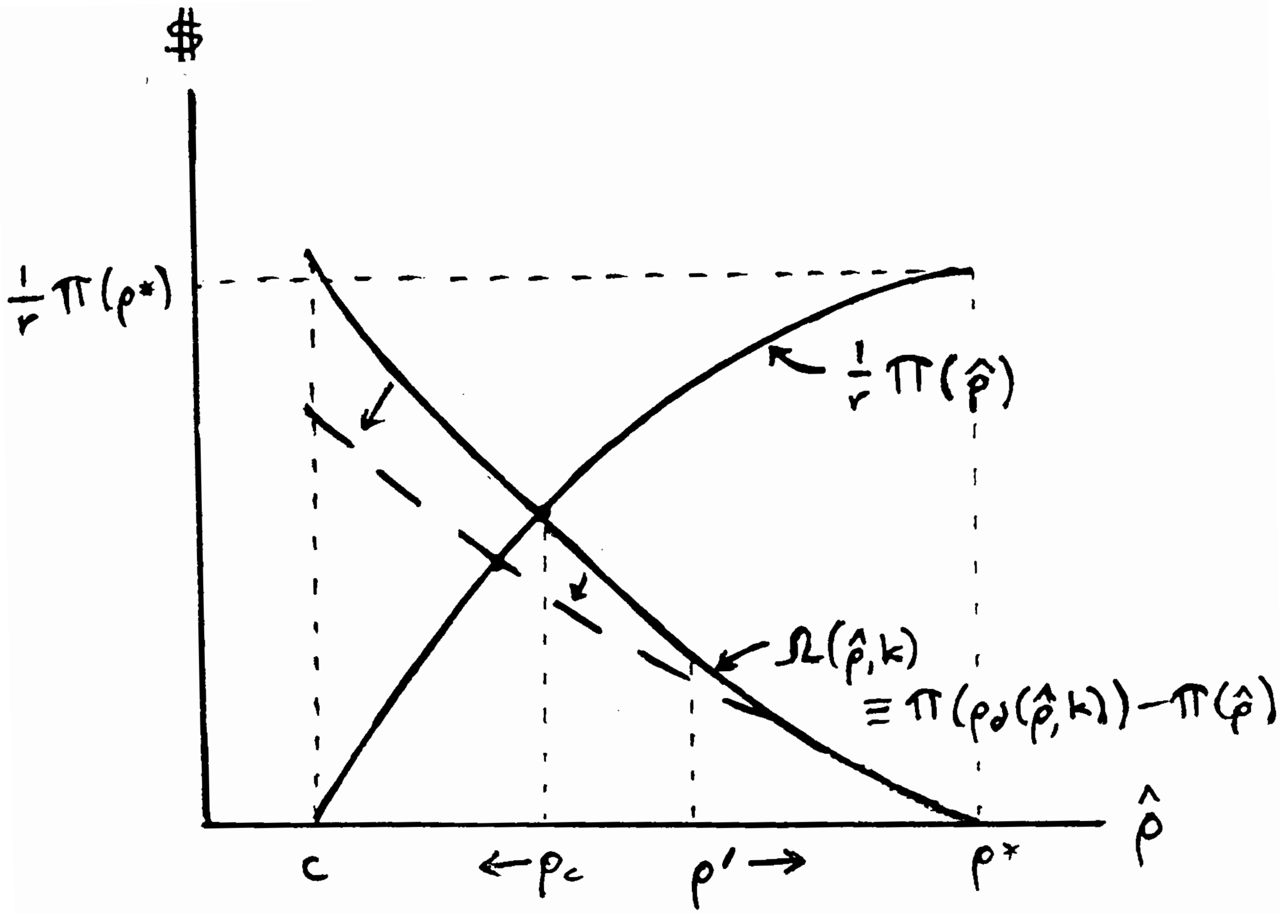

Figure 1

Derivation of Consumer-Optimal Single-Price Equilibria

Arrows show effect of decrease in k

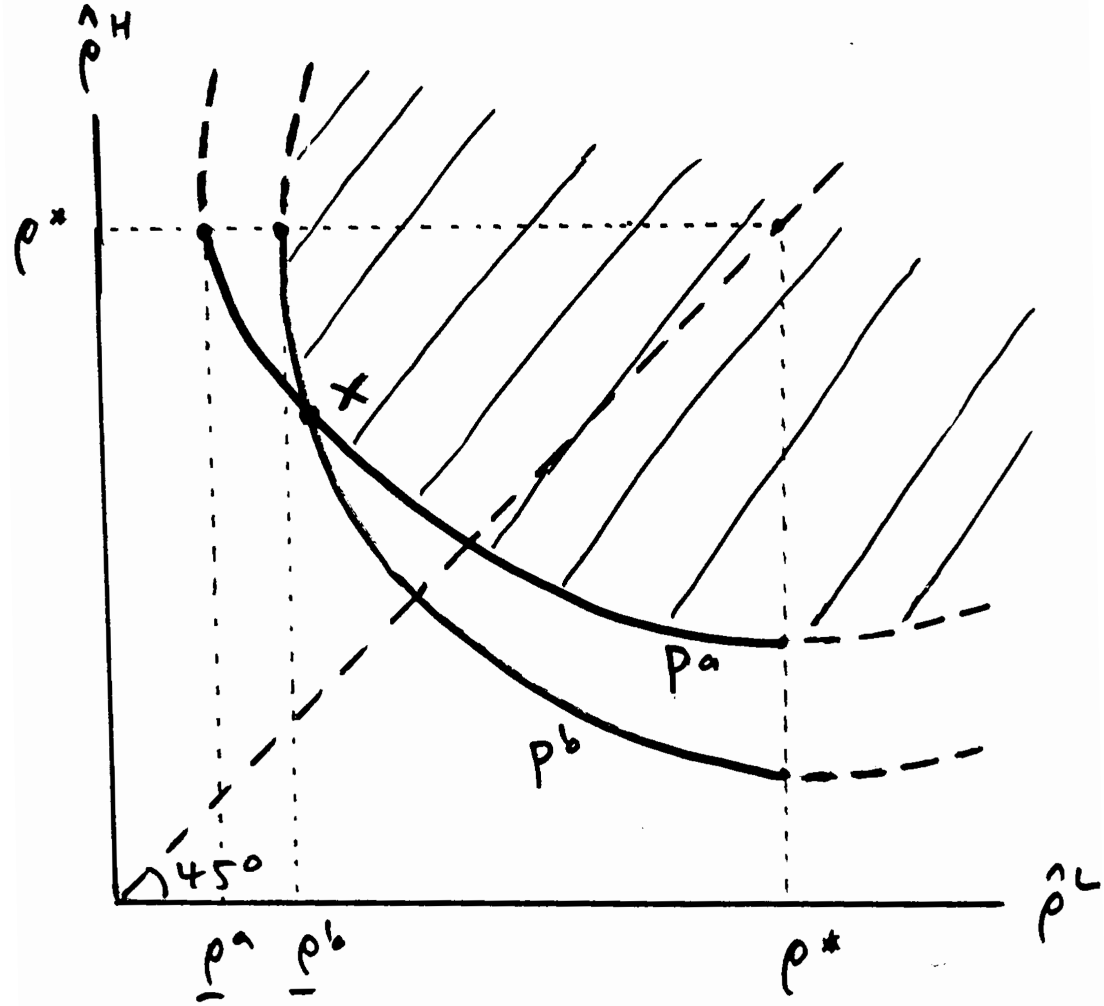

Figure 2

Curves P^a and P^b

shaded area gives set of
two-price equilibrium prices

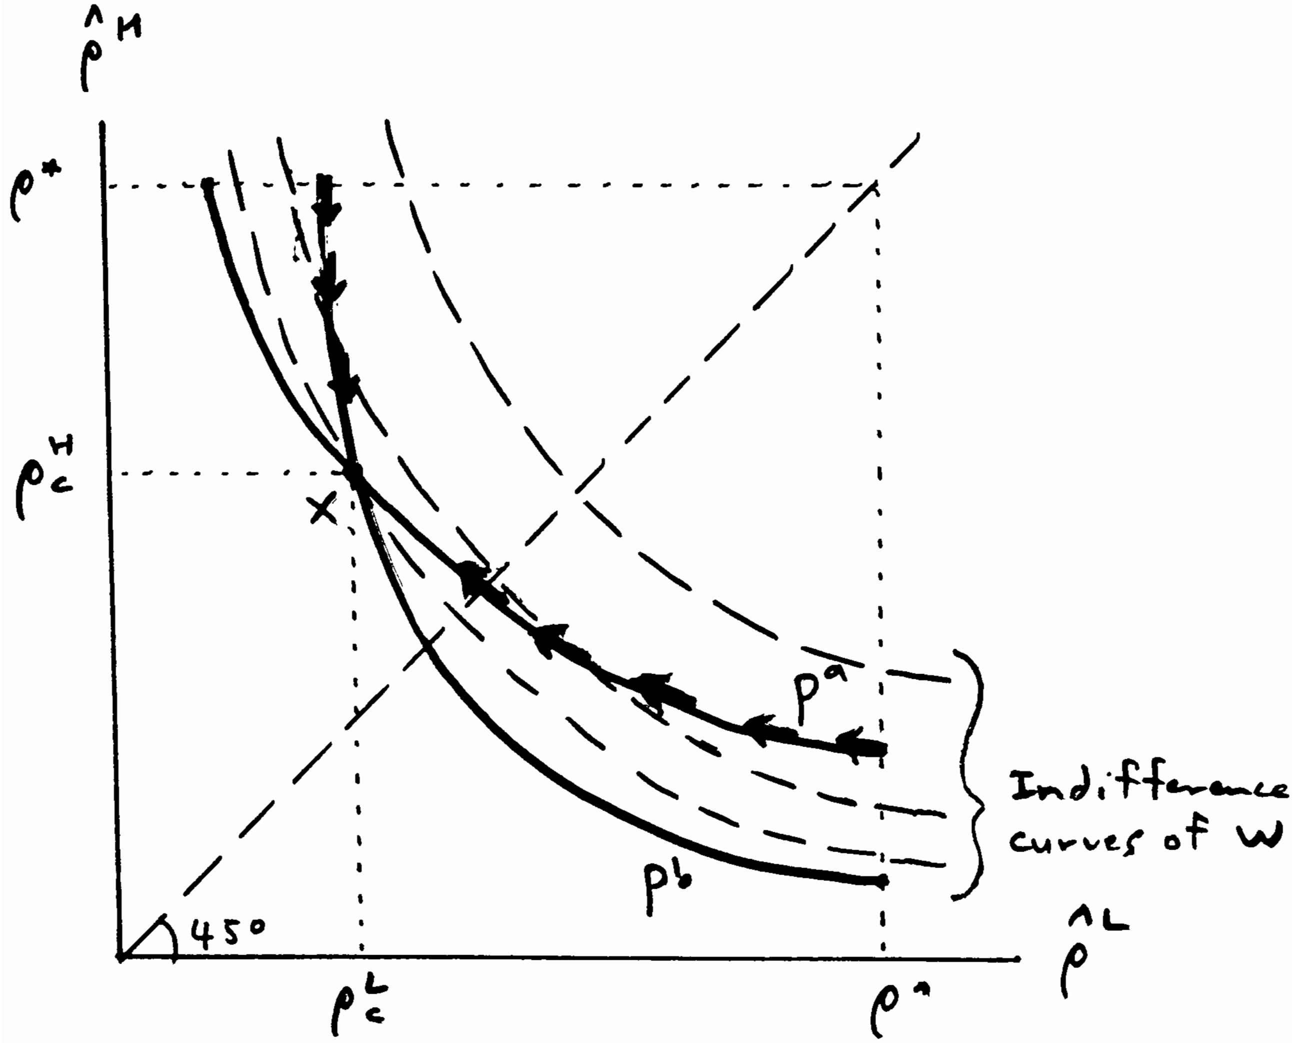

Figure 3

Derivation of Consumer-Optimal Two-Price Equilibria

Arrows show direction of increasing consumer preference