JOURNAL OF ECONOMIC THEORY 60, 140–158 (1993)

Search Market Equilibrium, Bilateral Heterogeneity, and Repeat Purchases*

ROLAND BENABOU

Massachusetts Institute of Technology, Cambridge, Massachusetts 02139 and National Bureau of Economic Research, Cambridge, Massachusetts 02138

Received November 12, 1990; revised March 24, 1992

We develop a general model of search market equilibrium with heterogeneous buyers and sellers. This framework unifies previous models with one-sided heterogeneity and clarifies many of their special properties. It easily accounts for price dispersion, active search, and the matching of buyer and seller types. It also extends to repeat purchases, once we embed it into a dynamic game with incomplete information. We formalize the inferences and strategies underlying equilibria where firms charge constant prices and customers patronize them loyally. We establish a general correspondence between such equilibria and single purchase markets. *Journal of Economic Literature* Classification Numbers: D83, L13.

INTRODUCTION

The following stylized facts seem descriptive of many markets: (a) there is significant diversity among buyers and sellers, allowing in particular less efficient firms to persist; (b) sellers with higher costs charge higher prices, and more of their customers have high opportunity costs of seeking alternatives; (c) for goods which are purchased repeatedly, buyers do not search each time but invest in an initial thorough search for a long-term supplier.

This paper develops a general model of search market equilibrium with heterogeneity among both buyers and sellers. It serves three purposes. The first is to account for the stylized facts described above. The second is to develop a unifying framework in which the special features of models with one-sided heterogeneity can be better understood and combined. The third is to establish general links between single and repeat purchase markets.

A market with identical consumers and firms is subject to Diamond's [6] paradox: as long as the search cost is positive, the unique equilibrium has all firms charging the monopoly price. The sequential search literature

* I am grateful to Patrick Bolton, Peter Diamond, Pierre Picard, and Jean Tirole for helpful comments. All errors are mine.

0022-0531/93 $5.00

has resolved this problem by allowing heterogeneity either in search costs or in production costs. Each approach, however, has drawbacks.

Models with buyer heterogeneity (Axell [1], Von zur Muehlen [15], Rob [13], Stiglitz [14]) generate search and price dispersion, with firms indifferent between all prices. Observation suggests that high and low price sellers generally differ, especially with respect to input costs. But more troublesome than mixed strategies is the fact that the nature of equilibrium—and especially price dispersion—depends critically on the assumed shape of the distribution of search costs near zero.[1] This type of assumption, lacking empirical content, seems like a fragile base on which to rest an explanation of equilibrium search and price dispersion.

Reinganum's [12] model of heterogeneity in firms' costs does not have these problems, since sellers' cost differences generate price dispersion for any value of buyers' common search cost. On the other hand, search plays a rather limited role. Sellers with cost below a critical level charge their monopoly price, as in Diamond [6], while those with higher cost bunch at consumers' reservation price; there is then no search in equilibrium.

These two types of models are clearly complementary. This paper combines them into a unifying framework which provides more adequate and robust explanations of equilibrium price dispersion, search, and entry.[2] It also extends the analysis to repeat purchases, modelled as a dynamic game with incomplete information between buyers and sellers. This sheds light in particular on the inference and strategies which lead firms to charge constant prices and customers to purchase loyally.

Section I presents the static model and derives a general characterization of equilibria as functional fixed-points. Section II studies repeat purchases and establishes an equivalence between Bayesian perfect equilibria with stationary outcomes and equilibria of the single purchase market. All proofs are gathered in appendix at the end of the paper.

I. Search with Bilateral Heterogeneity

1. *Sellers and Buyers*

There is a continuum of potential producers, with constant marginal costs distributed on $[\underline{c}, \bar{c}]$, $0 \leqslant \underline{c} \leqslant \bar{c} \leqslant +\infty$, according to the cumulative

[1] An atom, a positive density, a zero but increasing density or an interval of zero density at the origin lead to completely different results; see Rob [13].

[2] Its purpose is thus *not* to obtain price dispersion under minimal assumptions. Benabou [3] shows that Diamond's paradox can be resolved without any heterogeneity, provided that frictions affect not only buyers (search cost) but also sellers (costly price adjustment), even in arbitrarily small amounts.

distribution function (cdf) $G(c)$. There is a continuum of consumers, and each of them derives indirect utility $V(p)$ ($V' \leqslant 0$, $V'' \geqslant 0$) from buying the good at price p. Consumers know the support $[\underline{p}, \bar{p}]$ and distribution $F(p)$ of equilibrium prices in the market, but not what individual sellers charge. They get a first price quotation for free; all others require costly search. Buyers differ only in their search costs, which are distributed on $[\underline{\sigma}, \bar{\sigma}]$, $\underline{\sigma} < \bar{\sigma} \leqslant +\infty$, with cdf $Q(\sigma)$.

With little loss of generality, we assume that Q has a right-continuous density $q(\sigma) > 0$ on $(\underline{\sigma}, \bar{\sigma})$, and that $x(p) \equiv -V'(p) > 0$ within the equilibrium range. Since $x(p)$ is what consumers would buy from a monopolist, absent search opportunities, we call it "monopoly demand." We take the monopolist's problem to be well-behaved: letting $e_x(p) = -px'(p)/x(p)$, $p(1 - 1/e_x(p))$ is increasing. A firm with cost c has monopoly price $p_m(c) \leqslant +\infty$. We denote the left and right limits of a function ϕ at some point a as $\phi(a^-)$ and $\phi(a^+)$.

2. *Search*

There is a single period, during which prices remain fixed and all search takes place. We make the standard assumption that wealth effects from search costs are negligible. A consumer's optimal search rule is then given by a reservation price r, equating her cost σ and expected gain $\Gamma_F(r)$ from search:

$$\Gamma_F(r) \equiv \int_0^r [V(p) - V(r)]\, dF(p) = \int_0^r x(p)\, F(p)\, dp = \sigma \tag{1}$$

provided $\sigma < \Gamma(+\infty)$, so (1) has a solution. For $\sigma \geqslant \Gamma(+\infty)$, any offer is preferable to search. Thus in all cases the reservation price is

$$R_F(\sigma) \equiv \sup \left\{ r \geqslant 0 \mid \Gamma_F(r) = \int_0^r x(p)\, F(p)\, dp \leqslant \sigma \right\}. \tag{2}$$

When offered $p \leqslant R_F(\sigma)$, the consumer accepts and buys $x(p)$ units. The lowest and highest reservation prices will be denoted as $\underline{r}_F \equiv R_F(\underline{\sigma})$ and $\bar{r}_F \equiv R_F(\bar{\sigma})$, $\underline{r} \leqslant \bar{r} \leqslant +\infty$. The subscript F will be dropped when no confusion results.

3. *Demand.*

Let $1/\theta$ be the equilibrium density of operating firms (per consumer) in the market. We now derive their demand curve by aggregating consumers' search rules. Consider first those with $\sigma < \Gamma(+\infty)$. By (1), the density $v(r)$ of their reservation prices on $[\underline{r}, \bar{r}]$ is

$$v(r) = q(\Gamma(r))\, \Gamma'(r) = q(\Gamma(r))\, F(r)\, x(r). \tag{3}$$

Consumers with $R_F(\sigma) \in [r, r+dr]$ each have probability $F(r)$ of being successful in any single search. So by the law of large numbers, each firm is visited by $\theta \cdot v(r) \cdot dr$ of them on their first search, $\theta \cdot v(r) \cdot (1-F(r)) \cdot dr$ on their second search, etc.; hence a total of $\theta \cdot v(r)/F(r)\, dr = \theta \cdot q(\Gamma(r)) \cdot x(r) \cdot dr$ such visitors. Summing up those who accept an offer of p, plus the non-searchers ($r = +\infty$, or $\sigma > \Gamma(+\infty)$), yields the number of customers $N_F(p)$ and demand function $D_F(p)$:

$$D_F(p) \equiv x(p)\, N_F(p) \equiv \theta x(p) \left[\int_{\max\{p, \underline{r}\}}^{\infty} q(\Gamma(r))\, x(r)\, dr + \int_{\Gamma(+\infty)}^{\infty} q(\sigma)\, d\sigma \right]. \tag{4}$$

Because preferences affect reservation prices, $D_F(p)$ is not just the product of $x(p)$ and the demand curve from consumers with linear utility (contrary to the formula used by Axell [1]). We also see that for $p < \underline{r}$, $D_F(p)$ is simply proportional to monopoly demand $x(p)$, as in Reinganum [12], where $\underline{r} = \bar{r} = \bar{p}$. Finally, D_F is almost everywhere differentiable, and its kinks correspond to the left discontinuities of q; we denote its elasticity as $e_F(p)$.

4. *Pricing*

Since only firms with $c \leqslant \bar{r}$ can operate profitably, we have:[3]

$$1/\theta = G(\bar{r}) = G(R_F(\bar{\sigma})). \tag{5}$$

These firms maximize $\pi_F(p, c) \equiv (p-c)\, D_F(p)$ over $[c, \bar{r}]$. For $c < \bar{r}$, any solution p^* is interior.[4] Assuming for the moment that π is twice differentiable at p^*, the usual first- and second-order conditions must hold,

$$C_F(p) = c, \qquad C'_F(p) \geqslant 0, \tag{6}$$

where $C_F(p) \equiv p(1 - 1/e_F(p))$ is marginal revenue with respect to output, as a function of price: for any $p \in (0, \bar{r}_F)$,

$$C_F(p) = p \left[1 - \left[e_x(p) + \frac{q(\Gamma(p))\, x(p)}{\int_p^{\infty} q(\Gamma(r))\, x(r)\, dr + \int_{\Gamma(+\infty)}^{\infty} q(\sigma)\, d\sigma} \right]^{-1} \right]. \tag{7}$$

5. *Equilibrium*

In equilibrium, F must be consistent with each operating firm charging a solution $p(c)$ to (6). Suppose for the moment that C_F is increasing and

[3] Allowing for the case where $\bar{r}$ coincides with an atom of G, the most general form of (5) is: $G(\bar{r}^-) \leqslant 1/\theta \leqslant G(\bar{r}^+) = G(\bar{r})$. We shall assume that when they are indifferent between entering and staying out, firms choose to enter.

[4] This assumes $\bar{r} < +\infty$, for expositional simplicity; the theorems hold for $\bar{r} \leqslant +\infty$.

continuous (the standard case of marginal revenue falling continuously with output). By (7), $p(1-1/e_x(p)) \leqslant C_F(p) \leqslant p$, with equality if and only if, respectively, $p < \underline{r}$ and $p = \bar{r}$; moreover, $C_F(0^+) \leqslant 0$.[5] Therefore (6) defines for all $c \leqslant \bar{r}$ a unique optimal price $p_F(c) = C_F^{-1}(c)$, and

$$c \leqslant p_F(c) \leqslant p_m(c), \tag{8}$$

with equality if and only if, respectively, $c = \bar{r}$ and $c < C_F(\underline{r})$. Finally, since C_F is increasing, the resulting price distribution is Φ_F defined below.

Definition. For any $F\colon R_+ \to R_+$, define Γ_F as in (1), C_F as in (7), $\bar{r}_F = \sup\{r \geqslant 0 \mid \Gamma_F(r) \leqslant \bar{\sigma}\}$, and $\Phi_F\colon R_+ \to [0, 1]$ by

$$\begin{aligned} \Phi_F(p) &= \frac{G(C_F(p))}{G(\bar{r}_F)} \qquad &&\text{for} \quad p < \bar{r}_F \\ \Phi_F(p) &= 1 &&\text{for} \quad p \geqslant \bar{r}_F. \end{aligned} \tag{9}$$

In general, $C_F(p)$ may be discontinuous, where $q(\Gamma_F(p))$ is.[6] Optimality then requires that it cross the horizontal $C_F = c$ from below. Similarly, C_F need not be increasing. But since $F(p) = G(C_F(p)) \cdot G(\bar{r}_F)$ is non-decreasing, if $p_1 < p_2$ and $C_F(p_1) \geqslant C_F(p_2)$ then G must then be constant on $[C_F(p_2), C_F(p_1)]$, and F must be constant on $[p_1, p_2]$. Thus only a negligible set of firms have cost in $(C_F(p_2), C_F(p_1)]$, or charge prices in $(p_1, p_2]$. This will allow us to extend the above reasoning and derive a general characterization of equilibrium.

We shall call a function strongly quasi-concave if it is strictly quasi-concave and has no sadlepoint extremum. This property defines functions for which the first-order condition is sufficient for a global maximum.

Theorem I.1. *Let G be continuous. A cdf F on R_+ is an equilibrium price distribution, where almost all firms' profit functions are strongly quasi-concave if and only if it is a fixed point of the functional mapping $\Phi\colon F \to \Phi_F$.*

Proof. In Appendix.

For a monopolist, the first-order condition $p = 1/(1 - 1/e_x(p))$ is a one-dimensional fixed-point problem; the standard second-order condition is

[5] Again we focus the exposition on the case $\bar{r} < +\infty$ ($\Gamma(+\infty) > \bar{\sigma}$). Then as $p \to \bar{r}$, the search elasticity $e_s(p) \equiv px(p)\, q(\Gamma(p)) / \int_p^\infty q(\Gamma(r))\, x(r)\, dr \to +\infty$, so $C_F(p) \to \bar{r}$, if $q(\bar{\sigma}^-) > 0$. If $q(\bar{\sigma}^-) = 0$, the minor assumption that q is nonincreasing near $\bar{\sigma}$ ensures the same result, since $e_s(p) > 1/(\bar{r} - p)$. In the case $\bar{r} = +\infty$, as $p \to +\infty$ $e_s(p)$ becomes proportional to $px(p)$. So if $\limsup_{p \to +\infty} e_x(p) < 1$, $C_F(p)/p \to 1$, and $p_F(c)/c \to 1$; if $\liminf_{p \to +\infty} e_x(p) > 1$, $C_F(p)/[p(1 - 1/e_x(p))] \to 1$, and $p_F(c)/p_m(c) \to 1$.

[6] For instance, at $C_F(\underline{r})$ if σ is uniformly distributed on $[\underline{\sigma}, \bar{\sigma}]$, $\underline{\sigma} > 0$.

that at any solution, marginal revenue be increasing in price, making profits strictly quasi-concave. Theorem I.1 generalizes these requirements to the space of distributions on R_+. Interestingly, the global monotonicity of the cdf is equivalent, given $\Phi_F = F$, to the local monotonicity of marginal revenue $C_F(p)$ at almost every firm's solution to its first-order condition.

When G has discontinuities, a somewhat weaker result applies, as it could be that an atom of firms are indifferent between several prices.

THEOREM I.2. *For any distribution G, the condition $\Phi_F = F$ remains necessary for an equilibrium where almost all firms have strongly quasi-concave profits. It is also sufficient for F to define an equilibrium with quasi-concave profits.*

Proof. In Appendix.

The first result covers, for instance, the model of Carlson and McAfee [5], where a discrete distribution of firms face strictly concave profit functions. Theorem I.2 could be extended to equilibria where a positive measure of firms randomize over an interval of profit-maximizing prices (Axell [1], Von zur Muehlen [15], Rob [13], Stiglitz [14]), by looking at fixed points of a correspondence instead of a mapping.

6. *Discussion*

Figure 1 describes the essential features of an equilibrium.[7] Firms with cost above $\bar{r}_F$ cannot sell profitably and so stay out of the market. Firms with cost below $c_0 \equiv C_F(\underline{r}_F^-)$ can charge their monopoly price ($p_F(c) = p_m(c)$) without inducing any consumer to leave. All others are constrained by search to price below their monopoly level: $c < p_F(c) < p_m(c)$.

A discontinuity in q at some point σ (in the case of Fig. 1, at $\underline{\sigma}$) causes a similar discontinuity in marginal cost C_F at $R_F(\sigma)$ (here at $\underline{r}_F$). This leads to a *bunching* of prices by a whole segment $[c_0, c_1] = [C_F(\underline{r}_F^-), C_F(\underline{r}_F)]$ of firms.[8] Equivalently, a firm's price is insensitive to cost variations in the $[c_0, c_1]$ range; this is the source of the "rigidity" discussed by Stiglitz [14]. The combination of this bunching and the monopolistic behavior of firms with $c \leqslant c_0 = C_F(\underline{r}_F)$ is the central feature of the Reinganum [12] model.

A flat section in $C_F(p)$ yields a multivalued optimal price $[p^-(c_2), p^+(c_2)]$ for a point like c_2 on Fig. 1. If c_2 coincides with an atom of G, a positive mass of firms are *indifferent* between all prices in $[p^-(c_2), p^+(c_2)]$. This indeterminacy is the central feature of models with identical firms such as those of Axell [1], Von zur Muehlen [15], Rob [13], or Stiglitz [14].

[7] The figure is drawn for the case $\bar{r} < +\infty$. When $\bar{r} = +\infty$, $p_F(c)$ is asymptotic to c or to $p_m(c)$, depending on $e_x(p) \lesseqgtr 1$ as $p \to +\infty$ (see footnote 5).

[8] Only upward discontinuities are consistent with a well-behaved equilibrium.

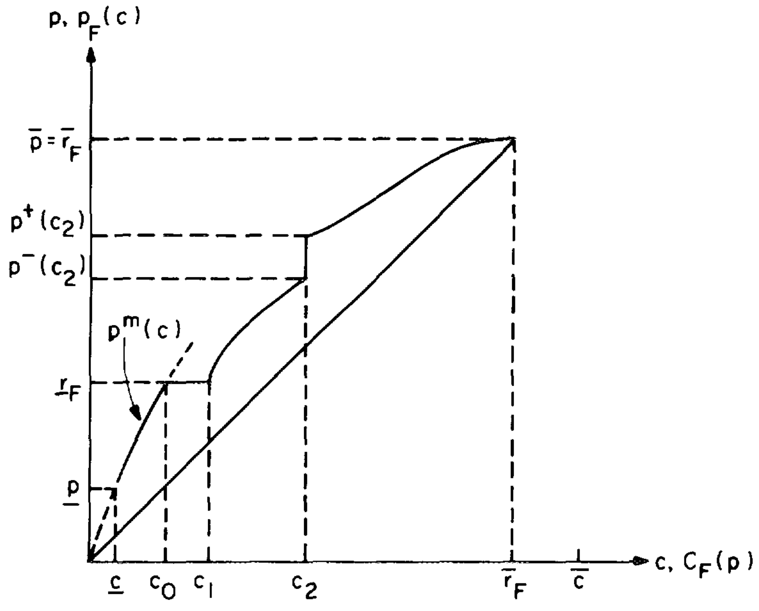

FIG. 1. The marginal revenue function $C_F(p)$ and the optimal pricing rule $p_F(c)$.

We shall not address here the issues of existence and uniqueness of equilibrium. Note from (7) that marginal revenue $C_F(p)$ need not be increasing for all F, so that in general Φ does not map cdf's into themselves. The monotonicity of $C_F(p)$ and $\Phi_F(p)$ must be built up from appropriate properties of the functions $V(p)$, $G(c)$, and $Q(\sigma)$. This is done in Benabou [4], where we derive several existence, uniqueness, and comparative statics results for uniformly distributed search costs. In particular, they generalize the model of MacMinn [10] and provide an extension of Reinganum [12] to the case of heterogeneous buyers (generating search in equilibrium) and entry.

II. Repeat Purchases

1. *General Issues*

Many goods, such as non-durables, are purchased repeatedly. Buyers do not search each time for a suitable seller, but rather invest in an initial search to find a long-term supplier. Job search in labor market shares the same feature. Implicit in repeat purchases is some inference of future prices from past ones: if yesterday's price was acceptable (respectively, too high) it can be expected that today's and tomorrow's will also be acceptable (respectively, too high); consumers can then (respectively, cannot) economize on search costs by coming back to the same seller.

The aim of this section is to formalize the inference and decision processes of agents in a repeat purchase market, in order to account for such long-term relationships and examine their consequences on equilibrium prices. McMillan and Morgan [11] also examine a market where prices are constant and consumers purchase loyally, but the two approaches are quite different. Their model explores a kind of price rigidity: identical firms are "locked" by consumers' repeat purchase behavior into an initial distribution of prices and clienteles, each yielding different profits. Their strategies form a Nash equilibrium in subgames at $t \geqslant 1$, but not in the full game starting at $t = 0$. They would also not be robust to an inflow of new buyers, which enables firms to rectify their customer base. We construct Bayesian perfect equilibria of a dynamic game between buyers and sellers, for any amount of consumer renewal.

2. *The Game*

Firms are indexed by $f \in [0, 1]$, with cost c^f; they set prices p_t^f in every period $t \in \mathbb{N}$. Buyers are indexed by $\sigma \in [\underline{\sigma}, \bar{\sigma}]$, as $Q(\sigma)$ is atomless. They live forever, have the same discount factor $\delta < 1$ as firms, and indirect utility $V(p)$ in each period. At the start of period $t = 0$, they each receive one price quotation for free; additional offers require search, with unit cost σ. In each period $t \geqslant 1$ they can costlessly return to previously encountered sellers, or sample new firms at random with cost σ. There is no limit on the number of searches which can be carried out within a period.

The cost of returning to a previously visited seller is an issue with important implications. With positive return costs, buyers' decisions would be different when faced with a given offer (e.g., $p = R_F(\sigma)$) and when considering whether to go back to a seller believed to charge that same price. As shown by Bagwell [2] for a monopolist with unknown cost, this gives firms an incentive to trick customers into coming back, then exploit them. The resulting nonstationary price strategies would be too difficult to incorporate here. We assume instead that return costs are zero. The search costs we consider thus correspond to resources required to find out about a seller (its location, the quality its merchandise, etc.). They are "inspection costs" rather than expenses incurred for every visit or purchase (trips or ordering costs).

3. *Equilibrium*

When setting its price, a firm does not know how many previous clients have decided to come back and check its price again. Buyers do not know a given firm's cost type, nor the price currently charged by previous suppliers. We look for Bayesian perfect equilibria (BPE) of this game of incomplete information between firms and buyers. We use Fudenberg and Tirole's [7] definition of BPE, as it is very similar to Kreps and Wilson's [9] sequential equilibrium which is not available for games with continua of types and actions. We thus require that: (i) a deviation by any player affects beliefs about its type only; (ii) beliefs obey Bayes' rule not only on the equilibrium path, but also starting from the new beliefs triggered by any deviation, and until the next event with zero subjective probability occurs.

The general game is quite complicated: firms' prices may vary over time, and consumers must now take into account entire expected price paths. We make two simplifications. First, when looking for equilibrium paths, we need only consider histories which do not involve simultaneous deviations by a positive measure of players (Gul, Sonnenschein, and Wilson [8]). Second, we restrict attention to BPE with *stationary prices* on the *equilibrium path* (strategies, however, are unrestricted). This not only keeps the analysis tractable, but also avoids a Folk-theorem multiplicity of equilibria.

A consumer's *initial* decision problem thus reduces to a once and for all search between the present values of constant price paths, distributed according to $F(p)$. By (2), the appropriate reservation price is

$$R_F^\delta(\sigma) \equiv R_F(\sigma(1-\delta)) \equiv \sup \left\{ r \geq 0 \;\middle|\; \int_0^r F(p)\, x(p)\, dp \leq \sigma(1-\delta) \right\}. \tag{10}$$

By the optimality principle, she keeps buying from any firm thus selected as long as its price remains constant, i.e., forever on the equilibrium path. Let $Q^\delta(\sigma) \equiv Q(\sigma(1-\delta))$ denote the distribution of scaled-down search costs, and q^δ its density. A constant price strategy p^f generates demand per period $D_F^\delta(p^f)$, given by (4) with q replaced by q^δ. If it is optimal, p^f must in particular maximize $\pi_F^\delta(p) = (p-c)\, D_F^\delta(p)/(1-\delta)$. Therefore:

THEOREM II.1. *Along the equilibrium path of any BPE with stationary prices*:

(i) *Each consumer σ initially searches with reservation price $R_F^\delta(\sigma)$ given by* (10), *then keeps buying from the same firm in all subsequent periods.*

(ii) *Each firm f charge $p_F^\delta(c^f)$, its optimal price in a single purchase equilibrium for production and search costs distributions $G(c)$ and $Q^\delta(\sigma)$.*

Consumers invest in a thorough first search, so as to benefit from a low price and economize on search costs in future periods. As a result, prices are lower, and price dispersion smaller, the more frequently the good is purchases.

We now turn to the converse result of Theorem II.1: any one-shot purchase equilibrium is the outcome of a BPE of the repeat purchase game, in which prices on the equilibrium path (but not outside) are stationary.

The equilibrium works as follows: as shown below, consumers patronize the same firm until it raises its price above their reservation level; if this occurs, they leave and never come back. Consider then firm f at time t, with information set H_t^f, and denote by P_{t-1}^f the highest price it ever charged in the past. If firm f ever exceeded its equilibrium price, it has no more customers with reservation price in $[p_F^\delta(c^f), P_{t-1}^f)$ and cannot attract any new ones (this will change when buyer renewal is introduced). It will therefore set $p_s = \min\{P_{t-1}^f, p_m(c^f)\}$ in all future periods. Conversely, if $P_{t-1}^f < p_F^\delta(c^f)$, some clients are not worth retaining; it will dismiss them by setting $p_s = p_F^\delta(c^f)$ from time t on. The optimal strategy at any H_t^f is, therefore,

$$\tilde{p}(H_t^f) \equiv \max[p_F^\delta(c^f), \min\{\max\{p_s^f \mid 0 \leqslant s \leqslant t-1\}, p_m(c^f)\}], \quad (11)$$

leading to constant prices. Letting $P_{-1}^f \equiv p_F^\delta(c^f)$ extends (11) to $t=0$.

Consider now buyer σ just after her nth search of period t. Her information set $H_{t,n}^\sigma$ consist of the "names" $I_{t,n}^\sigma$ of all firms previously encountered, and for each of them the history of prices observed when she was there. This includes her most recent and her highest observations at f, respectively $p_{t,n}^\sigma(f)$ and $P_{t,n}^\sigma(f)$. Consumer σ must form beliefs about firm f's cost type c^f. Given (11), a plausible assumption is that she believes with probability one most the recent price seen to be the firm's equilibrium price, i.e., that $c_f = C_F^\delta(p_{t,n}^\sigma(f))$. A price increase is thus taken as a signal of a higher cost than previously thought, and symmetrically for a price decrease. Given this belief and (11), consumer σ expects with probability one the price

$$EP_{t,n}^\sigma(f) = \min[P_{t,n}^\sigma(f), p_m(C_F^\delta(p_{t,n}^\sigma(f))] \geqslant p_{t,n}^\sigma(f) \quad (12)$$

in all future periods. If her last visit to firm f was at some $s \leqslant t-1$, she also believes that $EP_{t,n}^\sigma(f)$ was charged since then and is the current price. Note from (12) that consumers *never* expect price *decreases*, even though upward and downward deviations by a firm trigger completely *symmetric* revisions in their beliefs about its cost type. It is then optimal for them to leave and not come back if their reservation level is ever exceeded, as assumed earlier. This in turn ensures that firms never lower prices. To complete the construction of the equilibrium, let $J_{t,n}^\sigma = \{f \in I_{t,n}^\sigma \mid EP_{t,n}^\sigma(f) \leqslant R_f^\delta(\sigma)\}$ denote the set of firms where the price expected by consumer σ is below her reservation price, and let $K_{t,n}^\sigma = \arg\min\{EP_{t,n}^\sigma(f) \mid f \in J_{t,n}^\sigma\}$ denote the subset where it is minimal.

THEOREM II.2. *Let $G(c)$ and $Q^{\delta}(\sigma) \equiv Q(\sigma(1-\delta))$ generate an equilibrium $F(p)$ of the one-shot search market, with increasing pricing rule $p_F^{\delta}(c)$. The following strategies and beliefs then constitute a BPE of the dynamic game, with stationary path:*

(i) *At any information set H_t^f, firm f charges charges $\tilde{p}(H_t^f)$. It expects all consumers who bought from it at $t-1$ to visit it first in period t $(t \geq 1)$.*

(ii) *At any information set $H_{t,n}^{\sigma}$, consumer σ has beliefs distributed according to $G(c)$ and $F(p)$ over the costs and prices of firms $f' \notin I_{t,n}^{\sigma}$. She believes with probability one each firm $f \in I_{t,n}^{\sigma}$ to be of cost type $C_F^{\delta}(p_{t,n}^{\sigma}(f))$ and to have charged $EP_{t,n}^{\sigma}(f)$ in all periods $s \leq t$ when she did not visit it. Her strategy is given by the following program:*

(a) *If $J_{t,n}^{\sigma} = \varnothing$, she samples at random, then updates $H_{t,n}^{\sigma}$ to $H_{t,n+1}^{\sigma}$;*

(b) *If $J_{t,n}^{\sigma} \neq \varnothing$, she returns to any firm $f \in K_{t,n}^{\sigma}$ and buys there if:*

$$p_t^f \leq \min\{EP_{t,n}^{\sigma}(f') \mid f' \in J_{t,n}^{\sigma} \backslash \{f\}\}; \tag{13}$$

when $J_{t,n}^{\sigma} \backslash \{f\} = \varnothing$, the rhs is taken to mean $R_F^{\delta}(\sigma)$. If (13) does not hold, she pursues the program after using p_t^f to update $H_{t,n}^{\sigma}$ to $H_{t,n+1}^{\sigma}$.

Proof. In Appendix.

A consumer thus first goes back to the firm which she expects to have the lowest price $EP_{t,n}^{\sigma}(f) \leq R_F^{\delta}(\sigma)$. If its actual price p_t^f is still the lowest she knows of, she buys there again. If not, she goes to the next most attractive known firm, unless there is none where she expects a price below $R_F^{\delta}(\sigma)$. In that case, she starts sampling at random, until one is found.

The stationarity of the outcome relies importantly on consumers' infinite horizon and the possibility of costless return. For contrast, consider Bagwell's [2] model of introductory pricing, with only two periods and costly return. High cost firms have an incentive to initially charge low prices, so as to deceive consumers into incurring the return cost and exploit them in the last period. In our model this strategy is unprofitable: repeat customers discovering too high a price will leave without buying and will look for another seller, correctly expecting the current one to go on charging high prices.

While player's strategies are optimal at all information sets, close scrutiny reveals that off the equilibrium path buyers use a weakly dominated strategy (wds). Instead of buying from the firm believed with probability one to have the lowest price, they could first visit all known firms (at zero cost) to find out their actual prices. This is a minor caveat, as the players using a wds are not the ones signalling their types, i.e., not firms but consumers. Moreover, their out-of-equilibrium beliefs are quite reasonable (see proof of Theorem II.2). If buyers did keep informed of all previously visited sellers' prices, firms could "recall" lost customers through price cuts. An equilibrium with variable prices and consumers repeatedly switching between sellers would be intractable; an equilibrium with stationary prices would have the same outcome (by Theorem II.1) as the BPE constructed here.[9]

4. *Consumer Renewal*

Allowing for inflows of new buyers into the market is interesting for two reasons—in addition to realism. First, this is a desirable robustness property for an equilibrium. It prevents, for instance, the initial buyers' repeat purchase behavior from locking firms into prices which were not optimally chosen at $t=0$, as in McMillan and Morgan [11]. More generally, it allows firms to attract new customers by lowering their price and, in particular, to reverse in the long run the effects of a price increase, contrary to (11). We assume therefore that in every period each consumer leaves the market with probability $\mu \in (0, 1]$, while a measure $\lambda \geqslant 0$ of new buyers enter. Thus $\lambda/(1-\mu)$ measures the speed of consumer renewal, and λ/μ is the measure of customers in a steady-state market (on which we focus).

If firms charge constant prices, all cohorts have the same distribution of reservation prices. Along the equilibrium path a firm then has no incentive for intergenerational price discrimination. This suggests that the same type of BPE *path* as in Theorem II.1 should still prevail, with consumers searching only upon entering the market. What makes the problem in fact much more complicated is that *off* the equilibrium path, a firm does want to *discriminate between generations*. Whereas following a deviation to $p > p_F^{\delta}(c^f)$ it would previously not have benefited from lowering back its price (see (11)), it now has conflicting incentives to maintain it, because its existing customers have high reservation prices, and to lower it, so as to attract new clients in the optimal mix corresponding to its cost type. We show below that following any price history, the optimal strategy is bounded between those which correspond to the cases of no renewal ($\lambda = 0$), namely pricing according to (11), and complete renewal ($\lambda = +\infty$), namely charging the optimal static price.

LEMMA 1. *Assume that almost all firms f charge $p_F^{\delta\mu}(c^f)$ in every period and that consumers have the same strategies as in Theorem II.2, with δ*

[9] Also, when going back to any $f \in I_{t,n}^{\sigma}$ has cost $\varepsilon > 0$, the strategies and beliefs of Theorem II.2 form an ε-equilibrium. As $\varepsilon \to 0$, we obtain our BPE.

replaced by $\delta\mu$. For a firm at an information set H_t^f, the expected present value of profits from any sequence of prices $\{p_s \mid s \geqslant t\}$ is then[10]

$$\pi(\{p_s \mid s \geqslant t\} \mid H_t^f) \tag{14}$$
$$= \theta\lambda \sum_{s=t}^{+\infty} \delta^{s-t}(p_s - c^f)\, x(p_s) \sum_{\tau=0}^{+\infty} (1-\mu)^\tau N_F^{\delta\mu}(\max\{p_k \mid s-\tau \leqslant k \leqslant s\}),$$

where $p_k \equiv p_k^f$ for $k \leqslant t-1$. Moreover, there exists a sequence $\{p_s^*(H_t^f) \mid s \geqslant t\}$ maximizing this present value, with, for all $s \geqslant t$:

$$p_F^{\delta\mu}(c^f) \leqslant p_s^*(H_t^f)$$
$$\leqslant \max[p_F^{\delta\mu}(c^f), \min\{\max\{p_s^f \mid 0 \leqslant s \leqslant t-1\}, p_m(c^f)\}]. \tag{15}$$

Proof. In Appendix.

Note that if the equilibrium price $p_F^{\delta\mu}(c^f)$ was charged in all past periods, it remains optimal in the future. Moreover, firm f will never cut its price below $p_F^{\delta\mu}(c^f)$ in an attempt to attract new customers; but if it had exceeded $p_F^{\delta\mu}(c^f)$ earlier, it may go back down to that level in the long run. Yet if buyers take, as before, the last observation to be the equilibrium price for the firm's true cost type, they still never expect the price to decrease. So they still leave, not to come back, when their reservation price is exceeded. Lemma 1 thus serves as the foundation for the following result.

THEOREM II.3. *The results of Theorems* II.1 *and* II.2, *showing the equivalence between repeat purchase BPE with stationary paths and equilibria of the one-shot market, still hold with buyer renewal, provided that δ is replaced by $\delta\mu$ and firms' price strategy* (11) *by $p_s^*(H_t^f)$ defined in Lemma* 1.

Proof. In Appendix.

CONCLUSION

This paper has developed a general model of search market equilibrium with heterogeneous buyers and sellers, in which the special properties of many previous models are easily understood and combined. The analysis extends to repeat purchases, explaining the inferences and strategies which underlie equilibria, where firms charge constant prices and customers patronize them loyally.

[10] Equation (14) is written for a market which has been operating since $t = -\infty$. If it only starts at $t = 0$, (14) still holds, provided $p_k \equiv p_0^f$ for $k < 0$.

The next step is to incorporate stochastic shocks, which in real markets cause some prices to change and some buyers to switch suppliers. While such a model is bound to be complex, the methods and results derived here should provide a useful starting point.

Appendix

Proof of Theorems I.1 *and* I.2. We drop the subscript F from C_F, $\bar{r}_F$, etc. We denote by μ_G the measure induced by G over firms and define for any $c < \bar{r}$:

$$A(c) = \{p \in [0, \bar{r}) \mid C(p) \geqslant c\}; \qquad B(c) = \{p \in [0, \bar{r}) \mid C(p) \leqslant c\}. \qquad \text{(A1)}$$

Then $\pi(\cdot \mid c)$ is quasi-concave if and only if there exists a $p^* \in [0, \bar{r}]$ such that $[c, p^*] \subset B(c)$ and $[p^*, \bar{r}) \subset A(c)$. It is strictly quasi-concave if, moreover, $A(c) \cap B(c)$ contains no finite interval; p^* is then unique. It is strongly quasi-concave if in fact $A(c) \cap B(c) \subset \{p^*\}$; then $A(c) = [p^*, \bar{r})$.

A. *Necessity of the fixed point condition.* Consider an equilibrium where, for μ_G-all c, $\pi(\cdot \mid c)$ is strongly quasi-concave. For such a c, the unique optimal price is $p(c) = p^*(c)$, so $[p(c), \bar{r}) = A(c)$ and $\mu_G(\{c \mid p(c) \leqslant p\}) = G(C(p))$. Since only $c \leqslant \bar{r}$ participate, $F(p)$ must then equal $G(C(p))/G(\bar{r}) = \Phi_F(p)$.

B. *Sufficiency of the fixed point condition.* We showed in the text that if $\Phi_F = F$:

$$\text{if} \quad p_1 \leqslant p_2 \quad \text{and} \quad C(p_1) \geqslant C(p_2), \quad \text{then} \quad \mu_G((C(p_2), C(p_1)]) = 0. \qquad \text{(A2)}$$

Let us now show that for μ_G-all c, there exists a $p^* = p^*(c)$ such that $A(c) = [p^*, \bar{r})$; p^* is then necessarily unique. Indeed, if $A(c)$ is not of this form there must exist p_1, p_2, with $p_1 < p_2$ and $C(p_1) \geqslant c > C(p_2)$. Then (A2) implies $\mu_G((C(p_2), C(p_1)] = 0$, so $\mu_G((c - \varepsilon, c]) = 0$ for $\varepsilon > 0$.

Thus for μ_G-all c, $\pi(\cdot \mid c)$ is quasi-concave and reaches its global maximum at $p^*(c)$ (and perhaps elsewhere). Let all firms with $c \leqslant \bar{r}$ charge $p^*(c)$ and the others stay out. The fraction of prices no greater than p is then $G(C(p))/G(\bar{r}) = F(p)$, so this is an equilibrium. This concludes the proof of Theorem I.2.

To conclude the proof of Theorem I.1, we show that if G is continuous $\pi(\cdot \mid c)$ is *strongly* quasi-concave, for μ_G-all c: $\mu_G(\{c \mid \exists \hat{p} \neq p^*, \hat{p} \in A(c) \cap B(c)\}) = 0$. Clearly, we can focus on the case where $A(c) = [p^*(c), \bar{r})$. Then we must have $C(p) = c$ and $\hat{p} > p^*$, so either:

(a) $C(p)=c$ on $(p^*, \hat{p}]$. But $S \equiv \{c \mid \exists p_1, p_2: p_1 < p_2$ and $C(p)=c$ on $[p_1, p_2]\}$ cannot contain an interval of positive μ_G-measure: if $[c', c''] \subset S$, then for each $c \in [c', c'']$, $C(p)$ is constant on some $[p_1(c), p_2(c)]$, and so is $F(p)=G(C(p))$. So the non-decreasing function $F(\cdot)$ must assume each value in $[G(c'), G(c'')]$ over a non-trivial interval. This can only be if $G(c')=G(c'')$.

(b) $\hat{p}>p^*$, and there is some $p_1 \in (p^*, \hat{p})$ with $C(p_1)>c=C(\hat{p})$. By (A2), $\mu_G((c, C(p_1)]) = 0$. Since G is continuous, this implies $\mu_G([c, C(p_1)])=0$, or $\mu_G([c, c+\varepsilon)])=0$ for some $\varepsilon>0$. Q.E.D.

Proof of Theorem II.2. Given consumers' point expectations and zero return costs, it is clear that their proposed strategy maximizes expected utility. The formal verification is straightforward (see Benabou [4]). We now turn to firm's strategies. At any H_s^f, $s \geqslant t$, the distribution of reservation prices among a firm f's clientele is that of $t=0$, truncated below $P_{s-1} \equiv \max\{p_k \mid k \leqslant s-1\}$; we omit here the superscript f from prices. Its demand function is then $D_s^f(p)=x(p) \cdot N(\max\{p, P_{s-1}\})=x(p) \cdot N(P_s)$, where $N(\cdot) \equiv N_F^\delta(\cdot)$ is given by (4). At H_t^f, the firm thus chooses $\{p_s \mid s \geqslant t\}$ to maximize:

$$\pi(\{p_s \mid s \geqslant t\} \mid H_t^f)=\sum_{s=t}^{\infty} \delta^{s-t}(p_s-c)\, x(p_s)\, N(P_s), \tag{A3}$$

with $c=c^f$. Now, denoting $p(c)=p_F^\delta(c)$,

$$(p_s-c)\, x(p_s)\, N(P_s) \leqslant (p_s-c)\, x(p_s)\, N(p_s) \leqslant (p(c)-c)\, D(p(c)) \tag{A4}$$

with double equality if and only if $p_s=P_s=p(c)$. So if $P_{t-1} \leqslant p(c)$, the optimal strategy is to set $p_s=p(c)$; hence $P_s=p(c)$ for all future s. If $P_{t-1}>p(c)$, on the contrary, three cases must be distinguished:

(a) If $P_s \leqslant p_m(c)$, then $(p_s-c)\, x(p_s)\, N(P_s) \leqslant (P_s-c)\, x(P_s)\, N(P_s) \leqslant (P_{t-1}-c)\, x(P_{t-1})\, N(P_{t-1})$, since $(p-c)\, x(p)$ increases up to $p_m(c)$, while $(p-c)\, D(p)$ decreases beyond $p(c)$.

(b) If $P_s>p_m(c) \geqslant P_{t-1}$, then $(p_s-c)\, x(p_s)\, N(P_s) \leqslant (p_m(c)-c)\, x(p_m(c))\, N(P_s) \leqslant (p_m(c)-c)\, x(p_m(c))\, N(p_m(c)) \leqslant (P_{t-1}-c)\, D(P_{t-1})$, since $(p-c)\, D(p)$ decreases beyond $p(c)<P_{t-1}$.

In both cases, the most the firm can obtain in any period $s \geqslant t$ is $(P_{t-1}-c)\, D(P_{t-1})$; this is obtained by setting $p_s=P_{t-1}=P_s$ for all $s \geqslant t$.

(c) If $P_{t-1}>p_m(c)$, then $(p_s-c)\, x(p_s)\, N(P_s) \leqslant (p_m(c)-c)\, x(p_m(c))\, N(P_s) \leqslant (p_m(c)-c)\, x(p_m(c))\, N(P_{t-1})$.

The most the firm can obtain in any period is thus obtained by setting $p_s=p_m(c)$, $P_s=P_{t-1}$ for all $s \geqslant t$. This concludes the proof of the optimality of (11).

Finally, we consider beliefs. On the equilibrium path, they follow Bayes' rule and are fulfilled with probability one. Off the equilibrium path, they constitute "reasonable assessments" as defined in Fudenberg and Tirole's [7] BPE: (i) a firm's deviation only affects its customers' beliefs about its own cost type; (ii) until they again observe a zero (subjective) probability price, customers' beliefs about the firm's cost type and future prices are again consistent with Bayes' rule and firms' strategies (11). Q.E.D.

Proof of Lemma 1. Firm f's clientele at H_s^f consists of the buyers who joined at dates $s-\tau$, $0 \leqslant \tau \leqslant s$, have "survived" since, and have reservation price $R_F^{\delta\mu}(\sigma) \geqslant P_{s-\tau,s}^f \equiv \max\{p_k^f \mid s-\tau \leqslant k \leqslant s\}$. In each cohort the number of such consumers per firm is $\theta\lambda\, N_F^{\delta\mu}(P_{s-\tau,s}^f)$, with $N_F^{\delta\mu}(\cdot)$ given by (4). We now drop the indices f, δ, μ, F from c^f, P^f, $p_F^{\delta\mu}$, $N_F^{\delta\mu}$, etc. Summing up over all cohorts between 1 and s and adding the initial (time zero) measure λ/μ yields

$$N(p_s \mid H_s^f) = \theta\lambda \sum_{\tau=0}^{s-1} (1-\mu)^\tau N(P_{s-\tau,s}) + \theta(\lambda/\mu)^s (1-\mu)^s N(P_{0,s})$$

$$= \theta\lambda \sum_{\tau=0}^{\infty} (1-\mu)^\tau N(P_{s-\tau,s}),$$

with $P_{s-\tau,s} \equiv P_{0,s}$ for all $\tau \geqslant s$ (one can also view the market as having started at $t=-\infty$). Multiplying by $x(p_s)$, discounting, and summing up yields (14).

To construct a solution to the maximization of (14), we first solve the finite problems obtained by truncating the horizon at $T>t$, then let $T \to \infty$. Let H_t^f be given, so that it can be dropped from $\pi(\cdot \mid H_t^f)$. For any $T>t$, and any infinite sequence $\{p_s \mid s \geqslant t\}$, let $\pi_T(\{p_s\})$ denote discounted profits from t to T only. Since $\pi_T(\{p_s\})$ depends only on $\{p_s \mid t \leqslant s \leqslant T\}$, it trivially defines a function (also denoted as π_T) on R_+^{T-t+1}. The finite horizon problem consists in maximizing π_T on that space.

Claim 1. In the finite horizon problem, charging $p<p(c)$ is never optimal.

Indeed, for any sequence $\{p_s \mid t \leqslant s \leqslant T\}$, define $q_s = \max\{p_s, p(c)\}$ and $Q_{s-\tau,s} \equiv \max\{q_k \mid s-\tau \leqslant k \leqslant s\} = \max\{P_{s-\tau,s}, p(c)\}$, for all $t \leqslant s \leqslant T$ and $\tau \geqslant 0$. Then consider the following two cases.

Case 1. If $P_{s-\tau,s} < p(c)$. Since $(p-c)\,x(p)$ and $(p-c)\,D(p)$ are increasing up to $p(c)$, $(p_s-c)\,x(p_s)\,N(P_{s-\tau,s}) \leqslant (P_{s-\tau,s}-c)\,D(P_{s-\tau,s}) < (p(c)-c)\,D(p(c)) = (q_s-c)\,x(q_s)\,N(Q_{s-\tau,s})$.

Case 2. If $P_{s-\tau,s} \geqslant p(c)$. Then $(p_s-c)x(p_s)N(P_{s-\tau,s}) = (p_s-c)x(p_s)\,N(Q_{s-\tau,s}) \leqslant (q_s-c)\,x(p_s)\,N(Q_{s-\tau,s})$, with strict inequality if $p_s<q_s$, i.e., $p_s<p(c)$. Thus in both cases,

$$\pi_T(\{p_s\}) = \theta\lambda \sum_{s=t}^{T} \delta^{s-t} \sum_{\tau=0}^{\infty} (1-\mu)^\tau (p_s - c)\, x(p_s)\, N(P_{s-\tau,s})$$
$$\leqslant \theta\lambda \sum_{s=t}^{T} \delta^{s-t} \sum_{\tau=0}^{\infty} (1-\mu)^\tau (q_s - c)\, x(p_s)\, N(Q_{s-\tau,s}),$$

or

$$\pi_T(\{p_s\}) \leqslant \pi_T(\{\max(p(c), p_s)\}) \tag{A5}$$

with strict inequality if any p_s is smaller than $p(c)$.

Claim 2. In the finite horizon problem, $p > \max\{p(c), P_{t-1}\}$ is never optimal.

If $P_{t-1} \leqslant p(c)$, then by (A4) $p_s = p(c)$ for all $s \geqslant t$ maximizes π_T. If $P_{t-1} > p(c)$, redefine the sequences $\{q_s\}$ and $\{Q_{s-\tau,s}\}$ as $q_s = \min\{p_s, P_{t-1}\}$, $Q_{s-\tau,s} = \max\{q_k \mid s-\tau \leqslant k \leqslant s\} = \min\{P_{s-\tau,s}, P_{t-1}\}$. Then for any $s \geqslant t$ and $\tau \geqslant 0$, three cases are possible.

Case 1. If $P_{s-\tau,s} \leqslant P_{t-1}$, then $(p_s - c)\, x(p_s)\, N(P_{s-\tau,s}) = (q_s - c)\, x(p_s)\, N(Q_{s-\tau,s})$.

Case 2. If $P_{s-\tau,s} > P_{t-1} > p(c)$ and $p_s \geqslant P_{t-1}$, then either $P_{s-\tau,s} \leqslant p_m(c)$, implying $(p_s - c)\, x(p_s)\, N(P_{s-\tau,s}) \leqslant (P_{s-\tau,s} - c)\, x(P_{s-\tau,s})\, N(P_{s-\tau,s})$, or $P_{s-\tau,s} > p_m(c)$, implying $(p_s - c)\, x(p_s)\, N(P_{s-\tau,s}) \leqslant (p_m(c) - c)\, x(p_m(c))\, N(P_{s-\tau,s}) < (p_m(c) - c)\, x(p_m(c))\, N(p_m(c)) \leqslant (P_{s-\tau,s} - c)\, x(P_{s-\tau,s})\, N(P_{s-\tau,s})$. So in both cases, $(p_s - c)\, x(p_s)\, N(P_{s-\tau,s}) \leqslant (P_{s-\tau,s} - c)\, x(P_{s-\tau,s})\, N(P_{s-\tau,s}) \leqslant (P_{t-1} - c)\, D(P_{t-1}) = (q_s - c)\, x(q_s)\, N(Q_{s-\tau,s})$.

Case 3. If $\max\{p(c), p_s\} < P_{t-1} < P_{s-\tau,s}$, then $(p_s - c) x(p_s) N(P_{s-\tau,s}) = (q_s - c)\, x(q_s)\, D(P_{s-\tau,s}) \leqslant (q_s - c)\, x(q_s)\, D(P_{t-1}) = (q_s - c)\, x(q_s)\, D(Q_{s-\tau,s})$.

Thus, in every case, $(p_s - c)\, x(p_s)\, D(P_{s-\tau,s}) \leqslant (q_s - c)\, x(q_s)\, N(Q_{s-\tau,s})$, with strict inequality if and only if $P_{s-\tau,s} > P_{t-1}$. Hence,

$$\pi_T(\{p_s\}) \leqslant \theta\lambda \sum_{s=t}^{T} \delta^{s-t} \sum_{\tau=0}^{\infty} (1-\mu)^\tau (q_s - c)\, x(q_s)\, N(Q_{s-\tau,s}),$$

with strict inequality if and only if there exists some $s \geqslant t$ and $\tau \geqslant 0$ such that $P_{s-\tau,s} > P_{t-1}$, i.e., if there exists $k \geqslant t$ such that $p_k > P_{t-1}$. Thus,

$$\pi_T(\{p_s\}) \leqslant \pi_T(\{\min(p_s, P_{t-1})\}) \tag{A6}$$

with strict inequality if $p_k > P_{t-1}$ for some $k \geqslant t$. This concludes Claim 2.

Claim 3. In the finite horizon problem, $p > p_m(c)$ is never optimal.

If $p_m(c) \geqslant P_{t-1}$, this follows from Claim 2. If $p_m(c) < P_{t-1}$, redefine $q_s \equiv \min\{p_s, p_m(c)\}$, $Q_{s-\tau,s} \equiv \min\{q_k \mid s-\tau \leqslant k \leqslant s\} = \min\{P_{s-\tau,s}, p_m(c)\}$. Then, $(p_s - c)\ x(p_s)\ N(P_{s-\tau,s}) \leqslant (q_s - c)\ x(q_s)\ N(P_{s-\tau,s}) \leqslant (q_s - c)\ x(q_s)\ N(Q_{s-\tau,s})$, hence the result.

Define now the compact set $\Omega_t \equiv [p(c), \max\{p(c), \min(P_{t-1}, p_m(c))\}]$. Since π_T is continuous on Ω_t^{T-t+1}, it attains its maximum for some $\{p_{T,s} \mid t \leqslant s \leqslant T\}$, $p_{T,s} \in \Omega_t$; moreover, Claims 1 to 3 imply that $p_{T,s}$ maximizes π_T over all of R_+^{T-t+1}. We extend $\{p_{T,s} \mid t \leqslant s \leqslant T\}$ into an infinite sequence by setting $p_{T,s} \equiv c$ for $s > T$. Now fix $s \geqslant t$ and let $T \to \infty$; since the sequence $\{p_T, s \mid T \geqslant s\}$ is in Ω_t, there exists a subsequence $\{p_{T^k,s} \mid T \geqslant s\}$ converging to a limit $p_s^* \in \Omega_t$.

Claim 4. $\{p_s^* \mid s \geqslant t\}$ maximizes π; i.e., it is an optimal strategy from H_t^f.

Indeed, for any $\varepsilon > 0$ fix $T > t$ large enough to have $\theta\lambda(p(c) - c)\ D(p(c))\ \delta^{T-t+1} < \varepsilon\mu(1-\delta)/2$. Then,

$$\pi(\{p_s\}) = \pi_T(\{p_s\}) + \theta\lambda \sum_{s=T+1}^{\infty} \delta^{s-t} \sum_{\tau=0}^{\infty} (1-\mu)^\tau$$
$$\times (p_s - c)\ x(p_s)\ N(P_{s-\tau,s}) < \pi_T(\{p_s\}) + \varepsilon/2. \qquad \text{(A7)}$$

But, by definition of $\{p_{T,s}\}$: $\pi_T(\{p_s\}) \leqslant \pi_T(\{p_{T,s}\}) = \pi_{T'}(\{p_{T,s}\})$ for all $T' \geqslant T$, because $p_{T,s} = c$ for $s > T$. Similarly, by definition of $p_{T',s}$,

$$\pi_{T'}(\{p_{T,s}\}) \leqslant \pi_{T'}(\{p_{T',s}\}) = \pi(\{p_{T',s}\}) < \pi_T(\{p_{T',s}\}) + \varepsilon/2$$

by (A7) applied with $p_s = p_{T',s}$ for all $s \geqslant t$. Thus, finally,

$$\pi(\{p_s\}) < \pi_T(\{p_{T',s}\}) + \varepsilon \qquad \text{for all} \quad T' \geqslant T. \qquad \text{(A8)}$$

Now let T' take values T^k, with $k \to \infty$, so that each $p_{T',s} \to p_s^*$. Since the sum $\pi_T(\{p_{T',s}\})$ has only a fixed number $T - t + 1$ of terms, we can take limits in (A8) and obtain for all T large enough,

$$\pi(\{p_s\}) \leqslant \pi_T(\{p_s^*\}) + \varepsilon, \qquad \text{(A9)}$$

Finally, $\pi_T(\{p_s^*\}) \leqslant \pi(\{p_s^*\})$ since $p_s^* \geqslant p(c) \geqslant c$ for all s. Letting $\varepsilon \to 0$ concludes the claim. Q.E.D.

Proof of Theorem II.3. Firms' strategies $\{p_s^*(H_t^f)\}$ are optimal by construction. Consumers' problem is the same as in Theorem II.2, with effective discount rate $\delta\mu$. Their beliefs over firms' cost types and prices are also unchanged and obey Bayes' rule both on the equilibrium path and between two zero subjective probability events. Off the equilibrium path, firms' strategies may now induce them to expect non-constant prices; but it remains true that they never expect the price to fall below its current level. Therefore their optimal purchasing and return decisions are unchanged. Q.E.D.

REFERENCES

1. B. AXELL, Search market equilibrium, *Scand. J. Econ.* **79** (1977), 20–40.
2. K. BAGWELL, Introductory price as a signal of cost in a model of repeat business, *Rev. Econ. Stud.* **54** (1985), 365–384.
3. R. BENABOU, Search, price setting and inflation, *Rev. Econ. Stud.* **55** (1988), 353–373.
4. R. BENABOU, "Search Market Equilibrium, Bilateral Heterogeneity, and Repeat Purchases," CEPREMAP Working Paper 8806, 1988; revised, MIT mimeo, 1990.
5. J. CARLSON AND P. MCAFEE, Discrete equilibrium price dispersion, *J. Polit. Econ.* **91** (1983), 480–493.
6. P. DIAMOND, A model of price adjustment, *J. Econ. Theory* **3** (1971), 156–168.
7. D. FUDENBERG AND J. TIROLE, Perfect Bayesian equilibrium and sequential equilibrium, *J. Econ. Theory* **53** (1991), 236–260.
8. F. GUL, H. SONNENSCHEIN, AND R. WILSON, Foundation of dynamic monopoly and the Coase conjecture, *J. Econ. Theory* **39** (1986), 155–190.
9. D. KREPS AND R. WILSON, Sequential equilibria, *Econometrica* **50** (1982), 863–964.
10. R. MACMINN, Search and market equilibrium, *J. Polit. Econ.* **87** (1980), 851–858.
11. J. MCMILLAN AND P. MORGAN, Price dispersion, price flexibility and repeated purchasing, *Can. J. Econ.* **21** (1988), 883–902.
12. J. REINGANUM, A simple model of equilibrium price dispersion, *J. Polit. Econ.* **87** (1979), 851–858.
13. R. ROB, Equilibrium price distributions, *Rev. Econ. Stud.* **52** (1985), 487–504.
14. J. STIGLITZ, Competition and the number of firms in a market: Are duopolies more competitive than atomistic markets? *J. Polit. Econ.* **95** (1987), 1041–1061.
15. P. VON ZUR MUEHLEN, Monopolistic competition and sequential search, *J. Econ. Dynam. Control* **2** (1980), 257–281.