

Dynamic Pricing of Network Goods with Boundedly Rational Consumers*

Roy Radner and Arun Sundararajan

Leonard N. Stern School of Business, New York University

February 2006

Working Paper CeDER-06-03

Center for Digital Economy Research, New York University

Abstract

We present a model of dynamic monopoly pricing for a good that displays network effects. In contrast with the standard notion of a rational-expectations equilibrium, we model consumers as boundedly rational, and unable either to pay immediate attention to each price change, or to make accurate forecasts of the adoption of the network good. Our analysis shows that the seller's optimal price trajectory has the following simple structure: the price is zero when the product user base is below a specific threshold, and is chosen to keep user base stationary once this threshold demand level has been attained. We show that our prescribed pricing policy is robust to a number of extensions, which include the product's user base evolving over time, a fraction of consumers being sufficiently rational to make accurate adoption forecasts, and consumers basing their choices on a mixture of a myopic and a "stubborn" expectation of adoption. Our results differ significantly from those that would be predicted by a model based on rational-expectations equilibrium, and are more consistent with the pricing of network goods observed in practice.

*The authors thank Philipp Afeche, Toker Doganoglu, Joseph Harrington, Timothy Van Zandt, and workshop participants at Boston University, INSEAD, New York University, Stanford University, the University of Minnesota, the Sixteenth Workshop on Information Systems and Economics, the 2005 NBER/NSF Decentralization Conference, the 2005 North American Winter Meeting of the Econometric Society and the 2005 Kiel-Munich Workshop on the Economics of Network Industries.

1 Introduction

An important simplifying assumption made in economic models of network effects is that consumers are perfectly rational. This assumption is important in such models because the value to each consumer of a good or service that displays network effects (henceforth referred to as a "network good") is influenced by the consumption choices made by some or all other consumers. The solutions to such models are typically based on asserting the following "unboundedly" rational behavior: (a) consumers immediately react to each observable strategic decision made by a seller, and (b) having observed this decision, consumers form a common expectation of demand, and make their consumption choices unilaterally based on this expectation, which is then realized in equilibrium. This solution is commonly referred to as satisfying "fulfilled expectations", or as a "rational expectations equilibrium".

Some notion of perfect rationality is at the base of most current economic analysis, even though most researchers accept that agents are not in reality unboundedly rational. Such models continue to be used, perhaps because there is an implicit belief that the 'output' of analysis based on the approximation of unboundedly rationality agents is (reasonably) correct. In the specific case of network goods, however, many predictions of models based on unboundedly rational behavior do not appear to be a good description of reality. For example, it is believed that outcomes in markets for network goods are often path-dependent, and therefore, the dynamics of the adoption process and the path of choices made by a seller are important determinants of eventual outcomes. This would not be the case were consumers able to form rational expectations after each price change. It is therefore possible that managers, upon observing these aspects of real-world network markets, would hesitate to rely on the prescriptions of models that ascribe the extent of rationality, coordination and prediction accuracy that characterizes unboundedly rational behavior.

Our objective in this paper is to examine how the predictions of models of network effects change, if at all, under assumptions about consumer rationality that seem more realistic. We do so by presenting alternative models of demand for a network good, in which consumers are cognitively bounded, do not immediately react to every change in the seller's price, and furthermore, make their consumption choices based on a boundedly rational assessment of expected demand, which may depend on the current price, the current level of demand, and/or an exogenously specified "stubborn" expectation of equilibrium demand. We use these models of bounded rationality to

study the dynamic pricing problem for a monopoly seller of a network good. The adoption choices of the consumers continuously influence the rate at which demand adjusts over time, and the monopoly seller therefore chooses the price *trajectory* that maximizes her discounted stream of profits. The rate at which demand adjusts over time is also affected by a parameter λ which is proportional to the fraction of consumers in each period who "pay attention" to the current price.

Our first theorem shows that when consumers are myopic in their expectations, the monopolist's optimal pricing trajectory is generated by a target policy with the following properties: when current demand is below the target, the price is zero; when current demand is above the target, the price is the maximum possible; and when current demand is at the target, the price is chosen to keep demand stationary. The target could be interpreted as the level of adoption below which the monopolist invests in building a user base, and above which the monopolist profits from exploiting her installed base. The result thus prescribes an extreme form of penetration pricing that is not uncommon in markets for network goods.

This theorem also shows that the optimal demand target with myopic consumers is always strictly lower than the equilibrium level of demand predicted by a model with rational expectations. The difference between the target demand and the rational expectations equilibrium demand is a decreasing function of λ , and tends to zero as λ increases without bound, i.e., when all consumers react to price changes infinitely fast. We briefly discuss how the results of the theorem are affected when the distribution of consumer types is nonuniform, although we do not present this extension in detail.

Our subsequent results extend this theorem along two directions. First, we examine how the monopolist's optimal price trajectory varies when the population of consumers evolves over time. That is, in each period, a constant fraction of consumers is replaced by new ones, or there is an exogenous source of replacement in the population of potential customers, at a rate determined by a parameter c that is proportional to the fraction of consumers replaced in each period. Our second theorem establishes that the monopolist's optimal pricing trajectory continues to be generated by a target policy with the same properties as the one derived in Theorem 1, although with a strictly lower optimal demand target. Moreover, in this model, the price that keeps demand stationary at any desired level is progressively lower as c increases, and we discuss how this differentiates the effect of changes in the rate of replacement c , from corresponding changes in the rate of "attention"

λ .

In parallel, we establish the extent to which the prescribed price trajectory from our first theorem is robust to alternative models of bounded rationality. First, we generalize the result to the case in which the population of consumers contains both those who are myopic and those who are "fully rational." For this model, the results are qualitatively the same as for the case in which all consumers are myopic, but with a faster rate of adjustment. We then examine how the monopolist's optimal price trajectory varies when the expectation of demand formed by each consumer who pays attention is a weighted average of the myopic expectation and an exogenously specified "stubborn" expectation. Our final theorem establishes that, again, the monopolist's optimal pricing trajectory continues to be generated by a target policy with the same properties as the one derived in Theorem 1, with a lower target demand level. The target increases as consumers become less stubborn, eventually converging to the target demand level of the policy for purely myopic consumers.

We have organized the rest of this paper as follows. Section 2 describes our model of bounded rationality, our underlying discrete-time model, and the derivation of its continuous-time counterpart. Section 3 analyzes our "base" model with myopic consumers, derives the optimal price trajectory for this model, and contrasts its results with those predicted by the rational-expectations model. Section 4 extends the base model in three ways, incorporating in turn an evolving consumer base, a mixture of myopic and rational consumers, and a model of rationality in which consumers are "stubborn", and showing that the results of Section 3 are robust to each of these extensions. Section 5 concludes, and the Appendix contains those proofs not presented in the main body of the paper.

2 Overview of Our Model

2.1 Models of Bounded Rationality

We derive our continuous-time formulation as the limiting case of a discrete-time model. A network good is provided by a monopolist, who sells the good one period at a time. (Think of the good as a service.) The length of each period is h , and therefore time is indexed as $t = 0, h, 2h, \dots$ The monopolist announces a price $p(t)$ at the beginning of each period. We assume that the price, $p(t)$, is constrained to be nonnegative, and is bounded above (more on this later). A unit mass

of a continuum of consumers is indexed by a "type" parameter θ in the unit interval. Let $q(t)$ denote the mass of consumers who purchase the service during period t . A consumer of type θ who purchases the service during period t realizes a "net incremental utility" or "surplus" equal to $\theta q(t) - p(t)$. Thus θ is the marginal value to the consumer of the "network effect." Our basic assumption throughout the paper is that θ is *uniformly distributed on the unit interval*. However, in Section 3.3 we report some results for nonuniform distributions of θ .

In the process of deciding whether or not to purchase the service, the consumers are *boundedly rational*. The first aspect of this bounded rationality is that of *bounded attention*. In each period, a random fraction λh of consumers of each type "pay attention to" the price $p(t)$, where $0 < \lambda h < 1$. (Since in the continuous-time model we shall let h tend to zero, the last inequality will not constrain λ in an essential way.) Correspondingly, the remaining fraction $(1 - \lambda h)$ of consumers of each type do not respond to the monopolist's price announcement, and their choice remains unchanged from what it was during the previous period. Notice that an *equal* fraction λh of consumers of each type "pay attention" in each period, and that the magnitude of this fraction depends on the length of the interval h . Thus the average time between successive price checks by a consumer of any type is $(1/\lambda h)$. One might therefore also interpret λh as a "rate of adjustment" to changes in prices. However, the constraint $\lambda h < 1$ bounds this rate of adjustment in the discrete-time model.

The second aspect of bounded rationality in our model specifies how consumers who are paying attention form their *expectation of what the demand* $q(t)$ will be. Specifically, each consumer who notices the price $p(t)$ at the beginning of period t makes the same prediction, $q_E(t, h)$, of the total demand in period t . Therefore, a consumer of type θ who notices $p(t)$ will buy the good if and only if $\theta q_E(t, h) \geq p(t)$. In subsequent sections we explore the implications of a few simple models of expectations formation. In particular, we shall study a model of "myopic expectation," in which $q_E(t, h) = q(t - h)$.

2.2 A Continuous Time Model

Our analysis uses a continuous-time version of the discrete-time model described above, which is derived as a limiting case of the discrete-time model, as $h \rightarrow 0$. Our motivation for using a continuous-time model is two-fold. First, if consumers do not correctly forecast the demand in the coming period, then as noted above, *the discrete-time model has a built-in implicit minimum one-*

period lag of adjustment in demand, so the speed of adjustment is bounded. In the continuous-time model, the speed of adjustment is determined by the parameter λ , and may be taken to be as large as one likes. Second, a continuous-time model, although requiring more advanced mathematical methods, typically yields solutions with simpler dynamics (e.g., by avoiding the phenomenon of overshooting).

To begin, let

$$q_E(t) = \lim_{h \rightarrow 0} q_E(t, h),$$

and assume that $q_E(t)$ is well-defined. In the present paper, $q_E(t)$ will depend at most on the current demand and price, $q(t)$ and $p(t)$, respectively, although one can imagine plausible models in which it depends on the history of demand and price, as well (see, e.g., Radner and Richardson, 2003). We shall also assume that $p(t) \leq q_E(t)$. This is plausible, since if $p(t) > q_E(t)$ then the price would exceed every consumer's willingness-to-pay. (See further remarks about this assumption in Section 3.)

The resulting time-rate of change of demand is described in our first lemma. An intuitive understanding of how the discrete-time formulation is related to its continuous-time counterpart may help the reader interpret subsequent results more easily, and we therefore present the lemma's proof in the main body of the paper.

The third line of (1) in the Lemma is actually an *assumption*. In part, it reflects the fact (noted above) that, if $p(t) > q_E(t)$ then $p(t)$ exceeds every consumer willingness-to-pay. It could be argued that this would not affect the behavior of those consumers who are not currently "paying attention" to the price, and hence a price above $q_E(t)$ need not drive all consumers out of the market. The third line of (1) actually makes a stronger statement: if $p(t) > q_E(t)$ then every consumer will expect all the subscribers to unsubscribe immediately. It's as if a price above $q_E(t)$ were a "wake-up call." This has the effect of imposing on the monopolist the constraint that the current price must not exceed the current expectation of total demand. This discontinuous behavior can be thought of as an approximation to a more realistic model in which the parameter λ is itself some increasing function of the current price. In any case, for the model to be well-behaved some upper bound on price is required, and this approximation has the advantage of leading to a more tractable model.

In the third line of (1) we also impose the constraint that the price must be nonnegative. Other finite lower bounds could be imposed without changing the qualitative nature of our results.

Lemma 1 *If at time t (≥ 0) the demand and price are $q(t)$ and $p(t)$, respectively, then the time-rate of change of demand is specified by:*

$$q'(t) = \begin{cases} 0, & q(t) = 0, \\ \lambda \{Q[q_E(t), p(t)] - q(t)\}, & 0 < q(t) \leq 1, 0 \leq p(t) \leq q_E(t), \\ -\infty, & 0 < q(t) \leq 1, p(t) > q_E(t), \end{cases} \quad (1)$$

where

$$Q(x, p) \equiv 1 - \frac{p}{x}, \quad 0 < x \leq 1. \quad (2)$$

Proof. First, when $q(t) = 0$, the product is of no value to all consumers, which yields the first line of (1). Next, suppose the demand from consumers of type θ in period t is denoted by $w(\theta, t)$. Recall that a fraction λh of consumers of type θ notice $p(t)$, form a shared expectation of demand $q_E(t, h)$, and decide whether or not to adopt the product for period t . Therefore, if $\theta \geq [p(t)/q_E(t, h)]$, each consumer in this fraction λh adopts the product, and if $\theta < [p(t)/q_E(t, h)]$, then none of these consumers adopt the product. Since the remaining fraction $(1 - \lambda h)$ continue to do in period t what they were doing in period $t - h$, it follows that:

$$w(\theta, t) = \begin{cases} \lambda h + (1 - \lambda h)w(\theta, t - h), & \theta \geq p(t)/q_E(t, h), \\ (1 - \lambda h)w(\theta, t - h) & \theta < p(t)/q_E(t, h), \end{cases} \quad (3)$$

and therefore,

$$w(\theta, t) - w(\theta, t - h) = \begin{cases} \lambda h[1 - w(\theta, t - h)], & \theta \geq p(t)/q_E(t, h), \\ -\lambda h[w(\theta, t - h)], & \theta < p(t)/q_E(t, h). \end{cases} \quad (4)$$

Dividing both sides by h and letting h tend to zero yields the time rate of change of demand for consumers of type θ :

$$\frac{dw(\theta, t)}{dt} \equiv \lim_{h \rightarrow 0} \left(\frac{w(\theta, t) - w(\theta, t - h)}{h} \right) = \begin{cases} \lambda[1 - w(\theta, t)], & \theta \geq p(t)/q_E(t), \\ -\lambda[w(\theta, t)], & \theta < p(t)/q_E(t). \end{cases} \quad (5)$$

Recall that θ is uniformly distributed on the unit interval. Hence

$$q(t) = \int_0^1 w(\theta, t) d\theta, \quad (6)$$

it follows that

$$q'(t) = \int_0^1 \left(\frac{dw(\theta, t)}{dt} \right) d\theta, \quad (7)$$

so by (5) and (7),

$$q'(t) = \int_0^{p(t)/q_E(t)} -\lambda[w(\theta, t)]d\theta + \int_{p(t)/q_E(t)}^1 \lambda[1 - w(\theta, t)]d\theta, \quad (8)$$

which simplifies to

$$q'(t) = \lambda \left[1 - \frac{p(t)}{q_E(t)} \right] - \int_0^1 \lambda w(\theta, t) d\theta, \quad (9)$$

and using (6), this yields the second line of (1), which completes the proof of the second line of 1.

■

Note that the third line of (1) is equivalent to

$$0 \leq p(t) \leq q_E(t). \quad (10)$$

We also assume, for simplicity, that the marginal cost of providing the service is zero. The monopolist chooses the price trajectory $p(t)$ to maximize her total discounted profit (revenue),

$$\int_0^\infty e^{-rt} p(t) q(t) dt, \quad (11)$$

subject to (10), where $r > 0$ is her given discount rate, and $q(t)$ evolves according to (1).

2.3 Rational-Expectations Equilibrium

A standard alternative theory of consumer behavior in models of network effects is embodied in the concept of fulfilled-expectations, or rational-expectations equilibrium. In this section, we describe the rational-expectations equilibrium for our model that is optimal for the monopolist, since it is a natural benchmark for the results of the models with boundedly-rational consumers.

Imagine that, when faced with a price p , each consumer notices p , correctly predicts the total demand q at that price, and decides whether or not to subscribe on the basis of that prediction. This is consistent with consumers who are unboundedly rational, and is the behavior assumed by most models of network effects. Following standard terminology, we call such a pair (q, p) , a *rational expectations equilibrium* (REE). In what follows, we do not impose the constraint that $p \leq q$, although it will turn out that for the interesting REEs the constraint will be satisfied.

First, notice that a prediction of $q = 0$ made by all consumers is a correct prediction for any price p , and therefore, $(0, p)$ is an REE for any $p \in [0, 1]$. Next, for a prediction $q > 0$ to be correct,

it must satisfy

$$q = \left(1 - \frac{p}{q}\right), \quad (12)$$

since the expression on the right-hand side of (12) is the realized demand when the demand prediction is q [cf. (2)]. For $0 \leq p < 1/4$, this equation has two real solutions,

$$q = \frac{1 \pm (1 - 4p)^{1/2}}{2},$$

which are strictly positive and, incidentally, satisfy $p < q$. Alternatively, define the function $P(q)$ using (12) as

$$q = \left(1 - \frac{P(q)}{q}\right), \quad (13)$$

or

$$P(q) = q(1 - q). \quad (14)$$

Therefore, the pair $[q, P(q)]$ is an REE for each $q \in [0, 1]$.

We will make the standard assumption that if there are multiple REE's that correspond to a particular price, customers will coordinate on the one with the highest predicted demand. Under this assumption, define the optimal REE as the one that maximizes the monopolist's instantaneous profits pq . It follows that

$$q^* = \arg \max_q pqP(q), \quad (15)$$

$$p^* = P(q^*), \quad (16)$$

and computing the solution to the optimization problem (15) yields

$$q^* = (2/3), p^* = (2/9). \quad (17)$$

Furthermore, since all consumers notice each price change instantaneously and react rationally to the price change, demand adjusts to any new price instantaneously, independent of past prices/demand, and the demand-price pair at each instant is an REE, for any price trajectory the monopolist chooses. But there is a unique optimal REE; thus, there is no reason for the monopolist to vary her price over time. We have therefore shown:

Lemma 2 *When consumers are unboundedly rational, the monopolist's optimal price trajectory is $p(t) = p^*$, $q(t) = q^*$ for all $t \geq 0$, where (q^*, p^*) is the optimal rational-expectations equilibrium (17).*

3 Myopic Consumers

3.1 The Model

This section describes the monopolist's optimal price trajectory for a class of models of bounded rationality in which, in our underlying discrete-time model, every consumers' expectation of total demand during period t equals the actual demand in the previous period. Thus, in the discrete-time model,

$$q_E(t, h) = q(t - h). \quad (18)$$

Recall that θ is uniformly distributed on the unit interval. Corresponding to Lemma 1, the continuous-time approximation is

$$q_E(t) = q(t), \quad (19)$$

and therefore, if at time t the demand and price are $q(t)$ and $p(t)$, respectively, then the time-rate of change of demand is specified by

$$q'(t) = m(q(t), p(t)), \quad (20)$$

where:

$$m(q, p) = \begin{cases} 0, & q = 0, \\ \lambda \left[1 - \frac{p}{q} - q \right], & 0 < q \leq 1, 0 \leq p \leq q, \\ -\infty, & 0 < q \leq 1, p > q. \end{cases} \quad (21)$$

If one examines the first two lines of (21), it becomes apparent that each REE is also a stationary point of the demand-price process (and vice versa) when consumers are myopic, for any $\lambda > 0$. That is, (q, p) is an REE if and only if $m(q, p) = 0$. In subsequent sections, we therefore often refer to $P(q)$ in (14) as the "stay-where-you-are" price.

3.2 The Optimal Price Trajectory: A Target Policy

In this section, we establish that the monopolist's optimal price trajectory is generated by a policy that belongs to family that we call *target policies*. We shall analyze the maximization problem in (11) using the method of dynamic programming, in which the state variable is the current demand.

First, by Blackwell's Theorem, there is no loss in restricting our attention to *stationary* policies, i.e., policies for which the current price at any time is a function of the current state only:

$$p(t) = \mu[q(t)]. \quad (22)$$

Note that the function μ does not change in time. Of course, a policy is admissible only if the differential equation (1) has a unique solution starting from any initial state $q(0)$.

With some abuse of notation, we refer to the value of demand at time zero as q . The value of a policy μ at an initial state $q(0) = q$ is the corresponding profit,

$$V_\mu(q) = \int_0^\infty e^{-rt} \mu[q(t)] q(t) dt, \quad q(0) = q. \quad (23)$$

Define

$$V(q) = \sup_\mu V_\mu(q), \quad (24)$$

where the supremum is over all admissible policies μ . A policy is *optimal* if its profit attains the supremum at every state q .

Recall that the pair $[q, P(q)]$ is a stationary point for the demand-price process. Define the (stationary) *target policy* with target s ($0 < s < 1$) by

$$\pi(q) = \begin{cases} 0, & q < s, \\ P(q), & q = s, \\ q, & q > s. \end{cases} \quad (25)$$

If the initial state is less than the target s , then the demand will increase until it reaches the target, during which time the profit is zero. After that, the demand per unit time will remain at the target, and the profit will be $sP(s)$. Hence the total discounted profit will be

$$\int_T^\infty e^{-rt} sP(s) dt = e^{-rT} \left[\frac{sP(s)}{r} \right], \quad (26)$$

where T is the time at which demand reaches the target s . Note that there is a tradeoff between choosing a higher target and reaching it sooner. Similarly, if the initial state is greater than the target s , then the demand will decrease until it reaches the target, and the total discounted profit will be:

$$\int_0^T e^{-rt} [q(t)]^2 dt + \int_T^\infty e^{-rt} sP(s) dt, \quad (27)$$

where T is the time at which demand reaches the target s , and until that time $q(t)$ is determined by the simple linear differential equation

$$q'(t) = -\lambda q(t).$$

Let π^* be the optimal target policy. Our first theorem establishes that π^* is optimal among all policies, and therefore determines the seller's optimal price trajectory.

Theorem 1 *The optimal target policy π^* is optimal among all policies, and the optimal target is*

$$\sigma \equiv \frac{2\lambda}{3\lambda + r}. \quad (28)$$

Therefore, when current demand is below the target level σ , the monopolist charges a price of zero until it reaches its demand target σ , and then raises its price to the steady-state price,

$$P(\sigma) = \frac{2\lambda(\lambda + r)}{(3\lambda + r)^2}. \quad (29)$$

If current demand is above the target level, the monopolist charges the maximum possible price (equal to the current demand), until demand reaches the target level, at which point it lowers its price to the steady-state price in (29).

[Note: While the case in which the initial demand exceeds the target would typically managerially less relevant than the case in which it is less than the target, the characterization of the policy for $q > \sigma$ is necessary to compute the value function $V(q)$ in each state, so that we can establish that the policy is indeed optimal.]

Although Theorem 1 is a limiting case of both Theorems 2 and 3 (see Section 4), we nevertheless present a complete proof in Appendix A. A first managerial implication of Theorem 1 is the extent to which it differs from the corresponding REE solution. Rather than choosing a positive price and asserting that rational consumer expectations will induce instantaneous demand adjustment, the result indicates that the seller must pro-actively establish a user base. Furthermore, it prescribes an "extreme" penetration pricing policy – pricing at the lowest possible level until reaching a pre-specified demand target – as being the optimal price trajectory for establishing this user base. The target can be interpreted as the level of adoption below which the monopolist invests in building a user base, and above which the monopolist profits from her installed base. Notice that this pricing policy is optimal among all possible trajectories. Moreover, it is easily understandable, simple to




Figure 1: Illustrates how the demand target and steady-state price vary with changes in the "rate of attention" λ .

implement, and appears to be qualitatively consistent with those observed in practice for network goods.

The expression in (28) establishes that when $r = 0$, the optimal demand target σ is equal to the REE demand q^* . Furthermore, for all finite r ,

$$\lim_{\lambda \rightarrow \infty} \sigma = q^*, \quad (30)$$

and therefore, the REE outcome is a limiting case of the outcome in our model. For any finite rate of adjustment λ and positive discount rate r , (28) also establishes that the demand target σ is strictly *lower* than q^* .

Clearly, the optimal demand target in (28) increases in the "rate of attention" λ . However, the steady-state price is not monotonic in λ . As λ progressively decreases from ∞ (that is, as consumers become increasingly cognitively bounded), the steady-state price first increases and then subsequently decreases. This is depicted in Figure 1. More precisely, for rates of attention $\lambda > r$, the decrease in the demand target that a lowering of λ induces allows the seller to charge a higher price while still keeping its user base stationary at σ . However, for lower rates of attention ($\lambda < r$), the fact that a smaller and smaller fraction of customers are paying attention forces the seller to lower its price in order to maintain its desired target user base.

Despite the potential decrease in steady-state price $P(\sigma)$ as λ increases, it is straightforward to

verify that the seller's steady-state profits increase monotonically with λ :

Corollary 1 *The steady-state discounted profit,*

$$V(\sigma) = \frac{2\lambda^2(\lambda + r)}{(3\lambda + r)^3}, \quad (31)$$

is increasing in λ . Furthermore, if the seller's initial user base $q(0)$ is strictly lower than its demand target σ , the seller's discounted profit,

$$V(q_0) = \frac{1}{r} \left(\frac{1 - q_0}{1 - \sigma} \right)^{-\left(\frac{r}{\lambda}\right)} \sigma^2(1 - \sigma), \quad (32)$$

is also increasing in λ .

Thus, rather than being able to exploit the fact that its consumers do not pay attention to price changes, the seller is always adversely affected, and progressively more so as consumers become "more" boundedly rational (that is, as λ decreases). While we specify λ as an exogenous "rate of attention", it may be possible in practice for a seller to influence the rate of attention by making costly advertising investments. Studying this tradeoff is an interesting direction for future research.

Additionally, since the optimal policy is a target policy, the demand-price process eventually will "settle down" at the stationary point $[\sigma, P(\sigma)]$. As we have noted in Section 2.2, any stationary point of this process is a rational-expectations equilibrium (although not an optimal one). If one interprets the output of a static model of network effects as the steady-state of some underlying dynamic model (whose details have not been explicitly presented), this indicates that the use of static rational-expectations outcomes in this way is not entirely inconsistent with our model; only that the choice of the optimal REE is incorrect.

3.3 Nonuniform distributions of consumer type

In this subsection we comment briefly on the possible generalizations of the preceding model to cases with nonuniform distributions of the "type" parameter θ . (For a detailed account of these generalizations, see Radner and Sundararajan, 2005.)

First, in the preceding subsection it was shown that the steady-state demand in the monopolist's optimal price trajectory is strictly lower than q^* , the optimal REE demand, although the price that maintains this steady-state demand may be lower or higher than p^* , the optimal REE price. This

observation is more generally true in the following sense: *For any (sufficiently regular) customer type distribution, the optimal REE is not a steady state for an optimal policy in the model of myopic consumers.*

Second, if the cumulative distribution function of θ is strictly *concave*, one needs to admit "measure-valued controls," and the optimal policy is a "generalized target policy" with the *same* target as in (28), *independent of the actual distribution function*. When current demand is below the target, the price is zero; when current demand is above the target, the price is the maximum possible, and when current demand is at the target, the monopolist chooses the "mixture" between a price of zero and the maximum possible price that keeps demand stationary. In practice this would mean that, when the demand is in the neighborhood of the target, the price fluctuates irregularly between high and low values, which is suggestive of the phenomenon of frequent "sales."

We have also studied the case in which the cumulative distribution function of θ is *convex*, but thus far we have only partial results for this case. They suggest that, again, the optimal policy has a "target demand," but in the approach to that target from below, the price may rise gradually towards the steady-state price at the target (with a symmetric phenomenon when the approach to the target demand is from above.) In particular, we have a complete solution for a piecewise linear convex cumulative distribution function that has this property.

4 Extensions of the Model of Myopic Consumers

We now present three extensions to our model of myopic consumers. In each extension, we maintain the assumptions on bounded attention and the distribution of consumer types. Each extension establishes that the seller's optimal price trajectory retains the structure derived in Theorem 1, although the level of the demand target and steady-state price vary across the three cases.

4.1 An Evolving Consumer Population

The first extension we consider is a situation in which, rather than being static, the population of potential customers for the network good evolves over time. We return to the discrete-time model of Section 2 to characterize this precisely, while maintaining each of its other assumptions. Therefore, the length of each period is h , and time is indexed by $t = \{0, h, 2h, \dots\}$. A fraction ch

of consumers of each type (both adopters and non-adopters) is *replaced* in each period. That is, a fraction ch of existing consumers "leave" and an equal fraction ch of new consumers "arrive" and are added to the pool of potential customers (clearly, each customer in the latter fraction is not an adopter). The size of the total set of potential customers therefore remains constant, although it has a constant "rate of replacement", which is proportional to the parameter c .

Following this replacement at the beginning of each period, we continue to assume that an equal fraction λh of consumers of each type pay attention in each period, and an equal fraction $(1 - \lambda h)$ do not. Proceeding as in Section 2.2, the time-rate of change in demand in a continuous-time approximation of this discrete-time model is described in Lemma 4:

Lemma 3 *If at time t , the demand and price are $q(t)$ and $p(t)$, respectively, then the time-rate of change of demand is specified by $q'(t) = m(q(t), p(t))$, where,*

$$m(q, p) = \begin{cases} 0, & q(t) = 0, \\ \lambda \left[1 - \frac{p}{q} \right] - (\lambda + c)q, & 0 < q \leq 1, 0 \leq p \leq q, \\ -\infty, & 0 < q \leq 1, p > q, \end{cases} \quad (33)$$

The positive rate of consumer replacement therefore slows down the rate at which demand increases, for any price trajectory that causes an increase in demand. Next, proceeding as in Section 3.2, it follows that the "stay-where-you-are" price is

$$P(q) = q \left[1 - \left(1 + \frac{c}{\lambda} \right) q \right], \quad (34)$$

and a stationary target policy π with target σ ($0 < \sigma < 1$) is defined by

$$\pi(q) = \begin{cases} 0, & q < \sigma, \\ P(\sigma), & q = \sigma, \\ q, & q > \sigma. \end{cases} \quad (35)$$

Our second main result shows that the seller's optimal price trajectory continues to be generated by a target policy, although with a strictly lower demand target:

Theorem 2 *The seller's optimal price trajectory is generated by a target policy, and the optimal demand target is*

$$\sigma = \frac{2\lambda}{3(\lambda + c) + r}. \quad (36)$$




Figure 2: Illustrates how the "stay-where-you-are" pricing function varies with an increase in the rate of consumer churn c , and the corresponding impact of such an increase on the optimal demand target and steady-state price.

Therefore, when current demand is below the target level σ , the monopolist charges a price of zero until it reaches its demand target σ , and then raises its price to the steady-state price

$$P(\sigma) = \frac{2\lambda[(\lambda + c + r)]}{[3(\lambda + c) + r]^2}. \quad (37)$$

The proof of Theorem 2 is presented in the Appendix.

The comparative statics results for changes in λ of Section 3.2 continue to hold for this extension. An interesting contrast between "losing consumers" and "losing attention" is illustrated in Figure 2. Notice in (36) that as the rate at which consumers are replaced, c , increases, the demand target becomes lower, and this is similar to the effect of a decrease in the rate at which consumers pay attention to price changes. However, an increase in c also shifts the "stay-where you are" price curve *downward*, which decreases the seller's steady-state revenue. We state this formally as:

Corollary 2 *The steady-state price (37) and seller's revenue are strictly decreasing in the "rate of consumer replacement" c .*

These results imply that a shift in c that decreases the demand target to a specific level has a more adverse effect on a seller than a shift in λ that decreases the demand target to the same level. Put another way, the impact of an increase in the consumer replacement rate is, in a sense, more

detrimental to a seller than a corresponding change in the rate of attention.

4.2 Myopic and Rational Consumers

Our next extension examines a situation in which some consumers are myopic and some are fully rational. Again, we start with a discrete-time model and move to a continuous-time approximation. Formally, suppose that in a period of length h a fraction λh of the consumers myopically adjust their demands, and another fraction μh of the consumers can predict the total demand in the next period. This give rise to the difference equation

$$\begin{aligned} q(t) &= \lambda h \left[1 - \frac{p(t)}{q(t-h)} \right] + \mu h \left[1 - \frac{p(t)}{q(t)} \right] + (1 - \lambda h - \mu h)q(t-h), \\ 0 &< q(t-h), \\ 0 &\leq p(t) \leq \min\{q(t-h), q(t)\}. \end{aligned} \tag{38}$$

Hence

$$q(t) - q(t-h) = \lambda h \left[1 - \frac{p(t)}{q(t-h)} \right] + \mu h \left[1 - \frac{p(t)}{q(t)} \right] - (\lambda h + \mu h)q(t-h), \tag{39}$$

$$\frac{q(t) - q(t-h)}{h} = (\lambda + \mu) - \frac{\lambda p(t)}{q(t-h)} - \frac{\mu p(t)}{q(t)} - (\lambda + \mu)q(t-h). \tag{40}$$

Letting h tend to zero, we get

$$q'(t) = (\lambda + \mu) \left[1 - \frac{p(t)}{q(t)} - q(t) \right]. \tag{41}$$

Thus we get the same law of motion as in the "purely myopic" case ($\mu = 0$), except that λ has been replace by $(\lambda + \mu)$. Hence we see that in this model what matters is the total rate of adaptation, $(\lambda + \mu)$, not how the "adapters" are distributed between myopic and fully rational consumers. The pricing trajectory prescribed by Theorem 1 therefore continues to be optimal, with λ simply replaced by $(\lambda + \mu)$.

4.3 Myopic and "Stubborn" Consumers

Finally, we describes some properties of the monopolist's optimal price trajectory when consumers are something between being myopic and being "stubborn". Rather than basing their expectation of total demand in the next period on the current period's demand level, consumers who pay attention to the monopolist's price announcement partly base their prediction on a stubborn assessment, ω , of

the total demand for the good. The extent to which they base their expectation on ω is determined by a parameter γ , where $0 \leq \gamma \leq 1$, and

$$q_E(t, h) = \gamma q(t - h) + (1 - \gamma)\omega. \quad (42)$$

Proceeding as in Sections 2.1 and 3.2, it follows that

$$q_E(t) = \gamma q(t) + (1 - \gamma)\omega, \quad (43)$$

the law of motion is

$$m(q, p) = \begin{cases} 0, & q = 0, \\ \lambda \left[1 - \frac{p}{\gamma q + (1 - \gamma)\omega} - q \right], & 0 < q \leq 1, \quad 0 \leq p \leq \gamma q + (1 - \gamma)\omega, \\ -\infty, & 0 < q \leq 1, p > q, \end{cases} \quad (44)$$

and the "stay-where-you-are" price is

$$P(q) = (1 - q)[\gamma q + (1 - \gamma)\omega]. \quad (45)$$

A stationary target policy π with target σ ($0 < \sigma < 1$) is defined by

$$\pi(q) = \begin{cases} 0, & q < \sigma, \\ P(\sigma), & q = \sigma, \\ \gamma q + (1 - \gamma)\omega, & q > \sigma. \end{cases} \quad (46)$$

For a given γ and ω , let π^* be the optimal target policy, and denote its target as $\sigma(\gamma, \omega)$. Our final theorem confirms that the structure of the pricing policy prescribed in Section 3 is robust to this extension as well, although, independent of the value of ω , the demand target is always lower.

Theorem 3 (a) *The monopolist's optimal price trajectory is generated by the target policy with target $\sigma(\gamma, \omega)$.*

(b) *$\sigma(\gamma, \omega)$ is strictly increasing in γ , and has the following values at its end points:*

$$\sigma(0, \omega) = \frac{\lambda}{2\lambda + r}, \quad (47)$$

$$\sigma(1, \omega) = \frac{2\lambda}{3\lambda + r}. \quad (48)$$

(c) *$\sigma(\gamma, \omega)$ is strictly decreasing in ω .*

Notice that the demand target of Theorem 1 is a limiting case of the demand target above, when $\gamma = 1$. Parts (b) and (c) have a simple intuitive explanation. An increase in the installed base for a network good benefits the seller in two ways: through the direct increase in demand, and by increasing the willingness to pay of consumers. It is the latter property that increases the monopoly demand for the good beyond what a normal good would enjoy. Therefore, at any given stubborn expectation ω , a decrease in the weight γ placed on the current demand makes the good seem "less like" a network good, and more like a normal good with an exogenously specified value that is proportionate to ω , thus reducing the steady-state user base that the seller finds optimal. Correspondingly, for any given γ , an increase in ω reduces the fraction of perceived user value that is influenced by actual current demand, and increases the corresponding fraction influenced by the "stubborn" expectation. One might therefore expect outcomes that are qualitatively similar to those of Theorem 3 in the base model of Section 3.2 if, rather than being a pure network good as we have assumed, a fraction of the willingness-to-pay for the good is independent of the demand q .

5 Concluding Remarks

We have explored several variations of a model of optimal dynamic monopoly pricing of a network good with assumptions about consumer rationality that are more realistic than those embodied in the theory of rational expectations. Our results predict that when consumers cannot form correct demand expectations, or pay attention to every price announcement, a seller must establish a user base by pricing as low as is possible until reaching a desired stationary target demand level. This outcome is quite different from what is obtained from a model in which the process of demand adjustment is due to the correct formation of expectations by rational consumers after each price announcement. Moreover, our prescribed price path seems to be similar to those commonly observed in practice.

We have shown that our main results are robust to a number of extensions. During our analysis of these extensions, we contrast how the demand target and the price that keeps demand steady vary with changes in the rate at which myopic consumers pay attention to price changes, and the rate at which there is turnover in the potential consumer base. Similarly, we have examined how the target demand and the steady-state price vary with the extent to which these myopic consumers base their assessment of future demand on a stubborn expectation, and the level of

this expectation. While the extent to which each of these parameters characterize specific network goods and industries is an empirical question, managers who assess these parameters appropriately for their products can use our theory as a general basis on which to fine-tune the details of their pricing strategy.

Natural extensions to our theory would model competing network goods, or myopic consumers who make imperfect observations of demand. Additionally, we model the rate of attention and the rate of replacement as exogenous variables, although sellers may in fact be able to influence these by making advertising and branding investments. An interesting direction for future research would be to extend our model to permit investments of this kind. This may also indicate how such investments should vary over time, since the impact of a change in either parameter on the seller's profits depends on its timing.

Finally, building on recent models of local network effects (Sundararajan, 2005b, Tucker, 2004), the rate at which consumers pay attention to products may not be constant across the population, but may be influenced by the adoption decisions of other consumers that one is locally "connected" to. This represents an interesting extension to our model of bounded rationality, one that is especially pertinent to network goods, and a direction of research we hope to pursue in the future.

6 Bibliographic Remarks and References

Our results add to a broad theoretical and applied literature on network effects¹. The seminal papers, by Katz and Shapiro (1985, 1986) and by Farrell and Saloner (1985, 1986), and a large majority of the literature that followed, have focused more on analyses of oligopoly rather than monopoly pricing, and in contrast with our paper, almost always use the model of consumer behavior embodied in the concept of rational expectations. An exposition of the theory of rational expectations in economic analysis can be found in Radner (1982), and its use in defining "fulfilled-expectations" outcomes in the presence of network effects is described in Katz and Shapiro (1985).

The static model of network effects underlying our dynamic model is based on Rohlfs (1974), who provided the first model of monopoly pricing for a network good, and on the subsequent exposition by Economides (1996). The discussion of dynamic pricing closest to ours that we are

¹Comprehensive surveys of models of network effects can be found in Farrell and Klemperer (2004), and in Economides (1996).

aware of is by Dhebar and Oren (1985). Their formulation differs from ours in several respects, but in general it can be interpreted as incorporating a kind of bounded rationality on the part of the consumers. A special case of their model (Section 4 of their paper, with the parameter $\alpha = 0$) leads to a law of motion that is mathematically isomorphic to our benchmark case of purely myopic consumers with a uniform distribution of types (Section 3.2 above), and for that case they derive the optimal price trajectory. For their general model they discuss properties of the optimal price trajectory if the initial customer base is “small,” and indicate that the monopolist will eventually price to keep the demand at a steady state (what we would call a “target”), but they do not obtain a complete characterization of the optimal policy. (For a related discussion, see Dhebar and Oren, 1986).

Cabral, Salant and Woroch (1999) study the dynamic pricing of a durable network good in a two-stage model with rational consumers, where they illustrate how the presence of network effects may overturn Coasian dynamics and lead to first period pricing that is lower than second-period pricing. Fudenberg and Tirole (2000) model dynamic pricing by a monopolist who sells a network good to overlapping generations of consumers who live for two periods, although they assume perfect rationality on the part of their consumers. Related papers that study single-period monopoly price discrimination based on a model of rational demand expectations include those by Oren, Smith and Wilson (1982) and by Sundararajan (2004, 2005a).

Shared information systems often display network effects, and our model may thus inform the literature on the adoption of such systems. For example, Riggins, Kreibel and Mukhopadhyay (1994) model the two-stage adoption of an interorganizational system with positive and negative adoption externalities. While their model uses the standard notion of fulfilled expectations, they do discuss a case with myopic adopters. They show that subsidies are often necessary to induce adoption in the first stage, a result qualitatively similar to ours. Wang and Seidmann (1995) examine a related problem for the adoption of EDI in a two-sided network of buyers and suppliers, incorporating not just positive network effects from higher adoption, but negative (or “competitive”) externalities imposed by a buyer (supplier) on other buyers (suppliers) by their adoption; a similar tradeoff is modeled by Westland (1992) as well. Nault and Dexter (2005) provide a general model of pricing by a monopoly provider of a “network alliance” service, wherein the number of adopters and the investments made by these adopters each influence the demand for every adopter’s products.

They show that, when combined with an exclusivity arrangement with participating members, the provider's optimal commission level restricts membership in the alliance.

The bounded rationality of agents in our model leads to a demand adjustment process that is "viscous", and is similar in this regard to the model of Radner and Richardson (2003) and of Radner (2003). These papers, however, model a good of constant value, and do not study network effects – rather, the rate at which demand adjusts to price announcements by sellers varies in proportion to the magnitude of the difference between the announced price and each consumer's willingness to pay. A model of network effects with boundedly rational consumers that is closely related to ours is by Arthur (1989), who studies adoption choices between two competing durable network goods. In his model, myopic consumers make their choices based on the current market share of each good. He shows that over time, the market share of one of the goods will tend to 100%, though one cannot predict ex-ante which of the two goods it would be, and outcomes are path-dependent. He does not model the choice of price, dynamic or otherwise, instead implicitly assuming that prices are fixed and exogenously specified. An extension of our model of bounded rationality to one with two competing sellers of network goods may provide insight into whether his results continue to hold when sellers can strategically alter their prices over time, although this extension remains unsolved at this time.

References

1. Arthur, B., 1989. Competing Technologies, Increasing Returns and Lock-in by Historical Events. *Economic Journal* 99, 106-131.
2. Cabral, L., Salant, D., and Woroch, G., 1999. Monopoly Pricing with Network Externalities. *International Journal of Industrial Organization* 17, 199-214.
3. Dhebar, A. and Oren, S., 1985. Optimal Dynamic Pricing for Expanding Networks. *Marketing Science* 4, 336-351.
4. Dhebar, A. and Oren, S., 1986. Dynamic Nonlinear Pricing in Networks with Interdependent Demand. *Operations Research* 34, 384-394.
5. Economides, N., 1996, The Economics of Networks. *International Journal of Industrial Organization* 14, 673-699.
6. Farrell, Joseph, and P. Klemperer, 2004. Coordination and Lock-in: Competition with Switching Costs and Network Effects. <http://www.paulklemperer.org/>
7. Farrell, J. and G. Saloner, 1985, Standardization, Compatibility, and Innovation. *Rand Journal of Economics* 16, 70-83.

8. Fudenberg, D. and J. Tirole, 2000. Pricing a Network Good to Deter Entry. *Journal of Industrial Economics* 48, 373-390.
9. Gamkrelidze, R., 1978. *Principles of Optimal Control Theory*. Plenum, New York.
10. Katz, M. and C. Shapiro, 1985. Network Externalities, Competition and Contracting. *American Economic Review* 75, 424-440.
11. Kaufmann, R., McAndrews, J. and Wang, Y., 2000. Opening the "Black Box" of Network Externalities in Network Adoption. *Information Systems Research* 11, 61-82.
12. Nault, B. and Dexter, A., 2005. Membership and Incentives in Network Alliances. *IEEE Transactions on Engineering Management* (forthcoming).
13. Oren, S., Smith, S., and Wilson, R., 1982. Nonlinear Pricing in Markets with Interdependent Demand. *Marketing Science* 1, 287-313.
14. Riggins, F., Kriebel, C., and Mukhopadhyay, T., 1994. The Growth of Interorganizational Systems in the Presence of Network Externalities. *Management Science* 40, 984-998.
15. Rohlfs, J., 1974. A Theory of Interdependent Demand for a Communication Service. *Bell Journal of Economics* 10, 16-37.
16. Radner, R. 1982. Equilibrium Under Uncertainty. In K. Arrow and M. Intriligator, eds., *Handbook of Mathematical Economics* (North Holland, Amsterdam), 923-1006.
17. Radner, R. 2003. Viscous Demand. *Journal of Economic Theory* 112, 189-231.
18. Radner, R. and Richardson, T., 2003. Monopolists and Viscous Demand. *Games and Economic Behavior* 45, 442-464.
19. Radner, R. and Sundararajan, A., 2005. Bounded Rationality and Network Effects. Working Paper CeDER-05-08, Center for Digital Economy Research, New York University.
20. Seidmann, A., and Wang, E., 1995. Electronic Data Interchange: Competitive Externalities and Strategic Implementation Policies. *Management Science* 41, 401-418.
21. Sundararajan, A., 2004. Nonlinear Pricing and Type-Dependent Network Effects. *Economics Letters* 83, 107-113.
22. Sundararajan, A., 2005a. Network Effects, Nonlinear Pricing and Entry Deterrence. Available at <http://ssrn.com/abstract=382962>.
23. Sundararajan, A., 2005b. Local Network Effects and Network Structure. Available at <http://ssrn.com/abstract=650501>.
24. Westland, C., 1992. Congestion and Network Externalities in the Short-Run Pricing of Information Systems Services. *Management Science* 38, 992-1009

7 Appendix

Proof of Theorem 1

Consider an arbitrary target policy, π , and let s denote the corresponding target. Suppose that the initial state q is less than s . We shall first calculate the optimal target, s , starting from $q(0) = q$. Until it reaches s , $q(t)$ solves

$$q'(t) = \lambda[1 - q(t)]. \quad (49)$$

The unique solution to the differential equation in (49) with the initial condition $q(0) = q$ is:

$$q(t) = 1 - e^{-\lambda t}(1 - q). \quad (50)$$

As a consequence, if the initial state is $q < s$, the time T at which $q(T) = s$ solves

$$s = 1 - e^{-\lambda T}(1 - q), \quad (51)$$

and therefore

$$T = \frac{1}{\lambda} \log \left(\frac{1 - q}{1 - s} \right). \quad (52)$$

Therefore, under the policy π , the value function is:

$$V_\pi(q) = sP(s) \left(\int_T^\infty e^{-rt} dt \right), \quad (53)$$

which simplifies to:

$$V_\pi(q) = \frac{1}{r} \left(\frac{1 - q}{1 - s} \right)^{-\left(\frac{r}{\lambda}\right)} s^2(1 - s). \quad (54)$$

Equation (54) can be rewritten as:

$$V_\pi(q) = \frac{1}{r} \left[(1 - q)^{-\left(\frac{r}{\lambda}\right)} \right] \left[s^2(1 - s)^{\left(1 + \frac{r}{\lambda}\right)} \right]. \quad (55)$$

Differentiating (55) with respect to s yields:

$$\frac{dV_\pi(q)}{ds} = \frac{1}{r} \left[(1 - q)^{-\left(\frac{r}{\lambda}\right)} \right] \left[2s(1 - s)^{\left(1 + \frac{r}{\lambda}\right)} - \left[\left(1 + \frac{r}{\lambda}\right) s^2(1 - s)^{\left(\frac{r}{\lambda}\right)} \right] \right]. \quad (56)$$

For $0 < s < 1$, the right-hand side of (55) is strictly quasiconcave in s . Additionally, $\frac{dV_\pi(q)}{ds} = 0$ at $s = 0$ and $s = 1$, which are minima for which $V_\pi(q) = 0$ (In fact, both these statements are true for any function of the form $Kx^a(1 - x)^b$ for $a, b \geq 1$).

As a consequence, the value $\sigma \in [0, 1]$ that maximizes $V_\pi(q)$ with respect to s solves

$$2(1 - \sigma) = \left(1 + \frac{r}{\lambda}\right) \sigma, \quad (57)$$

which yields

$$\sigma = \frac{2\lambda}{3\lambda + r}. \quad (58)$$

The corresponding price is

$$P(\sigma) = \phi \equiv \frac{2\lambda(\lambda + r)}{(3\lambda + r)^2}. \quad (59)$$

(Note that when $r = 0$, this yields the rational expectations equilibrium quantity $q^* = 2/3$ and price $p^* = 2/9$. Also note that the price ϕ in (59) satisfies (??) for each $r \geq 0, \lambda > 0$.)

Correspondingly, if the initial state is $q > s$, until it reaches s , $q(t)$ solves

$$q'(t) = -\lambda q(t), \quad (60)$$

which corresponds uniquely to the demand trajectory:

$$q(t) = qe^{-\lambda t}, \quad (61)$$

and a similar computation yields the value function:

$$V_\pi(q) = \frac{1}{2\lambda + r} \left(q^2 + q^{-\frac{r}{\lambda}} s^{(2+\frac{r}{\lambda})} \left[\frac{2\lambda(1-s) - rs}{r} \right] \right) \quad (62)$$

which is also maximized with respect to s by the value of σ in (58).

Denote by π^* the target policy with target σ . We shall now show that the target policy π^* is optimal. For any given policy μ , if its value function is continuously differentiable, then the corresponding *Bellmanian Functional* is defined by

$$B_\mu(q, p) = pq - rV_\mu(q) + V'_\mu(q) m(q, p). \quad (63)$$

According to a well-known proposition, a policy μ is optimal if it satisfies the Hamilton-Jacobi-Bellman condition:

$$B_\mu(q, p) \leq 0 \text{ for all } q, p. \quad (64)$$

An alternative form for the last condition is

$$\mu(q) = \arg \max_p B_\mu(q, p). \quad (65)$$

This follows from the fact that, for all q ,

$$B_\mu[q, \mu(q)] = 0,$$

which is readily verified. (In fact, this identity is true for any stationary policy whose value function is C^1 .) Hence, from the above,

$$B_{\pi^*}[q, \pi^*(q)] = \pi^*(q)q - rV_{\pi^*}(q) + \lambda \left(1 - q - \frac{\pi^*(q)}{q} \right) V'_{\pi^*}(q) = 0. \quad (66)$$

It will be useful to define

$$G(q) \equiv q^2 - \lambda V'_\pi(q). \quad (67)$$

and write $B(q, p)$ in the form,

$$B(q, p) = pq - rV_\pi(q) + \lambda \left(1 - q - \frac{p}{q}\right) V'_\pi(q) \quad (68)$$

$$= -rV(q) + \lambda(1 - q)V'(q) + \frac{p}{q}G(q). \quad (69)$$

Thus $B(q, p)$ is linear in p , and the coefficient of p is $\frac{G(q)}{q}$. Hence

$$\arg \max_p B_\pi(p, q) = \begin{cases} 0, & \text{if } G(q) < 0, \\ q, & \text{if } G(q) > 0. \end{cases} \quad (70)$$

It will also be useful to define the *stay-where-you-are* policy by

$$p(t) = P[q(t)]. \quad (71)$$

With a slight abuse of notation, we denote this policy by P . With this policy, $q(t) = q(0)$ for all $t > 0$ (see (31)), and its value function is

$$V_P(q) = \frac{P(q)q}{r} = \frac{q^2(1 - q)}{r}. \quad (72)$$

Case 1. $0 < q < \sigma$: In this case $\pi(q) = 0$, and

$$B_\pi[q, \pi(q)] = -rV_\pi(q) + \lambda(1 - q)V'_\pi(q) = 0, \quad (73)$$

so

$$\lambda V'_\pi(q) = \frac{rV_\pi(q)}{1 - q}, \quad (74a)$$

$$G(q) = q^2 - \frac{rV_\pi(q)}{1 - q}, \quad (74b)$$

and from (72), it follows that

$$G(q) < 0 \Leftrightarrow V_\pi(q) > V_P(q). \quad (75)$$

Suppose that the monopolist uses the policy π for $0 \leq t < u$, and then switches to the "stay-where-you-are" policy P , from then on. Since her price is zero for $0 \leq t < u$, her resulting profit will be

$$g(u) \equiv e^{-ru}V_P[q(u)] = \left(\frac{1}{r}\right) e^{-ru}[q(u)]^2[1 - q(u)], \quad (76)$$

where $q(t)$ is determined by the differential equation $q'(t) = 1 - q(t)$ on the interval $[0, T]$, with $q(0) = q$. Note that

$$g(0) = V_P(q), \quad (77)$$

$$g(T) = V_\pi(q), \quad (78)$$

where, as before, T is the time at which $q(t)$ reaches the target σ under the policy π . Differentiating (76) with respect to u , and simplifying the resulting expression yields

$$g'(u) = \left(\frac{1}{r}\right) e^{-ru} [q(u)] [1 - q(u)] [2\lambda - (3\lambda + r)q(u)] > 0 \quad \text{for } 0 \leq u < T, \quad (79)$$

since

$$q(u) < \sigma = \frac{2\lambda}{3\lambda + r} \quad \text{for } 0 \leq u < T. \quad (80)$$

Hence, $g(u)$ is strictly increasing in u and so

$$V_\pi(q) = g(T) > g(0) = V_P(q), \quad (81)$$

and using (75), $B_\pi(q, p)$ is maximized at $p = 0$.

Case 2. $q > \sigma$: In this case $\pi(q) = q$. Using an analogous argument, we find that

$$\lambda V'_\pi(q) = \frac{-rV_\pi(q) + q^2}{q}, \quad (82)$$

which leads to a condition similar to (75),

$$G(q) > 0 \Leftrightarrow V_\pi(q) > V_P(q). \quad (83)$$

The analogous expression for g is

$$g(u) \equiv q \int_0^u e^{-rt} q(t) dt + e^{-ru} V_P[q(u)] \quad (84)$$

$$= q \int_0^u e^{-rt} q(t) dt + \left(\frac{1}{r}\right) e^{-ru} [q(u)]^2 [1 - q(u)], \quad (85)$$

where $q(t)$ is defined by the differential equation

$$q'(t) = -\lambda q(t), \quad q(0) = q \quad (86)$$

in $[0, T)$. Differentiating (85) with respect to u yields:

$$g'(u) = e^{-ru} \left([q(u)]^2 - [q(u)]^2 [1 - q(u)] + \frac{1}{r} (2q(u) - 3[q(u)]^2) q'(u) \right),$$

which simplifies to:

$$g'(u) = \frac{q(u)}{r} e^{-ru} [(3\lambda + r)q(u) - 2\lambda] q(u), \quad (87)$$

which is strictly positive, since

$$q(u) > \sigma = \frac{2\lambda}{3\lambda + r} \quad \text{for } 0 \leq u < T. \quad (88)$$

Therefore, $g(u)$ is strictly increasing in u and so

$$V_\pi(q) = g(T) > g(0) = V_P(q), \quad (89)$$

and therefore, $B_\pi(q, p)$ is maximized at $p = q$.

Finally, note that, from (74a) and (82),

$$V'_\pi(\sigma^-) = V'_\pi(\sigma^+) = V'_\pi(\sigma) = \frac{\sigma^2}{\lambda}, \quad (90)$$

$$G(\sigma) = 0, \quad (91)$$

so V_π is continuously differentiable for all q , and $B_\pi(\sigma, p)$ is independent of p . Hence B_π satisfies the Bellman Optimality Condition, which completes the proof.

Proof of Corollary 1

Differentiating both sides of (31) with respect to λ yields

$$\frac{dV(\sigma)}{d\lambda} = \frac{8r\lambda}{(r+3\lambda)^4}. \quad (92)$$

Similarly, while taking into account that σ is a function of λ , differentiating both sides of with respect to λ and simplifying yields

$$\frac{dV(q_0)}{d\lambda} = \frac{4(r+\lambda)}{(r+3\lambda)^3} \left(\frac{1-\sigma}{1-q_0} \right)^{\frac{r}{\lambda}} \text{Ln} \left(\frac{1-q_0}{1-\sigma} \right), \quad (93)$$

which is strictly positive when $(1-q_0) > (1-\sigma)$.

Proof of Lemma 3

Denote the demand from consumers of type θ in period t as $w(\theta, t)$. First, consider the fraction of consumers $(1-ch)$ who remain from the prior period. Recall that a fraction λh of these consumers of type θ notice $p(t)$, form a shared expectation of demand $q_E(t, h)$, and decide whether or not to adopt the product for period t . Therefore, if $\theta \geq [p(t)/q_E(t, h)]$, each consumer in the fraction $\lambda h(1-ch)$ adopts the product, and if $\theta < [p(t)/q_E(t, h)]$, then none of these consumers adopt the product. The remaining fraction $(1-\lambda h)(1-ch)$ continue to do in period t what they were doing in period $t-h$. Next, of the fraction ch of "new" consumers, a total fraction $(\lambda h)(ch)$ notice the product and adopt it if $\theta \geq [p(t)/q_E(t, h)]$, and a total fraction $ch(1-\lambda h)$ do not notice it, and therefore, do not adopt it, independent of the value of θ . It follows that:

$$w(\theta, t) = \begin{cases} \lambda h + (1-ch)(1-\lambda h)w(\theta, t-h), & \theta \geq p(t)/q_E(t, h), \\ (1-ch)(1-\lambda h)w(\theta, t-h) & \theta < p(t)/q_E(t, h), \end{cases}, \quad (94)$$

and therefore,

$$w(\theta, t) - w(\theta, t-h) = \begin{cases} \lambda h - [h(\lambda+c) - c\lambda h^2]w(\theta, t-h), & \theta \geq p(t)/q_E(t, h), \\ -[h(\lambda+c) - c\lambda h^2]w(\theta, t-h), & \theta < p(t)/q_E(t, h). \end{cases}. \quad (95)$$

Dividing both sides by h and letting h tend to zero yields the time rate of change of demand for consumers of type θ :

$$\frac{dw(\theta, t)}{dt} \equiv \lim_{h \rightarrow 0} \left(\frac{w(\theta, t) - w(\theta, t-h)}{h} \right) = \begin{cases} \lambda - [c + \lambda]w(\theta, t), & \theta \geq p(t)/q_E(t), \\ -[c + \lambda]w(\theta, t), & \theta < p(t)/q_E(t). \end{cases} \quad (96)$$

Proceeding as in Section 2.2, we compute $q'(t)$ as

$$q'(t) = \int_0^1 \left(\frac{dw(\theta, t)}{dt} \right) dF(\theta), \quad (97)$$

which simplifies to

$$q'(t) = \lambda \left[1 - F \left(\frac{p(t)}{q_E(t)} \right) \right] - [\lambda + c]q(t), \quad (98)$$

and using the fact that $q_E(t) = q(t)$ for myopic customers, the result follows.

Proof of Theorem 2

The law of motion can be rewritten as

$$m(q, p) = \begin{cases} 0, & q(t) = 0, \\ \lambda[1 - p/q - kq], & 0 < q \leq 1, 0 \leq p \leq q, \\ -\infty, & 0 < q \leq 1, p > q, \end{cases} \quad (99)$$

where

$$k \equiv 1 + \frac{c}{\lambda} \geq 1. \quad (100)$$

Now consider an arbitrary target policy with target s . Starting at $q < s$, until $q(t)$ reaches s , $q(t)$ satisfies the differential equation,

$$q'(t) = \lambda[1 - kq(t)], \quad q(0) = q, \quad (101)$$

whose solution is

$$q(t) = \frac{1}{k} [1 - e^{-\lambda kt} (1 - kq)], \quad (102)$$

and proceeding as in the proof of Theorem 1 yields the value function

$$V(q) = \frac{1}{r} s^2 [1 - ks]^{(1 + \frac{r}{k\lambda})} (1 - kq)^{-\frac{r}{k\lambda}}, \quad q < s \quad (103)$$

The RHS of (103) is strictly quasiconcave for $0 < s \leq 1$, and it is maximized at

$$\sigma = \frac{2\lambda}{3k\lambda + r}, \quad (104)$$

which is our candidate optimal target. Correspondingly, starting at $q > s$, the value function solves to being

$$V(q) = \frac{1}{2k\lambda + r} \left(q^2 + q^{-\frac{r}{k\lambda}} s^{(2 + \frac{r}{k\lambda})} \left[\frac{2\lambda(1 - ks) - rs}{r} \right] \right), \quad q > s, \quad (105)$$

which is also maximized in s at the value of σ in (104).

Now, the Bellmanian functional for the target policy with target σ is

$$B(q, p) = pq - rV(q) + m(q, p)V'(q), \quad (106)$$

which simplifies to

$$B(q, p) = p \left[q - \frac{\lambda V'(q)}{q} \right] - [rV(q) + \lambda - \lambda k V'(q)]. \quad (107)$$

Recalling the function $G(q)$ defined in (67)

$$G(q) = q^2 - \lambda V'(q), \quad (108)$$

it follows again from (107) that

$$\arg \max_p B(q, p) = \begin{cases} 0, & \text{if } G(q) < 0, \\ q, & \text{if } G(q) > 0. \end{cases} \quad (109)$$

Now, from below, differentiating both sides of (103) with respect to q and rearranging yields:

$$V'(q) = \frac{\sigma^2}{\lambda} \left(\frac{1 - k\sigma}{1 - kq} \right)^{[1 + \frac{r}{k\lambda}]}, \quad (110)$$

which in turn implies that,

$$G(q) = q^2 \left(1 - \frac{f_0(\sigma)}{f_0(q)} \right), \quad q < \sigma, \quad (111)$$

where

$$f_0(x) \equiv x^2 [1 - kx]^{[1 + \frac{r}{k\lambda}]} \quad (112)$$

However, the function $f_0(x)$ is maximized at $x = \sigma$, which implies that $f_0(q) < f_0(\sigma)$ for $q < \sigma$, and therefore, according to (111), $G(q) < 0$ for $q < \sigma$. A similar computation of $V'(q)$, which is omitted, verifies that $G(q) > 0$ for $q > \sigma$. Finally, one can verify that $V_1(\sigma-, \sigma) = V_1(\sigma+, \sigma)$, and this completes the proof of the theorem.

Proof of Theorem 3

(a) Consider an arbitrary target policy with target s . Proceeding as we did in Section 3.1, we shall first characterize the optimal target, starting from $q(0) = q$. We begin with the case of $q < s$. Until $q(t)$ reaches s , it satisfies the differential equation,

$$q'(t) = \lambda [1 - q(t)]. \quad (113)$$

As in Section 3.1, the value function for the target policy with target s is:

$$V(q, s) = \left(\frac{1 - s}{1 - q} \right)^\rho \frac{P(s)s}{r}, \quad (114)$$

where $\rho = \frac{r}{\lambda}$.

Using the expression (45) for $P(s)$, we have

$$V(q, s) = \frac{f(s)}{r(1 - q)^\rho}, \quad \text{where} \quad (115)$$

$$f(s) = (1 - s)^{\rho+1} [\gamma s^2 + (1 - \gamma)\omega s].$$

Hence the target that maximizes $V(q, s)$ is the value of s that maximizes $f(s)$. One verifies that

$$f'(s) = (1 - s)^\rho G(s), \quad \text{where} \quad (116)$$

$$G(s) = -\gamma(\rho + 3)s^2 + [2\gamma - (\rho + 2)(1 - \gamma)\omega]s + (1 - \gamma)\omega.$$

Note that $f'(s)$ and $G(s)$ have the same sign. Also, G is quadratic and concave, and $G(0) = (1 - \gamma)\omega > 0$. Hence f is maximized at the larger of the two roots of $G(s) = 0$. Call this root $\sigma(\gamma, \omega)$; it is the optimal target. Note that it is independent of the starting state, q .

We now show that, for $q < \sigma(\gamma, \omega)$, the target policy with target $\sigma(\gamma, \omega)$ is optimal among all policies. For this purpose, we abbreviate $\sigma(\gamma, \omega)$ to σ . The Bellmanian functional for this policy is

$$B(q, p) = pq - rV(q, \sigma) + \lambda \left[1 - \frac{p}{\gamma q + (1 - \gamma)\omega} - q \right] V_1(q, \sigma). \quad (117)$$

Differentiating with respect to p , we have

$$B_2(q, p) = q - \frac{\lambda V_1(q, \sigma)}{\gamma q + (1 - \gamma)\omega}. \quad (118)$$

From (115), we have

$$V_1(q, \sigma) = \frac{f(\sigma)}{\lambda(1 - q)^{\rho+1}},$$

and hence

$$B_2(q, p) = q - \frac{f(\sigma)}{(1 - q)^{\rho+1} [\gamma q + (1 - \gamma)\omega]}.$$

It follows that $B_2(q, p) < 0$ if and only if

$$\begin{aligned} (1 - q)^{\rho+1} [\gamma q^2 + (1 - \gamma)\omega q] &< f(\sigma), \text{ or} \\ f(q) &< f(\sigma), \end{aligned}$$

which is true for $q < \sigma$. This completes the proof of the optimality of the target policy with target σ in Case 1. The argument for Case 2, $q > \sigma$, is analogous, and is omitted. Finally, one can verify that $V_1(\sigma-, \sigma) = V_1(\sigma+, \sigma)$. This completes the proof of Part (a) of the theorem.

To prove Part (b), write $G(s)$ in (116) in the form

$$\begin{aligned} G(s, \gamma) &= \gamma g_b(s) + (1 - \gamma)g_a(s), \text{ where} \\ g_a(s) &= -(\rho + 3)s^2 + 2s, \\ g_b(s) &= -(\rho + 2)\omega s + \omega. \end{aligned} \quad (119)$$

Recall that $\sigma(\gamma, \omega)$ is the larger root of

$$G[s, \gamma] = 0.$$

A standard "comparative statics" calculation yields

$$\sigma_1(\gamma, \omega) = -\frac{g_b[\sigma(\gamma, \omega)] - g_a[\sigma(\gamma, \omega)]}{\gamma g'_b[\sigma(\gamma, \omega)] + (1 - \gamma)g'_a[\sigma(\gamma, \omega)]}. \quad (120)$$

Let σ_b be the positive root of $g_b(s) = 0$ (the other root is 0), and let σ_a be the root of $g_a(s) = 0$. Then

$$\sigma_b = \frac{2}{\rho + 3}, \quad \sigma_a = \frac{1}{\rho + 2}. \quad (121)$$




Figure 3: Illustrates the proof of part (b) of Theorem 3.

Note that

$$\frac{\sigma_b}{2} < \sigma_a < \sigma_b. \quad (122)$$

Note also that (1) $g_b(s)$ is decreasing and positive for $\frac{\sigma_b}{2} \leq s < \sigma_b$; (2) $g_a(s)$ is decreasing, and is negative for $\sigma_a < s \leq \sigma_b$, and (3) $\sigma_a < \sigma(\gamma, \omega) < \sigma_b$ for $0 < \gamma < 1$ (see Figure 3). Hence, by (120) and (122), $\sigma_1(\gamma, \omega) > 0$ for $0 < \gamma < 1$, which completes the proof of Part (b) of the theorem.

To prove part (c), first notice that, independent of the value of ω , part (b) establishes that for $0 < \gamma < 1$:

$$\frac{1}{2 + \rho} < \sigma(\gamma, \omega) < \frac{2}{3 + \rho}. \quad (123)$$

Also, from the second line of (116), σ is defined by

$$-\gamma(\rho + 3)[\sigma(\gamma, \omega)]^2 + [2\gamma - (\rho + 2)(1 - \gamma)\omega]\sigma(\gamma, \omega) + (1 - \gamma)\omega = 0. \quad (124)$$

Differentiating both sides of (124) with respect to ω and rearranging yields:

$$\sigma_2(\gamma, \omega) = - \left(\frac{(1 - \gamma)[(2 + \rho)\sigma(\gamma, \omega) - 1]}{2\gamma[(3 + \rho)\sigma(\gamma, \omega) - 1] + \omega[1 - \gamma](2 + \rho)} \right). \quad (125)$$

From (123), $(2 + \rho)\sigma(\gamma, \omega) > 1$, and thus both the numerator and the denominator of the expression in parentheses on the LHS of (125) are strictly positive. This completes the proof.