Dynamic Pricing of Network Goods with Boundedly Rational Consumers

Roy Radner and Arun Sundararajan
Stern School of Business, New York University

September 2004

1 Introduction

The theory of rational expectations is widely used to describe economic situations in which the outcome depends partly upon what people expect to happen. This idea is especially popular in models of network effects or externalities, wherein the value of a good or service (henceforth referred to as a "network good") to each consumer is influenced by the consumption choices made by some or all other consumers. In this literature (for instance, in the seminal papers by Katz and Shapiro, 1985, and by Farrell and Saloner, 1985), potential 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."

An important assumption underlying this solution is that consumers can in fact all "guess" the correct equilibrium level of demand. This assumption is consistent with a model of perfectly (or unboundedly) rational consumers. Some notion of perfect rationality is at the base of most current economic analysis, even though most of us 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 economic analysis based on the approximation of unboundedly rationality agents is (reasonably) correct.

In the specific case of network goods, however, the predictions of rational expectations models do not appear to be a good description of economic reality. For example, it has been observed that outcomes in markets for network goods are often path-dependent. This would not be the case were consumers able to form rational expectations; it also highlights the importance of the dynamics of the adoption process for eventual outcomes.

Our objective in this study is to present an alternative model of demand for a network good, in which consumers are not "rational enough" to be able to compute equilibrium demand. Instead, they make their consumption choices based on the current level of demand. The intensity of the network effect is heterogeneous across customers. Different customers therefore have a different marginal utility of total demand, which varies according to a given distribution, referred to as the distribution of consumer types. These instantaneous adoption choices of the consumers continuously influence the rate at which demand adjusts over time. A 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 an exogenous parameter k, which models the "extent" of consumer rationality. Stated differently, when k is higher, demand adjusts more rapidly, and eventually instantaneously as k tends to infinity, which is the case when consumers are unboundedly rational.

This paper is part of an ongoing program of research whose broad objective is to explore the conditions under which the assumption of unbounded rationality in economic models is a reasonable one.

2 Outline of our Model

We motivate the continuous-time formulation with a discrete-time version. A network good is provided by a monopolist, who sells the good one period at a time. (Think of the good as a service.) A unit mass of a continuum of consumers is indexed by a "type" parameter θ in the unit interval. Let F be the cumulative distribution function of θ, i.e., the fraction of consumers with type less than or equal to θ is $F(\theta)$. For simplicity, we assume F to be absolutely continuous. If the price is p, and a consumer expects the total demand (i.e., the mass of consumers) in a given period to be q, then the consumer will buy the good if and only if $\theta q \geq p$. The above description implies that the demand, q, is between zero and one. We also assume that the price, p, is constrained to be nonnegative, and is bounded above.

If consumers expect the total mass of customers in a period to be q then the total demand in that period will be

$$Q(q,p) = 1 - F\left(\frac{p}{q}\right). \tag{1}$$

Suppose that consumers are "myopic," in that they expect the total demand in a given period to be equal to the demand in the previous period. Then if the price is p, the change in demand will be

$$Q(q,p) - q = 1 - F\left(\frac{p}{q}\right) - q. \tag{2}$$

Accordingly, in the continuous time version, assume that if at time t (≥ 0) the demand and price are $q(t)$ and $p(t)$, respectively, then the time-rate of change of demand will be

$$q'(t) = m[q(t), p(t)], \tag{3}$$

where

$$m(q,p) = \begin{cases} 0, \quad q = 0, \\ k[1 - F(p/q) - q], \quad 0 < q \leq 1, 0 \leq p \leq q, \\ -\infty, \quad 0 < q \leq 1, p > q. \end{cases} \tag{4}$$

where $k > 0$ is a given parameter that measures the "speed of response" of the myopic decision makers. We restrict the monopolist's choice of price as follows:

$$0 \leq p \leq q. \tag{5}$$

The monopolist wants to choose a price trajectory $p(t)$ to maximize her profit. Assume, for simplicity, that her (marginal) cost of providing the service is zero; then her total discounted profit is

$$\int_0^\infty e^{-rt} p(t) q(t) dt, \tag{6}$$

where $r > 0$ is her given discount rate.

We shall analyze the maximization problem 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:

$$p(t) = \alpha[q(t)]. \tag{7}$$

Note that the function α does not change in time. Of course, a policy is admissible only if the differential equation (4) has a unique solution starting from any initial state $q(0)$.

The value of a policy α at an initial state $q(0) = q$ is the corresponding profit,

$$V_\alpha(q) = \int_0^\infty e^{-rt} \alpha[q(t)] q(t) dt. \tag{8}$$

Define

$$V(q) = \sup_{\alpha} V_{\alpha}(q), \tag{9}$$

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

3 Summary of Results

1. Rational Expectations Equilibrium is not Optimal. An alternative theory of consumer behavior is embodied in the concept of "rational-expectations equilibrium." Imagine that, when faced with a price p, each consumer correctly predicts the total demand at that price, and decides whether or not to subscribe on the basis of that prediction. Thus the total demand at that price must satisfy

$$Q(q,p) = q.$$

Following standard terminology, we shall call such a pair (q,p), a *rational-expectations equilibrium* (REE). Of course, for any price p, the pair $(0,p)$ is a REE; such a pair will be called *degenerate*. We shall call a REE *optimal* (for the monopolist) if it maximizes the product pq in the set of REEs. Thus, an implication of the REE hypothesis is that a monopoly market equilibrium must be an optimal REE. We shall show that, under a wide range of conditions, an optimal REE cannot be a steady state for an optimal policy in the preceding model with myopic consumers. In fact, under such conditions we can expect the REE demand to be larger, and the REE price to be smaller, than the respective demand and price in a (nondegenerate) steady state of an optimal policy.

Note that, since the speed-of-response parameter, k, is strictly positive, (q,p) is a REE if and only if

$$m(q,p) = 0,$$

i.e., q is a steady state of the demand process given the constant price p. For $q > 0$ define $P(q)$ implicitly by

$$m[q, P(q)] = 0. \tag{10}$$

Of course, this equation may have no solution or multiple solutions. In the latter case, take $P(q)$ to be the largest solution. Note that $[q, P(q)]$ is a REE. Let $\boldsymbol{Q}$ be the set of demands q such that $P(q)$ exists, and for q in $\boldsymbol{Q}$ define

$$v(q) = qP(q). \tag{11}$$

In what follows we make various assumptions about the regularity of the functions P and v. In subsequent sections we shall show that these assumptions are satisfied in a "robust" set of cases.

Define

$$\hat{q} = \arg\max\left\{v(q) | q \in \boldsymbol{Q}\right\},$$

and suppose that the usual first-order condition is satisfied at $\hat{q}$, namely,

$$v'(\hat{q}) = 0. \tag{12}$$

Theorem 1 *The optimal REE cannot be a steady state of an optimal dynamic price policy.*

Detailed proofs of our theorems are available on request.

Our first theorem has established that under fairly general assumptions about the distribution of types, the profit-maximizing rational expectations equilibrium is *never* a steady state of the optimal demand trajectory with boundedly rational consumers. Specifically, starting from the profit-maximizing rational expectations price, the monopolist can always improve her profits by charging the highest possible price for a finite period, and then switching to some other price trajectory.

2. Optimal Price Trajectory for Uniformly Distributed Types. We provide a complete solution of the optimal dynamic price problem for the special case in which θ is distributed uniformly on the unit interval; thus

$$F(\theta) = \theta, \quad 0 \leq \theta \leq 1; \tag{13}$$

the law of motion is

$$m(q,p) = \begin{cases} 0, & q = 0, \\ k[1 - (p/q) - q], & 0 < q \leq 1, 0 \leq p \leq q, \end{cases} \tag{14}$$

and the price path is constrained by

$$0 \leq p(t) \leq q. \tag{15}$$

We now consider a family of policies called *target policies*. Define the (stationary) target policy with target q^* $(0 < q^* < 1)$ by

$$\mu(q) = \begin{cases} 0, & q < q^*, \\ P(q^*), & q = q^*, \\ q, & q > q^*. \end{cases} \tag{16}$$

If the initial state is less than the target q^*, then the demand will increase until it reaches the target, during which time the profit is zero. After than, the demand will remain at the target, and the profit will be $q^* P(q^*)$. Hence the total discounted profit will be

$$\int_T^\infty e^{-rt} q^* P(q^*) dt, \tag{17}$$

where T is the time at which demand reaches the target q^*. Note that there is a tradeoff between reaching a higher target and getting there sooner. Let π be the "optimal target policy," i.e., the one that maximizes the last expression. This policy π *is optimal among all policies.*

Theorem 2 *The optimal target policy π is optimal among all policies, and the optimal target is*

$$\sigma \equiv \frac{2k}{3k + r}.$$

The target σ 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 her installed base. This theorem also shows that the optimal demand target with boundedly rational consumers is always strictly lower than the equilibrium level of demand predicted by a model with rational expectations, which is $\frac{2}{3}$. As the rate at which future profits are discounted tends to zero, the optimal demand target converges to the profit-maximizing rational expectations level of demand. Furthermore, the difference between the target demand and the rational expectations demand is higher when consumers are more boundedly rational, and converges to the rational expectations demand as the parameter k controlling the rate of demand adjustment tends to infinity.

3. Optimal Price Trajectory for Concave Type Distributions. We now extend the example above for any cumulative distribution F of consumer types that is *strictly concave.* While there is no optimal stationary policy in the classical sense, we show that there is an optimal *measure-valued* policy which is independent of F, and whose structure is very similar to that of the optimal target policy derived in Theorem 2.

Assume that the cumulative distribution function, F, of consumer types is strictly increasing and strictly concave on the unit interval. Recall that the law of motion is

$$m(q,p) = \begin{cases} 0, & q = 0, \\ k[1 - F(p/q) - q], & 0 < q \leq 1, 0 \leq p \leq q. \end{cases} \tag{18}$$

We extend the set of permissible controls to include those that specify a probability measure over the set of feasible prices at each time t. Therefore, at any time t, define a *generalized control* as a probability measure

$\psi(\cdot;t)$ on $[0,q(t)]$, and Ψ as the set of all such generalized controls. Analogously, let $\mathcal{M}(q)$ be the set of measures over $[0,q]$, with elements $M[\cdot;q]$, and let $\mu(\cdot;q)$ be a stationary generalized policy that maps the state q to a measure $M[\cdot;q] \in \mathcal{M}(q)$. If $V(q)$ is continuously differentiable, the optimal stationary generalized policy $\mu^*(\cdot;q)$ satisfies the Hamilton-Jacobi-Bellman condition:

$$\mu^*(\cdot;q) = \arg\max_{\mu \in \mathcal{M}(q)} \overline{B}(q,\mu), \tag{19}$$

where

$$\overline{B}(q,\mu) \equiv q\int p d\mu(p) - rV(q) + kV'(q)\left[1 - q - \int F\left(\frac{p}{q}\right) d\mu(p)\right]. \tag{20}$$

Let $\Phi(q)$ be the subset of measures in $\mathcal{M}(q)$ such that the probability is concentrated on the endpoints of the interval $[0,q]$, i.e., for which

$$M[\{q\},q] = \phi(q), \tag{21}$$

$$M[\{0\},q] = 1 - \phi(q)] \tag{22}$$

Finally, define $\Sigma(s,q)$ to be the subset of $\Phi(q)$ determined by functions $\phi_s(q)$ of the form

$$\phi_s(q) = \begin{cases} 0, & q < s; \\ (1-s), & q = s, \\ 1, & q > s. \end{cases} \tag{23}$$

Call a stationary generalized policy a *generalized target policy with target* s if $\mu(\cdot;q) \in \Sigma(s,q)$ for each q.

Theorem 3 *If the cumulative distribution of customer types F is strictly concave, then the optimal stationary generalized policy is a generalized target policy with target:*

$$\sigma = \frac{2k}{3k+r}. \tag{24}$$

According to this theorem, 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. This theorem therefore establishes that the optimal price trajectory suggested by considering uniformly distributed types is robust (in a sense) across a large class of distributions of customer types, as is the specific deviation from the rational expectations outcome that the example established.

4. Convex Distributions of Customer Types: A Conjecture. When the cumulative distribution function, F, of consumer types is continuously differentiable and strictly increasing on the unit interval, we have conjectured that the monopolist continues to choose a demand target lower than the rational expectations demand, but varies price gradually in the neighborhood of the demand target.

Conjecture 1 *If F is convex, and sufficiently close to uniform (in an appropriately chosen metric), then there is an optimal policy α with the following characteristics: there is a "target demand," σ, such the demand increases (decreases) to σ if the initial demand is less than (greater than) σ, and reaches σ in finite time. Once the demand reaches σ it is stabilized by charging the price $P(\sigma)$. If F is not uniform, then there is a nondegenerate interval $[c_0, c_1]$ such that σ is in the interior of the interval, and when the demand $q(t)$ is in the interior of the interval, the optimal price $\alpha[q(t)]$ is strictly between 0 and $q(t)$.*

We have a heuristic argument for this conjecture, but do not have a complete proof. This result would indicate of the extent of robustness of the example based on the uniform distribution.

4 Concluding Remarks

A characterization of the optimal price trajectory for more general customer type distributions is not available at this time. We propose to examine the equilibrium price trajectories chosen by sellers of competing network goods as our logical next step.