

R. Gibbens, R.Mason and Richard Steinberg

Internet service classes under competition

Article (Accepted version)
(Refereed)

Original citation:

Gibbens, R. and Mason, R. and Steinberg, Richard (2000) Internet service classes under competition. [IEEE journal on selected areas in communications](#), 18 (12). pp. 2490-2498. ISSN 0733-8716

DOI: [10.1109/49.898732](https://doi.org/10.1109/49.898732)

© 2000 [IEEE](#)

This version available at: <http://eprints.lse.ac.uk/23577/>

Available in LSE Research Online: September 2012

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final manuscript accepted version of the journal article, incorporating any revisions agreed during the peer review process. Some differences between this version and the published version may remain. You are advised to consult the publisher's version if you wish to cite from it.

Internet Service Classes Under Competition

Richard Gibbens, Robin Mason, and Richard Steinberg

Abstract—This paper analyzes competition between two Internet service providers (ISPs), either or both of which may choose to offer multiple service classes. In the model analyzed, a social planner who maximizes the total benefit from network usage and a profit maximizing monopolist will both form multiple service classes; but two networks competing to maximize profits will not. The reason is that a competition effect always outweighs a segmentation effect. Networks wish to offer multiple service classes in order to increase user benefits and hence charge higher prices. In doing so, however, they effectively increase the number of points in the service quality range at which they compete. Consequently, in any equilibrium competitive outcome, both ISPs offer a single service class. The analysis has particular implications for the Paris Metro pricing (PMP) proposal, which is considered in depth in this paper, since it suggests that PMP may not be viable under competition.

Index Terms—Congestion, differentiated services, Internet charging, multiproduct competition, Paris Metro pricing, quality of service.

I. INTRODUCTION AND MOTIVATION

IT IS widely recognized that Internet usage is subject to a “tragedy of the commons” [10], and that pricing of Internet resources is required to control congestion. Many proposals have been considered in the literature. The reader is referred, in particular, to [4], [9], [12], and [15]. One approach suggested by several authors is that a small number of service classes should be offered on the Internet. For example, Gupta, Stahl, and Whinston advocate a finite number of (perhaps four or less) priority classes. In Odlyzko’s “Paris Metro Pricing” (PMP) proposal [17], between two and four service classes would be generated by differential pricing on logically separate channels. (See below for more on this scheme.)

This paper assesses the viability of these service class proposals in general—and Odlyzko’s PMP in particular—when there are competing network providers. Three cases are referred to in the paper: the *social optimum*, where total benefit from network usage is maximized; *monopoly*, where a single network maximizes profit; and *duopoly*, where two networks compete to maximize their individual profits. In the model developed here, both a “social planner,” interested in maximizing the welfare

of users, and a profit maximizing monopolist will wish to use several service classes; duopolists will not.¹

The model is based on the particular proposal by Odlyzko. Some years ago on the Paris Metro, and up until quite recently on the suburban RER lines [18], users were offered a choice of travelling in first or second class carriages. The only difference was the price charged: both carriages had the same number and quality of seats, and obviously both reached the destination at the same time. The first class carriage was, however, more expensive, and consequently on average had fewer passengers in it. Those users with a strong preference for, e.g., obtaining a seat, were willing to pay the higher price; others, content to travel in what they would expect to be a more congested carriage, paid the lower second class fare.²

Odlyzko’s proposal is to apply the same scheme to packet based networks, such as the Internet. A network would be partitioned into separate logical networks, with different charges applied on each subnetwork. No guarantees of service quality would be offered; but, on average, networks charging higher prices will be less congested. Users will sort themselves according to their preferences for congestion and the prices charged on the subnetworks.

In order to assess the PMP proposal, we place the problem in a more general context of competition between networks who sell multiple products in the presence of negative externalities, i.e., congestion. There has been little analysis of multiproduct competition in the economic literature, and even less when congestion is present. The model involves two competing profit-maximizing networks who each may offer several service classes. Service classes are generated by forming “subnetworks,” differentiated by congestion level. The congestion on a subnetwork is determined by two factors: the fraction of the network’s total capacity allocated to the subnetwork; and the number of users on the subnetwork. A subnetwork with a low capacity and many users will have a high level of congestion. Quality is therefore demand-dependent, determined (in part) by the choices of networks’ prices.

Our main result is as follows. The unique outcome is that neither network subdivides its network. The intuition behind this can be understood by appealing to the analyzes by [2] and [24]. The desire to discriminate between users with heterogeneous valuations drives networks to charge different prices on separate subnetworks (i.e., produce multiple goods of different qualities).

¹See the Appendix for an analysis of the social optimum.

²There are many other examples of PMP. Reference [3] cites dentists in India who operate a “two queue scheme” in which people requiring dental treatment can pay a lower price and join a longer queue, or pay a higher price and join a shorter queue. Both sets of people are treated by the same dentist; but the second set receive treatment sooner. In the U.K., there is the choice to receive free (at the point of delivery) medical care via the National Health Service (NHS), or to pay for private treatment. While the patient sees the same doctor regardless of the option chosen, the waiting time is usually greater under the NHS.

Manuscript received October 15, 1999; revised April 15, 2000. The work of R. Gibbens was supported by a Royal Society University Research Fellowship; computing equipment was provided by the EPSRC under Grant GR/M09551. The work of R. Mason was supported by Alcatel Bell.

R. Gibbens is with the Statistical Laboratory, University of Cambridge, Cambridge, CB2 1SB, U.K. (e-mail: R.J.Gibbens@statslab.cam.ac.uk).

R. Mason is with the Department of Economics, University of Southampton.

R. Steinberg is with the Judge Institute of Management Studies, University of Cambridge.

Publisher Item Identifier S 0733-8716(00)09222-2.

This is a “segmentation effect.” Offsetting this is any increase in competition between networks which results from the use of another subnetwork (i.e., production of another good); this is a “competition” (in Champsaur and Rochet’s terms) or “expansion” (Shaked and Sutton) effect. In the model of this paper, these two effects interact so that either 1) there is no outcome that is an equilibrium³ under competition, or 2) when an equilibrium outcome exists, the competition effect always outweighs the segmentation effect, and both networks in a duopoly earn lower profits in any “multiproduct” outcome in which multiple service classes are offered than in the unique outcome in which each network offers a single service class.

The paper is structured as follows. The next section presents the model and discusses related work. Section III provides the main results of the paper concerning the outcomes of the model when there are two competing networks with fixed and equal capacities. Section IV presents numerical results to assess how allowing capacities to be freely and optimally chosen might change the conclusions of Section III. Section V concludes. Proofs are contained in the Appendix.

II. THE MODEL AND RELATED WORK

In this section, we present the basic framework and discuss related work. On joining network i , a user receives utility $U(\theta, i)$ per unit time, which has three components: a positive benefit V which is independent of which network he has joined; a disbenefit, depending on the degree of congestion on the network K^i , and his preference for (lack of) congestion θ ; and a disbenefit from having to pay a price p^i per unit time to the network for its services. Thus, the utility of a user with preference θ from joining a network i can be written as

$$U(\theta, i) = V - \theta K^i - p^i.$$

To simplify further, suppose that congestion on a network is simply the number of users divided by the capacity of the network:

$$K^i = \frac{Q^i}{C^i}$$

where Q^i is the mass of users on the network, and C^i is the network capacity. Three main assumptions are contained in this functional form for utility. First, the congestion function is linear. This may not be very realistic—it might be, for example, that the effect of congestion will be negligible at low usage levels, but will rise sharply beyond some critical level. The linear functional form captures, however, the relevant feature that congestion increases with the number of users, while maintaining tractability.⁴ Second, the utility function is separable in the various terms. In fact, there is no major loss of generality here: all that is required is that expenditure on network services is not too large relative to a user’s income (so that there are no wealth effects to be

considered). Finally, it is assumed that each user contributes equally to congestion, i.e., generates the same amount of traffic. This assumption is commented on further below.

Two further assumptions are made. First, all users join one and only one network. Second, a user pays a price per unit time for the right to be connected to network i ; hence, prices are *subscription-based* rather than *usage-based*. (In the latter case, the price paid would depend on the volume of data transmitted.) The first assumption seems more extreme than actually it is. In particular, it is compatible with users selecting networks on a per packet or call basis, provided that the decision as to which network to use is largely independent across the various packets and calls. In this case, the price charged by the network in the current model is price per packet or call. When the network choice decision is not independent across packets or calls, a fully dynamic analysis would be required; this is outside the scope of the present paper. The second assumption, that prices are subscription-based, is reasonable when users generate the same amount of traffic. When this is not the case, the networks might wish to use both subscription and usage charges. References [1] and [16] show that the degree of competition may be increased when both subscription- and usage-based pricing is allowed. These analyses suggest that including such two-part pricing might strengthen this paper’s results, since the same factors that will rule out multiple service classes in this model (see Section III) can lead to only subscription prices being charged in, e.g., [16]. There are no agreed results in this area, however; consequently, we ignore this issue in the rest of the paper.

With these assumptions, the profit of network i is $\pi^i = p^i Q^i$. (Costs of forming subnetworks, as well as any other costs, are for simplicity set to zero.)

Users differ in their preference for congestion. For example, those with elastic traffic will receive little disbenefit from congestion; they will have low values of θ . Users with inelastic traffic will be very sensitive to congestion, and will have high values of θ .⁵ To reflect the range of preferences in the population of users in the simplest manner, assume that there is a continuum of users whose θ parameters form a population distribution which is uniformly distributed on the interval $[0, 1]$. The uniform distribution is commonly used in economic models of competition between firms whose products are of different qualities; see, among many others, [22], [7], and [27]. The use of a uniform distribution is not crucial: as [2] shows, the important feature is that the density function is not too “irregular” (see [2, assumptions A.1 and A.1 bis]).⁶

A. Nash Equilibria

In this model, two networks compete to maximize individual profits. This assumption may not be as extreme as it appears. As far as users are concerned, choice is not necessarily limited by this restriction, since the two networks may offer multiple service classes. More importantly, there is good reason to suppose that, under certain circumstances, industries with congestion (and, more generally, industries in which firms sell products

³See Section II-A for a precise definition of equilibrium.

⁴In the monopoly case, it can be shown that the linearity assumption does not make a qualitative difference to the conclusions; see [3]. There is as yet no equivalent result when there is more than one network.

⁵Here, we use the terms elastic and inelastic in the sense introduced by [26]: a user with inelastic traffic is sensitive to congestion and hence delay; elastic is not.

⁶Taking the support to be $[0, 1]$ is simply a normalization.

of different qualities) may have very concentrated market structures (i.e., a small number of firms have a large market share); see the analyses of [22] and [23]. Nevertheless, the restriction to two networks is quite strong, and further work should generalize the model to allow for more networks to enter the industry. In the meantime, the current setting is the most transparent environment in which to study the effect of competition on the use of multiple service classes.

It is natural to assume that prices on the subnetworks are chosen after the decisions have been made regarding the number of subnetworks and the capacities of these subnetworks. Thus, a network has three decisions to make. First, it must choose the number of subnetworks to form. Second, it must set the capacity of its subnetworks. Finally, it must choose its prices. The non-cooperative game between networks may have, therefore, three stages. In Section III, for analytical convenience, capacities are treated as fixed; hence, there are two stages in the game in this section. The three stage game with chosen capacities is analyzed numerically in Section IV. The rest of the paper analyzes the outcomes of this game. For the basic game theoretic terminology and concepts, the reader is referred to [6].

In the two stage game, a network's strategy in the first stage is the number of subnetworks to form. A strategy for the network in the second stage consists of a pricing decision, taking the number of subnetworks as given.

This paper concentrates on pure strategies. An important finding of the paper is that equilibrium may fail to exist in the pricing stage of the game when networks form subnetworks. Note, however, that the results of [5] ensure that mixed strategy equilibria always exist in this model—that is, equilibria in which networks randomize over pure strategies (choose probability distributions over all prices) rather than a single price for sure.⁷ (A pure strategy is therefore a special case of a mixed strategy, where the probability distribution is degenerate.) Mixed strategy equilibria are not examined here, for several reasons. The first and most important is that the aim is to show that pure strategy equilibria may not exist, since the proof of this highlights the two key economic forces that are at play in the model (the segmentation and competition effects). Second, a standard criticism of mixed strategy equilibria is that they impose too large an informational burden on users; see, e.g., [6]. When choosing a network to join, users are faced not with certain price levels, but price distributions from which the final prices will be drawn. The complexity of this task is increased in the setting here, since quality is demand-dependent: in order for users to decide which subnetwork to join, they must be fully aware not only of the equilibrium strategies (probability distributions over prices) of networks, but also the choices of all other users. It is unlikely that actual users would be able to

⁷Hence, in this model, a mixed strategy in the pricing stage is a real-valued (distribution) function defined over R_+^n with a range in $[0, 1]^n$. In a simpler example, suppose that an individual has two actions open to her, A and B . A pure strategy specifies that one of the two actions should be played for sure in a particular situation; e.g., "choose A if another player chooses A ; otherwise choose B ." A mixed strategy specifies that an action be chosen with some probability; e.g., "choose A with probability p and B with probability $1-p$ if another player chooses A with probability q ; otherwise . . ." and so on. See [6] for more detail on mixed strategies.

perform this task.⁸ Third, the mixed strategy formulation in this setting does not entirely solve the problem of nonexistence of equilibrium. Randomization may generate an equilibrium; but once prices have actually been chosen (i.e., the randomness realized), the networks will adjust their prices. Hence, for the mixed strategy formulation to work, it must be that the networks are committed in some way to charge the prices that emerge from the randomization procedure. This is inconsistent with the setup in which prices are the networks' most flexible choice variable.

If a network is not divided into subnetworks, then the network's price is simply a scalar: $p \in \mathbf{R}_+$; if n subnetworks are formed, price is an n -vector: $\mathbf{p} \in \mathbf{R}_+^n$. Index the networks by I, II . Let the profit of firms I and II be denoted π^I and π^{II} , respectively; these profits are functions of the number of subnetworks formed and the prices charged. A Nash equilibrium in the second stage subgame is a pair of price vectors, $(\mathbf{p}^{I*}(n^I, n^{II}), \mathbf{p}^{II*}(n^I, n^{II}))$, such that for all n^I, n^{II} , and $\mathbf{p}^I(n^I, n^{II})$

$$\begin{aligned} & \pi^I(n^I, n^{II}, \mathbf{p}^{I*}(n^I, n^{II}), \mathbf{p}^{II*}(n^I, n^{II})) \\ & \geq \pi^I(n^I, n^{II}, \mathbf{p}^I(n^I, n^{II}), \mathbf{p}^{II*}(n^I, n^{II})) \end{aligned} \quad (1)$$

and for all n^I, n^{II} , and $\mathbf{p}^{II}(n^I, n^{II})$

$$\begin{aligned} & \pi^{II}(n^I, n^{II}, \mathbf{p}^{I*}(n^I, n^{II}), \mathbf{p}^{II*}(n^I, n^{II})) \\ & \geq \pi^{II}(n^I, n^{II}, \mathbf{p}^{I*}(n^I, n^{II}), \mathbf{p}^{II}(n^I, n^{II})). \end{aligned} \quad (2)$$

In other words, in a Nash equilibrium, no network has a unilateral incentive to change its strategy. We call a Nash equilibrium *conditionally subgame perfect* if the networks' pure strategies constitute a Nash equilibrium in every subgame in which a pure strategy equilibrium exists. This concept of conditional subgame perfection is motivated by the concept of subgame perfection, which allows mixed strategies as well. In our case, however, we rule out mixed strategies; and so it may be (and, in fact, will be) that no equilibrium exists in a particular subgame. The standard notion of subgame perfection, described in, e.g., [6] is modified to take account of this issue.

A strategy for a user is a choice of network to join, given the prices quoted by the networks. (If the user is indifferent between any two (sub)networks, his choice can be made randomly.) Any solution to this model will satisfy the following properties.⁹

Property 1: A (sub)network charging a higher price has lower congestion; e.g., in the two network case, if $p^I \geq p^{II}$, then $Q^I/C^I \leq Q^{II}/C^{II}$.

Property 2: Users who dislike congestion more will join a more expensive, less congested (sub)network; e.g., in the two network case, if $p^I \geq p^{II}$, then users with high θ will join network I ; in other words, there exists θ^* such that for $\theta \geq \theta^*$, $U(\theta, 1) \geq U(\theta, 2)$.

In the example with just two networks, the critical value θ^* identified in property 2 is the identity of the marginal user who is indifferent between network I and network II . When the two

⁸There are, however, justifications of mixed strategies that rely on the information incompleteness; see for example the famous purification theorem of [11].

⁹The proofs of these properties are straightforward, and so are omitted.

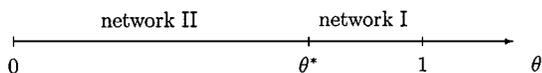


Fig. 1. The critical value of θ^* .

networks set prices p^I and p^{II} with $p^I \geq p^{II}$, and have capacities C^I and C^{II} , the mass of users who join network I is $1 - \theta^*$, and network II , θ^* (given the assumption of preferences uniformly distributed on the unit interval). This is illustrated in Fig. 1. Here θ^* is determined by the indifference relation $U(\theta^*, I) = U(\theta^*, II)$, which implies:¹⁰

$$V - \theta^* \frac{(1 - \theta^*)}{C^I} - p^I = V - \theta^* \frac{\theta^*}{C^{II}} - p^{II}$$

$$\theta^* = \frac{C^{II} + \sqrt{(C^{II})^2 + 4C^I C^{II}(C^I + C^{II})(p^I - p^{II})}}{2(C^I + C^{II})}. \quad (3)$$

B. Related Work

The model is related to the substantial body of existing work which examines charging schemes for congestible resources. The contribution of this paper is to analyze the equilibrium when multiple networks each offer one or more subnetworks. That is, the paper addresses the question of multiproduct duopoly with demand-dependent quality. None of the papers surveyed below has attempted a full analysis of this area.

The model used here is based on work by [3] and [19]. (In turn, these papers can be related to the theory of price competition with capacity constraints, e.g., [13].) Chander and Leruth show that a profit maximizing monopolist will charge the maximum number of different prices, and hence offer the maximum number of subnetworks with different qualities. By doing this, the monopolist segments the market, and can therefore extract more surplus from users and earn higher profits. Reference [19] studies duopoly outcomes with two firms, when each firm charges a single price, in an identical model setup to ours. (Part of their analysis is repeated briefly in the next section.) Reference [14] examines the duopoly case, and [21] oligopoly, again with each firm offering a single price. In contrast with these papers, we examine duopoly competition when both firms can charge more than one price, and hence offer more than one quality.

Our analysis is also related, therefore, to the literature on multiproduct competition. The majority of papers assume that the number and/or characteristics of products are fixed; [2] and [24] allow both the number, quality, and price of products to be chosen optimally. Both papers highlight two effects. Firms wish to offer a broad range of qualities, in order to segment the market. On the other hand, a firm lowers its profits when it offers qualities close to those of its rivals due to price competition; this effect drives firms toward a small quality range. The model in this paper differs from Champsaur and Rochet and Shaked and Sutton, since quality is determined by price, and costs play no part in the argument. Nevertheless, the same two effects will be important for competitive outcomes.

¹⁰Only the positive root is taken in the quadratic for θ^* , to ensure that a network attracts fewer users when it charges a higher price.

There is a small but growing literature that combines engineering with economics to provide a multidisciplinary analysis of quality of service provision on the Internet. For example, see [25] and [20]. These papers provide considerable detail on the engineering aspects of networks, and consequently adopt a more abstract approach to the economic analysis, concentrating on the existence and general inefficiency of noncooperative (Nash) equilibrium. In this paper, we concentrate on the economic aspects, constructing specific Nash equilibria to assess efficiency; consequently, the engineering aspects of the model are less detailed. These two approaches should be seen as complements—both are required to gain a complete understanding of the how service classes can and should be implemented on the Internet.

Of special relevance to our work is that of [28], who considers a problem of the supply of electricity when there is the possibility of excess demand. He confines his analysis to *priority rationing*, where customers are grouped into priority classes such that those with highest priority are supplied first. His model is therefore somewhat different to the one used in this paper; however, many of his conclusions have parallels in our setting.

In summary, many papers have considered charging and competitive outcomes when resources are congestible. A few papers have examined multiproduct competition when qualities are given. This paper combines these two approaches to address the question: will competing firms offer multiple products when quality is demand-dependent?

III. RESULTS

In order to derive analytical results, this section treats capacities as fixed and equal: $C^I = C^{II} = C$; and, when a network forms subnetworks, it assumes that the network splits its total capacity equally between the subnetworks. These assumptions will be relaxed in Section IV.

Our results, described below, confirm two results suggested by the work in [28] on priority rationing.

- 1) The only equilibrium when one network offers two classes (prices), while the other offers one, has the single-offer network's class lying between the two-offer network's classes. The two-class network earns lower profits in this equilibrium than in the single class, symmetric case. (By *symmetric*, we mean that the two networks offer the identical set of prices.)
- 2) No asymmetric equilibrium exists when both networks in a duopoly offer two or more classes.

In fact, the first result holds in a stronger form in this model: we show that *both* networks earn lower profits as a result of one of the networks introducing a second class. Moreover, unlike Wilson who uses numerical methods (which he himself concedes “are subject, of course, to the fallibility of numerical methods”), we derive our results analytically.

Proposition 1: Consider the case of fixed, equal capacities for the two networks. In the pricing subgame where neither network divides into subnetworks, there is a unique Nash equilibrium: the two networks charge the same price, $p^I = p^{II} = 0.5/C$ and have positive profits, $\pi^I = \pi^{II} = 0.25/C$.

In the equilibrium identified in the Proposition 1, the networks compete directly (both have the same capacity, charge the same

price, and offer the same level of congestion), and yet earn positive profits.¹¹ The level of equilibrium profits is a consequence of congestion. A cut in price by one network does not attract the entire market demand (as it would in a standard Bertrand model),¹² since as users defect from the high-price network, congestion rises on the low-price network. The fall in quality of the low-price network's good eventually stems the flow of users.

Proposition 2: Consider the case of fixed, equal capacities for the two networks. In the pricing subgame where (wlog) network I chooses to subdivide its network and network II does not, there exists a unique Nash equilibrium in pure strategies: network II charges price $p^{II} = 0.4766/C$ which lies strictly between network I 's two prices, $p_1^I = 0.5784/C$, $p_2^I = 0.4500/C$. The profits of the networks are: $\pi^I = 0.2455/C$ and $\pi^{II} = 0.2427/C$.

Proposition 3: Consider the case of fixed, equal capacities for the two networks. In the pricing subgame where (wlog) both networks choose to subdivide their networks into subnetworks, there exists no Nash equilibrium in pure strategies.

The proofs of Propositions 1, 2, and 3 appear in the Appendix. These propositions yield the following.

Proposition 4: The following strategy (expressed for network I , but symmetric for the two networks) constitutes the unique (up to the arbitrary prices p_1^I , p_2^I) conditionally subgame perfect equilibrium of the two-stage game.

First Stage: Do not subdivide into subnetworks.

Second Stage: If both networks have not formed subnetworks in the first stage, charge $0.5/C$; if network I has not subdivided its network, but network II formed two subnetworks, charge $0.4766/C$; if network I formed two subnetworks and network II did not subdivide its network, charge $0.5784/C$ and $0.4500/C$; and if both networks formed two subnetworks, charge any two nonnegative prices p_1^I , p_2^I .

Corollary 1: The unique equilibrium outcome is that neither network subdivides its network, and both charge the single price $0.5/C$.

Proposition 2 shows the outcome of the two economic forces—the segmentation and competition effects—that determine equilibrium in this model. The segmentation effect can be seen by considering the profit derivatives for the two networks when prices are all equal to $1/2C$, the level in the 2-network, 2-price equilibrium identified in proposition 1:

$$\begin{aligned} \left. \frac{\partial \pi^I}{\partial p_1^I} \right|_{1/2C} &= \frac{1}{12} \geq 0 \\ \left. \frac{\partial \pi^I}{\partial p_2^I} \right|_{1/2C} &= -\frac{1}{12} \leq 0 \\ \left. \frac{\partial \pi^{II}}{\partial p^{II}} \right|_{1/2C} &= 0. \end{aligned} \quad (4)$$

These derivatives imply that, from a starting point of equal prices, network I wishes to increase p_1^I and decrease p_2^I . Network II , on the other hand, will not wish to change

¹¹The requirement that all users join a network requires that the user with the highest dislike of congestion, with $\theta = 1$, should be willing to join a network. This user has utility equal to $V - (1 - \theta^*)/C - p^I$; for this to be nonnegative in equilibrium requires that $V \geq 1/C$.

¹²See, e.g., [27].

its price. (But of course, once network I changes its price, network II may wish to respond.) Network I —the segmenting network—wishes to offer a broader range of qualities.

Due to the competition effect, such a move lowers profit. In the equilibrium identified in Proposition 2, network I succeeds in segmenting that market: $p_2^I < p^{II} < p_1^I$. But in this case, it is as if the two networks compete in two places, rather than just one, since there are now two indifferent users (the one indifferent between network I 's cheaper, more congested subnetwork and network II ; and the one indifferent between network I 's more expensive, less congested subnetwork and network II). As a result, the networks compete more fiercely and the average price (weighted by market share) charged by the networks falls from $0.5/C$ in the two-network, two-price equilibrium in Proposition 1 to $0.4826/C$. Note that, although this difference is small, in practice it may be large, for several reasons. First, this model relies on a certain parameterization in order to show that profits decrease through introducing multiple service classes. With other parameterizations, the difference may be larger (although, of course, it could also be smaller). Second, for the sake of clarity, the cost of introducing multiple service classes has been ignored. If these are at all significant, then profits will be even lower and multiple service classes even less attractive to the networks.

These same two forces are at work in Proposition 3 and lead to nonexistence of equilibrium in pure strategies in the pricing subgame when both networks form two subnetworks. As in [28], the basic problem is that profit maximization drives the networks to move prices away from feasible levels. Consider, for example, the possible case in which $p_2^{II} < p_1^{II} < p_2^I < p_1^I$. Both networks gain from greater segmentation (other things equal), and so network I will decrease p_2^I and network II increase p_1^{II} until the case no longer holds. Instead, $p_2^{II} < p_2^I < p_1^{II} < p_1^I$; but then there are three indifferent users between the networks, rather than one, and so competition is increased. As a result, the networks will adjust prices to move away from this case. This process continues through all other possible cases and so an equilibrium in pure strategies does not exist.

IV. CHOICE OF CAPACITIES

Section III assumes that network capacities are fixed and symmetric. This assumption is convenient analytically; but, clearly, it is a strong restriction. In this section, numerical analysis is presented which suggests that the main conclusion—that multiproduct competition is not sustainable in a profit maximizing equilibrium—stands when networks are free to choose capacities, as well as prices.¹³

The game now has three stages. The number of subnetworks is chosen in the first stage; capacities are chosen in the second stage; and prices in the third stage. The Nash equilibrium is found for the third stage pricing sub-game, taking capacities and the number of subnetworks as given. Then the Nash equilibrium is found for the second stage subgame in capacities, taking into account the effect that capacity choice will have on optimal prices in the third stage pricing subgame. Finally, profits for different numbers of subnetworks are compared. The assumption of costless capacity is maintained in this section.

¹³Full details of the numerical analysis are contained in [8].

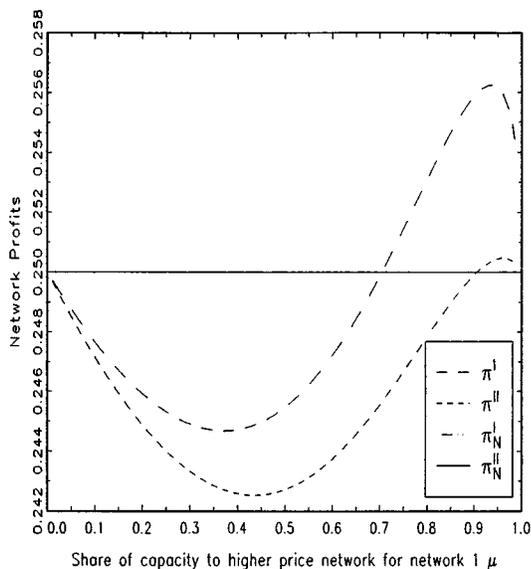


Fig. 2. The profit functions for $C^I = C^{II} = 1$.

In the simplest case of two networks each charging a single price, [19] proves that the Nash equilibrium in capacities is asymmetric. In the absence of costs of increasing capacity and when all users join one network, the network who charges the higher price has a strict incentive always to add to its capacity; the other network has a strict incentive always to decrease its capacity. While this result may appear counterintuitive, it is in fact the familiar result that firms have an incentive to differentiate themselves to the greatest extent; see, e.g., [22]. The reason is that, by making their services as different as possible, the two networks minimize price competition between them and hence maximize profits. In the absence of capacity costs, the one network increases its capacity, and hence its price, until the requirement is reached that the user with $\theta = 1$ receives nonnegative surplus from joining the network.

As in the previous section, it is helpful to consider first the 2-network, 3-price situation. There are two cases: 1) $p_1^{II} \leq p_2^I \leq p_1^I$ or $p_1^{II} \geq p_2^I \geq p_1^I$; and 2) $p_2^I \leq p_1^{II} \leq p_1^I$. Let the capacities of network I 's subnetworks be $C_1^I = \mu C^I$ and $C_2^I = (1-\mu)C^I$; network II 's capacity is C^{II} .

The analytical results of the previous section show that there is no equilibrium in the price stage in case 1 when $C^I = C^{II}$ and $\mu = 0.5$. Numerical investigation indicates that, for any fixed value of C^{II} , there are critical values of C^I and μ , denoted \underline{C}^I and $\bar{\mu}$, respectively, such that a pricing equilibrium exists only when $C^I \geq \underline{C}^I > C^{II}$ and $\mu \leq \bar{\mu}$. Furthermore, network I 's profits are decreasing in μ for all values of C^I , C^{II} and μ for which there exists a pricing equilibrium. Hence, the numerical analysis of this case suggests network I 's optimal choice of μ is zero, so that the 2-network, 3-price solution collapses to the 2-network, 2-price equilibrium: network I allocates all of its capacity to a single subnetwork, i.e., it does not subdivide.

Consider next the price stage under case 2, examined in Fig. 2. The figure plots profits with fixed total capacities $C^I = C^{II}$ ($=1$), and with μ varying between 0 and 1. It confirms the analytical results in Section III: when $\mu = 0.5$, profits when there

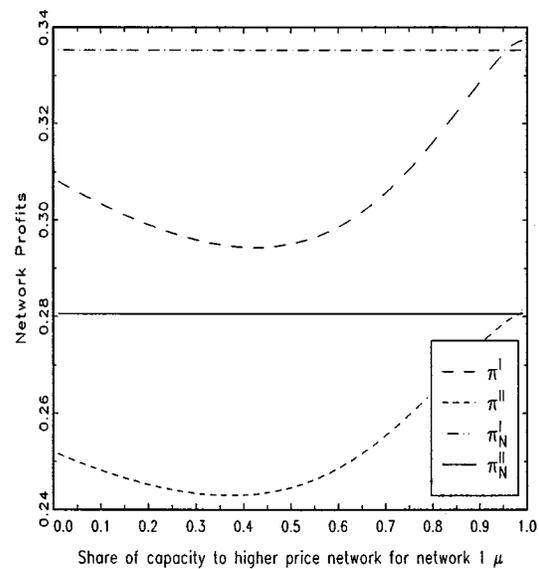


Fig. 3. The profit functions for $C^I = 1$, $C^{II} = 0.75$.

is PMP (denoted π^I and π^{II}) are, respectively, lower than those with no PMP (π_N^I and π_N^{II}). (Note that in this instance π_N^I coincides with π_N^{II} .) The figure shows that there is an interior optimal choice of μ , close to but less than 1. But for this value of μ , a computation shows that network II 's profits are decreasing in its capacity C^{II} . Network II will therefore lower its capacity; the effect of this is to make network I 's optimal choice of μ (still holding C^I fixed) the corner value of 1. This is illustrated in Fig. 3, which shows the profit functions for $C^I = 1$ and $C^{II} = 0.75$. With $\mu = 1$, the 2-network, 3-price solution collapses to the 2-network, 2-price equilibrium. This argument holds for any $C^I > C^{II}$. There is another case: suppose that C^I is much smaller than C^{II} ; in the calculations here, it is sufficient for $C^I = 0.5$ and $C^{II} = 1$. The numerical analysis then shows that case 3 is equivalent to case 1 (although the networks are reversed, with network 2 being the large, high-price network). Again, the 2-network, 3-price solution collapses to the 2-network, 2-price equilibrium.

A full numerical analysis of the two-network, four-price case has not been undertaken. Instead of calculating equilibrium prices and profits over a 4-dimensional grid of capacities, only certain values have been considered. Moreover, only the case in which $p_2^{II} \leq p_2^I \leq p_1^{II} \leq p_1^I$ has been examined. The aim is not to provide an exhaustive numerical analysis, but instead to indicate the sort of results that might emerge. The analysis suggests first that equilibria are possible only in certain ranges of networks' capacities, and second that even when a 4-price equilibrium is possible (i.e., even when profit-maximizing prices are feasible), the networks both earn higher profits in the 2-price solution.

This section has assessed whether the main result of Section III—that competing profit-maximizing networks will not sell multiple products—is robust to relaxing the assumption of fixed, equal, and symmetrically split capacities. This section offers substantial support for the analytical results. In all cases considered, the networks maximize profits by charging a single price and offering one network each.

V. CONCLUSION

This paper has developed a general analysis of an Internet pricing scheme for packet based networks under competition in which the networks are partitioned into logical subnetworks distinguished only by price level. The costs of increased competition as more subnetworks are introduced always outweigh the benefits from greater segmentation of the market. The intuition behind this is that the desire to segment drives the network with the high prices to lower one of its prices to straddle those of its rival; in turn, this leads to lower profits.

The complexity of the problem required that several assumptions be made, including: a uniform distribution of user preferences toward congestion, a linear congestion function, given and equal network capacities, and a fixed number of networks in the industry. Numerical analysis suggests that at least one of these assumptions—fixed capacity—may not be critical to the conclusions. Further work is required to assess the importance of the other assumptions. One area in particular seems promising. The current analysis suggests that there will be limited product differentiation in equilibrium, since each network offers only one service class. We have assumed that there is a fixed number of networks (i.e., two) in the network industry. The process of free entry may be, however, a mechanism by which a broad range of prices and qualities arises in equilibrium.

APPENDIX

The Social Optimum

We show that, given the model described in Section II, the social planner who aims to maximize the total welfare (user surplus) from network services will wish to use multiple products (subnetworks). It is assumed that he/she is able to allocate users to the most appropriate subnetwork. Consider the case where he/she splits the network, with total capacity of $2C$, into two subnetworks, each with capacity C . Let θ^* denote the critical value of the congestion parameter. Then total user surplus is

$$\begin{aligned} V_{SO} &= \int_0^{\theta^*} \left(V - \theta \frac{\theta^*}{C} \right) d\theta + \int_{\theta^*}^1 \left(V - \theta \frac{1 - \theta^*}{C} \right) d\theta \\ &= V - \frac{1}{2C} (2(\theta^*)^3 - (\theta^*)^2 - \theta^* - 1). \end{aligned}$$

The planner chooses θ^* to maximize V_{SO} . The first-order condition, $6(\theta^*)^2 - 2\theta^* - 1 = 0$, gives $\theta^* = (1 + \sqrt{7}/6)$. (The second-order condition for a maximum is satisfied.) Thus, he/she uses two subnetworks. Users with higher tolerance for congestion (i.e., lower θ s) are allocated to a network with a market share of around 61%. Users more sensitive to congestion are allocated to a network with market share of 39%. It is straightforward to extend this result to where the planner forms $n > 2$ subnetworks of equal capacity. In general, the planner will wish to form as many subnetworks as possible.

Proof of Proposition 1

The networks have equal and fixed capacities: $C^I = C^{II} = C$. There is a marginal user θ^* who is indifferent between the two networks; from equation (3), this user is given by $\theta^* =$

TABLE I
SOLUTIONS FOR PROOF OF PROPOSITION 2

	Solution 1	Solution 2
θ_1	0.7749	0.8014
θ_2	0.2899	0.4892
p_1^I	0.1887/C	0.5443/C
p_2^I	-0.2142/C	0.3623/C
p^{II}	-0.0170/C	0.4284/C

$(1 + \sqrt{1 + 8C(p^I - p^{II})})/4$. Suppose (p^I, p^{II}) is an equilibrium, where without loss of generality we assume $p^I > p^{II}$. Then $\pi^I = p^I(1 - \theta^*)$ and $\pi^{II} = p^{II}\theta^*$, where θ^* is given above. The first-order equilibrium conditions imply $p^I = (1 - \theta^*)(4\theta^* - 1)/C$ and $p^{II} = \theta^*(4\theta^* - 1)/C$. Since $p^I > p^{II}$, then $\theta^* > 1/2$, but then the first-order conditions imply $p^{II} > p^I$, a contradiction. The case $p^I = p^{II} = p$, say, has $\theta^* = 1/2$. In this case, the first-order conditions for an equilibrium imply $p = 1/2C$. The second-order conditions are easily verified.

Proof of Proposition 2

Suppose there exists a Nash equilibrium with $p_2^I \neq p_1^I$; without loss of generality, suppose $p_2^I < p_1^I$. Now consider the subcase where $p^{II} \leq p_2^I < p_1^I$. There are two marginal users: one, with congestion parameter θ_1 , is indifferent between joining network I 's higher priced subnetwork with capacity $C/2$ and joining network I 's lower priced subnetwork with capacity $C/2$; the other, with congestion parameter θ_2 , is indifferent between network II with capacity C and network I 's lower priced subnetwork. These users are defined by the indifference equations

$$V - \frac{2\theta_1(1 - \theta_1)}{C} - p_1^I = V - \frac{2\theta_1(\theta_1 - \theta_2)}{C} - p_2^I \quad (5)$$

$$V - \frac{2\theta_2(\theta_1 - \theta_2)}{C} - p_2^I = V - \frac{(\theta_2)^2}{C} - p^{II}. \quad (6)$$

Suppose there exist stationary solutions to the networks' profit maximization problems that are consistent with the construction. The profits are $\pi^I = p_2^I(\theta_1 - \theta_2) + p_1^I(1 - \theta_1)$ and $\pi^{II} = p^{II}\theta_2$. From (5) and the first-order conditions for stationary profit maximizing prices, it follows that $\theta_1 = (3 + \theta_2 + \sqrt{7(\theta_2)^2 - 6\theta_2 + 3})/6$.

Similarly, (6) and the first-order condition imply

$$\begin{aligned} &\frac{(1 - \theta_1)C - 2\theta_1(2\theta_1 - (1 + \theta_2))\theta_{11}}{\theta_{21}} + \frac{2\alpha\theta_2}{4\theta_1 - \theta_2 - 1} \\ &= 3(\theta_2)^2 - 2\theta_1\theta_2 \end{aligned}$$

where $\alpha(\theta_1, \theta_2) \equiv 4(\theta_1)^2 + 3(\theta_2)^2 - 12\theta_1\theta_2 - \theta_1 + 3\theta_2$. There are therefore two possible solutions in the feasible region (defined by $\theta_1 \geq ((1 + \theta_2)/2)$ and $\theta_2 \leq \theta_1 \leq (3/2)\theta_2$). The solutions are shown in Table I. The first solution can be ruled out immediately, since it involves negative prices. The second solution has positive variables, but violates the construction; in particular, $p^{II} > p_2^I$. Hence, there are no Nash equilibria in this case. If $p_2^I < p_1^I \leq p^{II}$, the proof proceeds similarly.

Now consider the subcase where $p_2^I < p^{II} < p_1^I$. There are two marginal users: one is indifferent between joining network I 's higher priced subnetwork and joining network II , with θ_1 ; the other is indifferent between network II and network I 's

lower priced subnetwork, with θ_2 . These users are defined by the indifference equations

$$V - \frac{2\theta_1(1 - \theta_1)}{C} - p_1^I = V - \frac{\theta_1(\theta_1 - \theta_2)}{C} - p^{II} \quad (7)$$

$$V - \frac{\theta_2(\theta_1 - \theta_2)}{C} - p^{II} = V - \frac{2(\theta_2)^2}{C} - p_2^I. \quad (8)$$

Profits of the two networks are $\pi^I = p_2^I\theta_2 + p_1^I(1 - \theta_1)$ and $\pi^{II} = p^{II}(\theta_1 - \theta_2)$. First-order conditions for stationary profit maximizing prices and the indifference equation yield two simultaneous nonlinear equations in the two unknowns, θ_1 and θ_2 :

$$\begin{aligned} &15\theta_1^3 + \theta_2^3 + 35(\theta_1)^2\theta_2 - 27\theta_1(\theta_2)^2 - 28(\theta_1)^2 \\ &\quad + 9(\theta_2)^2 - 23\theta_1\theta_2 + 14\theta_1 + 3\theta_2 - 2 = 0 \\ &\theta_1^3 + 15\theta_2^3 - 27(\theta_1)^2\theta_2 + 35\theta_1(\theta_2)^2 + 2(\theta_1)^2 \\ &\quad - 15(\theta_2)^2 + 11\theta_1\theta_2 - \theta_1 = 0. \end{aligned}$$

These equations have a unique solution in the relevant range that ensures that $p_2^I < p^{II} < p_1^I$. This is $\theta_1 = 0.8083$ and $\theta_2 = 0.2991$; equations (7) and (8) then give prices as $p_1^I = 0.5784/C$, $p_2^I = 0.4500/C$ and $p^{II} = 0.4766/C$, and so profits are $\pi^I = 0.2455/C$, $\pi^{II} = 0.2427/C$. These profits are both lower than the profits in the equilibrium where both networks offer only one price, derived in Proposition 1, *viz.*, $\pi = 0.25/C$. If $p^{II} = p_2^I = p_1^I$, then this solution, which was an equilibrium in Proposition 1, is not an equilibrium here, as easily verified from the first order conditions (4).

Proof of Proposition 3

Without loss of generality, assume that the highest price is charged by network *I* and that each network's 2nd price does not exceed its 1st. There are three cases: i) $p_2^{II} \leq p_2^I \leq p_1^{II} \leq p_1^I$, ii) $p_2^I \leq p_2^{II} \leq p_1^{II} \leq p_1^I$, iii) $p_2^{II} \leq p_1^{II} \leq p_2^I \leq p_1^I$.

Consider case i. The method of analysis is the same as for the cases where three prices are charged by the two networks—first, write down the indifference relations defining the marginal users, denoted θ_i , $i = 1, 2, 3$; then use the implicit function theorem to determine the derivatives of θ_i with respect to the prices; finally, calculate derivatives of profit functions. The indifference equations give

$$p_1^I - p_1^{II} = \frac{2\theta_1}{C} (2\theta_1 - 1 - \theta_2) \quad (9)$$

$$p_1^{II} - p_2^I = \frac{2\theta_2}{C} (2\theta_2 - \theta_1 - \theta_3) \quad (10)$$

$$p_2^I - p_2^{II} = \frac{2\theta_3}{C} (2\theta_3 - \theta_2). \quad (11)$$

The profits of the networks are $\pi^I = p_1^I(1 - \theta_1) + p_2^I(\theta_2 - \theta_3)$ and $\pi^{II} = p_1^{II}(\theta_1 - \theta_2) + p_2^{II}\theta_3$.

The proof will show that either network *I* will wish to lower p_2^I below p_2^{II} or, equivalently, network *II* will wish to raise p_2^{II} above p_2^I . Consider the derivatives $(\partial\pi^I/\partial p_2^I)$ and $(\partial\pi^{II}/\partial p_2^{II})$, evaluated at $p_2^I = p_2^{II}$, with p_1^I and p_1^{II} determined as stationary solutions. From (11), $p_2^I = p_2^{II} \equiv p_2$ implies that $2\theta_3 = \theta_2$. Setting the derivatives $(\partial\pi^I/\partial p_1^I)$ and

$(\partial\pi^{II}/\partial p_1^{II})$ equal to zero, and substituting the resulting expressions with the equality $2\theta_3 = \theta_2$ into the profit derivatives gives

$$\begin{aligned} \frac{\partial\pi^I}{\partial p_2^I} \Big|_{p_2^I=p_2^{II}} &= \frac{1}{2\theta_2} ((\theta_2)^2 - p_2C) \\ \frac{\partial\pi^{II}}{\partial p_2^{II}} \Big|_{p_2^I=p_2^{II}} &= \frac{\left(\begin{aligned} &4\theta_1\theta_2 - 12(\theta_1)^2\theta_2 - 2(\theta_2)^2 \\ &+ 12\theta_1(\theta_2)^2 - 6\theta_2^3 \\ &- Cp_2(2 - 4\theta_1 - 3\theta_2) \end{aligned} \right)}{4\theta_2(1 - 2\theta_1 - \theta_2)}. \end{aligned}$$

In order for an equilibrium to exist with $p_2^{II} \leq p_2^I \leq p_1^{II} \leq p_1^I$, it must be that

$$\frac{\partial\pi^I}{\partial p_2^I} \Big|_{p_2^I=p_2^{II}} \geq 0, \quad \frac{\partial\pi^{II}}{\partial p_2^{II}} \Big|_{p_2^I=p_2^{II}} \leq 0.$$

These inequalities together with (9) yield $\theta_1 \leq (1 + 4\theta_2 + \sqrt{1 - 4\theta_2 + 7(\theta_2)^2})/6$. (The negative root is ruled out by the requirement that $\theta_1 \geq \theta_2$.) This inequality must be combined with (9) and the requirement that $p_1^{II} \leq p_1^I$, *i.e.*, yielding: $\theta_1 \geq 1 + \theta_2/2$. This requires that $\theta_2 \geq \sqrt{2}/2$, and therefore that $\theta_3 \geq (\sqrt{2}/4)$ and $\theta_1 \geq 0.8536$.

The last stage of the proof shows that these inequalities are inconsistent with stationary solutions for p_1^I and p_1^{II} . The first-order conditions give

$$\begin{aligned} p_1^I &= \frac{\left(\begin{aligned} &4 + 90\theta_2 - 363(\theta_2)^2 + 289\theta_2^3 \\ &+ (4 - 118\theta_2 + 124(\theta_2)^2)\sqrt{1 - 4\theta_2 + 7(\theta_2)^2} \end{aligned} \right)}{9C(1 - 14\theta_2 + \sqrt{1 - 4\theta_2 + 7(\theta_2)^2})} \\ p_1^{II} &= \frac{\left(\begin{aligned} &-1 + 6\theta_2 - 12(\theta_2)^2 + 35\theta_2^3 \\ &+ (1 - 4\theta_2 + 11(\theta_2)^2)\sqrt{1 - 4\theta_2 + 7(\theta_2)^2} \end{aligned} \right)}{9C(-2 + 7\theta_2 + \sqrt{1 - 4\theta_2 + 7(\theta_2)^2})}. \end{aligned}$$

It is straightforward to show that, when $\theta_2 \geq (\sqrt{2}/2)$, $p_1^I \leq p_1^{II}$. When the inequality is strict (*i.e.*, $\theta_2 > (\sqrt{2}/2)$ and $p_1^I < p_1^{II}$), the solution violates the requirement that $p_1^{II} \leq p_1^I$. With equality, the solution collapses to $p_2^{II} = p_1^{II} = p_2^I = p_1^I$.

We conclude that in case i where $p_2^{II} \leq p_2^I \leq p_1^{II} \leq p_1^I$, either network *I* will wish to lower p_2^I below p_2^{II} , or network *II* will wish to raise p_2^{II} above p_2^I , or the inequalities will be weak, *i.e.*, the solution is the 2-network, 2-price case that arises as an equilibrium in Proposition 1; however, as in Proposition 2, this solution is not an equilibrium here.

Finally, although we have assumed here to be in case i, the other two cases proceed similarly. Note that in case ii, the subcase $p_2^I < p_2^{II} = p_1^I < p_1^{II}$ that arises as an equilibrium in Proposition 2 is not an equilibrium here. This is because when network *II* is able to charge two prices, its best response to network *I* charging (any) two prices is also to charge two prices. (This statement can be verified easily by examining the first-order conditions for prices that maximize network *II*'s profits when it has formed two subnetworks in the first stage.) In summary, it is always the case that one or another of the networks will wish to change the assumed order of prices.

ACKNOWLEDGMENT

The authors are very grateful to F. Kelly and D. Newbery for many helpful conversations; to J. Cave, M. Cripps, I. H. Lee, M. Miller, A. Odlyzko, S. Scholtes, and M. Waterson for comments; and to seminar participants in Bristol, Cambridge, and Warwick. Any errors are our responsibility.

REFERENCES

- [1] M. Armstrong and J. Vickers, "Competitive price discrimination," *Mimeo*, 1999.
- [2] P. Champsaur and J. C. Rochet, "Multiproduct duopolists," *Econometrica*, vol. 57, no. 3, pp. 533–557, 1989.
- [3] P. Chander and L. Leruth, "The optimal product mix for a monopolist in the presence of congestion effects: A model and some results," *Int. J. Industrial Org.*, vol. 7, no. 4, pp. 437–449, 1989.
- [4] D. Clark, "Internet cost allocation and pricing," in *Internet Economics*, L. W. McKnight and J. P. Bailey, Eds: MIT Press, 1997, pp. 215–252.
- [5] P. Dasgupta and E. Maskin, "The existence of equilibrium in discontinuous games, I: Theory," *Rev. Economic Studies*, vol. 53, no. 1, pp. 1–26, 1986.
- [6] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1991.
- [7] J. J. Gabszewicz, A. Shaked, J. Sutton, and J. F. Thisse, "Segmenting the market: The monopolist's optimal product mix," *J. Economic Theory*, vol. 39, no. 2, pp. 273–289, 1986.
- [8] R. J. Gibbens, R. Mason, and R. Steinberg, "Multiproduct competition between congestible networks," Working Paper WP 26/98, Judge Inst. Manage. Studies, Univ. Cambridge, Sept. 1998.
- [9] A. Gupta, D. O. Stahl, and A. B. Whinston, "A stochastic equilibrium model of internet pricing," *J. Economic Dynamics Contr.*, vol. 21, no. 4–5, pp. 697–722, 1997.
- [10] G. Hardin, "The tragedy of the commons," *Science*, vol. 162, pp. 1243–1248, 1968.
- [11] J. Harsanyi, "Games with randomly disturbed payoffs: A new rationale for mixed strategy equilibrium points," *Int. J. Game Theory*, vol. 2, pp. 1–23, 1973.
- [12] F. P. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: Shadow prices, proportional fairness and stability," *J. Operat. Res. Soc.*, vol. 49, pp. 237–252, 1998.
- [13] D. M. Kreps and J. A. Scheinkman, "Quantity precommitment and Bertrand competition yield Cournot outcomes," *Bell J. Economics*, vol. 14, no. 2, pp. 326–337, 1983.
- [14] I. Luski, "On partial equilibrium in a queuing system with two servers," *Rev. Economic Studies*, vol. 43, no. 3, pp. 519–525, 1976.
- [15] J. K. MacKie-Mason and H. R. Varian, "Pricing congestible network resources," *IEEE J. Select. Areas Commun.*, vol. 13, pp. 1141–1149, 1995.
- [16] R. A. Mason, "Simple competitive internet pricing," *Euro. Economic Rev.*, 2000.
- [17] A. Odlyzko, "Paris Metro pricing for the Internet," in *Proc. ACM Conf. Electronic Commerce (EC'99)*, ACM, 1999, pp. 140–147.
- [18] D. Owen, "Paris rail abolishes first-class travel," *Financial Times*, p. 2, Sept. 1, 1999.
- [19] A. de Palma and L. Leruth, "Congestion and game in capacity: A duopoly analysis in the presence of network externalities," *Annales d'Economie et de Statistique*, no. 15–16, pp. 389–407, 1989.
- [20] K. M. Park, M. Sitharam, and S. Chen, "Quality of service provision in noncooperative networks with diverse user requirements," *Decision Support Systems, Special Issue on Information and Computation Economics*, vol. 28, pp. 101–122, 2000.
- [21] D. Reitman, "Endogenous quality differentiation in congested markets," *J. Ind. Economics*, vol. 39, no. 6, pp. 621–647, 1991.

- [22] A. Shaked and J. Sutton, "Relaxing price competition through product differentiation," *Rev. Economic Studies*, vol. 49, no. 1, pp. 3–13, 1982.
- [23] —, "Natural oligopolies," *Econometrica*, vol. 51, no. 5, pp. 1469–1483, 1983.
- [24] —, "Multiproduct firms and market structure," *Rand J. Economics*, vol. 21, pp. 45–62, 1990.
- [25] S. Shenker, "Making greed work in networks: A game-theoretic analysis of switch service disciplines," in *Proc. SIGCOMM Symp. Commun. Architectures Protocols*, 1994, pp. 47–57.
- [26] —, "Services models and pricing policies for an integrated services Internet," in *Public Access to the Internet*, B. Kahin and J. Keller, Eds. Cambridge, MA: MIT Press, 1995, pp. 315–337.
- [27] J. Tirole, *The Theory of Industrial Organization*. Cambridge, MA: MIT Press, 1988.
- [28] R. B. Wilson, "Efficient and competitive rationing," *Econometrica*, vol. 57, no. 1, pp. 1–40, 1989.



Richard Gibbens was born in London, England, in 1962. He obtained the B.A. degree in mathematics, from the University of Cambridge, the Diploma in Mathematical Statistics, and the Ph.D. degree in 1988.

His doctoral thesis work, which was a collaboration between the University of Cambridge and British Telecom, was on the design and analysis of dynamic routing strategies in telecommunication networks, resulting in the development of the dynamic alternative routing strategy now in operation in the British Telecom trunk network. He has worked in the area of modeling of communication networks, mainly at the Statistical Laboratory, University of Cambridge, but was also a Visiting Consultant at AT&T Bell Laboratories, Murray Hill, NJ, during 1989.

Dr. Gibbens was awarded a Royal Society University Research Fellowship in 1993.



Robin Mason received the Masters and Ph.D. degrees in economics from the University of Cambridge.

He has held Research Fellowships at the Universities of Oxford and Cambridge. He is currently a Lecturer in the Economics Department at the University of Southampton. His work is on competition in and structure of network industries.



Richard Steinberg received the B.A. degree from Reed College, the Ph.D. degree in combinatorics and optimization from the University of Waterloo, and the MBA degree from the University of Chicago.

He has taught at the University of Chicago and Columbia University, and has served on the Technical Staffs of AT&T Bell Laboratories and GTE Laboratories. He is currently a Lecturer in Operations Management at the University of Cambridge in the Judge Institute of Management Studies.