Simple competitive Internet pricing

Robin Mason

Phil. Trans. R. Soc. Lond. A 2000 **358**, 2309-2318
doi: 10.1098/rsta.2000.0649

10.1098/rsta.2000.0649

THE ROYAL SOCIETY

Simple competitive Internet pricing

By Robin A. Mason

Department of Economics, University of Southampton, Highfield,
Southampton SO17 1BJ, UK (robin.mason@soton.ac.uk)

This paper aims to show that economics can make an important contribution to the question of how to price the Internet. It argues that any pricing scheme must be robust to strategic behaviour, in the sense that the scheme must be supportable as an equilibrium between firms or individuals. The paper analyses the implications of this restriction for interconnection agreements between networks, and for end-user pricing.

Keywords: Internet pricing; network effects; nonlinear pricing

1. Introduction

The aim of this paper is to show that economics can make an important contribution to the current debate about the appropriate way in which to price the public Internet. The paper will not maintain that the only relevant questions are economic ones. Instead, a multi-disciplinary approach is required in which engineering, statistics and economics are combined to devise feasible, stable and competitive charging schemes.

The argument proceeds by illustrating the economic issues that arise at two pricing interfaces. The first is between networks; the economic question of interest here is: what incentives do networks have to interconnect? To non-economists, this can seem a peculiar question: after all, what is the Internet without widespread connectivity? As the nature of the Internet continues to change from its original academic origin to a commercial market place, however, the issue of interconnection agreements between asymmetric networks is becoming increasingly relevant. The second is between networks and end-users; the question is: how will profit-maximizing networks price for network usage? There are many aspects to this question. This paper concentrates on two: simplicity and robustness to competition. Early economic proposals for Internet pricing were (rightly) criticized for their complexity. Subsequent work, in striving for simplicity, has neglected to consider whether the schemes proposed would be used by profit-maximizing networks competing for customers. These questions are clearest in the current market debate about whether usage prices should be employed at all.

The emphasis throughout the paper is the determination of pricing schemes that can be supported as an equilibrium† between firms or individuals. The restriction to equilibrium has real bite when the players involved are selfish, aiming rationally to

---

† The equilibrium concept used is Nash, or some refinement of Nash. If there is a set of strategies with the property that no player can benefit by changing his/her strategy while the other players keep their strategies unchanged, then that set of strategies and the corresponding pay-offs constitute the Nash equilibrium. In other words, no player has a unilateral incentive to move away from a Nash equilibrium, since, given what the other players are doing, each individual player is doing as well as it can.

maximize their own profit or utility; and when the number of players is relatively small. In such situations, any pricing scheme must be robust to strategic behaviour.

The rest of the paper is structured as follows. Interconnection agreements between networks are considered in § 2, the prices charged by networks to end-users are examined in § 3, and conclusions are drawn in § 4.

2. Network interconnection

In the early years of the Internet, networks operated a 'bill-and-keep' or peering system, in which no settlement payments were made. (See Srinagesh (1997) for a discussion of interconnection arrangements between packet-based networks making up the Internet.) Each network carried others' traffic without charge, the underlying assumptions being that either flows were roughly symmetric, or any other arrangement would stunt the growth of the Internet. The transition of the Internet from academic to commercial, large increases in traffic volumes, and the unequal development of networks have put this system under considerable stress.

In 1996, the extensive peering arrangements agreed under the Commercial Internet Exchange (CIX) started to dissolve. Large networks argued that they received little benefit, yet incurred substantial costs, from interconnection with small networks; this contrasted with the net benefits gained by the smaller networks from access to the customer base of the larger networks. Large networks began to apply pressure on smaller networks to change the relationship from peers to supplier–customer; instead of bill-and-keep, small networks would make settlement payments to larger networks. In 1997, UUNet, a large Internet Service Provider (ISP), informed 15 smaller ISPs that their peering arrangements would be cancelled; this was followed by UUNet's withdrawal from the CIX. At the same time, MCI and BBN, two other large ISPs, left the CIX agreement, meaning that three out the four largest networks in the US were no longer part of the CIX.† The larger networks continue to interconnect between themselves on a peering (no-settlement) basis.

The gulf between large and small networks has widened progressively with the consolidation taking place in the ISP industry. By November 1997, it was estimated that the US's four largest networks (UUNet, MCI, BBN and Sprint) accounted for between 85 and 95% of total backbone (i.e. core) Internet traffic, with the remaining volume carried by upwards of 40 other, small networks; see OECD (1998). There is a growing fear in the industry that large networks will use their size to limit competition in the ISP market by excluding smaller networks from interconnection agreements. See Crémer *et al.* (1998) for a consideration of the possible impact on the Internet of the recent merger between MCI and WorldCom.

In order to understand the economic incentives involved in interconnection, it is necessary to distinguish between two dimensions of service offered by a network. (Other dimensions might be relevant; but for current purposes, it is enough to concentrate on the broad features of networks which carry each others' traffic.) Internet networks may be *horizontally differentiated*: that is, when two networks are of equal size and charge the same price, some users prefer one network, and others the other.

---

† UUNet responded to criticism about its policy by publishing guidelines stating when UUNet was prepared to interconnect with a smaller network; the guidelines are reported in OECD (1998). UUNet reserved the right, however, to refuse interconnection with another network, even if that network meets the criteria laid down in the guidelines.

This can be for a variety of reasons. Some networks are regional in their coverage (although others are national); horizontal differentiation in this case is geographically based. More importantly, networks differ in the content and services provided on their networks. Moreover, networks often bundle Internet access with access to other communication networks (such as telephone and television); arguably cable access, with higher bandwidth but lower quality services, is horizontally differentiated from traditional twisted copper pair access, which has lower bandwidth but greater reliability. (See Foros & Hansen (1999) for a fuller discussion.) Networks are also *vertically differentiated*: when the networks offer the same content, are located at the same place, etc., and charge the same price, typically the larger network is preferred: each user gains greater benefit from belonging to a network with more members.† While all users value belonging to a large network, some value it more than others, i.e. users have heterogeneous preferences towards the vertical differentiation of networks.

Interconnection has two effects. First, it decreases the degree of vertical differentiation between networks of different sizes. In the limit when interconnection is perfect and costless, consumers are indifferent about which network they join (ignoring for the moment horizontal preferences), and the networks are pure price competitors. Competition between networks is therefore intensified by interconnection, and profits may fall as a result. The reverse side is that interconnection makes size less important for horizontally differentiated networks. When networks are not interconnected, the 'quality' of a network is entirely dependent on that network's size. Since a larger (higher quality) network is more profitable than a smaller (lower quality) network, the networks are more willing to undercut their rival's price in order to attract market share. Consequently, competition for market share is intense. Interconnection reduces the importance of the network's size, and consumers do not then mind joining a small network, since they receive the full benefit not only from other consumers on their network, but also from consumers on other networks. Hence, interconnection decreases competition.

Which effect dominates depends on the relative importance of the horizontal and vertical aspects. There are three possible cases. In the first, vertical aspects are relatively important, and the net effect of interconnection is to lower the profits of the networks. In the third, horizontal aspects are relatively important, and interconnection raises networks' profits. In the second, intermediate region, it will be seen below that the situation is asymmetric: there is a larger network and a smaller network; the larger network does not wish to interconnect, while the smaller network does wish to interconnect.

Figure 1 illustrates the result. The figure plots equilibrium profits π_i, prices p_i and size Q_i of two networks $i = 1, 2$, against a parameter β that measures the relative importance of horizontal and vertical aspects. When β is high (close to 1), horizontal aspects are most important; when β is low (close to 0), vertical aspects are most important. (The figure assumes that the function is linear in β. This is, of course, not the case; Mason (1999) contains the complete analysis.) The heaviest solid line gives

---

† When the utility that a user derives from consumption of the good increases with the number of other agents consuming the good, a 'positive network externality' is said to exist. Standard examples are physically connected computer networks and telecommunications systems. But the feature arises also in many other cases: for example, fans of live entertainment prefer big cities because the large market for entertainment assures a full variety of acts.

the level of these variables when the networks interconnect; the lighter solid line, and the heavier dashed line, give equilibrium variables when the networks do not interconnect. For $\beta < \beta_1$, both networks prefer not to interconnect. In this region, interconnection causes a net reduction in differentiation between the networks; consequently, competition is intensified by interconnection, and profits decrease. For $\beta_1 \leqslant \beta < \beta_2$, network 2 (the smaller network) prefers to interconnect, but network 1 does not. In this region, therefore, the networks have conflicting incentives towards interconnection. Finally, for $\beta \geqslant \beta_2$, both networks prefer to interconnect. In this third region, interconnection decreases the extent of competition between the networks.

At the same time, the networks' sizes become more equal as β increases. When $\beta = 0$, network 1 is significantly larger; by the time that β has increased to 1, the networks are symmetric. Note that the networks are not assumed to be asymmetric: it is inherent in this situation that the networks are of different sizes (except when β is exactly 1). Put more formally, in the unique Nash equilibrium (in pure strategies) between the profit-maximizing networks, the determined variables of prices, sizes and profits differ between the networks when $\beta < 1$.

3. End-user pricing

(*a*) *A survey*

A full survey of the many Internet end-user pricing proposals is not possible here (see, for example, Gupta *et al*. 1997; Kelly *et al*. 1998; Mackie-Mason & Varian 1997; Odlyzko 1997). This section looks at three schemes to illustrate the range. Two issues are particularly important. First, the (economic) incentives that users and networks face must be recognized; in particular, pricing schemes must be viable in a setting in which networks compete and maximize profits. Secondly, any scheme must be simple enough to implement at reasonable cost.

The 'smart market' of Mackie-Mason & Varian (1997) is perhaps the best known of the economic Internet pricing schemes. The proposal involves a zero usage price when network resources are not congested. At congested parts of the network, packets are prioritized based on the bids attached to them by users. Users whose packets are transmitted are not charged the amount that they bid, but rather the bid of the highest priority packet that is not admitted to the network. This $n + 1$-price, or Vickrey (1961), auction scheme has the well-known desirable features of

(i) provision of the right incentives for users to reveal their true willingness-to-pay for priority; and

(ii) generation of the socially optimal level of revenues for network expansion.

There are several criticisms of this scheme. The first is that it fails to take into account dynamic factors: often users are interested not only in instantaneous resource allocation, but also the allocation over the entire duration of a communication. For recent work on this question, see Crémer & Hariton (1999). Secondly, the smart market is generally viewed as being too complex to implement. The requirement

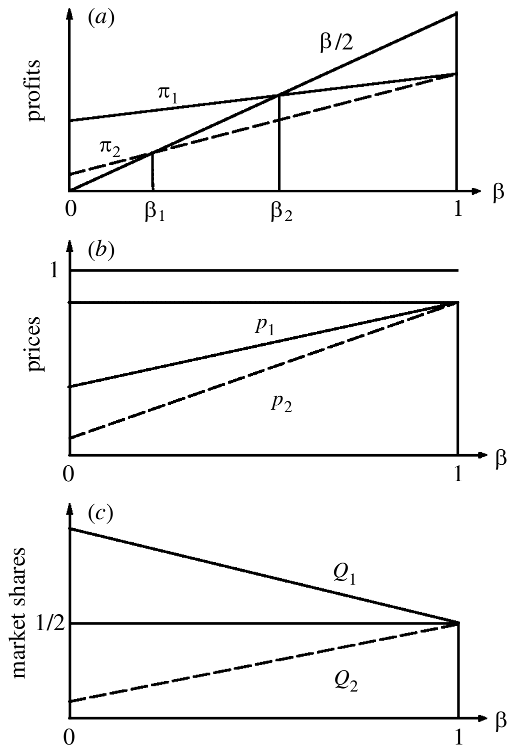

Figure 1. Illustrative network profits, prices and size.

that a bid is attached to every packet imposes large burdens on both users and already-congested resources (especially routers).†

Until 15 years ago, users of the Paris Metro were offered a choice of travelling in first or second class carriages. The only difference between the two carriages 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, for example, obtaining a seat were willing to pay the higher price; others, content to travel in a more congested carriage (on average), paid the lower second class fare.

Odlyzko (1997) has proposed that the same scheme be applied to the Internet. In the Paris Metro pricing (PMP) proposal, networks are partitioned into separate logical networks, with different usage charges applied on each sub-network. No guarantees of service quality are offered; but on average, networks charging higher prices are less congested. Users sort themselves according to their preferences for congestion and the prices charged on the sub-networks. The attraction of this scheme is its simplicity.

---

† The average packet size in TCP/IP, the suite of transport and application protocols used widely on the Internet, is 1600 bits. So a short email generates two or three packets; the PostScript version of this paper requires over 200 packets for file transfer; a 5 min telephone conversation generates around 1500 packets. Kelly *et al.* (1998) have proposed the 'proportionally fair pricing' scheme as a simpler alternative to the smart market that possesses several attractive features.

Even the simple PMP scheme is far removed from current practice. The question for ISPs is not what auction scheme should be used, but rather whether usage prices should be employed at all. The emerging market consensus for dial-up access to the Internet is a flat rate pricing structure: unlimited access for a fixed monthly fee. Leased-line access is charged according to the capacity of the line. The majority of users therefore face a marginal usage price of zero. There are, however, examples of usage pricing on the Internet. Users of NZGate, New Zealand's Internet gateway managed by Waikato University, are charged according to the volume of their traffic. Brownlee (1997) discusses the system, noting that the overheads of charging are significant. JANET, the UK academic network, charges users for incoming transatlantic traffic. The configuration of JANET means that measurement of traffic to and from North America is relatively inexpensive. Zero usage prices are more firmly established in the US.

(*b*) *An economic analysis*

There are two important economic considerations for networks maximizing profits. The desire to discriminate between consumers with heterogeneous valuations drives networks to form separate sub-networks (in Odlyzko's proposal) or to charge different prices (in a nonlinear pricing scheme). This is known as a 'segmentation effect': by differentiating a service or product, a firm can increase the gross surplus received by consumers, which can then be extracted to form profit. Offsetting this is any increase in competition between networks which results from the use of another sub-network or price; this is the 'competition' (in the terms of Champsaur & Rochet (1989)) or 'expansion' (Shaked & Sutton 1990) effect.

Gibbens *et al*. (1998) show that PMP may not survive in equilibrium between competing firms. In their model, the competition effect always outweighs the segmentation effect: both networks in a duopoly earn lower profits in any PMP equilibrium than in the non-PMP equilibrium. A more specific description follows.

(1) There is only one possible equilibrium when one firm offers two sub-networks, while the other offers one. Both firms earn lower profits in this equilibrium than in the non-PMP equilibrium.

(2) No PMP equilibrium exists when both firms in a duopoly offer two sub-networks. That is, given any combination of sub-networks and prices, one or other of the firms can increase its profit by changing its number of sub-networks and/or prices charged. The firms never settle in a stable (i.e. equilibrium) point in which both offer two sub-networks.

(These analytical findings are similar to the numerical results of Wilson (1989), who shows the same for priority supply classes of electricity.)

The same two economic effects come into play when considering flat rate (or 'buffet pricing') or nonlinear (e.g. a combination of a fixed plus usage) pricing. Traditional explanations of buffet pricing have focused on pure cost (see, for example, Nahata *et al*. 1999) or uncertainty (see, for example, Fishburn *et al*. 1997) factors. (A full survey of buffet pricing is not possible here; see Nahata *et al*. (1999) for many examples.) But strategic reasons are equally important.

The following model is used to demonstrate this. Consumers are able to buy a variable amount of a product from one of two firms. The firms charge up to two prices for their product: a fixed charge (like a membership or subscription fee) that is independent of the amount of the product bought by the consumer; and a per unit charge. The firms are horizontally differentiated, in the sense that when they charge the same prices, some consumers prefer one firm, other consumers prefer the other. Finally, there are positive network effects: each consumer values a firm's product more when that product is sold widely. To capture these features, suppose that consumers are distributed uniformly along the unit interval. Each consumer has a linear demand for a firm's product: with a usage price of p per unit, the consumer demands $1-p$ units of the product. (The fact that this number lies between 0 and 1 is just a normalization.) There are two firms, labelled 0 and 1, located respectively at 0 and 1 on the line. A consumer located at $0 \leqslant x \leqslant 1$ receives a utility from buying from firm i of

$$U(x,i) = V + t[ix + (1-i)(1-x)] + \tfrac{1}{2}(1-p_i)^2 + nD_i - f_i. \tag{3.1}$$

In this equation, V is a positive constant representing a common utility received by all consumers from either firm's product. The term $t[ix+(1-i)(1-x)]$ represents the 'transport cost' element of the consumer's utility: a consumer located at zero receives a utility component of t from buying firm 0's product; a consumer located at $x > 0$ receives a utility component of $t(1-x)$. Similarly, a consumer located at 1 (x) receives a utility component of t (tx) from buying firm 1's product. $t \geqslant 0$ is a parameter measuring the strength of this component of utility. (See below for further comment.) The third term is the surplus gained by the consumer from consumption of firm i's product when the usage price is p_i. It is straightforward to show (see, for example, Varian 1994) that the appropriate measure of this surplus is the area under the demand function $1-p_i$ and above the price line p_i.† This area is $\frac{1}{2}(1-p_i)^2$. The fourth term $nD_i \geqslant 0$ is a network effect term, where D_i is the total demand of consumers who buy from firm i, and $n \in [0,t]$ is a constant measuring the level of network effects. When n is high, there are strong positive benefits to each user on a network from other consumers' use of the network. When n is low (but still positive), these benefits are lower; this will be the case when there is congestion. Finally, f_i is the fixed price charged by firm i.

The two firms have identical production costs: a fixed cost k per customer (e.g. of installing a fixed link); a cost c per unit of demand (e.g. the marginal cost of transporting a fixed amount of data); and a fixed cost of m if a two-part tariff is employed, but not if a flat rate scheme is used (e.g. the cost of equipment to count packets on the Internet).

Note three things. First, the presence of a 'transport cost' captures horizontal differentiation: when the firms have equal demand and charge the same price, the consumers located near zero on the line prefer to buy firm 0's product, while consumers near the other end of the line prefer to buy firm 1's product. Secondly, the network effects depend only on firm 0's demand and not on firm 1's; in effect, the firm's products are assumed to be incompatible. At first, this may seem a strange assumption to make in the context of the Internet: after all, the Internet is nothing more than a 'network of networks', defined by widespread connectivity between

---

† To understand why, note that the consumer would have been willing to buy some units of the product at a higher price $\bar{p}_i > p_i$; call the demand at this price $\bar{q} \equiv 1 - \bar{p}_i$. Then at the lower price p_i, the consumer is buying the $\bar{q}$ units for less than (s)he would be willing to pay for them, and hence receives a surplus. Performing this comparison across all of the units bought by the consumer leaves the area stated in the text as the measure of total surplus.

multiple networks. As the previous section argued, however, the changing nature of the Internet is bringing interconnection agreements between asymmetric networks into question. In addition, if the firms' products in this model are interpreted as communication services (e.g. voice telephony or video conferencing), then lack of compatibility can arise through the use of proprietary standards in the communication applications. This section does not analyse the compatibility question. Instead, it assumes that compatibility is less than perfect; and, to simplify matters, that there is complete incompatibility. Allowing for partial compatibility would produce no substantial change in the model's results, but only complicate the calculations. Finally, it is assumed that positive network effects dominate, so that $n \geqslant 0$. Congestion in this model is equivalent to a reduction in n, but with n remaining non-negative.

The analysis is restricted to considering whether firms choose flat rates or two-part tariffs in equilibrium to maximize profits. This is in order to obtain the simplest possible assessment of whether it is possible for equilibrium to involve only volume-independent prices. There are two aspects to the restriction. First, even when prices do not depend on usage, they may nevertheless be nonlinear; for example, current Internet pricing schemes rarely depend on volume, but often are nonlinear in the maximum bandwidth available. Secondly, a full analysis would compare volume-dependent and volume-independent pricing schemes more generally (see, for example, Stole 1995; Armstrong & Vickers 1998; Rochet & Stole 1999).

The game has two stages. In the first stage, the firms choose what type of pricing scheme to employ (a flat rate or a two-part tariff) and the level of price(s) simultaneously. (The results are not sensitive to this order of choices; in fact, they are strengthened if the firms choose the type of pricing scheme before the price levels.) In the second stage, consumers choose which firm to buy from and how much to buy. The analysis concentrates on symmetric Nash equilibria in pure strategies. There are therefore two possible equilibria to consider: a flat rate equilibrium and a two-part tariff equilibrium.

Figure 2 illustrates the results of this model, showing the regions in (m, n) space in which different equilibrium outcomes hold.† The two-part tariff solution can be an equilibrium only if the fixed cost of charging usage prices is not too large: $m < \hat{m}$. Conversely, the flat rate solution can be an equilibrium only if the fixed cost of charging usage prices is sufficiently large $(m > \tilde{m})$.‡ In particular, the flat rate equilibrium does not exist when $m = 0$. In this case it can be shown that, when one firm charges a flat rate, the other firm's best response is always to choose a two-part tariff.

The figure also indicates that there may be zero, one or two possible equilibria, depending on the values of m and n, and the various critical levels. When an equilibrium exists, the flat rate is the unique equilibrium for 'high' values of m (greater than $\hat{m}$), and the two-part tariff the unique equilibrium for 'low' values of m (lower than $\tilde{m}$). For 'intermediate' values of m (between $\tilde{m}$ and $\hat{m}$), multiple equilibria can arise for sufficiently low n, while no equilibrium exists for sufficiently large n. Note, however, that the figure shows only the case in which $\hat{m} > \tilde{m}$. It is possible that $\hat{m} < \tilde{m}$, in which case multiple equilibria do not arise.

---

† In fact, the figure simplifies the situation substantially. Nevertheless, the figure shows some of the general features of equilibrium.

‡ This is true for the example illustrated in the figure; see Mason (2000) for the general condition.

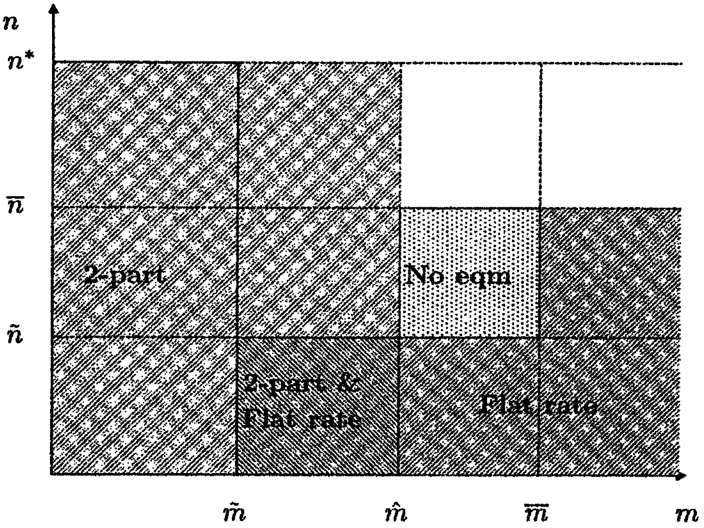

Figure 2. Equilibrium existence.

A particular example demonstrates the results even more starkly. Suppose that $t = 1, c = 0.04$ and $m = 0.003$. The absolute values of these variables are not important, only that m is an order of magnitude less than c.† Then for $n < 0.3026$, in the unique symmetric Nash equilibrium firms do not use two-part tariffs, but instead charge flat rates. Surprisingly, this suggests that the presence of congestion (i.e. a low n) makes it more likely that the unique equilibrium involves flat rates.

The explanation of this result uses again the segmentation and competition effects. A monopolist (weakly) prefers to use a nonlinear pricing scheme, since this increases consumers' gross surplus to be extracted as profit. But in a competitive market, firms must recognize that nonlinear prices can cause competition to be more fierce than when flat rates are used. With two-part tariffs, firms have two dimensions over which they can compete: they can cut both their fixed price and their usage price. In some situations, this latter competition effect can outweigh the segmentation effect. As a result, even a very small cost of implementing a pricing scheme with a usage component can make a large difference to the equilibrium outcome.

4. Conclusions

This paper has discussed some of the economic issues that arise in network interconnection and end-user pricing in the Internet. Previous work on Internet pricing has identified the need to have schemes that are 'efficient' (although this word is used in a variety of ways), stable (i.e. have desirable statistical properties), and simple. This paper has argued that these criteria need to be supplemented by the requirement of *equilibrium*: that interconnection agreements and pricing schemes be sustainable when firms maximize profit and individuals maximize utility.

---

† In this case, $\tilde{n} = 0.6803$, $\hat{n} = 0.7370$, $\bar{n} = 0.7406$, and $n^* = 0.7492$; and $\tilde{m} = 4.8 \times 10^{-6}$ and $\bar{m} = 0.1157$.

A relatively small number of economic principles have been discussed. In determining network interconnection, the relative importance of horizontal and vertical aspects proved to be a key determinant of whether networks wish to interconnect. In considering equilibrium pricing schemes, the analysis emphasized the segmentation and competition effects. These forces are always present and so, in this sense, are robust to various modelling assumptions. Of course, the exact balance between the forces depends on the situation in hand.

References

Armstrong, M. & Vickers, J. 1998 *Competitive price discrimination*. Mimeo.

Brownlee, N. 1997 Internet pricing in practice. In *Internet economics* (ed. L. W. McKnight & J. P. Bailey), pp. 77–90. Cambridge, MA: MIT Press.

Champsaur, P. & Rochet, J. C. 1989 Multiproduct duopolists. *Econometrica* **57**, 533–557.

Crémer, J. & Hariton, C. 1999 *The pricing of critical applications on the Internet*. Mimeo.

Crémer, J., Rey, P. & Tirole, J. 1998 *The degradation of quality and the domination of the Internet*. Mimeo.

Fishburn, P. C., Odlyzko, A. M. & Siders, R. C. 1997 Fixed fee versus unit pricing for information goods, competition, equilibria, and price wars. *First Monday* **2**(7).

Foros, O. & Hansen, B. 1999 *Competition and compatibility among Internet Service Providers*. Mimeo.

Gibbens, R., Mason, R. A. & Steinberg, R. 1998 Multiproduct competition between congestible networks. Working paper no. 9816, University of Southampton, Discussion Paper in Economics and Econometrics.

Gupta, A., Stahl, D. O. & Whinston, A. B. 1997 A stochastic equilibrium model of Internet pricing. *J. Econ. Dynam. Control* **21**, 697–722.

Kelly, F. P., Maulloo, A. & Tan, D. 1998 Rate control in communication networks, shadow prices, proportional fairness and stability. *J. Operation. Res. Soc.* **49**, 237–252.

Mackie-Mason, J. K. & Varian, H. 1997 Economic FAQs about the Internet. In *Internet economics* (ed. L. W. McKnight & J. P. Bailey), pp. 27–62. Cambridge, MA: MIT Press.

Mason, R. A. 1999 Compatibility between differentiated networks. Working paper no. 9909, University of Southampton, Discussion Paper in Economics and Econometrics.

Mason, R. A. 2000 Simple, competitive Internet pricing. *Eur. Econ. Rev.* (In the press.)

Nahata, B., Ostaszewski, K. & Sahoo, P. 1999 Buffet pricing. *J. Business* **72**, 215–228.

Odlyzko, A. 1997 A modest proposal for preventing Internet congestion. AT&T Labs Research. Mimeo.

OECD 1998 Internet traffic exchange: developments and policy. Report no. DSTI/ICCP/TISP(98). Paris: OECD.

Rochet, J. C. & Stole, L. 1999 *Nonlinear pricing with random participation constraints*. Mimeo.

Shaked, A. & Sutton, J. 1990 Multiproduct firms and market structure. *Rand J. Econ.* **21**, 45–62.

Srinagesh, P. 1997 Internet cost structures and interconnection agreements. In *Internet economics* (ed. L. W. McKnight & J. P. Bailey), pp. 121–154. Cambridge, MA: MIT Press.

Stole, L. 1995 Nonlinear pricing and oligopoly. *J. Econ. Management Strategy* **4**, 529–562.

Varian, H. 1994 *Microeconomic analysis*, 3rd edn. New York: W. W. Norton.

Vickrey, W. 1961 Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* **16**, 8–37.

Wilson, R. B. 1989 Efficient and competitive rationing. *Econometrica* **57**, 1–40.