PRIVATE NETWORKS PUBLIC OBJECTIVES
E. Noam and A. NíShúilleabháin (Editors)

Quality Choices in a Network of Networks

William Lehr

1. INTRODUCTION[1]

De-regulation, globalization, and rapid advances in technology are pushing us towards a world of virtual networks offering new, integrated services from a growing array of new service providers. The boundaries between public and private networks are blurring. Increasingly, we communicate via a patchwork mosaic of interconnected semi-autonomous sub-networks: a *network of networks*. In this new environment, the quality and reliability of our electronic communications infrastructure depends on the design decisions of each of the constituent sub-networks. Increasingly, these design decisions are being left to decentralized market forces. Can we rely on these forces to select the socially optimal level of quality?

Economists interested in the relationship between market structure and product quality have focused either on a monopolist's manipulation of quality to facilitate imperfect price discrimination[2] or oligopolistic quality differentiation to soften price competition.[3] The former research is applicable to the regulation of dominant carriers (e.g., AT&T or the Local Exchange Carriers (LECs)); while the latter may shed light on the effects of increased facilities-based competition (e.g., between AT&T and the other long distance companies, or between the LECs and alternative carriers such as Teleport or Metropolitan Fiber Systems). Further analysis of both of these circumstances is of great importance in assessing how the changes cited above will affect service quality. Both illustrate the potential inefficiency of market-based solutions: firms may either *under*- or *over*-invest in quality enhancements. Neither, however, considers what happens when control over the quality choice is decentralized. By assuming that the quality of *each* final good is chosen by a single agent, they ignore the coordination problem which confronts a network of networks.

This paper addresses this deficiency by presenting a model which shows how the fragmentation of network ownership and control creates a coordination problem which leads to inadequate investment incentives. The modelling perspective is quite abstract and ignores the messy details associated with engineering real-time, multimedia networks and the extremely difficult economic challenges of cost recovery and pricing in a broadband environment. The discussion of regulatory policies/politics is rather simplistic, and strong mathematical assumptions underlie the model's formulation (e.g., regarding the information available to decision-makers, the dimensionality of quality preferences, and the regularity of functions). The point of this exercise is to demonstrate how fundamental the basic coordination problem is. Even when much of the real world's complexity is stripped away, we are left with a difficult problem in collective choice which is not readily amenable to regulatory interventions. The principal lessons of this paper are cautionary. Although we should not expect a network of networks to adopt the first-best level of quality, traditional modes of direct intervention (via pricing regulations, penalties and/or minimum quality standards) may be worse. This conclusion may be obscured, but is unlikely to be reversed by a less abstract, more complex analysis.

In order to make the key features of the model and its results intelligible to non-economists, the discussion of technical details in the main body of the text is kept to a minimum.[4] Section 2 explains the formulation of the model in order to clarify the nature of the economic problem faced by a network of networks, and in order to relate the modelling abstraction to the real world situation. Section 3 presents the main conclusions which emerge from an analysis of the model. Section 4 addresses further extensions while Section 5 presents conclusions.

2. DESCRIPTION OF THE MODEL

The quality of communication services is inherently multidimensional in at least two important senses. First, we can measure quality in a variety of ways. For example, the quality of plain old telephone service depends on the level of line noise, the delay in establishing a connection, the probability of blocking or disruptions, and the flexibility/variety of service features. We may presume that all consumers would prefer higher quality along each of these dimensions, while differing in their willingness-to-pay for enhancements and in their willingness to trade off improvements along different quality dimensions. For example, voice telephony is intolerant of delays but (relatively) tolerant of bit errors, whereas certain data communications reflect opposite preferences. Thus, we should expect the willingness-to-pay for quality to vary systematically both across consumers and for each consumer based on the type of communication to be undertaken.

Appropriate product design and feature pricing strategies may exploit these multidimensional preferences in order to facilitate price discrimination. This can aid in recovering fixed costs and avoiding the deadweight losses associated with monopoly pricing (e.g., priority pricing to exploit variable tolerances for delay or special fees for data-reliable connections). Indeed, the desire to accommodate these diverse tastes is one of the forces which is fueling the emergence of semi-autonomous sub-networks. Since the focus of this paper is not pricing, let us ignore this type of multidimensionality and assume the quality of each sub-network can be represented by a scalar quality index, x_i.

The second sense in which quality is multidimensional arises because network services are a composite good. The effective *end-to-end* quality of service depends on the quality of each of the components which comprise the network. The quality of service in a network-of-networks will depend on the quality of each of the constituent sub-networks. Therefore, even if the quality of each sub-network may be described by a scalar index, x_i, the quality of the network of networks must be represented by a vector, $x \equiv (x_1, x_2, \ldots, x_N)$. To study the effects of decentralization of control over network design, I compare the choice of x when the quality of each of N interconnected sub-networks is chosen by a single agent with what happens when the sub-networks choose x_i for $i = 1,2,\ldots,N$ independently. Let us assume that the net benefits which accrue to the network designers of the ith sub-network can be described by a utility function, $U^i(x)$, and define the social welfare function as the sum of the N utility functions,

$$W(x) = \Sigma_{i=1,N} U^i(x). \tag{1}$$

This formulation permits great flexibility in identifying the sub-networks, while suppressing the distinction between service providers and consumers. It purposefully ignores the role of market prices and the competition for subscribers, implicitly assuming that the aggregate effects for each network are captured by the network utility functions. Several economists have noted how the positive externalities associated with larger networks may encourage the coalescence of smaller networks.[5] Subscribers with heterogeneous tastes on smaller networks may be willing to compromise in order to adopt a common, compatible technology. Often, however, formation of a single, centrally-controlled network is neither economically feasible nor desirable. Preferences may be sufficiently heterogeneous or privacy/control may be important enough to outweigh the positive externalities of forming a single network, even though interconnection remains desirable.

This paper begins with the assumption that we have a network of interconnected sub-networks. Perhaps the simplest example to consider is the case of N private corporate networks which are interconnected via the Public Switched Telecommunications Network (PSTN). Within-firm calls occur "on-net" and, presumably, the quality of these calls need depend only on the quality of the corporate network (i.e., office wiring, the PBX, etcetera).[6] The quality of "internet" calls between firms is likely to depend on the quality of *at least* the qualities of the originating, terminating, and interconnecting networks. In general, the quality of service experienced by subscribers on one corporate network may be affected by the quality of all the other networks which are interconnected.

The interdependence of network qualities may be quite complex, depending on what aspect of quality and type of network we are considering. For example, modem communications occur at the minimum baud rate supported along the transmission path. Only improvements in the speed of networks at the minimum rate permit faster communications. With back-up power arrangements, the reliability of each of network may be higher when interconnected if failure probabilities are uncorrelated.[7] In a packet network, the expansion of buffering at any node in the network may increase the likelihood of successful communications between all nodes.

The socially efficient quality choice, x^*, occurs when a single agent chooses the vector x so as to maximize total surplus, W(x). We may think of this as the *centralized* solution which would emerge if the sub-networks could costlessly coordinate their quality choices. If we exclude the possibility of external subsidies, then the centralized solution may require transfer payments among the sub-networks.

When control over quality choices is decentralized, we have an N-player game wherein each network attempts to set the quality of its network as a best reply to the independent choices of the other N-1 networks. A *decentralized* solution, x^0, is a Nash equilibrium to this game, wherein each network i chooses x_i^0 so as to maximize its net incremental willingness-to-pay for quality, $U^i(x)$. With suitable regularity conditions regarding the domain of x and the functional forms of the $U^i(x)$, we can guarantee existence of at least one centralized and one decentralized solution.[8] Additional assumptions allow us to assure that the solutions are in the interior of the domain of x and can be characterized by the following first order necessary conditions (FONC):[9]

Centralized Solution: x^* solves $\partial W/\partial x_i = \partial U^i/\partial x_i + \Sigma_{j\neq i}\partial U^j/\partial x_i = 0 \ \forall i \in N$ (2)

Decentralized (Nash) Solution: x^0 solves $\partial U^i/\partial x_i=0 \ \forall i \in N$ (3)

These two classes of solutions represent extremes. The centralized solution represents the ideal which could arise from perfect cooperation among the sub-networks. With efficient bargaining, perhaps during the stage when interconnection agreements are defined, it is conceivable that the sub-networks could agree to implement x^* and arrange side-payments among networks as necessary. In a sense, this is what occurred in the Bell Telephone network before divestiture. Ma Bell invested in improving the quality of network switches to optimize the benefits from both long distance and local exchange services. Improvements in intermachine trunking improved the quality of long distance services directly and local services indirectly by expanding the capacity for alternate routing. Ma Bell's cost/benefit analysis could take into account the effect on the combined long distance/local exchange market of changes in the quality of any constituent component.

At the other extreme, the decentralized solution presumes that perfect transfer payment schemes are impossible and assumes that agents care only about their private benefits. Although this is perhaps too pessimistic a view of how these decisions will be made, it seems overly optimistic to presume that the centralized solution will prevail. For example, consider the potential implications for antitrust and the costs of monitoring and enforcing the necessary transfer payments were the Baby Bells and AT&T to seek to reach such an agreement.

Without a further parameterization of the U^i it will be impossible to specify the magnitude of any welfare differences which may exist. The advantage of considering these two classes of solutions is their ability to highlight the coordination problem which arises when we decentralize network control.

3. COMPARISON OF THE CENTRALIZED AND DECENTRALIZED SOLUTIONS

A comparison of these two solutions yields four important conclusions:

(i) The decentralized solutions are all sub-optimal;
(ii) Symmetric quality is neither desirable nor likely;
(iii) The decentralized solution may provide either too much or too little quality; and,
(iv) Regulatory remedies may be difficult to implement.

Each of these conclusions is discussed at greater length below.

3.1. The decentralized solutions are all sub-optimal

From a comparison of the FONCs for the two solutions (Eqn 2 and 3), it is clear that x^* and x^0 will differ unless:

$$\Sigma_{j \neq i} \partial U^j / \partial x_i = 0 \ \forall i \in N. \tag{4}$$

This term measures the marginal externality imposed on other networks when network i improves its quality. In general, we would not expect this term to be zero and so should not expect a decentralized solution to be optimal. If the networks compete with each other and quality is costly to improve, then the externality might be negative. For example, the U^i might refer to profit functions for competing service providers or for corporations which use information technology as a source of competitive advantage (e.g., American Express improves its private network to offer more detailed customer billing services then VISA). Quality improvements on a competing network reduce own profits.

The networks may face a version of a Prisoners' Dilemma: when qualities differ, customers switch to the higher quality; but once qualities are matched, market shares are balanced. Although each firm's profits may be maximized if both could commit to low quality, it may be a dominant strategy to invest in high quality. Although the gains to consumers from having higher quality services may more than offset the reduction in industry profits, this need not be the case.

If the externality is negative, then the decentralized solution might result in excess quality wherein every network sets higher quality than is optimal. In this case, it may even be desirable to have *maximum* quality standards. The notion that there could be too much investment in quality may strike readers as implausible, yet corporate and political decision processes might impose just such a bias. After all, how often does one hear of a politician or corporate information officer (CIO) campaigning for lower quality and reliability (even when advocated cost-cutting will be expected to have this effect)? The costs of too little quality are apparent to everyone who uses the network, while awareness of the costs of too much quality may be apparent only to those few with detailed knowledge of the network. This may make the costs of errors in selecting the optimal quality level asymmetric for the decision-maker, inducing a bias in favor of erring on the side of excessive quality.

The externality would be positive if networks weakly benefited from improvements in the quality of other networks. For example, if other networks invest to reduce line noise or blocking probabilities, then internet callers to those networks may experience higher quality communications without bearing the costs of those improvements. In most cases, it seems likely that improvements in component quality lead to improvements in system quality, at least weakly. (The opposite might be true if quality improvements in one network lead to changes in traffic patterns which harmed the quality of another network or service.)

The above discussion should convince the reader that it is premature to presume *a priori* that we know the sign of the quality externality. Moreover, it is possible that the externality effects may be positive on some networks and negative on others. And, until we are able to sign the aggregate externality, we cannot determine which direction remedial policies should seek to move market outcomes should we seek to intervene. In spite of these difficulties, however, we can be reasonably certain that the externality is non-zero at x^*, and hence, the first term in Equation 2 is also nonzero, which corresponds to the FONC for the decentralized solution. Therefore, the decentralized solution will not be optimal.

It is also worth noting that generally the decentralized solution will not be unique, and when multiple solutions exist, these are likely to be Pareto rankable.[10] Let $x^{0\text{-Low}}$ and $x^{0\text{-high}}$ refer to the choices of x among the set of decentralized equilibria which minimize and maximize W(x), respectively. Thus, $W(x^*)-W(x^{0\text{ high}})$ is the minimum and $W(x^*)-W(x^{0\text{-low}})$ is the maximum welfare loss associated with decentralization. If there are multiple equilibria which are Pareto rankable, then there is the added coordination problem of selecting $x^{0\text{ high}}$ rather than $x^{0\text{-low}}$. For example, let us relax several of our assumptions regarding the form of U^i and suppose $U^i=\min(x_1,\ldots,x_N)$ such that the quality of communications depends on the minimum quality set by any network.[11] In that case every symmetric equilibrium with $x_i = x_j$, for $i,j \in N$ is a decentralized equilibrium, but only the one were $x_i=x\max$ for every network is socially efficient.

3.2. Symmetric quality is neither desirable nor likely

The previous example notwithstanding, we should not expect symmetric solutions for either the centralized or decentralized solutions to prevail in general. As noted earlier, customers are likely to differ in their willingness-to-pay for quality improvements based on differences in their traffic patterns. These heterogeneous tastes are reflected in differences in the U^i functions. If these differ, then it is reasonable to expect it to be optimal for networks to set different levels of quality and, in the market equilibrium, we should expect heterogeneous quality choices.

If all of the subscribers had identical tastes and started from symmetric positions (e.g., regarding their installed base), then we might expect all of the networks to agree to adopt the identical level of quality. Even with identical tastes, however, the symmetric solution may not be optimal if quality investments are a public good. For example, it may be optimal for only one of the networks to invest in a back-up power supply which could be made available to whichever network's primary power supply failed. Each of the networks may prefer that it not be selected as the one to make the investment and yet strictly prefer that at least one network makes the investment.

In the more typical case where subscribers tastes for quality differ, the heterogeneity in network quality will make it more difficult for policy-makers to effect welfare-improving quality improvements. For example, a uniform minimum quality standard will not be able to support the optimal solution since it cannot be higher than the minimum quality network in the centralized solution x^*, and thus will not be binding on networks which are supposed to set higher quality standards.

3.3. The decentralized solution may provide either too much or too little quality.

As noted earlier, unless we can specify the sign of the externality, we cannot say whether the welfare loss results from too much or too little quality. However, even if we assume that the externality is weakly positive for every network, we can only conclude that quality is too low on average.[12] By this we mean that it cannot be the case that *every* network sets quality too high in the optimal solution; however, it may be the case that *some* networks set too high quality in the decentralized solution.[13] This could occur if quality choices are *strategic substitutes* rather than *strategic complements*.[14] For example, in a packet switching network, investments to increase buffering capacity in one network may make it less desirable to increase buffering capacity in another network since it does not matter where packets are buffered. In this case, quality investments would be strategic substitutes. Alternatively, if the speed of one network is improved it may increase the marginal advantages from increasing the speed of other networks, or quality investment may be strategic complements.[15] When quality investments are strategic complements and the externality is positive, then we can be sure that the decentralized solution results in too little investment in quality. Although this seems to be the presumption with which most would-be regulators begin, it is instructive to note that this result depends on a number of important assumptions.

3.4. Regulatory remedies may be difficult to implement

Once we recognize that the market process will adopt a sub-optimal solution, it is reasonable to consider potential regulatory remedies. First, it should be obvious that there will exist a system of penalties and subsidies which, in theory, could support the efficient solution. Penalty/subsidy functions of the form:

$$T^i\,(x_i)=\Sigma_{j\neq i}U^j(x_i,x^*_{-i})-\Sigma_{j\neq i}U^j(x^*_i,x^*_{-i}) \qquad (5)$$

will cause each of the networks to internalize the externality and will support the efficient solution.[16]

Although this penalty scheme is theoretically plausible, it may be extremely difficult to implement in practice. The regulator would need to know each of the network U^i functions which would require detailed knowledge about both subscribers' willingness-to-pay and the costs of improving the quality of each network. Acquiring this information is likely to be quite expensive, if not impossible. The costs of implementation would represent a deadweight loss which might very well overwhelm the welfare loss which they were designed to correct.

Similarly, it may be possible to design a pricing scheme for interconnection or for terminating network messages which would support the optimal solution. Once again, however, these would require more detailed information than is likely to be available to regulators and there is no reason to expect the appropriate prices to arise naturally (e.g., via a tatonnement process). Traditional externality problems are often more readily amenable to price-based solutions because the externality is associated with a commodity good. In the model presented here, the incremental quality of different networks are not perfect substitutes which can be traded. It is important *which* network's quality is improved. The marginal externality associated with the quality of each network depends on traffic patterns which may be affected by pricing schemes, thereby further complicating the regulatory problem. Furthermore, pricing or penalty/subsidy schemes which do not accurately reflect network payoffs may result in an outcome which is worse then the worst decentralized solution, $x^{U\text{-low}}$.

Minimum quality standards (MQS) offer another regulatory alternative which are superficially appealing precisely because they are easy to implement and require less information. If the centralized solution is symmetric, then an MQS can support the optimal solution, but when tastes are heterogeneous and both centralized and decentralized solutions call for different network qualities, the effects of MQS are more complex.

If the decentralized equilibrium is unique, the externality is positive and qualities are strategic complements, then we know there exists an MQS which is welfare improving. One such MQS is equal to the minimum quality which is set by any network in the optimal solution. This has to be welfare improving because the above assumptions guarantee that each network chooses too low quality in the decentralized solution, so forcing the minimum quality network to assume its level will weakly increase the quality of every other network. Since we have assumed that the externality is positive, this benefits all networks in aggregate. An MQS which exceeds this level, however, may be worse than the worst decentralized solution.

Alternatively, if qualities are strategic substitutes then the gains from improving the quality of a lower quality network may be more than offset by reductions in the quality of a higher quality network. In this case, there may not exist a welfare enhancing MQS.

If there are multiple equilibria, then MQS which are not binding may be welfare improving in an ex ante sense. If the equilibria are locally separated, then an MQS which helps select a Pareto superior Nash equilibrium can yield important benefits and yet not be binding on any of the networks in the observed equilibrium. Therefore, it is wrong to presume that the existence of an MQS which is not binding has no effect. However, since the inferior equilibrium is not observed and we may not be sure that it would have been selected in any case, it may be difficult to assess the gains from such an MQS.

The difficulties of calculating the optimal MQS and the dangers of setting one so high that welfare actually declines should caution us against setting overly aggressive performance standards. On the other hand, the potential coordination benefits suggest that more modest standards may offer large gains which will be difficult to measure directly.

The generality of the preceding conclusions is both their greatest virtue and their greatest vice. They illustrate the fundamental nature of the coordination problem which results from decentralizing network control. In the next section, I present a very simple numerical example and discuss a variety of extensions to the basic model.

4. EXTENSIONS AND FURTHER QUESTIONS

The next two sub-sections present simple examples to further elucidate the preceding discussion, while the third sub-section presents initial speculations of how introducing uncertainty might further complicate matters.

4.1. Example #1: Linear Quality and Demand

Consider a situation where there are N private networks which are interconnected through the PSTN (the quality of which we will treat as a parameter) and let the U^i have the following form:

$$U^i = A_i x_i + B_i \Sigma_{j \neq i}(\alpha(x_i + x_j) + \beta x_p) - C x_i^{\delta} \qquad (6)$$

where A_i, B_i, C, α, β and δ are positive constants with $\delta > 1$. The first term corresponds to the value to subscribers of on-net calls, which increases as the quality of the network, x_i, increases. The second term corresponds to the value of internet calls. In this formulation, the quality of internet calls depends on the sum of the quality of the originating and terminating network, weighted by α, and the quality of the PSTN, given by x_p, and weighted by β. The (A_i, B_i)-parameters capture taste differences among the networks for on-net and off-net calling. The last term is the cost of increasing quality and δ is assumed greater than 1 to assure that costs are strictly convex. The FONC for the decentralized and centralized solutions are:

$$\partial U^i / \partial x_i = 0 = A_i + \alpha(N-1)B_i - \delta C x^{\delta-1} \qquad (7)$$

$$\partial W / \partial x_i = 0 = \partial U^i / \partial x_i + \alpha \Sigma_{j \neq i} B_j \qquad (8)$$

which yield the following unique solutions:

$$x_i^0 = \{[A_i + \alpha(N-1)B_i]/\delta C\}^{1/(\delta-1)} \text{ and } x_i^* = \{[A_i + \alpha(N-1)B_i + \alpha(N\hat{B} - B_i)\ ]/\delta C\}^{1/(\delta-1)} \qquad (9)$$

where $\hat{B} = (1/N)\Sigma_i B_i$. We know that the solution is unique since the U^i are strictly concave in x_i. Also, note that the decentralized solution provides too little quality ($x_i^0 < x_i^*$) since the externality is positive and quality is a strategic substitute, and that surplus is lower in the decentralized solution. Furthermore, note that unless the A_i and B_i are identical, asymmetric qualities are both optimal and a decentralized equilibrium. Finally, notice that the optimal and decentralized solutions do not depend on the quality of the PSTN and so changes in its quality, while welfare improving for the N private networks, would not influence their quality choices. These results are readily apparent by inspection of the two solutions.

Now, let us assume that costs are quadratic ($\delta=2$) and the networks are symmetric ($A_i=A$ and $B_i=B$). With these additional assumptions, the solutions simplify to:

$$x_i^0 = [A+\alpha(N-1)B]/2C \quad \text{and} \quad x_i^* = x_i^0+\alpha(N-1)B/2C \tag{10}$$

which allows us to obtain a simple expression for the magnitude of the welfare loss associated with decentralization:

$$W(x^*)-W(x^0) = \alpha^2N(N-1)^2B^2 \tag{11}$$

Notice that the welfare loss increases with the value placed on internet calling quite rapidly.

4.2. A second simple example: quality is the minimum of the network qualities

Now, consider a slightly different version of the above example with symmetric utilities and quadratic costs, where the U^i have the following form:

$$U^i = Ax_i + B\Sigma_{j\neq i}\min(x_i,x_j,x_p) - Cx_i^2 \tag{12}$$

In this example, the quality of internet calls is no longer the weighted sum of the qualities of the interconnecting networks, but rather the minimum of the quality experienced on the end-to-end path from the originating and terminating networks. With this formulation, the payoffs are no longer continuously differentiable, but it is still the case that the externality is weakly positive and that quality choices are weakly strategic complements.

In order to derive the solutions, notice that even if a network expected all of the other networks to choose zero qualities, it would choose to set $x_i=A/2C>0$ because on-net calling would still be valuable; and, no network would ever choose quality greater than $(A+(N-1)B)/2C$, even if all of the other networks and the PSTN had arbitrarily high qualities. Therefore, the optimal centralized solution is for all networks to choose the x_i equal to $A/2C$ if x_p is less than $A/2C$; and, otherwise choose x_i equal to the minimum of x_p and $(A+(N-1)B)/2C$. Any symmetric solution where x is between $A/2C$ and the minimum of x_p and $(A+(N-1)B)/2C)$ is a decentralized solution. Thus, in this example, the optimal solution is also always a Nash equilibrium; however, it is only when the quality of the PSTN is sufficiently low ($x_p<A/2C$) that we can be sure the two solutions will coincide and there will not be a welfare loss associated with decentralization of network control. Once the PSTN's quality gets sufficiently high, there is a potential coordination problem associated with selecting one of the symmetric low quality equilibria.[17]

Finally, note that as the quality of the PSTN increases, the losses from failing to coordinate on the efficient solution increases. In a sense, quality choices become more "strategic." Ignoring bypass issues, when x_p is sufficiently low, everyone agrees on the optimal solution and voluntary enforcement of the efficient outcome is easily achieved. As the network quality increases, however, externality issues become more important.[18] Whenever the quality-externalities become more important, both the efficiency losses from failing to enforce the centralized solution and the costs of enforcement[19] are likely to be larger. It is exactly under such circumstances that specialized institutional structures (e.g., voluntary standard setting bodies) may become important, especially if information asymmetries preclude effective use of more direct regulatory interventions (e.g., the government orders everyone to adopt q^*).

In the first example, quality investments are perfect substitutes -- improvements in any of the networks along the end-to-end path contribute to overall quality. The second example examines the case where quality is determined by the weakest link in the end-to-end chain. Which functional form better reflects reality is a question for the engineers, and surely depends on what one means by "quality." The first might correspond to a packet network where one's measure of quality is *average* end-to-end packet delay, while the second might be relevant if one is concerned with the *maximum* end-to-end delay. The former measure is important if one is interested in supporting asynchronous terminal sessions, while the latter is important if one wishes to support video or telephony.

4.3. Effect of Uncertainty

In all of the analyses above, I have assumed perfect information. As long as quality is *ex post* verifiable and it is possible to specify and enforce complete, contingent contracts, uncertainty should not represent a problem for implementation of the centralized solution. The presumption of ex post verifiability and enforcement may be reasonable for agreements governing quality attributes which are based on a large sample of (at least in principle) inexpensive observations. For example, delay until dial tone is received or line noise (bit error rates) seem amenable to low cost monitoring and successful contracting. Breach would be quickly detected. These sorts of situations are perhaps the easiest for regulators to address.

In contrast, network reliability -- interpreted as freedom from major disruptions -- offers a more difficult problem. Increased reliability is provided via back-up capacity. Since major network failures are (of necessity) a very infrequent occurrence, it may be much more difficult to reliably ascertain the quality of back-up systems ex ante. A moral hazard problem may arise if the probability of failure is low enough, the costs of unreliable back-up are great enough, and there are no criminal penalties available to deter breach ex ante. For example, in the absence of criminal penalties, a "fly-by-night" database network might agree to provide high-quality back-up services which would become operational in the event of a major system failure on connected sub-networks. The back-up provider could collect insurance premiums up front and declare bankruptcy if a major failure resulted in it breaching its contract.

From a modelling perspective, the effect of uncertainty would be likely to increase the multiplicity of equilibria since they would depend on participants' beliefs regarding true costs and benefits and each other's strategies. This seems likely to further complicate the coordination problem and thus the present analysis may be overly optimistic. On the other hand, if uncertainty is sufficiently severe we may be blissfully ignorant of just how far from the optimal outcome we are.

5. CONCLUSIONS

This paper presents a first step towards understanding how the fragmentation of control over our information infrastructure might affect the quality of service. The results suggest that we cannot remain comfortable that the market's *Invisible Hand* will guarantee a simple solution. In a network of interconnected sub-networks, each network will invest in improvements which will affect the interests of subscribers on other networks. The quality-externality may be positive or negative. There are likely to be a multiplicity of equilibria which will further complicate efforts to internalize the externality and make it resistant to simple regulatory remedies. Since tastes for quality are inherently heterogeneous, the optimal solution will involve heterogeneous qualities, which cannot be enforced with uniform quality standards. The present analysis assumes perfect information, but in more realistic situations, we might expect there to be significant uncertainty and asymmetric information which would hamper both public policy interventions and/or collective decision-making among the sub-networks. Although it is theoretically possible to support the optimal quality solution with centrally-administered subsidy/penalty schemes or pricing, implementing these in practice may be very difficult. Minimum quality standards, however, may be useful in precluding coordination on a Pareto inferior equilibrium, when there are multiple such equilibria.

Thus, the growing public concern regarding the effects on network quality of increased decentralization, deregulation and its attendant implications for industry structure appear warranted. Additional theoretical and empirical work is needed to help determine the magnitude of the potential threat. We need much better data on the costs and benefits of improving quality and reliability in our information infrastructure and how these are allocated among producers and consumers.

REFERENCES

Besanko, D., Donnenfeld, S. and White, L., "The Multiproduct Firm, Quality Choice and Regulation," *The Journal of Industrial Economics*, vol. XXXVI, no. 4, 411-429, June 1988.

Bulow, J., Geanakoplos, J. and Klemperer, P., "Multimarket Oligopoly: Strategic Substitutes and Complements," *Journal of Political Economy*, 93, 488-511, 1985.

Gabszewicz, J. and J. Thisse "Price Competition, Quality and Income Disparities," *Journal of Economic Theory*, vol 20, 340-359, 1979.

Heal, G., "Economics of Networks," draft CITI Private Network Conference paper, Winter 1992.

Laffont, J. and J. Tirole, "The Regulation of Multiproduct Firms, Part I: Theory," *Journal of Public Economics*, vol 43, 1-36, 1990.

Laffont, J. and J. Tirole, "The Regulation of Multiproduct Firms, Part II: Applications to Competitive Environments and Policy Analysis," *Journal of Public Economics*, vol 43, 37-66, 1990.

Motta, M., "Endogenous Quality and Coordination of Decisions," Center for Operations Research and Econometrics Discussion Paper #9152, Universite Catholique de Louvain, Belgium, October 1991.

Schmalensee, R., "Market Structure, Durability, and Quality: A Selective Survey," *Economic Inquiry*, vol 17, 177-196, April 1979.

Shaked, A. and J. Sutton, "Relaxing Price Competition Through Product Differentiation," *Review of Economic Studies*, vol 49, 3-1, 1982.

Spence, M., "Monopoly, Quality and Regulation," *Bell Journal of Economics*, vol 6, no 2, 417-429, Autumn 1975.

ENDNOTES

[1] I would like to thank Stan Besen, Niel Stolleman, Glen Woroch and participants in the Columbia Institute for Tele-Information's conference on *Private Networks and Public Objectives* on May 15, 1992, for helpful comments. All remaining errors are my own.

[2] Besanko, D., Donnenfeld, S. and White, L., "The Multiproduct Firm, Quality Choice and Regulation," *The Journal of Industrial Economics*, vol. XXXVI, no. 4, June 1988, 411-429, Laffont, J. and J. Tirole, "The Regulation of Multiproduct Firms, Part I: Theory," *Journal of Public Economics*, vol 43 (1990a) 1-36, Laffont, J. and J. Tirole, "The Regulation of Multiproduct Firms, Part II: Applications to Competitive Environments and Policy Analysis," *Journal of Public Economics*, vol 43 (1990b) 37-66, Schmalensee, R., "Market Structure, Durability, and Quality: A Selective Survey," *Economic Inquiry*, vol 17, April 1979, 177-196 and Spence, M., "Monopoly, Quality and Regulation," *Bell Journal of Economics*, vol 6, no 2, Autumn 1975, 417-429.

[3] Gabszewicz, J. and J. Thisse (1979) "Price Competition, Quality and Income Disparities," *Journal of Economic Theory*, vol 20, 340-359, Motta, M., "Endogenous Quality and Coordination of Decisions," Center for Operations Research and Econometrics Discussion Paper #9152, Universite Catholique de Louvain, Belgium, October 1991 and Shaked, A. and J. Sutton (1982) "Relaxing Price Competition Through Product Differentiation," *Review of Economic Studies*, vol 49, 3-14.

[4] Interested readers are encouraged to see the working paper on which the present chapter is based entitled "Network Quality Choices in a Network of Networks," January 1992, CITI Working Paper #521.

[5] Heal, G., "Economics of Networks", draft CITI Private Network Conference paper, Winter 1992.

[6] In a world of virtual networks, building networks and Centrex services, even on-net calls may be routed across facilities whose quality is controlled independently.

[7] Assume networks 1 and 2 provide mutual emergency back-up power such that either network fails only if both networks fail. This agreement improves the reliability of both networks if the probability failures are either negatively or uncorrelated; however, if they are positively correlated then reliability may decline. The latter could occur if failures propagate across networks.

[8] To guarantee existence of a socially optimal *centralized* solution, x^*, define the domain of x such that the quality of sub-network i is a real scalar such that $0 \leq x_i \leq xmax < \infty$; and let each of the $U^i(x)$ be continuous on the domain of x. To guarantee the existence of a *decentralized* or Nash solution, further assume that the U^i are concave in own-quality, x_i.

[9] To allow characterization of the solutions by FONC, assume that the U^i are twice continuously differentiable on x. To guarantee that the solutions x^* and x^N are interior, assume the following boundary conditions hold for each i: (i) $U^i(0,x_{-i}) \geq 0$ (normalization); (ii) $U^i(xmax,x_{-i}) < 0$ (guarantees never optimal nor equilibrium for any network to set $x_i = xmax$); and, (iii) $U^i_i(0,x_{-i}) > 0$ (guarantees never optimal to set quality of network i to zero). These assumptions are sufficient, but not necessary to guarantee the existence of interior solutions.

[10] Even if the centralized solution is not unique, all solutions will yield the same total welfare. Assumption that U^i is strictly concave in x_i is sufficient to guarantee uniqueness of decentralized solution.

[11] This example does not satisfy the assumptions made above regarding the form of U^i, (e.g., U^i is not twice differentiable and $U^i(xmax,x) > 0$). It is useful because it clearly demonstrates the problem of multiple equilibria.

[12] Formally, assume that $\partial U^j/\partial x_i \geq 0 \ \forall j \neq i$ and $\forall i$, $\exists\, j \neq i$ such that $\partial U^j/\partial x_i > 0$ so that for each network there is at least one other network which strictly benefits when network i improves its quality.

[13] The proof is by counter-example.

[14] Besanko, D., Donnenfeld, S. and White, L., "The Multiproduct Firm, Quality Choice and Regulation," *The Journal of Industrial Economics*, vol. XXXVI, no. 4, June 1988, 411-429.

[15] Formally, x_i and x_j are strategic complements or substitutes if $\partial U^i/\partial x_i \partial x_j$ and $\partial U^j/\partial x_i \partial x_j$ are both weakly positive or negative, respectively.

[16] Note that $\partial U^i/\partial x_i + \partial T^i/\partial x_i = \partial W/\partial x_i = 0$ at x^*.

[17] Although an MQS can easily support the efficient outcome, this is not especially interesting since the example is so simple.

[18] Free-riding is not really an issue in the symmetric case, but it could be in the asymmetric case where the B_i differ.

[19] One might expect monitoring/enforcement costs to be larger when the private gains from deviating from q^* are larger according to the old adage "where there's a will, there's a way.."