

Innovations and Upgrades in Virtualized Network Architectures

Paul Laskowski

Universidade Técnica de Lisboa / UC Berkeley

John Chuang

UC Berkeley

ABSTRACT

The difficulty of making changes to the internet architecture has spawned widespread interest in virtualized testbeds as a place to deploy new services. Despite the excitement, uncertainty surrounds the question of how technologies can bridge the gap from testbed to global availability. It is recognized that no amount of validation will spur today's ISPs to make architectural changes, so the testbed itself must somehow provide global availability. We investigate whether a virtualized architecture that is widely offered by commercial ISPs would support the adoption of new services or upgrades to the infrastructure, and whether ISPs would ever support such an architecture. According to our economic analysis, the answer depends critically on how money flows to network and service providers. If the virtualized network inherits the market structure prevalent on the internet today, which we call *network-gatekeeper*, investment levels are likely to be poor. On the other hand, we identify two superior market types, *mix-and-match* and *service-gatekeeper*, which can improve incentives to invest in services, and even in network upgrades. We discuss how these market types may be implemented.

Categories and Subject Descriptors

C.2.4 [Computer-Communication Networks]: Distributed Systems; J.4 [Social And Behavioral Sciences]: Economics

General Terms

Design, Economics, Theory, Legal Aspects

Keywords

Virtualization, Technology Adoption, Supply Chains, Evolution, Innovation, Clean-Slate Design.

1. INTRODUCTION

The gulf between the internet we have and the internet we want seems wider today than ever before. For decades, researchers have identified ways to improve the network, calling for such things as quality of service, multicast capability, mobility support, and IPv6. More recently, several large-scale research projects have been tasked with redesigning the internet from scratch, further expanding our view of potential capabilities [1][6][15]. Unfortunately, even as our aspirations have risen, the network in use by the vast majority of users has remained stagnant, and even ossified [3][12]. Many studies have noted the internet's resistance to evolution [3][5][8][10][11][12][13][16][17]. Internet service providers are the gatekeepers to any

architectural change, and we now accept that they lack the economic incentives to adopt new technologies [10][13].

Faced with the impossibility of changing the internet directly, attention has recently shifted to large-scale experimental testbeds as an alternate place to deploy new technologies. These flexible platforms can be thought of as meta-networks, capable of simulating a wide-variety of architectures. Several testbeds are now in development, including the Global Environment for Network Innovations (GENI), a collaboration among 29 initial research teams that may cost \$350 million and connect 100,000 users [6].

Despite the excitement, uncertainty surrounds the question of how technologies can eventually bridge the gap from a testbed to global availability by end-users. In a canonical paper, Peterson et al. point out that no amount of validation on a separate testbed will induce today's ISPs to implement significant changes to the internet architecture [12]. The obstacles preventing adoption are simply not technological, but rather economic. This means that if new technologies are to reach a global population of users, the testbed itself must provide global availability.

What are the chances, then, that today's ISPs (or a new set of ISPs) could be convinced to support GENI, or any similar architecture, expanding it to global proportions? Moreover, would such a network, spanning a great number of providers, be more evolvable than today's internet? To answer these questions, this study will analyze the economic incentives surrounding virtualized testbeds.

1.1 Three Supply Chains

Alarmingly, for all the resources invested in developing the technical aspects of testbeds, rather little is known about their economic properties. In the worst case, ISPs may agree to support a virtualized architecture, but insist on retaining their current role as gatekeepers, selecting which protocols to allow in their domains. If nothing else is done to alter the nature of payments, the virtualized architecture will act as nothing more than a technical lubricant, allowing ISPs to install new technologies with greater ease. By all accounts, today's lack of evolvability has nothing to do with technical barriers, so this is unlikely to break the impasse.

As these observations suggest, the behavior of a testbed depends crucially on who selects the technologies that run on the network. We will therefore embed this choice as a central feature of our network model.

A typical design for a virtualized architecture centers around a *virtualization interface*. This is an abstraction that hides the implementation of the infrastructure below, and allows multiple technologies to run on top of it. In keeping




Figure 1: Structure of a Virtualized Architecture

with previous work [5][9][17], we will refer to the architecture below the virtualization interface as *network or infrastructure*, and the architecture above the interface as *services*. As shown in Figure 1, we imagine that the network is operated by a set of network providers (NPs), and the services are operated by a set of service providers (SPs, see [17] for a technical discussion of how such an architecture may be implemented).

In order to make use of a service, an end-user will need both a network provider and a service provider, and each will need money to finance its cost of operation. Examining our abstraction, we see that there are three canonical ways that money can flow through this system. First, as depicted in Figure 2.A, the end-user can pay the network provider, who passes on a payment to the service provider. Economic viability also dictates that the network provider gets to select the service provider (if the service provider were selected by the user, it could charge as much as it wanted from the network provider).

We will call this arrangement a network-gatekeeper market, since a new service cannot be deployed without the consent of the network provider. A good example of a network-gatekeeper market involves cable television. In this case, the network provider is a cable operator, and the service providers are cable television networks. An end-user is not allowed to contract directly with television networks to purchase channels she is interested in. Instead, the channel lineup is controlled by the cable operator, who purchases content on the user's behalf.¹

In many cases, the network provider and the service provider in a network-gatekeeper market are the same firm. Able to choose any service provider it wants, it should come as no surprise that a network provider will often plump for itself, and avoid having to split profits with another firm. For example, Ratnasamy et al. consider the possibility of network evolution through competing network layer protocols [13]. The IP layer network can certainly be considered as a service in our model, and some users might prefer, say, IPv6 over IPv4. The choice is made by network providers, however, who generally choose a single protocol that they can supply themselves.

Similarly, end-users have no say in what protocol is used for their traditional (non-VoIP) telephone service, selection

¹ Advertising revenue can be thought of as a payment from end-users directly to cable networks, so this example does not fit the network-gatekeeper market perfectly.

of IP level routes, IP level multicast transmissions, or cell phone transmissions. All of these may be understood as network-gatekeeper markets, in which the network provider itself provides any available protocol. In principle, an entrepreneur would be free to form a service provider to provide an alternate protocol, but she would have to convince network providers to purchase it (or even accept it for free). Without their cooperation, users can have no access to the new technology.

Of course, our entrepreneur would be interested to see if her service could be implemented at the application layer, bypassing the need for ISP cooperation. For example, telephone service can be implemented as a network layer service, or as an application over IP. In the latter case, it is end-users who choose a VoIP service provider, and there is little network providers can do to interfere. Indeed, this is the idea behind our second canonical market type.

As depicted in Figure 2.B, an end-user can independently select a network provider and a service provider, and pay each of them directly. We will call this scenario a mix-and-match market. Such markets are common for services that can run on top of the internet, including online stores, multiplayer games, and photo-sharing sites.

In many cases, service providers in a mix-and-match market do not charge users directly, but rather earn revenue from advertising. Online newspapers, file-sharing networks, and social networking applications are likely to operate in this way. While this is an interesting strategy from a business perspective, we will set it aside for the current study. For our purposes, advertising and direct payments have similar properties, and we will not worry about the distinction.

One final possibility remains for the flow of money through a virtualized architecture, which we will call a service-gatekeeper market. As depicted in Figure 2.C, an end-user may contract directly with a service provider, who selects and pays a network provider for use of the underlying infrastructure.

An example of a service-gatekeeper market may be found in unbundling-based ISPs – providers that do not own their own cables, but purchase connectivity from an underlying network provider. Of course, such ISPs often have to compete against the underlying network's own IP service, so that network may not want to offer the unbundling-based ISP a good price. The situation may be different however,




Figure 2: Monetary Flows in a Virtualized Architecture

if a regulator forbids vertical integration between network and service providers. That is the situation we will focus on in this study.

It is worth noting how our three monetary paths correspond to the economic notion of supply chains. In this analogy, a service provider and a network provider each produce one of two inputs that must be combined to make a final product. If we were to reverse the direction of each arrow in Figure 2, they could represent the direction in which these inputs were sold, much like a supply chain diagram. For example, in the network-gatekeeper market, the service provider sells a service input to the network provider, who assembles both inputs into a complete product that can be sold to the end-user. Although supply chains have been thoroughly studied in economics, a virtualized architecture is rather unique in the flexibility we have to rearrange the order of production.

We should admit that the market types we have described are somewhat idealized; many services we observe in today’s industry do not align perfectly with any of the three. Nevertheless, we believe that the three markets we have identified are natural representatives of the broader space. In the spirit of clean-slate design, we will focus on these three in order to gain a broad view of how our design choices affect economic outcomes.

The market types that exist in today’s internet are there more by accident than by design; they are the result of founding-era engineering decisions molded by tussles among network players. As we design our next generation network, however, we have an opportunity to choose how network services should be financed. Even casual observation suggests that this decision can have an immense, if not definitive, effect on how evolvable the architecture is. Many of the services that are financed by a mix-and-match market – social networking sites, search engines, file-sharing networks – are renowned for the rate at which new innovations replace old ones. Meanwhile, the network-layer architecture, financed by a network-gatekeeper market, is stagnant and ossified [3][12].

Motivated by these observations, this study will build an economic model to analyze our three market types, paying particular attention to the incentives firms have to invest in service innovations, and to upgrade the infrastructure. The rest of this paper is organized as follows: In Section 2, we will set up our economic model. In Section 3, we will proceed to analyze each of our three market types. In Section 4, we will discuss implementation issues and conclude the paper.

2. MODEL SETUP

Our modeling framework is designed to meet three main criteria. First, it must be capable of representing all three market types. Second, it must allow consumers to switch gradually among products, not all together as in traditional price competition. This “softness” in the consumer demand

is what allows network providers to elevate prices above marginal cost – a prominent feature of today’s market. Third, we want our model to be tractable, yielding closed-form solutions.

To the best of our knowledge, no previous framework in economics is capable of meeting these criteria. On a high level, our strategy will be to represent the end-user market with Cournot, or quantity, competition. On the other hand, we will use price competition to represent exchanges between firms. This distinction captures the notion that firms are more discerning buyers than end-users. Firms are more likely than consumers to carefully compare alternate products and review their choices regularly, for example.

While these games are standard throughout economics, one major obstacle prevents us from applying them directly. Recall that a mix-and-match market allows both network providers and service providers to sell directly to end-users. We would therefore like to model both exchanges using Cournot competition. Classic Cournot games cannot combine in this fashion, however – if both types of providers choose quantities, there is no guarantee that the total quantities of network and service will match!

In response, our main analytic achievement is to reformulate the classic Cournot theory so that it extends naturally to the mix-and-match market. We present our full theory in the appendix, using an axiomatic approach.

2.1 Measuring investment

We begin with a duopoly of network providers, $\{NP_1, NP_2\}$, and a set of m service providers, $\{SP_1, \dots, SP_m\}$. Let x_i be the quantity of network i , and y_j be the quantity of service j . Our assumptions will ensure that $\sum x_i = \sum y_j$. We assume that prices are given by an inverse demand function, $\mathbf{t} = \{t_{ij}\}$, where $t_{ij}(x_1, x_2, y_1, \dots, y_m)$ is the price of the combination of network i and service j as a function of quantities. For each market type, we will begin with a base case in which demand is equal for all product combinations,

$$t_{ij} = 1 - \sum_{k=1}^2 x_k = 1 - \sum_{k=1}^M y_k, \text{ for all } i, j. \quad (1)$$

This linear form is standard in economics. Our investment model draws heavily upon the game-theoretic literature on innovation [14]. In accordance with this lineage, we will refer to investment opportunities in services as *innovations*. On the other hand, when an investment improves the quality of the network, we will yield to common parlance and refer to it as an *upgrade*.

We represent an investment opportunity by its strength, $\alpha \in [0, 1)$, and assume that it enhances demand by αb , where b is the user base. For example, if NP_i invests in an upgrade of strength α , demand for product combinations involving network i increases by αx_i ,

$$t_{ij}^{new} = t_{ij} + \alpha x_i, \text{ for all } j \quad (2)$$

This form assumes that innovations and upgrades are subject to network effects. This is the “hard case” for deployment, because a technology does not have much value if only a portion of providers adopt it. This implies that a small group of adopters cannot gain much competitive advantage over their rivals (such competitive advantages are used to support adoption in [4][10][13]). On the other hand, the technology will achieve its maximum value if all providers adopt it, but then no firm has any competitive advantage so profits may not rise by much.

Following a standard technique from economics, let π^L be the profit a firm makes in the base case, and $\pi^H(\alpha)$ be its profit if it invests in an innovation of strength α . Define the *incentive to invest*, I , as the difference between these,

$$I(\alpha) = \pi^H(\alpha) - \pi^L \quad (3)$$

We will write this as I^{NP} for a network provider and I^{SP} for a service provider. The incentive to invest is the most a rational firm would be willing to invest to deploy the new technology [10][14].

3. ANALYSIS OF MARKET TYPES

To give some sense of scale to the values we are about to compute, it helps to have a benchmark for comparison. A good one to use is the profit and incentives faced by a monopolist, M , which provides both the network and the service. In the base case, the monopolist faces demand, $t_M = 1 - x_M$, with well-known maximum profit $\pi_M = 1/4$. With a service innovation or network upgrade, demand becomes $t_M = 1 - (1 - \alpha)x_M$. The maximum profit may be

computed as $\pi_M = \frac{1}{4(1-\alpha)}$, so the incentive to invest is,

$$I^M(\alpha) = \frac{1}{4(1-\alpha)} - \frac{1}{4} = \frac{\alpha}{4(1-\alpha)} \quad (4)$$

By comparing the performance of our three market types to this benchmark, we will effectively measure how much investment can be attributed to the presence of competition.

3.1 The Network-Gatekeeper Market

Our network-gatekeeper game proceeds as follows:

- Stage 1: Service providers select prices $\{s_j\}$
- Stage 2: Network providers select service providers
- Stage 3: Network providers choose quantities, (x_1, x_2)

If NP_i selects SP_j , its profit is $x_i(t_{ij} - s_j)$. If SP_j is selected by NPs choosing total quantity y_j , its profit is $s_j y_j$.

First, we ask what happens when a service innovation of size α is offered to SP_1 . To simplify the expression in Claim 1, we assume $\alpha \leq 3/4$. This only leaves out the very largest innovations, those that at least quadruple the potential consumer market. If the provider invests, demand for product combinations that include SP_1 increases to

$$t_{i1} = 1 - \sum y_j + \alpha y_1 \quad (5)$$

Claim 1: For $\alpha \leq 3/4$, the incentive to invest in service innovations in the network-gatekeeper market is

$$I^{SP}(\alpha) = \frac{15\alpha - 16\alpha^2}{54(1-\alpha)^2}, \quad (6)$$

which is never more than 10% greater than $I^M(\alpha)$

The full computation involves an extensive case-by-case analysis; a proof sketch is provided in the appendix. The service-gatekeeper market does not provide substantially more incentive for service innovations than the monopolist benchmark. In fact, the monopolist actually enjoys greater incentive for large α .

Next, we ask what happens when a network upgrade of size α is offered to NP_1 . If the provider invests, demand for product combinations that include NP_1 increases to

$$t_{1j} = 1 - \sum x_i + \alpha x_1 \quad (7)$$

Claim 2: The incentive to invest in network upgrades in the network-gatekeeper market is,

$$I^{NP}(\alpha) = \begin{cases} \frac{(15-16\alpha)\alpha}{9(3-4\alpha)^2}, & \alpha < 1/2 \\ \frac{5+4\alpha}{36(1-\alpha)}, & \alpha \geq 1/2 \end{cases} \quad (8)$$

The proof is given in the appendix. This incentive is plotted as the bottom line in Figure 3, scaled by the monopolist benchmark. For small innovations, a monopolist enjoys higher incentives to invest, since it earns the maximum possible profit from the higher technology. Meanwhile, the firm in our network gatekeeper does not have enough market share to gain much value from the upgrade. As α increases, though, its incentive to upgrade rises quickly as it steals market share from its rival. There is a discontinuity at $\alpha = 1/2$, as the rival network provider is driven completely from the market. Above this level, the graph appears to drop, but that is an artifact of normalizing by the monopolist benchmark. In fact, the absolute incentive continues to increase for stronger upgrades.

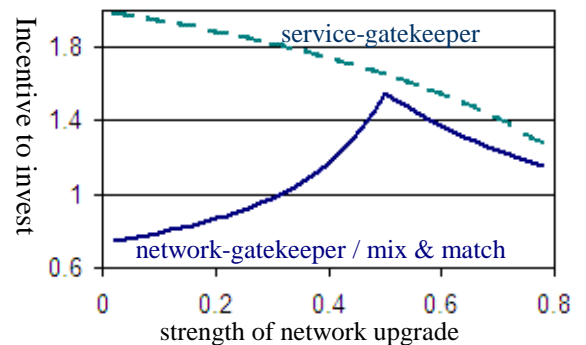


Figure 3: Incentive to upgrade network across three market types, relative to monopolist benchmark

3.2 The Mix-and-Match Market

To model the mix-and-match market, we apply our new theory of Cournot competition, in which firms select *quantity restrictions*. Details may be found in the appendix.

Stage 1: Each NP_i and each SP_j selects quantity restrictions \hat{x}_i and \hat{y}_j , respectively.

Claim 3: In the mix-and-match market, the incentive to invest in a service innovation is given by,

$$I^{SP}(\alpha) = \begin{cases} \frac{4\alpha}{(3-\alpha)^2}, & \alpha < 1/2 \\ \frac{4}{25(1-\alpha)}, & \alpha \geq 1/2 \end{cases} \quad (9)$$

The proof is given in the appendix. This incentive is plotted as the middle line in Figure 4. For nearly all values of α , the mix-and-match market yields a higher incentive to invest than the network-gatekeeper market. The difference is the most dramatic for small innovations, where the mix-and-match market offers a 60 percent improvement.

Intuitively, this market is better at aligning actors with technologies. In the network-gatekeeper market, a service must be sold to multiple network providers, after which it essentially has to compete against itself. The mix-and-match market largely avoids this situation, clearing the way for services to compete against each other.

Our model reveals a tension between short-term efficiency and long-term development. The mix-and-match market allows a service provider to profitably invest in a wider range of innovations than the network-gatekeeper market. Of course, there are technologies that are valuable or inexpensive enough that they would be deployed in either market. For those cases, it can be shown that the mix-and-match market results in lower social welfare and lower consumer surplus than the network-gatekeeper market.

Claim 4: The incentive to invest in a network upgrade in the mix-and-match market is the same as in the network-gatekeeper market,

$$I^{NP}(\alpha) = \begin{cases} \frac{(15-16\alpha)\alpha}{9(3-4\alpha)^2}, & \alpha < 1/2 \\ \frac{5+4\alpha}{36(1-\alpha)}, & \alpha \geq 1/2 \end{cases} \quad (10)$$

The proof follows a similar procedure to that of Claim 2, and we omit it. The mix-and-match market maintains the Cournot competition between network providers, so their incentives to upgrade do not change in our model. There is one important difference between the two markets, however, that our model cannot detect. This comes into play when there are multiple service classes.

Intuitively, different types of services may vary in terms of how valuable they are and the size of their potential user base. In a network-gatekeeper market, network providers

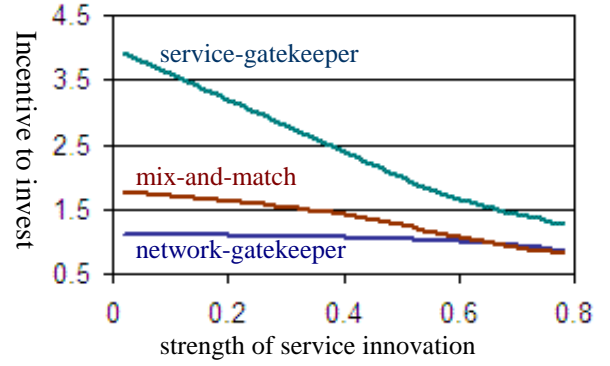


Figure 4: Incentive for service innovation across three market types, relative to monopolist benchmark
can tailor the supply of each service independently to maximize profits. For example, a network operator may try to sell more telephone service than IP service. In a mix-and-match market, however, a network provider can only choose a single quantity of the network input. The relative amounts of each service are determined by the competition between service providers.

Unable to discriminate, network providers are likely to face lower profits in the mix-and-match market. The incentive to upgrade the network, though, may increase or decrease, depending on how an upgrade affects each service.

3.3 The Service-Gatekeeper Market

Our service-gatekeeper market game proceeds as follows:

Stage 1: Network providers select prices (r_1, r_2)

Stage 2: Service providers select network providers

Stage 3: Service providers choose quantities, \mathbf{y} .

If SP_j selects NP_i , it earns profit $y_j(t_{ij} - r_i)$. If NP_i is selected by SPs with total quantity x_i , it earns profit $r_i x_i$.

Claim 5: The incentive to invest in a service innovation in the service-gatekeeper market is,

$$I_{SP}(\alpha) = \begin{cases} \alpha, & \alpha < 1/2 \\ \frac{1}{4(1-\alpha)}, & \alpha \geq 1/2 \end{cases} \quad (11)$$

The proof follows a similar structure to that of Claim 3 and we omit it. This incentive is plotted as the top line in Figure 4. For all α , the service-gatekeeper market provides substantially more incentive to innovate than the other two – more than twice as much for small values of α .

Moreover, unlike the mix-and-match market, this extra incentive to innovate does not come at the expense of welfare or consumer surplus. Because the network providers do not act to elevate prices, it can be shown that both measures are higher in the service-gatekeeper market.

Intuitively, the service-gatekeeper market makes competition between the network providers very strict. As a result, they do not elevate prices or restrict the amount of network available to consumers. A service provider with an

innovation can therefore sell it to the maximum possible consumer population. Like the mix-and-match market, the service-gatekeeper market allows services to compete directly against each other. In addition, though, it awards all of the reward for a service innovation to the service providers, rather than sharing some of it with the network providers. These factors make the service-gatekeeper market the clear winner in terms of service innovation.

The success of this market depends, in part, on the separation between the network and service markets. If the NPs are vertically integrated with service providers, they may drive out third-party SPs and turn that market into a duopoly. This suggests that a service-gatekeeper market is best paired with regulation to limit vertical integration.

One might worry, of course, that this extra service innovation comes at the expense of network upgrades. Indeed, this is a common theme in today’s network policy debates. The next claim shows that this is not the case.

Claim 6: The incentive to invest in a network upgrade in the service-gatekeeper market is

$$I^{NP}(\alpha) = -1 + 1/\sqrt{1-\alpha} \quad (12)$$

A proof sketch is given in the appendix. This incentive is plotted as the top line in Figure 3. Surprisingly, for all α , the service-gatekeeper market offers more incentive for network upgrades – almost three times as much for small α .

The service gatekeeper market offers the most incentive to invest in network upgrades, even though network providers earn zero profit in the base case. Because the service-gatekeeper market makes network competition much stricter, consumers will flock to the provider that can offer a better deal, offering a strong reward for investment.

It is a feature of our service-gatekeeper model that only one network provider can earn a profit at a time. This is because we have assumed only one class of service. We can imagine that with multiple classes, service providers might arrange themselves into different networks, some seeking a low price, perhaps, while others seek the highest possible performance. In fact, we may speculate that over time, providers may try to distinguish themselves by tailoring their networks to particular classes of service. By diversifying in this way, network providers can all earn a positive profit, while also enabling a broader range of services on the network.

4. DISCUSSION

If next-generation architectures are to move from a community of researchers to global deployment, their behavior will increasingly be governed by economic forces. In this study, we have argued that the direction of money flow has a decisive effect on innovation and upgrade in a virtualized architecture. If we do not make market structure a part of our design, there is a risk that ISPs will transplant the network-gatekeeper market from today’s internet onto the virtualized architecture. We have seen that this market

supports by far the least service innovation, so the end result might be a new network that replicates the stagnation in today’s network layer.

By contrast, we have identified two alternate market types that make better choices. A mix-and-match market can increase incentives for service innovation by up to 60%. That figure rises to over 300% for a service-gatekeeper market. This market also dramatically enhances investments to upgrade infrastructure, and improves welfare and consumer surplus. Unfortunately, network providers may oppose these alternate markets, as they may expect less profit, at least in the base case. The success of these markets may thus depend on involvement from regulators.

Though we are proposing substantial changes to the way networks and services are financed, we believe that implementing a virtualized architecture with a mix-and-match or service-gatekeeper market is within reach. Existing proposals for next-generation architectures can be adapted to support these market types. The Cabernet architecture, in particular, features a connectivity layer that abstracts multiple network providers underneath, and allows multiple services to coexist on top [17]. It thus fulfills the role of the virtualization interface in Figure 1. To support a service-gatekeeper market, the connectivity layer would further have to transmit price and traffic information between network providers and service providers.

One of the benefits of the service-gatekeeper market is that diverse service providers can select the best networks for their specific requirements. We therefore expect services to thrive when there are as many choices for the underlying network as possible. This suggests a synergy between a service-gatekeeper market and multi-homing, which allows network devices to connect to more than one access network. Some mobile phones already do this today, switching seamlessly between cellular and wifi networks.

Multi-homing would enable the benefits of a service-gatekeeper market to extend throughout the network, all the way to the end-user. Without multi-homing, service providers cannot select a user’s first hop, which really becomes the user’s choice. Even so, the service-gatekeeper market continues to enhance incentives for service innovation – at worst, it just becomes equivalent to a mix-and-match market. As the amount of multi-homing on the network increases, however, service providers start to select among access networks, and the access market becomes stricter, yielding greater incentives for both service innovation and network upgrade.

5. ACKNOWLEDGEMENTS

We thank Deirdre Mulligan, Jennifer Rexford, Pamela Samuelson, Jeff Ubois, David Marvit, Pedro Pereira, Jens Grossklags, and Yuri Takhteyev for their comments and insight. This work is supported in part by the National Science Foundation under ITR award ANI-0331659 and by Fujitsu Labs of America.

6. REFERENCES

- [1] 100x100 Clean Slate Design Project. <http://100x100network.org/>
- [2] Afegan, M. & Wroclawski, J. On the Benefits and Feasibility of Incentive Based Routing Infrastructure. *ACM Workshop on Practice and Theory of Incentives in Networked Systems (PINS)* (2004).
- [3] Computer Science and Telecommunications Board, National Research Council. *Looking Over the Fence at Networks: A Neighbor's View of Networking Research*. The National Academies Press. 2001.
- [4] Estan, C., Akella, A., & Banerjee, S. Achieving Good End-to-End Service Using Bill-Pay. *HotNets* (2006).
- [5] Feamster, N., Gao, L., & Rexford, J. How to lease the Internet in your spare time. *SIGCOMM Comput. Commun. Rev.*, p. 61-64, Jan 2007
- [6] GENI: Global Environment for Network Innovations. <http://www.geni.net/>.
- [7] Jin, Y., Sen, S., Guerin, R., Hosanagar, K., & Zhang, Z.-L. Dynamics of Competition between Incumbent and Emerging Network Technologies. *ACM NetEcon* 2008.
- [8] Joseph, D., Shetty, N., Chuang, J., & Stoica, I. Modeling the Adoption of new Network Architectures. *CoNEXT* 2007.
- [9] Lakshminarayanan, K., Stoica, I., & Shenker, S. Routing as a Service. UC Berkeley EECS Technical Report No. UCB/CSD-04-1327, 2004.
- [10] Laskowski, P. & Chuang, J. Network Monitors and Contracting Systems: Competition and Innovation. *ACM SIGCOMM* (2006).
- [11] Laskowski, P. & Chuang, J. A Leap of Faith? From Large-Scale Testbed to the Global Internet. *TPRC* 2009.
- [12] Peterson, L., Shenker, S., & Turner, J. Overcoming the Internet Impasse through Virtualization. *HotNets* (2004).
- [13] Ratnasamy, S., Shenker, S., & McCanne, S. Towards an Evolvable Internet Architecture. *ACM SIGCOMM* (2005).
- [14] Scotchmer, S. *Innovation and Incentives*, Cambridge, MA: MIT Press, 2004.
- [15] Stanford Clean Slate Internet Design Program. <http://cleanslate.stanford.edu/>
- [16] Turner, J.S. & Taylor, D.E. Diversifying the Internet. *IEEE GLOBECOM* (2005).
- [17] Zhu, Y., Zhang-Shen, R., Rangarajan, S., & Rexford, J. Cabernet: Connectivity Architecture for Better Network Services. *Proc. Workshop on Re-architecting the Internet*, (2008).

7. APPENDIX

A new theory of Cournot competition

We begin with just a single type of provider to illustrate our strategy. In the classic linear Cournot game, a set of n firms select quantities, $\{x_i\}_{1 \leq i \leq n}$, and prices are determined by a demand function, $t = N - \sum x_i$. Here, N represents the maximum population of buyers. We will focus, however, not on how much of the market the firms supply, but rather on how much of the market they decide to leave unsupplied. Define the quantity restriction for firm i , \hat{x}_i , to be how many fewer customers the firm supplies than it would have if all customers were divided equally: $\hat{x}_i = N/n - x_i$. This form is chosen so that the total quantity restriction is the number of unsupplied consumers:

$$\sum \hat{x}_i = N - \sum x_i. \quad (13)$$

Because there is a one-to-one correspondence between quantities and quantity restrictions, our game is fully equivalent when firms choose quantity restrictions instead of quantities.

If this change of decision variables strikes the reader as somewhat artificial, rest assured that it may be considered temporary. Quantity restrictions will give us a natural way to join two Cournot games together to form a two-sided market. In the end, though, we will see that each side faces a regular Cournot game in Nash equilibrium, and the game may be rephrased in terms of quantities.

Moving on to our network setting, assume a duopoly of network providers, $\{NP_1, NP_2\}$, and a set of m service providers, $\{SP_1, \dots, SP_m\}$. Let x_i be the quantity of network i , and y_j be the quantity of service j . Furthermore, let $\mathbf{x} = (x_1, x_2)$ and $\mathbf{y} = (y_1, \dots, y_m)$ be the vectors of all such quantities. We now assume a demand function $\mathbf{t} = \{t_{ij}\}$, where $t_{ij}(\mathbf{x}, \mathbf{y})$ is the price of the combination of network i and service j as a function of quantities. Let N be the maximum population of buyers (technically, the maximum total quantity for which total profit calculated through \mathbf{t} may be non-negative).

We now extend Cournot competition to a two-sided market as follows: Each network provider, NP_i , and service provider, SP_j , simultaneously chooses a quantity restriction, labeled \hat{x}_i and \hat{y}_j , respectively. We will let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_n)$ and $\hat{\mathbf{y}} = (\hat{y}_1, \dots, \hat{y}_m)$ be the vectors of all quantity restrictions. We then compute quantities, $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{y} = (y_1, \dots, y_m)$ as,

$$x_i = \frac{N - \sum \hat{y}_j}{n} - \hat{x}_i, \quad y_j = \frac{N - \sum \hat{x}_i}{m} - \hat{y}_j \quad (14)$$

This form ensures, as before, that the total quantity restriction is the total number of unsupplied consumers,

$$\sum \hat{x}_i + \sum \hat{y}_j = N - \sum x_i = N - \sum y_j. \quad (15)$$

We will assume that our demand function is separable into a network component and a service component, as,

$$t_{ij} = u_i(\mathbf{x}) + v_j(\mathbf{y}), \quad (16)$$

for some functions u_i and v_j (Note that this separation will not be unique). This assumption means that there are no complementarities between network technologies and service technologies. We will need it to ensure that prices remain continuous in supply.

Let the price for network provider NP i be $r_i(\hat{\mathbf{x}}, \hat{\mathbf{y}})$, and the price for service provider SP j be $s_j(\hat{\mathbf{x}}, \hat{\mathbf{y}})$. In classic Cournot competition, there is a unique set of prices such that demand is fulfilled and consumers are indifferent between purchasing all products. Unfortunately, the same will not hold in our two-sided market, because given any set of prices that fulfills demand, any amount of money can be subtracted from the network providers' prices and added to the service providers' prices without affecting the final prices seen by consumers. We will therefore need more assumptions to identify the most natural set of prices. Our approach will use the following four:

1: Demand is fulfilled, $t_{ij}(\mathbf{x}, \mathbf{y}) = r_i(\mathbf{x}, \mathbf{y}) + s_j(\mathbf{x}, \mathbf{y})$.

2: Prices on one side depend only on the total restriction on the other side, not individual actions. That is, $r_i = r_i(\hat{x}_1, \dots, \hat{x}_n, \sum \hat{y}_j)$, and $s_j = s_j(\sum \hat{x}_i, \hat{y}_1, \dots, \hat{y}_m)$.

3: When the domain is restricted to a fixed supply, $(\mathbf{x}^*, \mathbf{y}^*)$, prices are linear in $\hat{\mathbf{x}}, \hat{\mathbf{y}}$.

4: If either side of the market does not restrict quantity at all, the maximum total profit for that side is zero. That is, for any $\hat{\mathbf{y}}$, $\max_{(\hat{x}_1, \dots, \hat{x}_n) | \sum \hat{x}_i = 0} \sum x_i r_i = 0$, and for any $\hat{\mathbf{x}}$, $\max_{(\hat{y}_1, \dots, \hat{y}_m) | \sum \hat{y}_j = 0} \sum y_j s_j = 0$.

Theorem 1: There exists a unique set of price functions, $\{r_i\}$ and $\{s_j\}$, such that assumptions 1-4 are fulfilled.

Proof: First, assume that the price functions, $\{r_i\}$ and $\{s_j\}$, fulfill assumption 1. Then for any two network providers, $i \neq i'$, and an arbitrary service provider, j , we may write $r_i - r_{i'} = t_{ij} - s_j - t_{i'j} + s_j = u_i + v_j - u_{i'} - v_j = u_i - u_{i'}$, so r_i must take the form, $r_i = u_i + c$, where $c(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is a real-valued function of the decision variables. Then $s_j = t_{ij} - r_j = u_i + v_j - u_i - c = v_j - c$. Conversely, given any function c , the price functions, $r_i = u_i + c$, $s_j = v_j - c$ will fulfill assumption 1, since $r_i + s_j = t_{ij}$. We therefore need only show that there exists a unique c , such that the resulting price functions fulfill assumptions 2,3, and 4.

Next, assume that a function $c(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ and the resulting price functions fulfill assumption 2. We know changes to $\hat{\mathbf{y}}$ that preserve $\sum \hat{y}_j$ do not affect $r_i = u_i(\mathbf{x}) + c$. Therefore, such changes must also leave c unchanged. Similarly, changes to $\hat{\mathbf{x}}$ that do not affect $\sum \hat{x}_i$ cannot affect c , so we may write $c = c(\sum \hat{x}_i, \sum \hat{y}_j)$. Conversely, it is easy to see that for any c of this form, the resulting functions $r_i = u_i + c(\sum \hat{x}_i, \sum \hat{y}_j)$, and $s_j = v_j - c(\sum \hat{x}_i, \sum \hat{y}_j)$ fulfill assumption 2. We therefore need only show that there exists a unique $c(\sum \hat{x}_i, \sum \hat{y}_j)$ such that the resulting price functions fulfill assumptions 3 and 4.

Next, assume that a function $c(\sum \hat{x}_i, \sum \hat{y}_j)$, and the resulting price functions fulfill assumption 3. Fix a supply vector, $(x_1^*, \dots, x_n^*, y_1^*, \dots, y_m^*)$, and note that this also fixes $e = \sum \hat{x}_i + \sum \hat{y}_j$. Since within this domain, $r_i = u_i + c$, and $s_j = v_j - c$ are linear, and u_i and v_j are fixed, c must be linear, and we may write the relation,

$$c(\sum \hat{x}_i, \sum \hat{y}_j) = \frac{c(e, 0) \sum \hat{x}_i + c(0, e) \sum \hat{y}_j}{e}. \quad (17)$$

Thus, c is completely determined by its intercepts. Conversely, given any intercepts, $c(\sum \hat{x}_i, 0)$ and $c(0, \sum \hat{y}_j)$, it is easy to see that the function formed by linearly extending it in the manner of (17) results in price functions that fulfill assumption 3. We therefore need only show that there exist unique intercept values for c , such that the resulting price functions fulfill assumption 4.

Finally, for any $\sum \hat{y}_j$, with $\sum \hat{x}_i = 0$, the total profit to the network providers is,

$$\sum x_i r_i = \sum x_i (u_i + c(0, \sum \hat{y}_j)) = \sum x_i u_i + nc(0, \sum \hat{y}_j) \quad (18)$$

with a maximum of $\max \sum x_i u_i + nc(0, \sum \hat{y}_j)$. Note that this maximum is zero for a unique value of $c(0, \sum \hat{y}_j)$. Similarly, for any $\sum \hat{x}_i$, with $\sum \hat{y}_j = 0$, there is unique value of $c(\sum \hat{x}_i, 0)$ such that the maximum profit to the service providers is zero. These are the unique intercept values such that the resulting price functions fulfill assumption 4, which completes the proof. \square

We have constructed our two-sided Cournot game to utilize the most general demand function possible. Our only assumption was that the range of supply for which demand is positive is bounded. To further demonstrate why our

extension is a natural one, we may now restrict to the simpler case we will actually use. Let demand t_{ij} be linear, and suppose that there is a particular product combination that is always “best”: For any fixed total supply, selling this product by itself attains the maximum possible demand. We will call this the *single best product assumption*. Without loss of generality, we assume that this demand-maximizing product comes from NP_1 and SP_1 .

We may (uniquely) choose a separation of our demand function, $t_{ij}(\mathbf{x}, \mathbf{y}) = u_i^{NP}(\mathbf{x}) + v_j^{SP}(\mathbf{y})$ so the second term is zero whenever the demand-maximizing product is the only one being sold: $v_1^{SP}(y_1, 0, \dots, 0) = 0$ for all y_1 . u_i^{NP} can be interpreted as the maximum consumer demand for a given set of network quantities, \mathbf{x} , achieved by matching all networks with service 1, $u_i^{NP} = t_{i1}(x_1, \dots, x_n, \sum x_k, 0, \dots, 0)$. We let the slope of demand for the demand-maximizing product be $-\beta$, so $u_1^{NP}(x_1, 0, \dots, 0) = \beta(N - x_1)$.

Lemma 7: Under linear demand and the single best product assumption, network provider prices are $r_i = u_i^{NP}(x_1, \dots, x_n) - \beta \sum \hat{y}_j$.

Proof: From the proof of the previous theorem, we know that $r_i = u_i + c$, where c is linearly determined by its intercepts according to (17). Furthermore, we know that $c(0, \sum \hat{y}_j)$ is the unique value such that the maximum profit to the network providers over $\sum \hat{x}_i = 0$ is zero. By assumption, this maximum profit is attained by selling network 1 by itself,

$$\begin{aligned} \max_{\sum \hat{x}_i = 0} x_i r_i &= u_1^{NP}(N - \sum \hat{y}_j, 0, \dots, 0) + c(0, \sum \hat{y}_j) \\ &= \beta \sum \hat{y}_j + c(0, \sum \hat{y}_j) \end{aligned} \quad (19)$$

Setting this to zero yields $c(0, \sum \hat{y}_j) = -\beta \sum \hat{y}_j$.

Similarly, $c(\sum \hat{x}_i, 0)$ is the unique value such that the maximum profit to the service providers over $\sum \hat{y}_j = 0$ is zero. Again, this maximum profit is attained by selling service 1 by itself,

$$\begin{aligned} \max_{\sum \hat{y}_j = 0} y_j s_j &= v_1^{SP}(N - \sum \hat{x}_i, 0, \dots, 0) + c(\sum \hat{x}_i, 0) \\ &= 0 + c(\sum \hat{x}_i, 0) \end{aligned} \quad (20)$$

Setting this to zero, we have $c(\sum \hat{x}_i, 0) = 0$.

Extending c according to (17) yields, $c(\sum \hat{x}_i, \sum \hat{y}_j) = -\beta \sum \hat{y}_j$, so, $r_i = u_i^{NP} + c = u_i^{NP} - \beta \sum \hat{y}_j$, as required. \square

Alternately, let $u_i^{SP} = u_i^{NP} - \beta(N - \sum x_i)$ and $v_j^{SP} = v_j^{NP} + \beta(N - \sum y_j)$ to obtain a new separation of demand, $t_{ij}(\mathbf{x}, \mathbf{y}) = u_i^{SP}(\mathbf{x}) + v_j^{SP}(\mathbf{y})$, such that

$u_i^{SP}(x_1, 0, \dots, 0) = 0$, and $v_j^{SP}(\mathbf{y}) = t_{1j}(\sum y_j, 0, \dots, 0, \mathbf{y})$ is the maximum consumer demand for service provider quantities \mathbf{y} . This yields an analogous expression for service provider prices:

Lemma 8: Under linear demand and the single best product assumption, service provider prices are $s_j = v_j^{SP}(y_1, \dots, y_m) - \beta \sum \hat{x}_i$.

The proof is symmetric to that of the previous lemma. Taken together, these lemmas suggest that the two sides of the market interact by imposing a straightforward price on each other. Since the two sides are in a vertical arrangement, this is exactly the type of relationship we naturally expect. A further example will clarify this point.

Example 1: If there is exactly one network provider and one service provider with linear demand, the two-sided Cournot game is isomorphic to a price competition game, in which NP_1 and SP_1 simultaneously select prices r_1 and s_1 , respectively, and demand is given by the inverse demand function, $t_{1,1}^{-1}(r_1 + s_1)$.

Proof: Define bijection $\phi: \mathbb{R} \rightarrow \mathbb{R}$, $\phi(\hat{x}) = \beta \hat{x}$ which takes quantity restrictions to prices. Then for strategy point (\hat{x}, \hat{y}) , the two-sided Cournot game yields prices $r_i = u_i^{NP}(x_i) - \beta \sum \hat{y}_j = \beta(\hat{x} + \hat{y}) - \beta \hat{y} = \beta \hat{x}$, $s_1 = v_1^{SP}(y_1) - \beta \sum \hat{x}_i = \beta(\hat{x} + \hat{y}) - \beta \hat{x} = \beta \hat{y}$. And since demand is fulfilled, we know $t_{1,1}(M - \hat{x} - \hat{y}) = r_1 + s_1 = \beta(\hat{x} + \hat{y})$. Meanwhile, our bijection yields the same set of prices: $\phi(\hat{x}) = \beta \hat{x}$, $\phi(\hat{y}) = \beta \hat{y}$, and demand in the price competition game is computed as $t_{1,1}^{-1}[\beta \hat{x} + \beta \hat{y}] = N - \hat{x} - \hat{y}$. Since both prices and quantity are the same as in the two-sided Cournot game, payoffs are also unchanged. Therefore, ϕ is an isomorphism. \square

This example shows that with monopolies on both sides of the market, firms may simply select the prices for their inputs, which are then added together. In particular, the combined product will be double-marginalized in Nash equilibrium. This behavior is exactly what we would expect for monopolists in a vertical arrangement, and further validates our method for joining Cournot games into a two-sided market.

Proof sketch of Claim 1: Suppose that SP_1 selects price s_1 , and consider the ensuing subgame. Write H for an NP 's action of selecting SP_1 , and L for the action of selecting another service provider. Let π_{ab} , $a, b \in \{H, L\}$ be an NP 's profit in the subgame if it selects action a, and the other NP selects action b. In general, it can be shown that an NP 's profit always decreases when the other NP upgrades to the new technology, so $\pi_{HH} < \pi_{HL}$.

Suppose that LL is a Nash equilibrium in this subgame. Then $\pi_{LL} \geq \pi_{HL}$, which implies $\pi_{LL} > \pi_{HH}$. Recall that our solution concept requires the NPs to choose the equilibrium with the highest total profit. Thus, even if HH is another Nash equilibrium, LL will be chosen. Hence, if SP_1 is to sell any service at all, LL cannot be an equilibrium.

In subgame LL, the network providers receive prices, $t_{1j} = t_{2j} = 1 - x_1 - x_2$. This is a classic Cournot competition, with the well-known equilibrium,

$$x_1 = x_2 = 1/3, \quad (21)$$

yielding profits $\pi_{LL} = 1/9$.

If NP_1 plays H and NP_2 plays L, NP profits are,

$$\begin{aligned} \pi_{NP_1} &= x_1(t_{11} - s_1) = x_1(1 - (1 - \alpha)x_1 - x_2 - s_1), \\ \pi_{NP_2} &= x_2 t_{2j} = x_2(1 - x_1 - x_2) \end{aligned}, \quad (22)$$

with first-order conditions,

$$\begin{aligned} 0 &= -2(1 - \alpha)x_1 + 1 - x_2 - s_1, \\ 0 &= -2x_2 + 1 - x_1 \end{aligned}, \quad (23)$$

which can be solved to give,

$$x_1 = \frac{1 - 2s_1}{3 - 4\alpha}, x_2 = \frac{1 - 2\alpha + s_1}{3 - 4\alpha}, \quad (24)$$

yielding profits,

$$\pi_{HL} = (1 - \alpha) \left(\frac{1 - 2s_1}{3 - 4\alpha} \right)^2, \pi_{LH} = \left(\frac{1 - 2\alpha + s_1}{3 - 4\alpha} \right)^2 \quad (25)$$

This holds as long as $x_2 > 0$, which can be confirmed for the range of α we are investigating. For SP_1 to sell anything, we must have $\pi_{HL} \geq \pi_{LL}$, or

$$s_1 \leq \frac{1}{2} - \frac{3 - 4\alpha}{6\sqrt{1 - \alpha}} \quad (26)$$

In subgame HH, NP profits are,

$$\begin{aligned} \pi_{NP_1} &= x_1(t_{11} - s_1) = x_1(1 - (1 - \alpha)(x_1 + x_2) - s_1) \\ \pi_{NP_2} &= x_2(t_{21} - s_1) = x_2(1 - (1 - \alpha)(x_1 + x_2) - s_1) \end{aligned}, \quad (27)$$

With equilibrium,

$$x_1 = x_2 = \frac{1 - s_1}{3(1 - \alpha)} \quad (28)$$

Yielding profits,

$$\pi_{HH} = \frac{(1 - s_1)^2}{9(1 - \alpha)} \quad (29)$$

We may check that (26) implies $\pi_{HH} > \pi_{LH}$, so as long as SP_1 sets its price low enough to attract one NP, the other will follow. We may also check that setting s_1 below this maximum value only lowers SP_1 's profit, so it is, in fact, the equilibrium value. SP_1 's profit is then

$$\pi_{SP_1} = \frac{15\alpha - 16\alpha^2}{54(1 - \alpha)^2}. \quad (30)$$

Since service providers earn zero profit in the base case, this is also SP_1 's incentive to invest. \square

Proof of Claim 2: In the base case, the service providers are identical and in price competition, so they will compete prices down to zero. The network providers then receive prices, $t_{1j} = t_{2j} = 1 - x_1 - x_2$. This is a classic Cournot competition, with the well-known equilibrium,

$$x_1 = x_2 = 1/3, \quad (31)$$

and equilibrium profits for the SPs and NPs are,

$$\pi_{SP_j} = 0, \pi_{NP_1} = \pi_{NP_2} = 1/9. \quad (32)$$

After a network upgrade, the service providers remain identical, and will compete prices down to zero, $s_j = 0$ for all j . Network provider profits are then,

$$\begin{aligned} \pi_{NP_1} &= x_1(1 - (1 - \alpha)x_1 - x_2) \\ \pi_{NP_2} &= x_2(1 - x_1 - x_2) \end{aligned}, \quad (33)$$

With first-order conditions,

$$x_1 = \frac{1 - x_2}{2(1 - \alpha)}, x_2 = \max\left(\frac{1 - x_1}{2}, 0\right). \quad (34)$$

These can be solved for the equilibrium,

$\alpha < 1/2$	$\begin{aligned} x_1 &= \frac{1}{3 - 4\alpha}, x_2 = \frac{1 - 2\alpha}{3 - 4\alpha} \\ \pi_{NP_1} &= \frac{1 - \alpha}{(3 - 4\alpha)^2}, \pi_{NP_2} = \frac{(1 - 2\alpha)^2}{(3 - 4\alpha)^2} \end{aligned}$
$\alpha \geq 1/2$	$x_1 = \frac{1}{2(1 - \alpha)}, x_2 = 0, \pi_{NP_1} = \frac{1}{4(1 - \alpha)}, \pi_{NP_2} = 0$

Subtracting NP_1 's profit in the base case ($1/9$) from π_{NP_1} yields the required incentive to upgrade the network. \square

Proof of Claim 3: For a service innovation of strength α , we may compute,

$$\begin{aligned} u_i^{NP} &= 1 - (1 - \alpha) \sum x_k \\ v_1^{SP} &= 1 + \alpha y_1 - \sum y_k \\ v_j^{SP} &= 1 - \sum y_k, j \neq 1 \end{aligned} \quad (35)$$

Service providers with the base technology receive prices,

$$s_j = v_j^{SP}(\mathbf{y}) - (1 - \alpha) \sum \hat{x}_i = 1 - \sum y_j - (1 - \alpha) \sum \hat{x}_i \quad (36)$$

If $y_1 < 1 - (1 - \alpha) \sum \hat{x}_i$, the supply of the base service technology will rise during free entry until its price is zero:

$$\sum_{j \neq 1} y_j = 1 - y_1 - (1 - \alpha) \sum \hat{x}_i \quad (37)$$

The price of service 1 will then be,

$$s_1 = 1 + \alpha y_1 - \sum y_j - (1 - \alpha) \sum \hat{x}_i = \alpha y_1, \quad (38)$$

which is strictly increasing in y_1 . Therefore, y_1 can never fall in this range in equilibrium.

If $y_1 \geq 1 - (1 - \alpha) \sum \hat{x}_i$, then the base service technology will receive a negative price for any positive supply, and all base technology service providers will exit the market. SP_1 will then receive price and profit,

$$\begin{aligned} s_1 &= v_1^{SP} - (1 - \alpha) \sum \hat{x}_i = 1 - (1 - \alpha)(y_1 + \sum \hat{x}_i) \\ \pi_{SP_1} &= y_1 s_1 = y_1 (1 - (1 - \alpha)(y_1 + \sum \hat{x}_i)) \end{aligned} \quad (39)$$

This attains its maximum at $y_1 = \frac{1 - (1 - \alpha) \sum \hat{x}_i}{2(1 - \alpha)}$.

For, $\alpha \geq 1/2$ this falls within the permissible range of quantities, $y_1 \geq 1 - (1 - \alpha) \sum \hat{x}_i$, and is the equilibrium value of y_1 . For $\alpha < 1/2$, π_{SP_1} decreases throughout the permissible range, and so SP_1 will set quantity just high enough to drive all other service providers from the market,

$$y_1 = \begin{cases} 1 - (1 - \alpha) \sum \hat{x}_i, & \alpha < 1/2 \\ \frac{1 - (1 - \alpha) \sum \hat{x}_i}{2(1 - \alpha)}, & \alpha \geq 1/2 \end{cases} \quad (40)$$

Next, network provider profits are,

$$\begin{aligned} \pi_{NP_i} &= x_i r_i = x_i (u_i^{NP}(\mathbf{x}) - (1 - \alpha) \sum \hat{y}_j) \\ &= x_i (1 - (1 - \alpha) \sum x_k - (1 - \alpha) \sum \hat{y}_j) \end{aligned} \quad (41)$$

which gives two first-order conditions,

$$1 - (1 - \alpha)x_i - (1 - \alpha) \left(\sum x_k + \sum \hat{y}_j \right) = 0, \quad i \in \{1, 2\} \quad (42)$$

These can be solved to give,

$$x_1 = x_2 = \frac{1 - (1 - \alpha) \sum \hat{y}_j}{3(1 - \alpha)} \quad (43)$$

Combining (40) and (43), we may derive SP_1 's profit,

$$\pi_{SP_1} = \begin{cases} \frac{4\alpha}{(3 - \alpha)^2}, & \alpha < 1/2 \\ \frac{4}{25(1 - \alpha)}, & \alpha \geq 1/2 \end{cases} \quad (44)$$

Since SP_1 earns zero profit in the base case, its profit in this equilibrium is also its incentive to invest. \square

Proof sketch of Claim 6: We need to introduce a technical assumption to prove this claim. When $1 > r_1 > r_2$, the service providers will have two possible equilibria, one in which they all select network 1, and one in which they all select network 2. Because free entry drives service provider profits to zero in either configuration, the SPs are unable to distinguish between these equilibria. We will assume that firms choose the equilibrium that a finite number of SPs, m , would prefer.

For NP prices r_1, r_2 , SPs receive prices,

$$t_{1j} = 1 - (1 - \alpha)x_1 - x_2 - r_1, \quad t_{2j} = 1 - x_1 - x_2 - r_2 \quad (45)$$

It can be checked that both networks cannot be supplied in equilibrium, or an SP will want to switch from network 2 to network 1. Total SP profit if all select network 1 may be

computed as $\sum \pi_{SP_j} = \frac{(1 - r_1)^2}{(1 - \alpha)(m + 1)^2}$, and if all select

network 2, $\sum \pi_{SP_j} = (1 - r_2)^2 / (m + 1)^2$. Price competition

yields the equilibrium, $r_1 = 1 - \sqrt{1 - \alpha}$, $r_2 = 0$. Free entry drives the final price down to $t_{1j} = r_1$, so NP_1 makes profit,

$r_1(1 - (1 - \alpha)r_1) = -1 + 1/\sqrt{1 - \alpha}$, as required. \square