

Spillovers and Network Neutrality

Christiaan Hogendorn*

January 2010

* Christiaan Hogendorn is an Associate Professor of Economics at Wesleyan University, Middletown, CT. His research focuses on market structure and competition in deregulated infrastructure industries. His current projects study network neutrality and other types of “open systems” regulation, and vertical dis-integration in several types of infrastructure industries. He was formerly a Member of Technical Staff at Bell Laboratories and a visiting scholar at the Columbia Institute for Tele-Information (CITI). The author gratefully acknowledges financial support for this work provided by Google.

Table of Contents

<u>1 INTRODUCTION</u>	<u>1</u>
<u>2 ECONOMIC APPROACHES TO EFFICIENCY</u>	<u>3</u>
2.1 Efficiency – Marginal and Inframarginal	4
2.2 Spillovers – Externalities, Complementarities, and Appropriability	5
<u>3 GENERAL PURPOSE TECHNOLOGY</u>	<u>7</u>
<u>4 NETWORK EFFECTS</u>	<u>9</u>
4.1 Direct Network Effects	10
4.2 Indirect Network Effects	11
4.3 Cross-Platform Indirect Network Effects	12
4.4 Network Neutrality and Network Effects	12
4.5 Estimates of Network Effects	13
<u>5 INNOVATION</u>	<u>14</u>
<u>6 CONCLUSION</u>	<u>17</u>
<u>APPENDIX</u>	<u>18</u>
<u>REFERENCES</u>	<u>21</u>

1 Introduction

The Internet's architecture was designed to be fundamentally neutral among different types of content, applications, and devices. Last-mile facilities-based broadband Internet access service providers (hereafter, "ISPs") provide end-users with access to the Internet, but they are expected to route all traffic in a non-discriminatory manner. They do not charge Internet content or application providers to reach end-users, and they are expected to route traffic without regard to what that traffic contains, who it is from, or where it is going.

This paper discusses an important underlying economic justification - namely, positive spillovers - for "network neutrality" regulation that aims to prevent ISPs from engaging in discriminatory practices. As in many regulatory contexts, an important reason for intervention in the market is that private decisions are incompletely reflected in private costs and benefits. Rather, there are *spillover* costs and benefits that do not accrue to those making the decisions, and thus the decisions may not be optimal from society's point of view. This paper argues that ISPs engaging in discriminatory practices would be a significant, discrete change in the economic configuration of the Internet. As such, the loss of positive spillovers would be much greater than, for example, a change in pricing of existing Internet services.

Suppose an ISP offered content providers the ability to have their traffic prioritized over other providers' traffic. Or suppose the ISP preferred one service or one set of applications over others or chose to block or slow a service or set of services. Clearly these would be large changes from current practice, but the question is why they should be a public policy concern. The reason is that the ISP would choose its preferred services or its prices for service classes based on the *static, private* valuations of those services.

That is, the ISP will base its preferred services and fees on what it can *privately appropriate from* the content provider's *current* (i.e. *static*) revenue

from its content (which could be current advertising revenue or current subscription or product prices to consumers).

There are many standard economic results that suggest such preferences and pricing will be efficient, but they require that the static private values be consistent with *dynamic, public* values. That is, the preferences and prices set by the ISP would have to take into account the *future (dynamic)* potential of the content, and they would have to include any *public* values beyond the pure private revenue of the content provider.

Why would there be a divergence? First, the Internet is a *general purpose technology (GPT)*, which means that it is an input into a wide range of uses across the economy. This means that the applications of content providers often have public values in excess of their private value.

Second, the Internet, as the name suggests, is a network of networks, which means that *network effects* are rampant throughout all Internet-associated products. These take time to grow and they accrue to users beyond the content provider, so the dynamic, public value is likely to be higher than the static, private value.

Third, the Internet is an *innovation-spawning technology*, so that almost all of its content is rapidly changing and developing, adding new value. This means that the static value of content may be much less than the dynamic value, since successful development will cause a future increase in value.

All of these effects are present to a small degree with an ordinary product, for example an industrial part. But they are greatly heightened in the context of the Internet by these three features. As I explain below, the full value to society of an open Internet and maintaining network neutrality is likely to be considerably higher than the value that an ISP would consider when setting prices and policies that would degrade the quality of some sources of content relative to others.

I also discuss the wide body of literature that examines these features in a number of areas. While a precise accounting of Internet spillovers would be difficult and is beyond the scope of this paper, the literature around various technologies can provide helpful indicators. This survey of the literature includes the following findings:

- GPTs are a crucial factor in economic growth. The spillovers from GPTs, and from information technology (IT) in particular, are very large and affect entire economies. A recent study by economists at the Institute for Economic Research in Munich recognizes that broadband Internet access's benefits are rooted in its general purpose nature, and finds that a 10 percentage-point increase in broadband penetration raises a country's annual per-capita growth by 0.9-1.5 percentage points (Czernich et al. 2009).
- Several research papers focusing on network effects within just a single network estimate that consumers value a 10% increase in the installed base of compatible users of a network the same as a 5% decrease in price. Thus, if discrimination were to 'balkanize' the Internet and *reduce* the size of many network communities using various applications by similar amounts, the costs in terms of lost surplus from network externalities could be very large.
- In several studies on R&D investment, estimated innovation spillover returns to society can be 2-3 times higher than private returns.

2 Economic Approaches to Efficiency

In this section we briefly discuss some standard terminology and approaches to whether market outcomes produce the highest possible social welfare. The discussion is illustrated using a supply-and-demand diagram in the Appendix.

Consumer surplus refers to the amount of welfare that goes to consumers of a product. It is the difference between the maximum price a consumer would be willing to pay and the actual price he pays in the marketplace. For free content, consumer surplus is the value, translated into monetary terms, that the consumer derives from viewing or reading the content.

Producer surplus is the difference between the price of a product and the variable cost of production. For a great many Internet-related products, the variable cost is close to zero, since duplication of digital content is essentially costless. Even more tangible items like bandwidth and computer parts typically have quite low variable costs. Thus, most of the revenue of Internet-related firms is producer surplus. This has to be weighed against the fixed costs of development of content or products, which may be quite high.

2.1 *Efficiency - Marginal and Inframarginal*

Standard economic theory predicts that all agents, both producers and consumers, will engage in production or consumption up to the point where the marginal benefit (the surplus from selling or consuming one more unit) is equal to the marginal cost (the incremental variable cost of producing or purchasing one more unit). This means that all possible consumer and producer surplus from a market is achieved.

In addition to this standard marginal result, there is also the surplus that accrues to *inframarginal* producers and consumers. For these agents, the market price is far above or below what they need to give them the incentive to produce or consume a marginal unit. Therefore, as long as there are no other impediments to the market, these agents will certainly be in the market and achieve their surpluses. Although the marginal agents determine the price in a competitive market, most of the surplus, and thus most of the social welfare, accrues to the *inframarginal* agents.

2.2 *Spillovers - Externalities, Complementarities, and Appropriability*

The case for government policy-making that affects markets often rests on *spillovers*, generally some type of direct effect of economic actions that is not mediated through the price system. Because spillovers are so central to justifying policy, this section discusses the nature of Internet spillovers from an economic theory perspective. The following sections then turn to three general categories of specific examples of spillovers.

An *externality* usually refers to an exception to the marginal efficiency result. It is a case where, on the margin, too much or too little of an activity is done in a private market because the price of the good signals only the *private marginal benefit* or *private marginal cost* and not some additional social benefit or cost. (This is the approach pioneered by Pigou (1924) that is taught in most economic textbooks.) The concept of externality is usually taken to imply that an agent does indeed have *control* over the activity, but that “the effect produced is not a deliberate creation but an *unintended* or *incidental* by-product of some otherwise legitimate activity.” (Michan 1971, pg. 2)

In the Internet context, externalities are often positive: for example, the contributors to Wikipedia create positive externalities to other Internet users by providing useful information, but they are not given monetary compensation for their efforts. In some cases it is possible to internalize an externality by finding a way to charge for a service that would otherwise be available for free. The most common method on the Web is to sell advertising; a second, less-common method is to charge a subscription fee and protect the website with a password. Usually when an externality is internalized, it is then simply referred to as an effect (see Leibowitz and Margolis (1994) for discussion of this in the network context).

In most economic literature, an externality only creates an inefficient market outcome if it meets two conditions. First, it must be possible to find an

alternative pricing or advertising scheme that fixes the externality without prohibitively high transactions costs. Second, the pricing must actually change behavior on the margin relating to the externality. (Papanderou 1994) For example, computer networks may increase employee productivity at a firm, but a system for ISPs to charge for these effects would have prohibitively high transactions costs. Also, given limits in individual employees' ability to discern how the network affects productivity, it is not clear that such a system could be behavior affecting.

Nonetheless, there is a strand of economic thinking on this type of spillover effect that is important for understanding policies like network neutrality. This thinking goes back to Dupuit (1844) who was looking at *indivisible* decisions such as whether or not to build a bridge. He noted that while bridge tolls might be efficient on the margin, the large inframarginal surplus created by a bridge might make it socially desirable even if the tolls did not cover the cost.

One important implication of these *inframarginal externalities*, or more simply, *complementarities*, is that certain configurations of economic activity may be more desirable for an industry or a nation than others. This logic has been used both with respect to individual industries and whole economies of developing countries - for early and more recent reviews see Scitovsky (1954) and Azariadis and Stachurski (2005).

The other area where complementarities are important is innovation. Arrow (1962) emphasized that innovators often create large surplus (both on the margin and inframarginally) that is not easily *appropriable*. Since the surplus accrues to those other than the innovator, there is not enough incentive to create the innovation in the first place. More recently, Lipsey, Carlaw, and Bekar (hereinafter "LCB") emphasized that the majority of spillovers caused by GPTs are not marginal positive externalities but instead what they term *technological complementarities*. The important difference between externalities and complementarities is that externalities can create inefficient prices, whereas

complementarities only matter when considering discrete policy changes or outright prohibitions of certain activities. If, for example, an ISP decided to disable e-mail service, then all the complementarities that come from people using e-mail would be lost, and this would be a cost of the change in policy.

With this background, we are now equipped with concepts to think about the effects of discriminatory practices by ISPs. In addition to the private effects on ISPs, there would be various types of marginal externalities and inframarginal complementarities. These need to be accounted for in policymaking. In the following sections, we identify three main reasons why these effects may be large for the Internet: the fact that the Internet is a general purpose technology, the presence of network effects, and the importance of innovation.

3 General Purpose Technology

The Internet is used in a huge number of different ways throughout the economy. This feature is shared with other key technologies like electricity, railroads, and the internal combustion engine. Such technologies have been given the name *General Purpose Technologies* (Bresnahan and Trajtenberg 1995) and they have profound effects on the whole economy due to their pervasive uses. A recent book *Economic Transformations: General Purpose Technologies and Long Term Economic Growth* (LCB 2005) discusses the nature of GPTs and their effects.

LCB define a GPT as “a single generic technology, recognizable as such over its whole lifetime, that initially has much scope for improvement and eventually comes to be widely used, to have many uses, and to have many spillover effects.” (98) We deal with the “scope for improvement” in section 5 on Innovation. In this section we concern ourselves with the “widely used” issue, and the resultant spillover effects.

The spillovers from GPTs, and from information technology (IT) in particular, are very large and affect entire economies. Jorgenson and Stiroh (1999) estimate that one sixth of the United States' productivity growth from 1990-96 was attributable to IT. Czernich et al. (2009) find that an increase of 10% in broadband penetration increases annual GDP growth by 0.9-1.5 percentage points. Jorgenson et al. (2008) show that U.S. productivity growth in the early 2000s was based on a wide variety of industries adopting new forms of IT in production. Indeed, most research on economic growth and GPTs suggest that economies need GPTs in order to grow (LCB, Jovanovic and Rousseau 2005, Frischmann and Lemley 2006)

The important point here is that the range of uses is so large, and the effects so diffuse, that it is very difficult to appropriate all of the benefits - see OECD (2009) for a long but still partial listing of these benefits. The wide range of a GPT may not result in an externality on the margin, but it certainly creates large inframarginal complementarities. In this respect, spillovers from the Internet, and from other GPTs, are very different from spillovers from special purpose shared facilities like a swimming pool (Frischmann and van Schewick 2007). For these special purpose facilities, the users directly appropriate almost all the benefits, for example, by personally enjoying swimming. In contrast, general purpose facilities create spillovers that impact others - for example, consider a new application for efficient voice communication, which someone decides to use for distance learning. The business that created the tool may receive a benefit, in the form of monetary compensation, and the consumers receive a benefit from this activity. In addition, there is a benefit to society as a whole from a more edified population. Frischmann and van Schewick call these infrastructure effects, and note that "... infrastructure effects do not necessarily increase users' willingness to pay for access to the infrastructure resource, and therefore, cannot be appropriated by the network owner through its pricing of the infrastructure good" (20).

Now consider the providers of consumers' Internet access. By charging consumers a subscription fee, they receive compensation for use of the Internet and capture some value from the consumer. But they do not capture the full value of consumers' use of this communication tool -- neither the full value in terms of the consumers themselves, and particularly not the spillovers that come from the consumers putting their Internet access to productive uses, and communicating or transacting with someone else. "If there is consumer surplus in the second transaction—and there always is—that consumer surplus is external to the original transaction, because neither the original buyer nor the original seller can capture it." (Frischmann and Lemley 2006, pg. 117).

Since Internet access causes such large spillovers, one might ask whether ISPs would do everything they could to stimulate the spillovers in hopes of capturing at least some of the surplus to themselves - Farrell and Weiser (2003) call this internalizing complementary externalities or ICE ("complementary externalities" are another term for spillovers). They go on to show that unfortunately, there are several reasons that an ISP's self-interest may not be compatible with the public interest, and ICE may fail. The reasons an ISP might block or degrade certain services in violation of ICE include (i) it helps the ISP charge different prices to different customers in order to increase profits, (ii) it may make it harder for other firms to enter the market, either as ISPs or as application providers, (iii) there may be bargaining problems that prevent internalization, (iv) firms may not fully realize the benefits of all externalities, and (v) if Internet applications have other spillover benefits, it may make it easier to capture (but in the process reduce) those spillovers.

4 Network Effects

For many technologies, and especially for the Internet, the value of the network as a whole rises as the number of users increases. Furthermore, this is

also true of many applications on the Internet, which are themselves networks that connect users – whether for e-commerce, socializing, information retrieval, and so forth. Katz and Shapiro (1985) identified two ways this could happen, either a direct effect from user to user or an indirect effect operating through complementary goods.

4.1 Direct Network Effects

When the effect is direct, the consumers get value directly from communicating with each other. The most famous early example of this was the telephone network, and many of the basic Internet protocols also involve direct communication – e-mail and chat for example. There has been a resurgence of Internet applications designed to create direct network effects, often under the label “Web 2.0.” These allow users to directly share profiles and pictures (Facebook), current activities (Twitter), favorite news stories (Digg), speak to each other (Voice over IP, e.g., Skype,) and so on.

Regardless of the specific application, all direct networks share the feature that users desire to communicate with one another, and the more possible pairs or groups of users, the more value the entire system has. Thus, when a new user joins such a network, all of the other users receive a positive spillover. Sometimes the network can encourage joining with introductory pricing or other incentives, in which case the spillover may be internalized into a network effect, while in other cases it may remain an externality.

The reverse of a network effect can be called a nonuser negative network effect, so that when a user either leaves or does not join a network, there is a loss of what would otherwise be the positive network effect. (Nagler 2009) Any impediment to a user joining the network therefore has nonuser negative effect on all users. For example, suppose an ISP introduces a traffic management practice that forces users to pay extra for a voice application like Skype. Then some users will not feel the benefits of joining are sufficient to justify paying the

fee, and they will not join. But this does not create a loss for that consumer alone. All of the other users of the network will also receive less value, even if they choose to pay the fee, because the direct network effect has been reduced.

4.2 *Indirect Network Effects*

With an indirect network effect, users also value the system more as the total number of users rises, but through a different mechanism. There are also content providers who choose to offer various products on the system, and these are what users want to use. The more users, the more content providers will provide the products, and therefore the more value to the users.

Indirect network effects occur at several levels in the Internet. At the broadest level, the more Internet users there are, the more Internet-related applications will be developed, but this is more of an innovation effect, which we discuss below. But at a narrower level, many of the applications developed for the Internet are themselves platforms with their own user and content provider communities. For example, standards like html and xml have this property - the more users with web browsers or other applications that use the latest versions of these standards, the more incentive for programmers to develop new applications that use them. The same is true for some websites - for example, the more users of Facebook, the more incentive for developers to create new Facebook applications.

Again, each user creates a spillover effect beyond his or her own benefit from joining the system and using various products. If a user is prevented or dissuaded from being part of the system due to certain types of traffic being prohibited or delayed, then the effect spills over into less new product development, hurting both the business of the developers and the value of other users.

4.3 *Cross-Platform Indirect Network Effects*

A cross-platform indirect network effect is a type of indirect network effect that happens through the (partial) compatibility of different platforms. (Hogendorn and Yuen, 2009) For example, suppose another Internet user joins a Web-based financial information service like Pageonce. This will increase the incentive for Pageonce to improve all of its product offerings, including the one available on the Apple iPhone, since the protocols used to provide the Web-based service are partially compatible with those used to provide the iPhone service. The improved iPhone version of Pageonce generates additional value for iPhone users. Thus, we have an indirect network effect that creates a positive spillover for iPhone users even though it originated with a web-based user.

Because the Internet is at the heart of a constellation of information and communication technologies, these cross-platform indirect network effects can be quite important. They also suggest a further cost to traffic management that differentially affects certain types of services. If these services are partially compatible across platforms, then there could be a nonuser negative effect outside of the Internet on other ICT platforms.

4.4 *Network Neutrality and Network Effects*

If ISPs were to prioritize or degrade service for certain applications, certain websites, etc., there could be a reduction in both types of network effects. The question is whether an individual ISP would have incentives that do not align with social incentives. There are two reasons this might happen.

First, network effects, particularly direct ones between users, have the same appropriability problems related to the many uses of GPTs. Again, there are small benefits distributed broadly across a large number of users. If an ISP cannot appropriate these surpluses to itself, it may block or degrade

applications that generate them without regard to the inframarginal surplus being lost.

Second, there is a large literature suggesting that competing platforms will not privately choose to be compatible with one another as much as socially optimal. (Farrell and Saloner 1985; Chou and Shy 1993; Crémer, Rey, and Tirole 2000, etc.) Blocking websites and applications, degrading them, and offering multiple classes or types of services are all forms of incompatibility that are presently not available or very insignificant. If these measures were increased, then both types of network effects could be inefficiently lost.

4.5 *Estimates of Network Effects*

No one knows the total value of network effects emanating from the Internet. Since the Internet is a network of networks, and since many of the applications that run on the Internet are themselves networks, the total number of networks is huge. There is a well-developed body of research that estimates the value of network effects on just one network. A summary of this work, which is surveyed in Birke (2009), is presented in Table 1 in the appendix.

The first thing to note is that consumers value network effects a lot. Several of the papers study how much consumers value a 10% increase in the installed base of compatible users of a network. While the methods used vary, it seems that consumers typically equate this to about a 5% decrease in price of the network good. Thus, for example, if a change like network discrimination reduced the compatible installed base by, say 30%, the consumer would be harmed the same as if the networked good (say, a social network website) had raised its price by 15%.

Only one of the studies directly attempts to measure the uninternalized network externalities from compatibility. Ohashi (2003) examined the well-known example of VHS and Betamax video cassettes in the 1980s. He estimates that if, hypothetically, these had been compatible in 1986, there would have been

indirect network externalities equal to almost \$10 per US household. Multiplied across many millions of consumers, this is particularly substantial.

While these measures do not directly value the network effects of the Internet, they do suggest that the values are very large, especially when multiplied over the many networked applications that run on the Internet. If discrimination were to 'balkanize' the Internet and thus reduce the size of network communities using various applications, the costs in terms of lost surplus from network externalities could be very large.

5 Innovation

The most dramatic source of spillovers from the Internet is innovation. Firms and other types of users throughout the economy make the Internet an input into new activities, and these in turn add value for society.

At the broadest level, the Internet is used to increase efficiency and productivity in a wide array of industries. This diffusion of the Internet throughout the economy takes time, as firms produce new products and find improved ways to make existing products, using the Internet as an input. A bank, for example, can offer financial services with web-based or cell-phone-based updates. A trucking firm can use online GPS tracking to improve scheduling and delivery times. A firm like eBay may take a very old activity - running an auction house - and completely redefine it by using the Internet as an input. These are examples of extended technological complementarities of the Internet, where the Internet is the key driver, but the spillover takes place in another industry, like trucking or auctions. Diffusion like this is common with all GPTs - electricity for example took a long time to fully penetrate the economy, reorganizing everything from factory floor layouts to the hours of shopping and working. (David 1990)

Internet applications themselves are a large source of innovation, and they usually piggyback on one another. In other words, the Internet is a platform for many innovations, some of which are platforms for innovations as well. For example, consider the http on the World Wide Web itself. http is an innovation built on top of the Internet that in turn ushered in countless additional technologies - starting with websites themselves, to search engines, social networking, and so on. Some of these tools might also be platforms for follow-on innovation; for instance, people use Google Maps to create “mash-ups” that combine, for instance, maps with housing data, and developers create applications that operate on top of Facebook. In this way, innovation online is often “recursive” (Zittrain).

One form of “innovation” is less dramatic, but at least as important to economic growth. This is the simple use of existing Internet applications by new businesses and consumers. Every time a firm or consumer adopts an existing Internet technology, it faces its own unique problems and develops its own unique solutions. For example, a firm may adopt an Internet-based travel expense voucher system for its employees. Obviously many other firms have already done this, but since each firm has slightly different needs, flexible technologies are key to a successful adoption. LCB note that just because there is a “blueprint” for implementing a technology, this does not include all the “tacit knowledge” that goes with the blueprint. Thus, adoption of existing Internet technologies becomes a form of innovation, where each firm has to solve its own problems and implement the technology in its own way. The ability of a firm to accept and implement spillovers from another firm is called absorptive capacity (Cohen and Levinthal 1989), and firms that are more open to new technology and more experienced in developing it themselves are generally better at absorbing spillovers.

In absorbing an Internet application, firms face a great deal of risk. There is risk relating to the costs of the project, the benefits of the project, and also the ongoing value of the application in relation to other Internet applications. This

relation to other applications is important because most firms try to reuse their applications to perform multiple tasks; this is called technological convergence by Rosenberg (1976).

Arrow (1962) first noted the extensive appropriation problems in settings like this. The chain of innovation is so extensive that it is very unlikely that the owner of any one innovation fully appropriates all of the innovation's surplus. Indeed, society would not want him or her to do so, since the follow-on uses are often not predictable. For example, when Tim Berners-Lee invented the Web, he could never have predicted the wide array of innovations to which it would give rise. Thus, whatever the apparent, immediate impact of discriminating among applications today, the full consequences of that cannot be predicted.

As in the previous two cases, the difference between externalities and complementarities is important. In many cases, innovation may not involve an externality on the margin. Probably more common, there are discrete innovation decisions that are indivisible. Either the innovator has access to crucial inputs or access to the market, or he or she does not. If not, then potential inframarginal surplus is lost since the innovation is not created.

Discrimination by ISPs would also make it less likely that developers would use new inputs, since there would be uncertainty as to whether the input would work for all users and whether it might be arbitrarily degraded in the future.

The size of spillovers from innovation have been studied by several authors and are surveyed in Lemley and Frischmann (2006). Details are presented in Table 2 in the appendix. Generally the findings are that innovators do in fact earn high returns to the capital they invest, which is not surprising since innovation is a very risky activity. The estimated spillover returns to society as a whole are much higher, in many studies 2 or 3 times higher. The fact that these social returns are so much higher than most financial returns should not be surprising, since the financial returns to any investment only reflect the surplus

the investor can appropriate to himself or herself, whereas the returns to society reflect all the spillovers discussed above.

6 Conclusion

The Internet's benefits to society and to economic growth are enormous, and this paper has emphasized that most of these benefits come in the form of spillovers - benefits that accrue to third parties in the rest of the economy rather than directly to the ISP or its customers. Spillovers occur beyond the standard buyer-seller transaction represented by a consumer subscribing to an ISP. As with other GPTs, the key to maintaining the value of the spillovers is to maintain the open and general purpose character of the Internet connection itself. If ISPs reduce the essentially open nature of the Internet, they can cause considerable harm to society.

Appendix

Figure 1: A Market with Positive Externalities

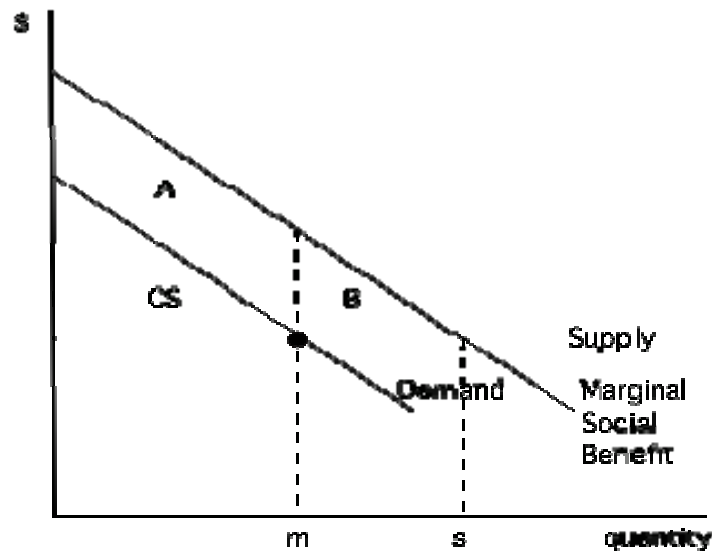


Figure 1 shows a typical market with positive externalities. As discussed in the main text, these might come from general purpose technology, network effects, or innovation. They show up in a conventional supply-and-demand diagram through the marginal social benefit curve, which lies above the demand curve and indicates that additional surplus goes to consumers other than those who actually purchase the product.

If the good is sold at a market clearing price, equilibrium price and quantity settle at the black dot and the quantity produced and consumed is m . As a result, those who purchase the good receive consumer surplus CS and in addition other members of society receive surplus A .

The conventional economics literature emphasizes the inefficiency of quantity m . In fact, the social surplus exceeds costs at that point, and the efficient quantity is s . If quantity s were produced, additional surplus B would be generated.

What we have emphasized in this paper is the surplus shown by area CS and the “inframarginal externality” shown by area A . Since these occur under a free market, they are usually not of policy concern. However, in this case, as discussed in the main text, discrimination by ISPs against certain applications could reduce those applications’ spillovers and network effects. As a result, such discrimination would reduce both CS and A . Since area A does not even accrue to customers of ISPs, it is especially vulnerable to this reduction.

Table 1: Empirical Research on Network Effects

Authors	Topic	Spillover
Pashigian, Gould (1998)	Shopping mall rents	72–87% reduction in rent for large stores in shopping malls, controlling for sales per square foot.
Rauch (1993)	Industrial park land prices	417–947% greater price appreciation in successful industrial parks compared with surrounding land.
Brynjolfsson, Kemerer (1996)	Spreadsheet software	10% increase in installed base associated with a 7.5% increase in price.
Gandal, Kende, Rob (2000)	Effect of CD titles on CD player sales	10% increase in CD titles is equivalent to a 5% price cut in CD player prices.
Madden, Coble-Neal, Dalzell (2004)	Mobile phone adoption	10% increase in subscribers associated with 1.5% to 3.3% increase in growth rate of subscriptions.
Ohashi (2003)	VHS vs. Betamax VCRs	10% increase in installed base associated with a 9.6% increase in market share of one type of VCR. 1986 value of network externalities: \$29 per US household. 1986 value of having VHS and Betamax be hypothetically compatible: \$9.73 per US household.
Gowrisankaran, Stavins (2004)	Banks adopting ACH electronic payments	10% increase in share of banks adopting ACH associated with approximately 4–9% higher probability of a bank adopting ACH.
Miller, Tucker (2007)	Electronic health records	State privacy laws associated with 25% reduction in probability of a hospital adopting electronic records.
Prieger, Hu (2006)	Video game adoption	10% increase in game titles equivalent to a 10% decrease in console price.
Clements, Ohashi (2005)	Video game adoption	10% increase in game titles equivalent to a 3.6% decrease in console price.

Table 2: Empirical Research on Innovation Spillovers

Authors	Method	Results
Mansfield (1977)	Survey firms, use results to estimate demand curve	25% private return, 56% social return
Deng (2008)	Patent citations	Spillover value of \$0.6 to \$1.2 million per citation.
Jones, Williams (1998)	Average of 6 other studies of social return to R&D.	Social return 27% within industry, 100% across whole economy.
Goto, Suzuki (1989)	Data on Japanese R&D	Return from spillovers 3.14 times larger than return within an industry.
Bernstein, Nadiri (1989)	Data on US industries	Social rate of return 1.30–2.23 times private rate of return.
Bloom, Schankerman, Van Reenen (2005)	Data on US firms	Social rate of return 3.5 times private rate of return.

References

- Arrow, K. Economic welfare and the allocation of resources for invention. Richard Nelson, ed., *The Rate and Direction of Inventive Activity* (Jan 1962), 609-625.
- Bernstein, J., and Nadiri, M. Research and development and intra-industry spillovers: An empirical application *The Review of Economic Studies* 56, 2 (Jan 1989), 249-267.
- Birke, D. The economics of networks: A survey of the empirical literature. *Journal of Economic Surveys* 23, 4, 762-793.
- Bloom, N., Schankerman, M., and Reenen, J. V. Identifying technology spillovers and product market rivalry. *CEP Discussion Paper No 675* (2005).
- Bresnahan, T. F., and Trajtenberg, M. General purpose technologies - 'engines of growth'? *Journal of Econometrics* 65 (1995), 83-108.
- Brynjolfsson, E., and Kemerer, C. Network externalities in microcomputer software: An econometric analysis of the *Management Science* (Jan 1996).
- Chou, C., and Shy, O. Partial compatibility and supporting services. *Economics letters* 41 (Jan 1993), 193-197.
- Clements, M., and Ohashi, H. Indirect network effects and the product cycle: Video games in the us, 1994-2002. *Journal of Industrial Economics* (Jan 2005).
- Cohen, W., and Levinthal, D. Innovation and learning: The two faces of r d. *The Economic Journal* 99, 397 (Sep 1989), 569-596.
- Cremer, J., Rey, P., and Tirole, J. Connectivity on the commercial internet. *Journal of Industrial Economics* (Jan 2000).
- Czernich, N., Falck, O., and Kretschmer, T. Broadband infrastructure and economic growth. *CESifo Working Paper ...* (Jan 2009).
- David, P. The dynamo and the computer: An historical perspective on the modern productivity paradox. *The American Economic Review* 80, 2 (May 1990), 355-361.
- Deng, Y. The value of knowledge spillovers in the us semiconductor industry. *International Journal of Industrial Organization* 26 (Jan 2008), 1044-1058.

- Dupuit, J. On the measurement of the utility of public works. *Annales des Ponts et Chaussées* 8 (Nov 1844).
- Farrell, J., and Saloner, G. Standardization, compatibility, and innovation. *The RAND Journal of Economics* (Jan 1985).
- Goto, A., and Suzuki, K. R & d capital, rate of return on r & d investment and spillover of r & d in *The Review of Economics and Statistics* (Jan 1989).
- Gowrisankaran, G., and Stavins, J. Network externalities and technology adoption: lessons from electronic payments. *RAND Journal of Economics* (Jan 2004).
- Hogendorn, C., and Yuen, S. Platform competition with 'must-have' components. *Journal of Industrial Economics* 57, 2 (May 2009), 294-318.
- Jones, C., and Williams, J. Measuring the social return to r&d*. *Quarterly Journal of Economics* (Jan 1998).
- Jorgenson, D., Ho, M. S., Samuels, J. D., and Stiroh, K. Industry origins of the american productivity resurgence. *Interdisciplinary Information Sciences* 14, 1 (Jan 2008), 43-59.
- Jorgenson, D., and Stiroh, K. Information technology and growth. *The American Economic Review* 89, 2 (May 1999), 109-115.
- Jovanovic, B., and Rousseau, P. L. General purpose technologies. *Handbook of Economic Growth* (Jan 2005), 1181-1224.
- Katz, M., and Shapiro, C. Network externalities, competition, and compatibility. *The American Economic Review* 75, 3 (Jun 1985), 424-440.
- Lemley, M., and FRISCHMANN, B. Spillovers. *Columbia Law Review* 100, 2 (Jan 2006).
- Liebowitz, S., and Margolis, S. Network externality: An uncommon tragedy. *The Journal of Economic Perspectives* 8, 2 (Apr 1994), 133-150.
- Lipsey, R., Carlaw, K., and Bekar, C. *Economic Transformations: General Purpose Technologies and Long-term ...* 2005.
- Madden, G., Coble-Neal, G., and Dalzell, B. A dynamic model of mobile telephony subscription incorporating a network effect. *Telecommunications Policy* (Jan 2004).

- Mansfield, E., Rapoport, J., Romeo, A., Wagner, S., and Beardsley, G. Social and private rates of return from industrial innovations*. *The Quarterly Journal of Economics* 91, 2 (May 1977), 221-240.
- Miller, A., and Tucker, C. Privacy, network effects and electronic medical record technology adoption. *Proceedings of WEIS* (Jan 2007).
- Mishan, E. The postwar literature on externalities: an interpretative essay. *Journal of Economic Literature* (Jan 1971).
- Nagler, M. Network externalities, mutuality, and compatibility. *papers.ssrn.com* (2009).
- OECD. Network developments in support of innovation and user needs. *OECD Working Party on Communication Infrastructures and Services Policy* (2009).
- Ohashi, H. The role of network effects in the us vcr market, 1978-1986. *Journal of Economics & Management Strategy* (Jan 2003).
- Papandreou, A. *Externality and Institutions*. 1998.
- Pashigian, B., and Gould, E. Internalizing externalities: the pricing of space in shopping malls. *The Journal of Law and Economics* (Jan 1998).
- Pigou, A. C. *The economics of welfare* (2nd ed.). 1924.
- Prieger, J., and Hu, W. An empirical analysis of indirect network effects in the home video game market. *papers.ssrn.com* (2006).
- Rauch, J. Does history matter only when it matters little? the case of city-industry location. *The Quarterly Journal of Economics* (Jan 1993).
- Rosenberg, N. *Perspectives on technology*. Cambridge (UK) (Jan 1976).
- Scitovsky, T. Two concepts of external economies. *The Journal of Political Economy* (Jan 1954).
- Trajtenberg, M., and Bresnahan, T. General purpose technologies: "engines of growth"? *Journal of Econometrics* (Jan 1995).
- Zittrain, J. *The Future of the Internet-And How to Stop It*. 2009.