

LAW, SOCIAL WELFARE, AND NET NEUTRALITY

Boston University School of Law
Law & Economics Paper No. 16-08

Forthcoming, Net Neutrality Symposium, Review of Industrial
Organization, 2016

February 23, 2016

Keith N. Hylton
Boston University School of Law

This paper can be downloaded without charge at:

<http://www.bu.edu/law/faculty-scholarship/working-paper-series/>

Law, Social Welfare, and Net Neutrality

Keith N. Hylton*

January 2016

(forthcoming, Net Neutrality Symposium, Review of Industrial Organization, 2016)

Abstract: The net neutrality norm generates wealth transfers from one type of internet content provider to another. In theory, these transfers might be socially desirable, and could be justified on the basis of informational externalities similar to those that have been identified to justify the fair use doctrine in copyright law. However, in practice, the conditions that justify fair use in the copyright context do not appear to hold in the settings in which the net neutrality principle operates. Moreover, the internal subsidization required by net neutrality generates a transfer from the relatively poor to the relatively rich. The potential welfare gains that might come from controlling anticompetitive abuse or government coercion through implementation of the policy can be achieved by alternative policies with less harmful consequences.

Keywords: net neutrality, vertical integration, anticompetitive exclusion, telecommunications regulation, subsidy, fair use

* William Fairfield Warren Distinguished Professor, Boston University; Professor of Law, Boston University School of Law, knhylton@bu.edu. I thank Michelle Connolly, Thomas Hazlett, David Sappington, Elliot Maxwell, Roger Noll, and Chris Yoo for very helpful comments. I have also benefited from suggestions of participants at the Tech Policy Institute's Conference on Net Neutrality, January 8, 2016, in Washington D.C.

Net neutrality is an appealing term. After all, what might be the conceptual antithesis of net neutrality: net discrimination? The word discrimination suggests an evil purpose, while neutrality suggests immediately an open, fair process. The term alone conditions one to be supportive of the concept.

Net neutrality consists, as far as I can tell, of a regulatory norm that is as simple as it is appealing. Providers of content transported through the networks of broadband internet firms cannot be charged different prices by the broadband firms for the service of transporting their content.¹ Thus, a firm that provides a highly demanded type of content that absorbs much of the capacity of the broadband network cannot be charged for the additional congestion and wear-and-tear associated with the transporting of its content.

If there were no differences between the consumers of various services – for example, if everyone consumed the same information services from the internet – net neutrality would truly be neutral in effect. It would not permit differential pricing, or discrimination, to adversely affect any providers of content. It would not differentially impact any consumer, since every consumer is identical by hypothesis.

But there are differences between consumers of internet services, and consumers demand different services, which impose dissimilar costs on broadband firms.² Consequently, net neutrality requires some consumers to subsidize the consumption of others. In this sense, net neutrality is not neutral at all: it forces A to pay for the consumption of B. Viewed from this perspective, net neutrality is a form of differential pricing.

The basic economics of net neutrality seems similar to that of toll bridges.³ Charging all users of the bridge the same amount may force some users to subsidize others. This has moral hazard and adverse selection implications. Consumer welfare declines, unless there is an efficiency case for a regulation that requires internal subsidization of one group of consumers by another group of consumers (cross-subsidization).⁴ I will consider the possible welfare cases for such internal

¹ As Becker, Carleton and Sider explain, the net neutrality principle has devolved into four specific requirements: broadband providers are “prohibited from: (1) prioritizing traffic and charging differential prices based on the priority status; (2) imposing congestion-related charges; (3) adopting business models that offer exclusive content or that establish exclusive relationships with particular content providers; and (4) charging content providers to access the Internet based on factors other than the bandwidth supplied.” (Becker, Carlton, Sider, at 498)

² On the variation in network services and costs, see Yoo (2013).

³ The similarity is that in both contexts, one observes large up front capital expenses to construct a facility, and the incremental cost of using the facility is comparatively small. The economics of bridges is a topic with an extensive literature. See, e.g., Hotelling (1938); Minasian (1979). The difficult problem of optimal supply analyzed in the earlier literature has important implications for the treatment of different types of user, and the extent to which one can determine when the pricing of bridge services inefficiently advantages one user at the expense of another user. I will oversimplify and avoid most of these difficulties here.

⁴ I use the term cross-subsidy as an equivalent to “internal subsidy”, as in Posner (1971). This is a popular, though admittedly non-rigorous use of the term. In my treatment, a cross subsidy or internal subsidy results from a regulation (such as a net neutrality requirement) when a firm is forced to change its pricing structure, as a result of the regulation, in a manner that benefits one group of consumers and disadvantages another group of consumers. The concept of cross-subsidy is difficult and would require a much more extensive treatment than offered here to deal with adequately in economic terms. A rigorous treatment of the concept of cross-subsidy is provided in Faulhaber (1975).

subsidization here, drawing on familiar analogies in the law. After that, I will consider alternative arguments for the net neutrality norm based on distributional concerns, potential anticompetitive abuse, and government coercion.

I conclude that the case for the net neutrality norm is weak. The potential efficiency justifications for the policy are speculative, and unsupported so far by the evidence at hand. The distributional consequences of the policy are undesirable. The policy effects a transfer of resources from the less advantaged to the more advantaged without any significant offsetting welfare gain. Wherever there is a speculative gain that might come from the policy, such as controlling anticompetitive abuse or government coercion, the gain can be achieved by an alternative policy with less harmful consequences.

1. The Bridge Analogy

Consider a toll bridge with two types of user: ordinary cars and heavy trucks. The difference between the two is that the heavy trucks are more costly to the bridge owner to service. They impose a greater degree of congestion on bridge traffic. Congestion reduces the flow of traffic over the bridge, thereby increasing the average cost of the service. In addition, the trucks impose a greater degree of depreciation (wear and tear) on the bridge. A profit-maximizing bridge owner would adopt a system of discriminatory Ramsey prices, charging a mark-up above user-specific marginal cost that is inversely related to the elasticity of user demand.⁵ Assuming the demand elasticity of cars is no greater than that of trucks – cars consisting of a larger percentage of leisure travelers are more likely to seek cheaper though more time-consuming routes – the owner would charge higher prices to trucks than to cars.

A regulatory authority that seeks to maximize welfare subject to a given level of profit promised to the bridge owner would also choose a system of discriminatory Ramsey prices, though not precisely the same as those chosen by the unregulated monopolist. The mark-up above user-specific marginal cost would be somewhat less, but would still vary inversely with the elasticity of user demand.⁶

Suppose, however, the regulator requires the bridge owner to charge the same price to both cars and trucks. Under the conditions assumed, cars would, in effect, finance an internal subsidy for trucks. As a result, trucks would tend to use the bridge more than if the cross-subsidy were not present, and cars would use the bridge less than if the cross-subsidy were not present. The additional use by trucks would increase congestion costs, driving up the cost of the service and increasing depreciation costs.

Charging different usage prices to the cars and the trucks would permit the bridge owner to internalize to truck owners the additional costs imposed by the trucks. This, in turn, would discourage the trucks from excessive use – that is, from imposing a marginal cost of \$1 on the bridge owner and other users when the marginal benefit to the truck owner from the particular

⁵ Here I refer to the simplest form of Ramsey pricing that results from profit maximization by a multi-product monopolist. For simplicity, I assume that the user demand functions are independent of each other. For a straightforward presentation, see Tirole (1989).

⁶ See generally Baumol and Bradford (1970).

use is only \$.50. If the charge varied with the intensity of the use, it would encourage the truck owners to consider the congestion costs and the miles of wear and tear imposed in each relevant time period. The higher charge would also have the sorting effect of inducing some truck owners to forgo using the bridge and to find some other route. Over time, charges might encourage technological innovation toward trucks that carry the same freight while imposing lower congestion and depreciation costs.

As a consequence of charging separate prices, the bridge owner would reduce congestion and depreciation costs, and be able to pass those cost savings on to consumers in the form of lower general prices (for an equivalent unit of service) for use of the bridge, and this in turn would increase the total consumption of the services offered by the bridge.

Admittedly, there are some cases where the bridge owner might choose not to charge differential prices. Perhaps the differences in service costs are minor, and the administrative costs of differential pricing exceed the efficiency gains. Alternatively, perhaps trucks provide the greatest source of demand for new bridge capacity. Foresighted bridge owners would therefore be reluctant to tax a major source of industrial capacity growth. In these cases, the bridge owner may choose not to impose differential pricing even if completely free to do so.

The bridge analogy seems to apply straightforwardly to the net neutrality problem. Net neutrality is equivalent to prohibiting the bridge owner from using differential pricing, and generates similar costs. Some providers of internet content, such as Netflix, impose extraordinary congestion costs as a result of the internal subsidy from consumers of other internet services. Hence, permitting the network owner to price differentially can and probably would enhance consumer welfare.⁷ To the extent that heavy use of the service has a depreciation effect (electrical components suffer wear and tear from use), similar costs are imposed.

2. Is There a Case for Cross-Subsidization?

The question of interest is whether there is a case for cross-subsidization? As a general matter, it is not difficult to make an argument for cross-subsidization in the presence of significant externalities, but the question is whether such a justification is applicable in the net neutrality context.

Return to the bridge story. If trucks provide a positive externality to all bridge users, then a policy of cross-subsidization might be socially optimal. Suppose, for example, the trucks provide a good that benefits other bridge users (or the consumers of those other bridge users) through a mechanism that fails to charge a price for those external benefits. Vaccines, for example, provide a beneficial externality to consumers who do not consume the vaccines. A subsidy toward the trucks carrying those vaccines might be optimal. If most of the trucks carry vaccines, a subsidy funded by cars and benefitting trucks in general might be optimal.

⁷ See, e.g., Becker, Carlton, and Sider (2010); Peitz, Martin and Schuett, Florian (2015). In my example in the text, if the privately-chosen prices are relatively close to the socially optimal Ramsey prices, then unregulated monopoly would enhance welfare relative to the equal-pricing rule.

Does this argument for cross-subsidization have any application to the internet? If the services that impose the greatest congestion costs on the internet also provide a positive externality to all consumers of internet services, then the cross-subsidization required by net neutrality might be socially optimal. The informational services that are often cited for providing important positive externalities to society are, in essence, educational. Educational services enhance the stock of human capital, and thereby help to make members of society more productive, or otherwise advance science and culture.

The fair use doctrine of copyright law, for example, imposes a subsidy funded by one type of information service for the benefit of another type of information service: from the holder of a copyright to an infringer of the copyright. Under the fair use doctrine, an agent who copies copyrighted material can avoid a finding of infringement if the use is deemed fair. Fair use is determined by several factors, such as the extent and purpose of the copying. One key factor supporting a finding of fairness is copying for educational purposes. Educational purposes can be understood broadly to include information that advances culture. Satire, for example, arguably has an educational purpose, because it is through satire that the public becomes aware of the shortcomings or flaws of certain theories or works of art. For example, Alice Randall's *The Wind Done Gone* provides an account, deemed a parody by one court, of the fictional events from Margaret Mitchell's *Gone With the Wind*, told from the eyes of slaves rather than from the eyes of slaveholders. In a society matured to the stage that its prior belief that one particular race should be given the legal power to hold another race in slavery now seems comically unjust, Randall's retelling, from an upside-down perspective, of a famous work of art embodying the now-discredited view of social order has significant educational value. By subsidizing such iconoclastic works of literature, fair use doctrine probably enhances society's welfare.⁸

Cross-subsidization of the sort required under the net neutrality principle is therefore not foreign to the common law, which for the most part strongly protects property rights, nor necessarily harmful to social welfare. The question is whether this fair-use based theory of cross-subsidization has any justificatory power when brought to the net neutrality debate.

If the providers of heavily demanded internet content were providing educational services, or informational services that generally enhance the stock of human capital, then the net neutrality norm might lead to a socially optimal regulatory regime. However, as far as I am aware, this is not the case. Netflix, for example, streams popular movies and television shows through the internet, using up an extraordinary amount of pipe capacity as it does so. Some of the shows are critically acclaimed, but I doubt that any of them could be accurately described as educational in any significant sense. The most popular Netflix program, *House of Cards*, offers a realistically cynical view of the inner workings of American politics, but such depictions of the political system have been part of western literature for a long time and probably do little to significantly enhance the stock of human capital. To be sure, Shakespeare also gave us realistically cynical portraits of the inner workings of politics, but the similarity between Shakespeare and *House of Cards* is probably equivalent to the similarity between a Rembrandt and a typical piece of

⁸ Gordon (1982).

critically-acclaimed modern art: one truly educates and significantly enhances culture while the other mainly reflects the current tastes of the median consumer.

Economides and Tag (2012) offer the most prominent externality justification for network neutrality to date, though it is different from the fair-use based theory just described. Instead of externalities across different types of consumer, the external effect may be present in the form of network externalities between consumer and content provider.⁹ Network effects may work in both directions across the platform, with more consumers generating more content, and more content generating more consumers. The broadband network provider may tend to overcharge the content provider relative to the social optimum. If the network externality suggested by Economides and Tag is sufficiently strong, implementation of the network neutrality norm might result in an optimal outcome.

Although the Economides and Tag result is rather weak – showing that under special conditions net neutrality might be socially optimal – there are still reasons to question its applicability to the real world. The assumptions required for the result are strong; including, for example, identical consumers and inelastic demand. However, much of the economic controversy surrounding network neutrality stems from the heterogeneity of consumers and the existence of some degree of demand elasticity for network services. The bridge analogy with which this paper begins assumes these features. Moreover, the empirical evidence regarding the importance of provider-to-consumer network effects is speculative at this stage,¹⁰ and the theory may not capture all of the relevant externalities (Becker, Carlton, Sider, 2010).

3. Distributional Considerations

On distributional grounds, cross-subsidization of informational services can be harmful to social welfare, at least if you adopt a Rawlsian perspective,¹¹ which disfavors wealth transfers from the materially advantaged to the relatively disadvantaged. A Rawls-inspired welfare function would incorporate the distribution of resources as a component of social welfare.¹² A more equal distribution of resources reduces the frequency of instances where individuals are born into, or somehow find themselves through no fault of their own, in impoverished households. Thus, institutions that reduce the risk of being placed into such households enhance society's welfare if members of society are sufficiently risk averse, and conversely institutions that increase this risk reduce society's welfare. This perspective suggests efficiency can be decomposed into general

⁹ Nicholas Economides and Joacim Tåg (2012).

¹⁰ The provider-to-consumer network effect, which is important for Economides and Tag, can be examined in the context of newspapers to get a sense of its plausibility. The argument is similar to saying that more newspapers leads to more newspaper readers. While it is highly plausible that more readers will tend to generate more newspapers, it is not immediately plausible that more newspapers will generate more readers. In addition, it is the *content* provided by the newspapers that might plausibly generate more readers, not the *number* of newspapers. Thus, an enhancement of quality (content) coupled with a reduction in quantity might have a stronger feedback effect on the number of consumers than a simple increase in quantity.

¹¹ Rawls (1971).

¹² See, e.g., Posner (1981, at 59). Of course, one could say that this is just an ordinary social welfare function that takes risk-aversion into account. See *id.*

efficiency (or “wealth maximization”) and distributional efficiency concerns. I will focus on the distributional efficiency concern here.

On the distributional effects of imposing cross-subsidies in information services, consider the British Broadcasting Corporation (BBC). Because of the greater educational value of its programming, the BBC generally has a stronger claim to subsidization than does Netflix, yet the BBC subsidy is regressive in distributional terms. It reduces the price of a service consumed mainly by the wealthy and compels the poor, through its government funding, to pay part of the cost of providing those services to the rich.

Similarly, to the extent the net neutrality norm requires the subsidization of Netflix, it imposes a regressive tax on society. The Netflix programs are less high-brow than those of the BBC, and as a result have a greater appeal to the average middle-income consumer, but the poorest of the poor are unlikely to be among the heavy consumers of Netflix. A recent study of the demographics of Netflix users notes that “when compared to the average U.S. adult, recent Netflix subscribers are: (a) Gender: Fairly even gender split, 49% male, 51% female; (b) Age: 37% are Millennials (18-34), which is 23% more likely than average; (c) Residential status: slightly more likely to be adults who still live with their parents, although the majority (61%) of recent subscribers are homeowners; (d) Suburban living: 50% live in the suburbs, which is slightly more than average (+9%); (e) Household income (HHI): 42% have a HHI of \$50K and under, 35% have a HHI of \$50K-\$100K, and 23% have a HHI over \$100K; (f) Education: They are slightly more likely than average (+11%) to say their highest level of education is an associate or bachelor’s degree; (g) Parental status: 24% more likely to be parents and 52% more likely to have school-aged children living with them; (h) Household: 38% more likely to have 4 or more people living in their current household.”¹³ These factors suggest that the average Netflix consumers is wealthier than the average U.S. citizen. On distributional grounds, the net neutrality norm therefore operates as a regressive tax by compelling the materially less-advantaged to support the consumption of a relatively advantaged group.

As a general rule, a norm prohibiting forced wealth transfers from the relatively disadvantaged to the relatively advantaged appears to have a stronger appeal on welfare grounds than the net neutrality norm. The policy against regressive transfers is likely to improve social welfare, when distribution of wealth is taken into account, while the net neutrality norm degrades distributional efficiency. If the net neutrality norm enhanced general efficiency – that is, enhanced society’s wealth – then it might be desirable in spite of its negative distributional effects. But the general efficiency argument for net neutrality is ambiguous, lacking empirical support currently, and unlikely to be proven valid in the end when the empirical evidence has been fully explored. Hence, the net neutrality norm at present offers an unambiguous reduction in distributional efficiency coupled with a likely negative impact on general efficiency.

The norm against regressive coerced transfers operates generally, while the net neutrality norm operates within a narrow subset of society. This suggests an important difference between the two norms, with distributional implications. The policy disfavoring regressive transfers is

¹³ <http://civicscience.com/wp-content/uploads/2015/05/CivicScience-Netflixs-New-Users-May-2015-Final.pdf>.

unlikely to serve effectively as a veil for concealing the aims of some faction interested in transferring wealth to itself, because it operates across a wide spectrum of society. A narrow faction that appeals to the antiregressiveness principle purely out of self-interest – that is, to justify efforts to expropriate wealth from others – could easily find itself losing as a result of the principle just as much as it gains. The net neutrality norm, by contrast, offers a narrowly tailored argument that, by seeming to invoke a lofty goal, obscures self-interested motives. The beneficiaries of net neutrality form a concentrated interest group, and therefore would benefit from a norm that conceals an effort to persuade government to transfer wealth in its favor. Given the net neutrality norm’s usefulness as an instrument that simultaneously facilitates and obscures expropriative activity in the political process, it should be regarded with some suspicion from a distributional efficiency perspective.

4. Vertical Integration

One reason the bridge analogy may seem incomplete as stated previously is that there may be some settings in which the ability to price differentially among bridge users might be used in an anticompetitive manner. Suppose one truck service purchases the bridge and charges high prices to rival truckers to use the bridge. This scenario is often described as the essential facilities problem in antitrust, where the classic case, *United States v. Terminal Railroad*,¹⁴ involved a consortium of railroads that purchased a bridge (actually, the rail terminal facilities connected to the bridge) that had to be used by competing railroads. In the internet network setting, a vertically integrated network owner (integrated into content) might use differential pricing to harm competitors in the content market.¹⁵

I have already established that differential pricing could encourage efficient (in a second-best sense) consumption of bridge services, and the same is true of general platforms such as a broadband network. Differential pricing might also provide the most efficient method of recoupment of the costs of constructing a platform, such as a bridge or a broadband network.¹⁶ The net neutrality norm is therefore inadvisable because of its efficiency costs. However, net neutrality might be defended still, because it prevents anticompetitive price discrimination among platform users.

There are reasons to question this argument for net neutrality. First, net neutrality goes further than necessary. The antitrust laws already exist for regulating anticompetitive conduct, and they attempt to regulate with a much finer brush than the net neutrality rule. An antitrust court would take efficiencies into account in any analysis of a complaint against a vertically-integrated platform owner on antitrust grounds. The net neutrality principle ignores efficiencies.

Second, the market provides disincentives to predatory abuse in the vertical integration setting. If there are no efficiencies resulting from vertical integration, the vertically-integrated bridge owner will be punished, to some degree, by capital markets as investors observe the risk that the vertically integrated owner might reduce the market value of the bridge in order to provide an

¹⁴ 224 U.S. 383 (1912).

¹⁵ See, e.g., Hazlett and Wright (2012).

¹⁶ On Ramsey pricing and broadband networks, see Yoo (2013, at 580).

advantage to its own content flowing over the bridge. This may seem speculative, at first, but examples have been observed. Keurig, a maker of coffee brewing devices famously associated with the small plastic cups (K-cups) used for making coffee, recently abandoned its effort to lock out the products of rival K-cup makers from compatibility with its brewing device. When Keurig attempted to switch to a new brewing device that made rival K-cups technologically incompatible, the firm's stock price dropped so quickly that it reversed course and announced that it would continue marketing a device that could accommodate rival cups. Moreover, this reversal happened despite the fact that Keurig had virtually no significant competition in the market for its brewing devices.

The Keurig experience suggests that a firm that owns a facility which rivals must access to provide service to customers may suffer a severe penalty in the capital markets for attempting to exclude rival service providers. The value of the facility is determined by the overall quality and variety of services that the facility can provide. Locking out rival service providers may reduce the value of the facility. Hence, predatory exclusion is not always a costless exercise, even in the absence of the threat of competition or antitrust penalties, and in the case of Keurig would have been permitted by the capital market to occur only if it were efficient.

Broadband firms face the same capital market constraints as the makers of coffee brewing devices. If content exclusion were an efficient strategy for broadband firms, capital markets would steer them toward it. But remarkably few instances that potentially could be described as anticompetitive content exclusion have been observed.¹⁷

Finally, direct competition with other platform providers constrains anticompetitive abuse. While high-speed broadband networks tend to face few competitors in most American markets,¹⁸ the existence of geographic markets in which some consumers have access to more than one high-speed network suggests that the threat of competitive entry exists in the broadband market. Google's fiber service exists in three cities and promises to expand to more,¹⁹ and other technological novelties are in the development phase.²⁰ Moreover, the most heavily demanded content providers, which create the greatest congestion costs, also drive a substantial share of the demand for high-speed networks – an external effect which broadband providers have strong incentives to consider. An anticompetitive move to disadvantage such content, merely to provide a short-term boost in content owned by the vertically-integrated network, would risk a much larger loss in market share in the future as competitive substitutes develop. The currently proposed merger between AT&T and DirecTV is partially premised on the argument that the combined entity will be able to take advantage of scale economies to invest in a more expansive broadband network. If these statements are at least partially true, they signal to incumbent firms

¹⁷ See Hazlett & Wright (2012, at 782).

¹⁸ Most people have only one high-speed broadband provider, see <http://consumerist.com/2014/12/18/govt-report-true-high-speed-broadband-competition-in-the-u-s-remains-largely-nonexistent/>.

¹⁹ Google Fiber gigabit service is expanding, but only in three cities now, see <http://consumerist.com/2015/10/29/shortlist-of-future-google-fiber-cities-keeps-getting-longer-three-more-cities-added/>.

²⁰ And there are new ideas for bringing high-speed internet, such as small satellites: <http://consumerist.com/2015/06/11/the-guy-who-started-tesla-wants-to-shoot-the-internet-to-you-from-space/>.

that anticompetitive abusive treatment of content providers is at best a short-term strategy with great long-term risks.

Perhaps the bigger issue in the competitive landscape for broadband is the role of Schumpeterian, or dynamic, competition. It happens that at any given moment, the market for high-speed broadband can be defined in a manner that seems to imply that most consumers are vulnerable to mistreatment by firms with monopoly power. Network quality improves through substantial capital investments, and short-run market power provides the primary incentive for such investments. On the content side, technology continually generates new types of content or methods of enhancing existing content, in turn creating more demand for higher quality broadband networks. The recent history of the market appears to one of chicken-and-egg evolution with advances in networks supporting advances in content, and advances in content creating demand for more sophisticated networks.

Yet, if every time a substantial investment in network quality occurs, regulatory agencies treat the resulting picture of high concentration, when the market is defined by selecting the highest quality broadband service, as evidence of a lack of competition and therefore a justification for price regulation, then incentives to invest will be dulled. Indeed, the factions that demand price regulation will always be able to find an empirical justification given the economic relationship between quality, investment, and market concentration. This generates the paradoxical result that as the potential for significant network quality improvements increases – resulting in more choices for consumers among different levels of service quality – the threat of burdensome regulation based on a static vision of competition increases too.

5. Government Coercion

Another reason the bridge analogy may seem incomplete is that it leaves out the problem of government coercion as a potential justification for the net neutrality norm. The government coercion problem provides net neutrality proponents with perhaps their strongest argument, but it is far from compelling in the end.

The threat of government coercion in the broadband service market arises in two settings; one potentially desirable from a social welfare perspective, and the other almost certainly undesirable. The potentially desirable form of coercion arises as a response to internet piracy. Content owners have sued internet service providers, seeking a more active effort by the network providers to block pirate websites or sellers of knock-off products.²¹ Content owners invariably argue that the internet service providers have not done enough to ensure that copyrights or trademarks are not regularly infringed through their networks. In the case of trademarks, the networks typically make available cheap knock-offs of established brands that either directly infringe the established trademark, or exhibit sufficient similarity to have the same effect (trade dress infringement).

²¹ *MGM Studios, Inc. v. Grokster, Ltd.*, 545 U.S. 913 (2005); *BMG Rights Management (US) LLC et al. v. Cox Enterprises Inc. et al.*, case number 1:14-cv-01611, in the U.S. District Court for the Eastern District of Virginia.

As a general matter, broadband networks should not be held strictly liable for copyright or trademark infringement. The networks provide a valuable social good by enhancing the flow of information in society. The positive spillover benefits provided by broadband networks probably outweigh the harms to copyright and trademark owners. Still, the remaining question is whether the owners of broadband networks have an obligation to monitor the theft of intellectual property and just how much effort such an obligation would entail on their parts.

The net neutrality norm enters here with an obvious application. Broadband network owners may prefer something like a network neutrality rule as a barrier to any claims of liability for copyright or trademark infringement. A broadband network could assert the neutrality norm as a justification for treating all content providers alike, and for reluctance to differentially treat any of its content providers. A network neutrality rule enacted into law might provide a powerful preemption defense to any broadband network sued for facilitating infringement of intellectual property.

In this setting, the net neutrality norm provides an overly broad principle that stands in the way of designing an optimal regulatory scheme. To monitor and control the infringement of intellectual property, broadband providers may have to adopt more aggressive measures than have been adopted to date. Courts applying standard negligence principles are in an ideal position to determine the proper balance between risk and precaution on the part of broadband providers. A net neutrality regulation, asserted in an effort to preempt a court from applying negligence principles, would effectively remove courts from an area of regulation in which they have traditionally operated, and in which they may be able to provide an appropriate solution to conflicting interests.

The other form of government coercion I referred to earlier, which is not socially desirable, is state-directed control of internet content. The net neutrality norm arguably provides the benefit of enabling broadband networks to cite the norm (or regulation) as a justification for refusing to carry out the censorship aims of a particular government. Here, the norm appears to be unnecessary in a society in which the government is constrained to protect free speech, and wholly ineffective in a society in which the government is not so constrained. In the United States, the First Amendment to the U.S. Constitution should form a sufficient legal basis for any broadband provider to resist efforts on the part of the U.S. government to control content, especially of a political nature, on the internet. It is not clear what a net neutrality norm could provide beyond the protection already provided by the constitutional guarantee of free speech. If anything, a net neutrality norm would likely be harmful in preventing the government from establishing, through either judicial or regulatory standards, optimal regulations to control abuses such as copyright infringement.

What if, instead of the government directing the network to constrain speech, the network itself attempts to constrain speech? For example, suppose a network announces that it will impose a special charge on content providers who intend to transmit speech the network regards as offensive to some particular person or group? The neutrality principle would arguably promote free expression in this special case by preventing the network firm from imposing discriminatory viewpoint-based charges. However, there are several reasons to doubt the desirability of the

neutrality norm even in this scenario. First, the neutrality norm goes too far by prohibiting all discriminatory charges rather than limiting its prohibition to viewpoint-based discriminatory charges. As in the scenarios considered earlier (anticompetitive abuse, preemption), the neutrality rule prohibits desirable forms of discrimination in an effort to root out potentially undesirable forms of discrimination. Second, competitive pressures (for customers, for capital) will constrain the incentives of network firms to discriminate on the basis of viewpoint. Indeed, if viewpoint discrimination were attractive to consumers, the heavy-handed internet censoring observed in China would not be met by massive efforts to circumvent it. Third, large broadband firms have valuable contracts with federal and state governments,²² and would surely put these contracts at risk by engaging in viewpoint-based discrimination of content providers. These last two considerations suggest that broadband firms will be disincentivized by the market to engage in viewpoint-based discrimination among content providers.

In a government that is not constrained (through both law and actual enforcement) to protect free speech, such as that of China, a net neutrality norm would be the feeblest of barriers against government coercion. The content providers themselves would face the threat of punishment if they were to violate the government's censorship rules. Broadband providers might or might not comply with a rule of net neutrality, and it would make no difference to the underlying problem of censorship by the government. Indeed, the net neutrality norm might facilitate the government's censorship, by preventing broadband providers from taking any actions that might counteract the government's direct control over content providers. The net neutrality norm could then be paraded ostentatiously by the censoring government as a sign of its relative enlightenment.

6. Conclusion

In theory, implementation of the net neutrality norm might be socially desirable, and could be justified on the same grounds as the fair use doctrine in copyright law. However, in practice, the conditions that justify fair use in the copyright context do not appear to hold in the settings in which the net neutrality principle operates. Moreover, the cross-subsidization required by net neutrality generates a transfer from the relatively poor to the relatively rich. For every potential social gain that might be provided by the neutrality policy, an alternative, narrower policy exists that would be at least as effective and less likely to have harmful side effects. Efficiency and equity considerations provide no support for the net neutrality norm.

²² For example, on Comcast's government business, see <http://business.comcast.com/ethernet/industry-solutions/government>.

References

- Baumol, W. J., & Bradford, D. F. (1970). Optimal Departures from Marginal Cost Pricing. *The American Economic Review*, 60(3), 265-283.
- Becker, G. S., Carlton, D. W., & Sider, H. S. (2010). Net Neutrality and Consumer Welfare. *Journal of Competition Law and Economics*, 6(3), 497-519.
- Economides, N. & Tåg, J. (2012). Net Neutrality on the Internet: A Two-sided Market Analysis. *Information Economics and Policy*, 24(2), 91-104.
- Faulhaber, G. R. (1975). Cross-Subsidization: Pricing in Public Enterprises. *The American Economic Review*, 65(5), 966-977.
- Gordon, W. J. (1982). Fair Use as Market Failure: a Structural and Economic Analysis of the Betamax Case and its Predecessors. *Columbia Law Review*, 82(8), 1600-1657.
- Hazlett, T. W., & Wright, J. D. (2012). The Law and Economics of Network Neutrality. *Indiana Law Review*, 45(3), 767-840.
- Hotelling, H. (1938). The General Welfare in Relation to Problems of Taxation and of Railway and Utility Rates. *Econometrica*, 6(3), 242-269.
- Minasian J. R. (1979). Indivisibility, Decreasing Cost, and Excess Capacity: The Bridge. *Journal of Law & Economics*, 22(2), 385-397.
- Peitz, M. & Schuett, F. (2015). Net Neutrality and Inflation of Traffic. TILEC Discussion Paper No. 2015-006. Available at SSRN: <http://ssrn.com/abstract=2573466> or <http://dx.doi.org/10.2139/ssrn.2573466>.
- Posner, R.A. (1981). *The Economics of Justice*.
- Posner, R. A. (1971). Taxation by Regulation. *The Bell Journal of Economics and Management Science*, 2(1), 22-50.
- Rawls, J. (1971). *Theory of Justice*.
- Tirole, J. (1980). *The Theory of Industrial Organization*.
- Yoo, C. S. (2013). Is There a Role for Common Carriage in an Internet-Based World? *Houston Law Review*, 51(2), 545-608.