NBER WORKING PAPER SERIES

WEAK VERSUS STRONG NET NEUTRALITY

Joshua S. Gans

Working Paper 20160
http://www.nber.org/papers/w20160

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2014

During the course of this research I worked at Microsoft Research. All views remain my own and Microsoft did not have pre-approval of the publication of this paper. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Weak Versus Strong Net Neutrality
Joshua S. Gans
NBER Working Paper No. 20160
May 2014
JEL No. D04,D42,D43,K2,L1,L12,L13

ABSTRACT

This paper provides a framework to classify and evaluate the impact of net neutrality regulations on the allocation of consumer attention and the distribution of surplus between consumers, ISPs and content providers. While the model provided largely nests other contributions in the literature, here the focus is on including direct payments from consumers to content providers. With this additional price it is demonstrated that the type of net neutrality regulation (i.e., weak versus strong net neutrality) matters for such regulations to have real effects. In addition, we provide support for the notion that strong net neutrality may stimulate content provider investment while the model concludes that there is unlikely to be any negative impact from such regulation on ISP investment. Counter to many claims, it is argued here that ISP competition may not be a substitute for net neutrality regulation in bringing about these effects

Joshua S. Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto ON M5S 3E6
and NBER
joshua.gans@gmail.com

1 Introduction

Net (or Network) Neutrality has been a concept discussed in the United States for over a decade (Wu, 2003; Lee and Wu, 2009). It relates to a principle that Internet Service Providers (ISPs) or any other network operator on the Internet should not be able to create conditions under which consumers would be induced to favour some content providers (CPs) over another.[2] At times, this has evoked the notion that CPs should not pay for access to consumers while at others it is that CPs should not face charges that are discriminatory with respect to one another or related to the quality of the experience their consumers receive. Most recently, this issue has emerged with Netflix (a streaming video supplier) paying Comcast (a large US ISP) for a higher quality route or 'fast lane' to its consumers. However, the notion of content-based price discrimination has also emerged with respect to deals with some ISPs (especially outside of the US) to allow consumers to access services such as Wikipedia[3] or Facebook without charge on mobile devices and AT&T considering offering that option too if a CP pays them directly.[4]

The debate has focussed on numerous dimensions. From the principle of neutrality expressed above to concerns that restricting the ability of ISPs to engage in price discrimination or offer different quality products might impact adversely on infrastructure investment. Finally, there is a concern that ISP discrimination options may lead to barriers to CPs themselves to invest; particularly, in new services. This may come because of a lack of affordable access, consumers shying away from lower quality services or from hold-up whereby ISPs increase charges as CPs become more successful. This is of particular concern where ISPs themselves are integrated CPs perhaps in competing non-Internet services such as cable television or media production.

This paper is designed to bring some clarity to the debate. First of all, rather than speculate on what rules might enshrine certain principles of behaviour – namely, can we have a commercial Internet whereby consumers do not internalise ISP pricing when choosing how to allocate their attention? – here I evaluate these issues using economic

[2] Where 'content providers' could include websites, applications or other services that make use of the Internet as a means of interacting and supplying consumers.
[3] http://wikimediafoundation.org/wiki/Mobile_partnerships (accessed 20 May 2014).
[4] See, https://gigaom.com/2014/01/06/att-launches-sponsored-data-inviting-content-providers-to-pay-consumers-mobile-data-bills/ (accessed 20 May 2014).

criteria. Specifically, how do various price regulations impact on consumers' actual allocation of attention and its comparison to a socially optimal allocation? And how do various regulations impact on ISP or CP investment in their activities? Thus, this paper is as much about the effectiveness of proposed regulations as it is about their expected impact.

Second, while there is an existing formal literature that asks questions similar to those just specified, that literature has several difficulties that have made applying it to this debate difficult; let alone comparing alternative claims made in those debates. For one, the starting point for that literature is to assume that CPs do not have a direct monetary relationship with consumers. With respect to recent streaming video issues, this is an unreasonable assumption. Thus, here I start with the situation where CPs set prices to consumers directly.[5] In addition, that existing literature assumes that consumers do not pay ISPs for services that involve different charges depending upon what type of content they access. While this is an accurate depiction of the initial period of the US commercial Internet, it is not an accurate depiction elsewhere or for the mobile Internet. While ISPs have not engaged in content-specific charging to consumers per se, they have imposed download caps that, for bandwidth intensive content, impact on consumer choices as to what content to consume. Consequently, these prices now and perhaps into the future will impact on the consumers allocation of attention. For that reason, it is explicitly considered here.

The final aspect where the existing formal literature falls short is that it is complex in its analysis. As is standard in industrial organisation in academic economics, there is a push towards generality of the analysis but, despite abstracting away from consumer charging of the forms described above, the formal models have not been conclusive in their recommendations and, indeed, they have identified a plethora of effects that have made building intuition difficult. This is not to say those models are unimportant or unrealistic. Far from it. My contention here is that a simpler, less realistic model can demonstrate some first-order insights regarding the net neutrality debate. However, the analysis also confirms that such regulations have benefits and costs. The

[5] For example, in Cheng et al (2009), Choi (2010), Economides and Tag (2007), Choi and Kim (2010) and Reggiani and Valetti (2011) there is no direct CP charging to consumers and CP earn all revenue from other means such as advertising.

aim here is to more clearly identify these rather than make a definitive policy recommendation per se.

To this end, this paper provides an unashamedly simple model with a single ISP, two CPs and a consumer. There is no nuance on demand from the consumer (they can allocate their attention to one CP or another). There is no detailed analysis of network priority rules. And there is no close examination of alternative pricing structures and download caps – everything is comprised of a set of simple transfers for services. However, in the process, I believe a number of useful insights can be gathered.

First, in this model with a flow of payments between different entities, we can distinguish between *weak net neutrality* (whereby content-based price discrimination is outlawed in single relationships) and *strong net neutrality* (whereby it is outlawed altogether). In so doing, I demonstrate that only strong net neutrality leads to real effects; in particular, over the allocation of profits between the ISP and content providers. Moreover, it does not impact at all on the consumer allocation of attention that remains socially optimal regardless of regulatory regime. This is because the price between CPs and consumers mediates this and other payments ultimately do not distort it. Consequently, *neutrality regulation is itself neutral in important ways*.

Second, where there is no price between consumers and content providers, certain forms of neutrality regulation can lead to real effects on the allocation of consumer attention. In this case, strong net neutrality can cause consumers to choose a socially sub-optimal allocation. This is because the socially optimal allocation includes both the consumer's value on the service and the CPs earns from advertising and the like. In the absence of sufficient pricing flexibility, that latter component does not factor into consumer choice, even indirectly, and hence, they may, under some circumstances, choose a sub-optimal allocation. Moreover, it is shown that this reduces surplus for all parties.

Third, when net neutrality regulation is effective (i.e., under strong net neutrality), then there is a transfer for surplus from the ISP to content providers and, moreover, the surplus retained by content providers is related to their marginal value in the market. Hence, we can expect that dynamic efficiency, in terms of content provider investment, will be enhanced by strong net neutrality rules. At the same time, it is demonstrated that

net neutrality regulation (of any form) *does not* impact on the incentives of ISPs to improve network quality. Consequently, this suggests that claims that price flexibility is required to ensure such investment takes place does not stand up in this model. Finally, it is demonstrated that 'fast lanes' as a means of counteracting strong net neutrality will only do so imperfectly. There is no first-order prediction that CPs will be unable to capture the surplus they would if a fast lane option did not exist but, at the same time, ISPs will be able to capture the increment from providing a higher quality service. That said, if ISPs were to be compelled to offer fast lanes at the same price as slow lanes, there is an increased likelihood that CPs will appropriate their full marginal value if the value they provide to consumers is complementary with network quality (as it may be for streaming video).

Finally, I examine what happens when there is ISP competition. Not surprisingly, such competition allows consumers to appropriate more surplus. However, more subtly, because consumers operate at the edge of the Internet and are, in effect, monopoly suppliers of their own access, net neutrality regulations simply have no effect. The consumer's chosen ISP becomes their agent in propagating that monopoly power and so CPs appropriate no surplus. This suggests that claims that competition amongst ISP would eliminate the importance of the issue are more nuanced than is generally thought.

2 Baseline Model

2.1 Set-up

Consider a simple model with one ISP, one representative consumer and two content providers (CP1 and CP2). The content providers give value to consumers of $v_1 > v_2$. Suppose also that the content providers earn other revenue so that $v_1 + r_1 > v_2 + r_2$. Here the content providers may earn revenue from advertisers rather than Internet consumers. We will assume content providers have no other costs related to the usage of their service but we will consider, below, that they have fixed costs of entering or being in business.[6] Let's suppose that the consumer chooses content from one provider only;

[6] With additional notion, it would be relatively straightforward to incorporate usage related content provider costs (e.g., royalty payments to publishers and movie studios) into the analysis.

perhaps because of limited attention in a given time period. We will assume the ISP costs associated with bringing content to consumers are c regardless of which content is chosen. Below we will consider what happens if there is a service that causes this cost to vary.

2.2 Definitions

The ISP charges consumers a connection fee that can, potentially, be contingent on the content chosen (that is, (t_1,t_2) where t_i is the content associated with CPi).[7]

The ISP can also charge content providers an access or transit price. Once again, this can be different from different content providers (that is, (a_1,a_2)). Finally, content providers can directly charge consumers for the use of their services (p_1,p_2). It is assumed throughout that these prices are not regulated. The flow of payments is depicted in Figure 1.

Figure 1: Flow of Payments

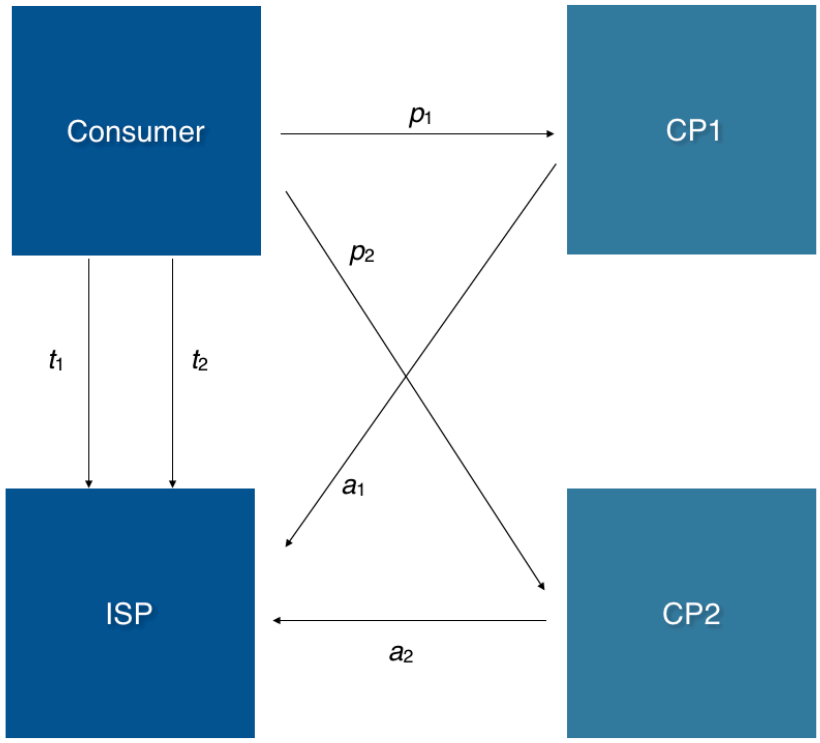

[7] While there are many ways this might occur, a common practice is for ISP to charge on the basis of usage. Thus, if CP1 is a video streaming service, $t_1 > t_2$ may represent those additional usage fees. The specification here allows for this and any other basis of charging consumers in a manner that causes them to perceive a different charge associated with different content.

Given this specification of prices, we can define the following.

Definition (Strong net neutrality). *The ISP cannot discriminate in its pricing to either CPs or consumers. That is, they must set* $t_1 = t_2 = t$ *and* $a_1 = a_2 = a$.

Definition (Weak content provider net neutrality). *The ISP cannot discriminate in its price to CPs; that is,* $a_1 = a_2 = a$.

Definition (Weak consumer net neutrality). *The ISP cannot discriminate in its price to consumers; that is,* $t_1 = t_2 = t$.

Definition (No regulation). *The ISP can charge content-contingent prices to consumers* (t_1, t_2) *and to content providers* (a_1, a_2).

To date, as mentioned in the introduction, analyses of net neutrality have not made a distinction between strong and weak net neutrality.[8] These analyses, therefore, examine moves from weak to strong net neutrality rather than moves from no regulation to a form of net neutrality.

Here it will be shown that the distinction matters. In particular, we will evaluate the market outcomes under different regulations and see whether the consumer ends up choosing the CP with the highest value. We will also examine the division of rents between CPs and the ISP under each scenario. In each case, we assume that in the market, the ISP sets (a_1, a_2) and (t_1, t_2) prior to the CPs setting (p_1, p_2).

2.3 No regulation

Under no regulation, the ISP is free to engage in content-based price discrimination as it sees fit. The following proposition characterises the equilibrium outcome.

Proposition 1. *Under no regulation, the consumer selects the socially optimal content provider (i.e., CP1) with content providers and the consumer earning zero surplus and the ISP earning* $v_1 + r_1 - c$.

All proofs are in the appendix. Intuitively, the content providers compete for the attention of the consumer. At that stage, the consumer will be concerned with the price charged by the content providers and also the different charges the ISP might set depending upon

[8] For example, in Cheng et al (2009), Choi (2010), Economides and Tag (2007), Choi and Kim (2010) and Reggiani and Valetti (2011) there is no direct CP charging to consumers and CP earn all revenue from other means such as advertising.

which provider is chosen. Similarly, the content providers' prices will, in turn, depend on the charges they must pay the ISP. The proof of the proposition demonstrates that by setting $t_i + a_i = v_i + r_i$ the content providers will price in such a way that the consumer ends up choosing the content provider with the highest $v_i + r_i$ and, in equilibrium, will that value will be passed through to the ISP.

2.4 Weak net neutrality

Now we turn to consider weak net neutrality where content-based price discrimination is prohibited either to content providers or consumers. The following proposition is demonstrated.

Proposition 2. *Under either weak content provider net neutrality or weak consumer net neutrality, the consumer chooses the socially optimal content provider (i.e., CP1) but the allocation of surplus remains the same as Proposition 1.*

This result says that weak net neutrality (regardless of its form) is neutral. That is, the only thing that may change as we move from no regulation to a form of weak net neutrality is that prices change. The payoffs to each party and the choices (in particular, for the consumer) do not change. Thus, the precise same level of static welfare results as the no regulation case.

The neutrality of weak net neutrality mirrors other results in the two-sided markets literature when there are prices set between all relevant parties and would arise in a much more general model than that presented here (Gans and King, 2003). It implies that weak net neutrality regulations are unlikely to be effective but also that a monopoly ISP could choose not to engage in content-based price discrimination along a given dimension and still achieve maximum profits. In other words, undertakings to engage in weak net neutrality could be freely made.

This also suggests avenues for empirical examination. For instance, it is not unreasonable to suppose that engaging in content-based price discrimination to either consumers or content providers involves transaction costs to the ISP. Thus, a monopoly ISP may minimise those costs by engaging in only one type of content-based price discrimination but that political and other pressures across different markets may influence which side of the market is chosen. Consequently, we would expect to see

content-based price discrimination related to historical norms in pricing perhaps related to vertical integration between ISPs and content providers.

3 Non-Neutral Results

In the previous section, we demonstrated that the outcomes under no regulation and either form of weak net neutrality were themselves neutral. That is, while there may be impacts of the regulation on the prices charged between ISPs, CPs and consumers, the choices and, ultimately, the payoffs to each agent are unchanged. Here we examine situations where, compared with no regulation, there may be non-neutral outcomes.

3.1 Strong net neutrality

Under strong net neutrality, the ISP is not permitted to engage in content-based discrimination of any form. Thus, $a_1 = a_2 = a$ and $t_1 = t_2 = t$. The following proposition demonstrates that regulation of this form is not neutral compared with other regulations.

Proposition 3. *Under strong net neutrality, the consumer chooses the socially optimal content provider (i.e., CP1) but earns zero surplus as does CP2. CP1 earns profits of* $v_1 + r_1 - (v_2 + r_2)$ *while the ISP earns profits of* $v_2 + r_2 - c$.

Thus, while (static) efficiency is maintained under strong net neutrality, the ISPs profits fall to $\min\{v_1 + r_1, v_2 + r_2\} - c$ while the chosen CP's profits, say, CP1, becomes $p_1 + r_1 - a = v_1 + r_1 - (v_2 + r_2)$. In this set-up, consumer surplus remains at zero. Consequently, a move from no regulation (or weak net neutrality) to strong net neutrality does not change overall (static) surplus generated but it does shift the appropriation of that surplus, in part, from ISPs to CPs.

Another way of seeing this is to note that the profits of CP1 are:

$$p_1 + r_1 - a_1 = v_1 - v_2 - (t_1 - t_2) + a_2 - r_2 + r_1 - a_1$$

If these are to be zero that implies that:

$$a_1 + t_1 - (a_2 + t_2) = v_1 + r_1 - (v_2 + r_2)$$

Thus, you can see that when the ISP can flexibly choose all of its prices or, at the very least, control the sums $a_i + t_i$, as it can under weak net neutrality, it is possible the ISP to extract all of the content provider profits. By contrast, under strong net neutrality

$a_1 + t_1 = a_2 + t_2$ so that CP1's profits must be equal to $v_1 + r_1 - (v_2 + r_2)$. Thus, the difference, $a_1 + t_1 - (a_2 + t_2)$ is a proxy for the expected transfer resulting from strong net neutrality.

Strong net neutrality is effective in the sense that it can change outcomes in the market. In the simple model here, it serves to shift surplus from the ISP to the content provider. Below, we see how this impacts on investments that each may make ex ante. However, intuitively, the content provider chosen by the consumer would not retain surplus under no regulation because the ISP could price in such a way that that surplus was extracted; effectively treating that content provider as higher cost than its rival. Under strong net neutrality, that extraction becomes impossible because the playing field between content providers is level. This allows the superior content provider to appropriate its marginal contribution to the market.

3.2 Outcomes with a missing price

Thusfar, we have assumed that consumers pay content providers directly to access their services. Of course, much of the Internet does not rely on that model and the services are free to consumers. Consequently, it is instructive to analyse the case where CPs do not charge consumers; i.e., $p_1 = p_2 = 0$. They instead rely exclusively on r_i for revenue or other benefits. This also corresponds to the case where consumers cannot charge CPs to give their attention. In this case, the consumer will choose CP1 over CP2 if:

$$v_1 - t_1 \geq v_2 - t_2$$

and there is no separate bidding from CPs for consumers. Importantly, as is demonstrated here, consumers do not themselves care about r_i in making their choices which gives rise to the possibility of distortions.

In this setting, we can demonstrate the following:

Proposition 4. *Suppose that* $v_1 + r_1 > v_2 + r_2$. *Then, if* $v_1 \geq v_2$ *then, regardless of regulatory regime, the consumer chooses the socially optimal content provider (i.e., CP1) and the ISP earns* $v_1 + r_1 - c$. *However, if* $v_1 < v_2$ *then, under no regulation or weak content provider net neutrality, the same outcome results while, under weak consumer net neutrality or strong net neutrality, the consumer chooses CP2 and the ISP earns* $v_2 + r_2 - c$.

When there is a missing price (in this case, the price between content providers and consumers), preventing content-based price discrimination in the charges paid by consumers to the ISP, has real effects (regardless of the regulations on the content provider fees to the ISP). Put simply, under weak consumer net neutrality and strong net neutrality, the ISP cannot use pricing to guide the consumer's choice of content provider – either directly or indirectly (through access fees paid by the content provider). The absence of that pricing signal means that consumers only choose their content provider on the basis of their own value and not any additional returns (i.e., advertising) made by the content provider. If consumer value and those additional returns are positively correlated, the consumer will choose the socially optimal content regardless of regulatory regime. However, if they are negatively correlated, this will not happen. In this case, the ISPs returns will be reduced but, as the proposition demonstrates, there is no corresponding transfer of surplus to content providers. Thus, net neutrality regulations are Pareto sub-optimal.

It is this possibility that suggests that net neutrality regulations are not innocuous in terms of allocative efficiency. When there is already a missing price in the market, further suppression of price signals can be welfare reducing. This might also arise if, for some reason, content from different content providers had different costs to the ISP in delivering that content to consumers. That is, for the same attention allocation of the consumer, some content may be more bandwidth intensive than others. Note, however, that for this to arise, the business model of content providers must not include direct charging to consumers. Consequently, the impact of net neutrality regulations is not independent of assumptions on business model type opening up an avenue for further empirical investigation of the welfare impacts of such regulations.[9]

[9] Choi, Jeon and Kim (2013) also identify content provider business model as interacting in important ways with net neutrality regulations. However, they examine different dimensions of business models than the ones discussed here.

4 Dynamic Consequences

Thusfar, the model has analysed static considerations centred around the consumers' choice of content provider. However, many of the discussions with regard to the impact of net neutrality have been on the investments that are made by content providers and ISPs. Here we examine those dynamic issues.

4.1 Content Provider Investments

The previous analysis demonstrated that when content providers can charge consumers directly, the only regulation that results in a change in their payoffs is strong net neutrality. Thus, moving from any other regime to strong net neutrality, increases the profits of the content provider that attracts consumer attention, say CP1, from 0 to $v_1 + r_1 - (v_2 + r_2)$; their marginal contribution to surplus. This is precisely the payoff that would normally be associated with a competitive content provider market. By contrast, in the absence of strong net neutrality, that marginal surplus is appropriated by the ISP. Hence, it is clear that if stimulating content provider investments is desirable then so is strong net neutrality.

When there is a missing price, the analysis becomes more complex. If revenue from other sources (e.g., advertising) is positively correlated with consumer value for the content of each content provider, then regardless of the regime, the ISP appropriates all surplus and the CPs earn 0. Thus, in this case, net neutrality regulation would not be effective in assisting content provider investments. If revenue from other sources is negatively correlated with consumer value, then weak consumer net neutrality and strong net neutrality can have real effects. However, there is no corresponding transfer of rents from the ISP to CPs. Hence, there is no expected impact on content provider investments.

4.2 Network quality

The main argument raised as to why net neutrality regulation of any form would be a poor outcome is that it would deter incentives for ISPs to invest in capacity and improve network quality. As already demonstrated, only if there is strong net neutrality is there an impact on ISP payoffs and thus, other forms of regulation would be neutral in

this regard. The question is whether the impact on payoffs also has a corresponding impact on ISP incentives to invest in network quality.

We consider a simple representation of such investment. Suppose that, at a sunk cost of $C > 0$, the ISP can improve network quality and hence, consumer value realised by Δ. We will assume that such investment is socially desirable with $\Delta > C$.

When the content provider can charge consumers directly, under no regulation and weak net neutrality, the ISP's profits are $v_1 + r_1 - c$ without an investment in network quality and $v_1 + \Delta + r_1 - c - C$ with that investment. Thus, they appropriate the full social returns from that investment and would be expected to undertake it. By contrast, under strong net neutrality, ISP profits are $v_2 + r_2 - c$ without and $v_2 + \Delta + r_2 - c - C$ with the investment. Consequently, it can be seen that there is no change in the ISPs incentives to invest in network quality as a result of a change in regulatory regime. A similar outcome can be demonstrated for the case where there is a missing price as, in that case, even if the consumer chooses the non-welfare maximising content, the ISP appropriates the increment to consumer value from its investment in quality by an increase in charges to consumers.

4.3 Fast lanes

A recent form of investment or, specifically, product introduction from ISPs are so-called 'fast lanes.' These are products available to content providers to speed up their service to the ISP's customers. Concerns have been raised that such products are a way of circumventing pricing regulation rules for ISPs with respect to content providers – that is, circumventing the effects of weak content provider net neutrality or strong net neutrality. Here I evaluate those claims.

Suppose that the ISP offers a fast lane to content providers at a fee of A and a slow lane to them at a fee of a. Note that the fees, A and a, are not contingent on content. If CPi chooses a fast lane, the value to consumers is Δ_i greater than if they choose a slow lane.[10] It is assumed that $\Delta_1 > \Delta_2$. It is also assumed that, if a fast lane is used, it costs the ISP C rather than c.

[10] Choi and Kim (2010) provide a detailed model of how these different qualities may arise using network prioritization rules.

The timeline is only changed slightly with this set-up. First, the ISP announces the consumer and content provider fees (for both lanes). Then the consumer chooses an ISP. Then the content providers simultaneously announce their prices to consumers and whether they have chosen a fast or slow lane. Finally, consumers choose their content provider.

Given this, we can prove the following result:

Proposition 5. *Under weak content provider net neutrality, the consumer selects the socially optimal content (CP1) and there is an equilibrium where CP1 chooses a fast lane. Neither the consumer nor the content providers earn any surplus while the ISP earns* $v_1 + \Delta_1 + r_1 - C$. *Under strong net neutrality, the consumer selects the socially optimal content (CP1) while CP1 chooses a fast lane. Neither the consumer nor CP2 earn any surplus while the ISP earns* $v_2 + \Delta_1 + r_2 - C$ *and CP1 earns* $v_1 + r_1 - (v_2 + r_2)$.

With regard to weak content provider net neutrality, the result in the proposition is not too surprising. Put simply, content-based fees from the ISP at the consumer level can do the work of the regulation at the content provider level (as in Proposition 2).

The strong net neutrality result shows that, while the fast lane does not change surplus appropriated by CP1, the increased value from the fast lane itself is appropriated by the ISP. As in the result in the previous section, this demonstrates that strong net neutrality will not change the ISP's incentives to invest in higher quality broadband. However, here we have assumed that some the benefit to the consumer that comes from higher quality broadband is complementary with the content provider (i.e., that $\Delta_1 > \Delta_2$). This means that there is no incentive for the content provider to invest in content that will exploit the fast lane, if it is provided.

Alternatively, what would happen if the ISP was compelled to offer a single lane. In this situation, it can be easily shown that while nothing changes under weak net neutrality, under strong net neutrality the ISP earns $v_2 + \Delta_2 + r_2 - C$ while CP1 earns $v_1 + \Delta_1 + r_1 - (v_2 + \Delta_2 + r_2)$. Thus, the content provider's incentives are aligned with their marginal impact on surplus but the ISP only cares about investing in quality to the extent that it improves quality across a variety of content providers. It is not difficult to imagine situations where this may be socially more or less efficient.[11]

[11] Some of these trade-offs are analysed by Choi and Kim (2010). See also Hermalin and Katz (2007).

5 Competition amongst ISPs

As a final exercise, it is instructive to consider the role of the monopoly ISP here. What happens if there is ISP competition? Specifically, suppose there are two identical ISPs who announce their charges to consumers and content providers prior to the consumers choosing their ISP and, subsequently, their content.

First, observe that the price structure of the content providers would not change as a result of this and, in particular, in the unregulated case, if an ISP managed to sign up a consumer, then the net cost of serving a consumer that chose CP*i* would be $c - a_i$. Thus, the higher an ISP can make each a_i, the lower those costs can be; something that will matter with ISP competition.

Second, this suggests that if a_i were equal to $v_i + r_i$, there would be potential for the ISP to transfer surplus from either CP to the consumer. As a CP's profits are $p_i + r_i - a_i$ then this outcome could only arise if $p_i = v_i$. The problem is that the CPs compete for the consumer but also set the price to the consumer. For CP2, as the inferior provider, this means that the lowest they can set their price is $p_2 = a_2 - r_2$. Thus, an ISP could set $a_2 = v_2 + r_2$ resulting in $p_2 = v_2$; interestingly, the ISP is motivated, in order to compete for the consumer, to create conditions whereby the competition between CPs is softened.

Given this, in competition for the consumer, CP1 sets its price at $p_1 = v_1 - v_2 - (t_1 - t_2) + p_2 = v_1 - (t_1 - t_2)$. Thus, the surplus accruing to the consumer is $v_1 - t_1 - p_1 = -t_2$. So, by setting $t_1 = t_2 = c - a_1 = c - r_1 - v_1$, an ISP will potentially be able to bid successfully for a consumer with the consumer appropriating the total surplus generated.

As would be expected, consumers appropriate all of the surplus when there is ISP competition. The reason that content providers appropriate none of the surplus is that it is the consumer that chooses the ISP and the content provider is forced to accept the terms of that ISP in accessing the consumer. Importantly, it can readily be seen that net neutrality regulations would not change this outcome – at least for this specialised model. It is already the case that, under no regulation, the ISP implements weak consumer net neutrality.

What about strong net neutrality? If this were required, it is easy to demonstrate that the ISP could simply set $a = a_1 = a_2 = v_1 + r_1$ and the same outcome would result. To see this, note that, in this case, $p_1 = v_1 - v_2 - (t_1 - t_2) + v_1 + r_1 - r_2$ with surplus to the consumer of $v_1 - t_1 - p_1 = -v_1 - r_1 + v_2 + r_2 - t_2$. For this surplus to be $v_1 + r_1 - c$, we must have $t_2 = -2v_1 - 2r_1 + c + v_2 + r_2$. Thus, $t = t_1 = t_2 = -2v_1 - 2r_1 + c + v_2 + r_2$. So we have $p_1 = 2v_1 - v_2 + r_1 - r_2$, $p_2 = v_1 + r_1 - r_2$ so that $t + p_1 = -r_1 + c$. In this case, the consumer appropriates the entire surplus of $v_1 + r_1 - c$.[12]

It is interesting that having ISP competition does not, in fact, result in either more surplus to content providers or to some effectiveness of net neutrality regulations. In reality, this result is strong because of the special nature of the model. Here the ISP controls terms on both sides of the market – consumer and content provider. Thus, because the consumer chooses one ISP, the ISP has a monopoly over access to that consumer. Moreover, the ISP chooses to exercise that monopoly power by softening competition amongst CPs and extracting rents. In competition, those monopoly rents are passed to the consumer in the firm of negative consumer access fees for broadband. Net neutrality regulations only mean that those consumer access fees are more negative but otherwise the allocation of surplus in the market is unchanged.

While it is possible to conceive of negative consumer fees for broadband as rebates for signing on to an ISP, many might consider this outcome unrealistic. Thus, it is instructive to consider what happens if we assume that all prices must be positive and so t_1 and t_2 must be at least 0. Suppose, therefore, that t_1 is 0. Then, the profits of CP1 would be $p_1 + r_1 - a_1 = v_1 - v_2 - (0 - t_2) + a_2 - r_2 + r_1 - a_1$. Can it be the case that, like the situation with negative prices, the profits of CP1 can be 0? This implies that $a_1 - (a_2 + t_2) = v_1 + r_1 - (v_2 + r_2)$. However, if $a_1 = v_1 + r_1$ and $a_2 + t_2 = v_2 + r_2$, then $p_1 = v_1$ and consumer surplus would be 0. An ISP would be able to compete for a consumer by instead ensuring that p_1 were as low as possible. This implies setting $a_2 = r_2$ so that $p_2 = 0$ and also t_2 as low as possible (i.e., 0). Thus, $p_1 = v_1 - v_2$ leaving CP1 with profit

[12] This provides some support for Becker, Carlton and Sider's (2010) conjecture that ISP competition makes net neutrality regulation ineffective. However, below if prices are constrained to be non-negative, then some net neutrality regulation is effective.

$v_1 - v_2 + r_1 - a_1$ and consumer surplus of v_2. So the consumer is indifferent as to the level of a1 and so the ISP can set $a_1 = v_1 + r_1 - v_2$ and earn a profit. Thus, as in the flexible price case, the non-negative price case leaves content providers with no surplus.

In this situation, what is the impact of net neutrality regulations? First of all, note that weak consumer net neutrality is implemented in equilibrium regardless. Second, if the price charged by the ISP to content providers must be equal, then $a = a_1 = a_2 = r_2$. Thus, under both weak content provider net neutrality and strong net neutrality, CP1 receives profit of $v_1 - v_2 + r_1 - r_2$ while the 'winning' ISP receives profit of $r_2 - c$ and the consumer receives surplus of v_2. Thus, compared with the monopoly ISP case, when there are non-negative prices, a broader range of net neutrality regulations are potentially effective in transferring surplus from ISPs to content providers.

As noted above, what drives some of these results is that the ISP controls both sides of the market. Alternatively, what if the content provider had its own ISP? In this case, that ISP might charge the content provider to access the Internet as a whole. However, between the two ISPs, they will negotiate the link between the consumer and the content provider. This is the *interconnection fee* that the content provider's ISP will pay the consumer's ISP for transit. In this situation, if there is competition between ISPs on both sides of the market, we would expect that some surplus would accrue to the consumer while others would accrue to the content provider.[13] In either case, net neutrality regulations would likely not change the outcomes here.

6 Conclusions

The economics of net neutrality is subtle. The existence of a network of pricing relationship means that regulating any one set of prices may not have real effects. Consequently, the form and extent of net neutrality regulation matters in terms of what might be effective in reallocating surplus or changing the allocation of attention of consumers. This paper has provided a simple model to illustrate effects, some of whom were already present in the literature, but to place them in a broader pricing context.

[13] Choi, Jeon and Kim (2013) explore interconnected and competing ISPs.

That said, the simplicity of the model no doubt abstracts away from other mechanisms. My claim here is that the ones identified in this paper will remain first order. In particular, when it is effective net neutrality can shift surplus from ISPs to content providers with consequent impact on content provider investments. By contrast, it is difficult to find a mechanism in the paper whereby net neutrality regulation would have an adverse impact on ISP investment. While other papers have found an impact this has been in the context of mitigating effects that have been largely supressed here.

The hope now is that by providing a complete framework to analyse the impact of net neutrality, empirical researchers may be able to find experiments and other opportunities to measure and predict the likely impact of such regulations.

7 Appendix: Proof of Propositions

7.1 Proof of Proposition 1

Working backwards, given (t_1,t_2), the consumer will choose CP1 over CP2 if:

$$v_1 - t_1 - p_1 \geq v_2 - t_2 - p_2$$

The CPs will set their consumer prices, in effect, to bid for the consumer's single unit of attention. Consequently:

- If $v_1 + r_1 - t_1 - a_1 \geq v_2 + r_2 - t_2 - a_2$, $p_2 = a_2 - r_2$ and $p_1 = v_1 - v_2 - r_2 - (t_1 - t_2) + a_2$ (CP1 is used)
- If $v_1 + r_1 - t_1 - a_1 \leq v_2 + r_2 - t_2 - a_2$, $p_1 = a_1 - r_1$ and $p_2 = v_2 - v_1 - r_1 - (t_2 - t_1) + a_1$ (CP2 is used)

Thus, there are two cases:

- If $v_1 + r_1 - t_1 - a_1 \geq v_2 + r_2 - t_2 - a_2$, then the ISP's profits are $t_1 + a_1 - c$. If it wants to achieve this outcome, then $t_1 + a_1 \leq v_1 - v_2 + (r_1 - r_2) + t_2 + a_2$; which also implies that $p_1 + r_1 \geq a_1$. For the consumer to participate, however, $v_1 - t_1 - p_1 \geq 0 \Rightarrow v_2 + r_2 \geq t_2 + a_2$. Thus, the ISP sets $t_2 + a_2 = v_2 + r_2$ and $t_1 + a_1 = v_1 + r_1$ and earns profit of $v_1 + r_1 - c$. Both consumer and CP surpluses are zero.
- If $v_1 + r_1 - t_1 - a_1 \leq v_2 + r_2 - t_2 - a_2$, the same prices $t_2 + a_2 = v_2 + r_2$ and $t_1 + a_1 = v_1 + r_1$ maximise ISP profits.

Notice that, in each case, $p_i = a_i - r_i$ while the sum of consumer and CP access charges are constant. However, (static) social welfare is maximized as the ISP wants to structure prices at the point of indifference to maximise $v_i + r_i - c$. To see this, given p_i, the consumer will choose CP1 over CP2 if:

$$v_1 - t_1 - a_1 + r_1 \geq v_2 - t_2 - a_2 + r_2 \text{ or}$$

$$v_1 + r_1 - (v_2 + r_2) \geq t_1 + a_1 - (t_2 + a_2)$$

Thus, if $t_2 + a_2 = v_2 + r_2$ and $t_1 + a_1 = v_1 + r_1$, the consumer's choice will align with the (static) socially optimal outcome.

7.2 Proof of Proposition 2

We begin with content provider net neutrality. In this case, the ISP can set a single access price, *a*, to each CP. Note that since it was $a_i + t_i$ that mattered for everything under no regulation, then regardless of *a*, the ISP will choose (t_1,t_2) to achieve the same outcome as above. That is, the prices charged by the CPs to the consumers will be determined by:

- If $v_1 + r_1 - t_1 - a \geq v_2 + r_2 - t_2 - a$, $p_2 = a - r_2$ and $p_1 = v_1 - v_2 - r_2 - (t_1 - t_2) + a$ (CP1 is used)

- If $v_1 + r_1 - t_1 - a \leq v_2 + r_2 - t_2 - a$, $p_1 = a - r_1$ and $p_2 = v_2 - v_1 - r_1 - (t_2 - t_1) + a$ (CP2 is used)

Again, there are two cases:

- If $v_1 + r_1 - t_1 \geq v_2 + r_2 - t_2$, then the ISP's profits are $t_1 + a - c$. If it wants to achieve this outcome, then $t_1 \leq v_1 - v_2 + (r_1 - r_2) + t_2$; which also implies that $p_1 + r_1 \geq a$. For the consumer to participate, however, $v_1 - t_1 - p_1 \geq 0 \Rightarrow v_2 + r_2 \geq t_2 + a$. Thus, it sets $t_2 + a = v_2 + r_2$ and $t_1 + a = v_1 + r_1$ and earns profit of $v_1 + r_1 - c$. Both consumer and CP surpluses remain at zero.
- If $v_1 + r_1 - t_1 \leq v_2 + r_2 - t_2$, the same prices $t_2 + a = v_2 + r_2$ and $t_1 + a = v_1 + r_1$ maximise ISP profits.

Thus, the ISP profits are exactly the same as would arise under no regulation as are consumer and CP surplus that remain at zero.

For the case of consumer net neutrality, the outcome is the same and for the same reasons. In this case, the ISP price to consumers had to be non-discriminatory (i.e., $t_1 = t_2 = t$), then if (a_1, a_2) could be chosen freely, the same outcome as under no regulation would emerge as the prices that maximise ISP profits are $t + a_2 = v_2 + r_2$ and $t + a_1 = v_1 + r_1$.

7.3 Proof of Proposition 3

Given this, in bidding for the consumer's attention:

- If $v_1 + r_1 - t - a \geq v_2 + r_2 - t - a$, $p_2 = a - r_2$ and $p_1 = v_1 - v_2 - r_2 + a$ (CP1 is used)
- If $v_1 + r_1 - t - a \leq v_2 + r_2 - t - a$, $p_1 = a - r_1$ and $p_2 = v_2 - v_1 - r_1 + a$ (CP2 is used)

Note that, in this case, consumers are not induced to choose one content provider over the other on the basis of any price directly under the control of the ISP. By contrast, under weak net neutrality, there was always a way for the ISP to charge content-based prices to impact on the consumer's choice.

This means that ISP profits are $t + a - c$ and it sets $t + a = \min\{v_1 + r_1, v_2 + r_2\}$ to maximise those profits. This is because if $v_1 + r_1 \geq v_2 + r_2$, $p_2 = a - r_2$ and $p_1 = v_1 - v_2 - r_2 + a$. Thus, if the ISP sets $t + a = v_1 + r_1$, if the consumer chooses CP1, then the consumer receives $v_1 - p_1 - t = v_2 + r_2 - a - t$ whereas if the consumer chooses CP2, it receives $v_2 - p_2 - t = v_2 + r_2 - a - t$. In either case, because $v_1 + r_1 \geq v_2 + r_2$, the consumer's surplus is negative. Thus, it can be seen that the maximum price the ISP can set is $t + a = v_2 + r_2$.

7.4 Proof of Proposition 4

We begin by not making any assumption regarding the comparison between v_1 and v_2. Under no regulation, ISP profits are maximised by setting charges to consumers and

content providers so that consumers make the socially optimal decision. It is readily apparent that this is achieved with $t_i = v_i$ and $a_i = r_i$. Thus, profits are the same as in the no regulation case above.

Under weak content provider net neutrality, $a_1 = a_2 = a$, the ISP can set $t_i = v_i$ so that the consumer chooses the optimal CP. To see this, suppose that $v_1 + r_1 > v_2 + r_2$. Then if $r_1 \geq r_2$ by setting $a = r_1$ and $t_1 = v_1 - \varepsilon$, $t_2 = v_2$ where ε is arbitrarily small the ISP will receive the maximum profits of $v_1 + r_1 - c$. However, if $r_1 < r_2$, then, with these same prices, the ISP can receive maximum profits as it can guide the consumer towards CP1 and extract the other rents from CP1. Thus, the same outcome as no regulation occurs.

Now suppose that $v_1 \geq v_2$. Under weak consumer net neutrality, when t is non-discriminatory, as it is under strong net neutrality, the ISP cannot guide the consumer at all. In this case, the consumer chooses CP1 and so the ISP sets $t = v_1$ and $a_i = r_i$. It should be readily apparent that this same outcome will arise under strong net neutrality with $t = v_1$ and $a = r_1$.

If $v_1 < v_2$ then, so long as it can set different prices to consumers, the ISP can set $t_1 = v_1 - \varepsilon$, $t_2 = v_2$ to guide the consumer to the correct CP leaving a little surplus with the consumer for choosing CP1. Thus, under no regulation and under weak content provider net neutrality, the socially optimal outcome arises.

By contrast, under weak consumer net neutrality or strong net neutrality, the ISP has no mechanism by which it can guide the consumer to choose CP1. The consumer will always choose CP2 and so the ISP maximises profits by charging $t = v_2$ and $a = r_2$.

7.5 Proof of Proposition 5

Under weak content provider net neutrality, suppose that the ISP sets $t_i + a = v_i + r_i$ and $t_i + A = v_i + \Delta_i + r_i$. This implies that $A - a$ is a constant. In addition, suppose that $A - a = \Delta_1$. This implies that CP2 will never choose the fast lane. Consequently, $p_2 = a - r_2$ and $p_1 = v_1 + \Delta_1 - v_2 - r_2 - (t_1 - t_2) + a$ so that the consumer chooses CP1. In this case, $p_1 = v_1 + \Delta_1 - t_1 = A - r_1$ so that CP1 earns zero surplus. Note that CP1 gains no advantage if it chooses the slow lane and thus, this is an equilibrium outcome.

Under strong net neutrality, content-based price discrimination is not possible in any form. In this case, suppose that, as in Proposition 4, the ISP sets $t + a = v_2 + r_2$ and, therefore, that $t + A = v_2 + \Delta_1 + r_2$. Consequently, CP2 will never choose the fast lane and $p_2 = a - r_2$ and $p_1 = v_1 + \Delta_1 - t$. In this case, the consumer receives $v_1 + \Delta_1 - p_1 - t = v_2 - t = a - r_2 = 0$. However, CP1 receives

$p_1 + r_1 - A = v_1 + \Delta_1 - t + r_1 - A = v_1 + r_1 - (v_2 + r_2)$ and would no more if it did not choose the fast lane. Finally, the ISP receives $v_2 + \Delta_1 + r_2 - C$.

8 References

Becker, G.S., D.W. Carlton and H.S. Sider (2010), "Net Neutrality and Consumer Welfare," *Journal of Competition Law and Economics*, 6 (3), pp.497-519.

Cheng, H.K., S. Bandyopadhyay and H. Guo (2009), "The Debate on Net Neutrality: A Policy Perspective." *Information Systems Research*.

Choi, J.P. (2010), "Optimal Tariffs and the Choice of Technology: Discriminatory Tariffs vs. The 'Most Favored Nation' Clause." *Journal of International Economics,* Vol. 58, pp.560–579.

Choi, J.P., D-S. Jeon and B.C. Kim (2013), "Net Neutrality, Business Models and Internet Interconnection," *mimeo*., UNSW.

Choi, J.P. and B.C. Kim (2010), "Net Neutrality and Investment Incentives," *RAND Journal of Economics,* 41, pp.446–471.

Economides, N. (2008), "Net Neutrality, Non-Discrimination and Digital Distribution of Content through the Internet," *A Journal of Law and Policy for the Information Society,* 4, pp.209–233.

Economides, N. And B. Hermalin (2012), "The economics of network neutrality," RAND Journal of Economics, 43 (4), pp.602-629.

Economides, N. and J. Tag (2012), "Network Neutrality on the Internet: A Two-Sided Market Analysis," Information Economics and Policy, 24 (2), pp.91-104.

Gans, J.S. and S.P. King (2003), "The Neutrality of Interchange Fees in Payment Systems," *Topics in Economic Analysis and Policy*, 3 (1), Article 1.

Hermalin, B.E. and M.L. Katz (2007), "The Economics of Product-Line Restrictions with an Application to the Network Neutrality Debate." *Information Economics and Policy,* 19, pp. 215–248.

Lee, R. S. and T. Wu (2009), "Subsidizing creativity through network design: Zero-pricing and net neutrality," *Journal of Economic Perspectives*, 23, 61–76.

Reggiani, C. and T. Valletti (2011), "Net neutrality and innovation at the core and at the edge," mimeo., Imperial.

Wu, T. (2003), "Network Neutrality, Broadband Discrimination." *Journal on Telecommunications and High Technology Law,* 2, pp.141–178.