NBER WORKING PAPER SERIES

NET NEUTRALITY, PRICING INSTRUMENTS AND INCENTIVES

Joshua S. Gans
Michael L. Katz

Working Paper 22040
http://www.nber.org/papers/w22040

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2016

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w22040.ack

Net Neutrality, Pricing Instruments and Incentives
Joshua S. Gans and Michael L. Katz
NBER Working Paper No. 22040
February 2016
JEL No. D4,D42,D43,L1,L12,L13

ABSTRACT

We correct and extend the results of Gans (2015) regarding the effects of net neutrality regulation on equilibrium outcomes in settings where a content provider sells its services to consumers for a fee. We examine both pricing and investment effects. We extend the earlier paper's result that weak forms of net neutrality are ineffective and also show that even a strong form of net neutrality may be ineffective. In addition, we demonstrate that, when strong net neutrality does affect the equilibrium outcome, it may harm efficiency by distorting both ISP and content provider investment and service-quality choices.

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

Michael L. Katz
UC, Berkeley
Haas School of Business
545 Student Services #1900
Berkeley, CA 94720-1900
katz@haas.berkeley.edu

1 Introduction

In 2015, the U.S. Federal Communications Commission (FCC) adopted new Net Neutrality rules that would apply to both fixed and wireless Internet Service Providers (ISPs).[1] These rules are examples of what Gans (2015) described as *weak net neutrality*: the rules prohibit ISPs from discriminating on the basis of content in setting their charges to content providers for exchanging traffic with the ISPs' end users (consumers). Gans (2015) found that, when content providers charge for their services and an ISP can discriminate on the basis of content in setting its charges to consumers, weak net neutrality will have no impact on any agent's payoffs. He argued that stronger net neutrality rules are required if the policy is to have an impact. Specifically, he examined *strong net neutrality:* rules that prohibit content-based discrimination in pricing to both content providers *and* consumers.

In this paper, we identify an error in Gans (2015)'s conclusion regarding the non-neutrality of strong net neutrality. This insight gives rise to several further considerations and a more nuanced result on strong net neutrality. The paper proceeds as follows. After describing the model in Section 2, we demonstrate in Section 3 that, when the ISP can choose its access and termination rates subject only to non-discrimination requirements, it has sufficient instruments to extract the entire social surplus.

Section 4 then considers two situations in which the ISP can no longer extract all of the social surplus: (a) when there is a regulatory cap on the IPS's access charges, and (b) when consumers have heterogeneous preferences with respect to content providers. We first show that, when net neutrality takes the form of binding limit on the fees the ISP chargers to content providers, this policy meaningfully affects the equilibrium outcome but can result in inefficiencies, including the excessive provision of quality by content providers. We then consider the impact of consumer heterogeneity on these results (Gans, 2015, had only a single, representative consumer). We demonstrate that strong net neutrality can harm total surplus by inducing the ISP to exclude low-value content providers from the market. This result is a variant of the standard trade-off inherent with heterogeneous consumers and an inability to engage in price discrimination: the ISP compares setting a high price, which excludes low-value transactions, with setting a low price that results in greater sales. We also examine the effects of

[1] Federal Communications Commission, *In the Matter of Protecting and Promoting the Open Internet, Report and Order on Remand*, Declaratory Ruling, and Order, rel. March 12, 2015.

net neutrality on investment incentives.[2] Gans (2015) found that strong net neutrality can enhance content providers' incentives to invest in quality, but will have no effect on ISP investment. We demonstrate that the conclusion with respect to content providers' investment incentives extends to the case of consumers with heterogeneous preferences with respect to quality, but the conclusion with respect to the ISP's investment incentives does not—there are circumstances under which strong net neutrality regulation reduces the ISP's investment incentives.

Section 5 presents a brief concluding remarks.

2 Model Set-Up

We utilize the simple set-up of Gans (2015). There is a single ISP, two content providers (CP1 and CP2), and a unit mass of a single type of consumers (until Section 4, where we relax that assumption). In terms of notation, we assume the following:

- v_i is the value a consumer derives consuming content from provider *i*.
- r_i is the revenue content provider *i* derives from other sources (*e.g.*, advertising).
- t_i is the transmission price charged by the ISP to a consumer who purchases content from provider *i*.
- a_i is the access price the ISP charges provider *i* for having its content delivered to the consumer.
- p_i is the price provider *i* charges a consumer for content.

We consider two different sets of assumptions with regard to the timing of price setting. Gans (2015) assumed that the ISP sets $\{a_i, t_i\}_{i\in 1,2}$ prior to either content provider's setting its price, p_i. In addition to considering such *sequential* pricing, we also examine setting in which $\{a_i, t_i, p_i\}_{i\in 1,2}$ are all set *simultaneously*.

Following Gans (2015), we define the following policies:

[2] There has been considerable debate regarding the effects of net neutrality regulation on ISP and CP investment incentives; for a review see Greenstein et.al. (2016). They note that an important factor in determining how net neutrality impacts on investment incentives has to do with the alleviation and pricing of congestion on networks. Depending on the model of queuing and on the business model of ISPs, net neutrality regulation can both detract and promote investment in infrastructure to alleviate congestion.

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

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

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

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

As mentioned in the Introduction, the FCC's recently imposed rules are properly characterized as weak content provider net neutrality.

We make the following assumptions about demands and costs. Because of limited attention, a consumer selects content from at most one provider. The values a consumer derives from content consumption satisfy $v_1 > v_2$, and the sums of consumer value and other content provider revenue (*e.g.*, fees from advertisers) satisfy $v_1 + r_1 > v_2 + r_2$. We normalize the marginal cost of content provision to zero. The ISP incurs a marginal cost of *c* to deliver content to a consumer, regardless of which content provider is chosen.

3 Sufficient Instruments

Gans (2015) showed that the ISP can use discriminatory pricing based on the identity of the content provider to extract the entire surplus available in the market. Proposition 3 of Gans (2015) proposed that only strong net neutrality—which prevents all forms of price discrimination based on the identity of the content provider—changes the equilibrium set of choices and payoffs and allows a content provider to enjoy positive surplus. However, the proof of the proposition contained an error—one that is instructive in evaluating source of the ISP's hold-up power. Here, we will demonstrate that the hold-up power arises because of the number of pricing instruments at the ISP's disposal relative to the number of consumer types (specifically, in the baseline model there is only one consumer type).

3.1 Strong Net Neutrality can be Neutral

Under strong net neutrality, the ISP has only two prices that it can set, t and a; that is, both the ISP's transmission price and access price are independent of content chosen. Assuming that the ISP sets its prices before the content providers set theirs, the following proposition holds:

Proposition 1: *Under strong net neutrality, every consumer chooses the socially optimal content provider (i.e., CP1) but earns zero surplus, as do CP1 and CP2. The ISP earns profits* $v_1 + r_1 - c$.

> **Proof:** Suppose that the ISP sets $t = v_1$ and $a = r_1$. The highest price that CP1 can charge and still makes sales is $p_1 = 0$. CP2 cannot make sales at any positive price and cannot earn non-negative profits at any price consumers would be willing to pay. The ISP earns profits of $v_1 + r_1 - c$, which clearly is the upper bound on its profits given consumer and content-provider rationality. **QED**

The equilibrium surplus to each agent under Proposition 1 is the same as that which arises in the absence of regulation or under either form of weak net neutrality.[3]

Proposition 1 demonstrates that, under strong net neutrality, the ISP maximizes its profit by identifying CP1 as the efficient choice for consumers and choosing prices that extract all of the surplus when a consumer chooses CP1. Because the ISP's prices are independent of the content provider's identity and CP2 offers less potential total surplus, CP2 is excluded from the market; even if CP2 set a price equal to 0, a consumer would earn negative surplus given the ISP's prices. The error in Gans (2015) lay in an implicit assumption that both content providers must be active in equilibrium, so that a consumer would never suffer negative surplus from choosing either content provider.[4]

This error helps to identify a potentially important difference between strong and weak net neutrality. Although the equilibrium payoffs are the same under all three forms of net neutrality in this simple model, there is a sense in which the ISP has a set of more robust

---

[3] See Gans (2015).

[4] This error also effects Proposition 5 of Gans (2015). That proposition should now read:

> *If* $\Delta_1 \geq C_1$ *then absent net neutrality or under either weak content provider net neutrality or strong net neutrality, the consumer selects the socially optimal content (CP1) while CP1 chooses a fast lane. Neither the consumer nor the content providers earn any surplus while the ISP earns* $v_1 + \Delta_1 + r_1 - C$.

The corrected proposition demonstrates that, regardless of the regulatory regime, the fast lane does not alter the surplus appropriated by CP1, although it does change the surplus that the ISP receives. It follows that CP1 does not have incentives to invest in ways that exploit the possibility of generating greater consumer benefits through use of a fast lane.

strategies available under weak net neutrality. Specifically, under weak net neutrality, the ISP could also rely on a pricing strategy that allowed both content providers to be active in equilibrium: the ISP can set either t_i or a_i in a discriminatory fashion to satisfy $t_i + a_i = v_i + r_i$ for both *i*. Under weak net neutrality, the ISP need only make content-provider-specific pricing decisions without choosing to exclude any content provider. Consumers then choose the efficient content provider in equilibrium. Because the ISP's strategy under strong net neutrality hinges on all consumers' choosing the same content provider, it would appear unlikely that this result will fully extend to setting which consumers have heterogeneous content preferences. We confirm this intuition in Section 4.2 below.

3.2 Invariance to Pricing Timing

It is useful briefly to explore whether the neutrality of strong net neutrality is sensitive to changes in assumptions regarding the timing of pricing decisions. To see that it is not, suppose, instead that each p_i is set prior to the ISP's setting its prices. We will now demonstrate that:

Proposition 2: *If the content providers set their prices, $\{p_1, p_2\}$, prior to the ISP's setting its prices, then under any regulatory regime there exist equilibria in which consumers choose the socially optimal content provider (i.e., CP1) but earn zero surplus as do CP1 and CP2. The ISP earns profits of $v_1 + r_1 - c$.*

> **Proof:** Consider the following strategies. The content providers choose prices such that $v_1 - p_1 > v_2 - p_2$ and ISP sets $t = v_j - p_j$ and $a = p_j + r_j$ for any CP prices such that $v_j - p_j \geq v_k - p_k$ for $j \neq k$. Observe that the ISP's strategy states how it will set its prices contingent on the prices set by the CPs. Given the CPs' prices, CP2 makes no sales. At any price at which CP2 would make positive sales, its profit per consumer would be $p_2 + r_2 - a = p_2 + r_2 - p_1 - r_1 < v_2 + r_2 - v_1 - r_1 < 0$. Hence, CP2 cannot benefit from deviating from the candidate equilibrium. Given the ISP's strategy, CP1 also cannot deviate to earn positive profits. The ISP fully extracts the maximum surplus that can be realized given the cost and demand conditions. It follows that, given CP and consumer rationality, there is no profitable deviation for the ISP. **QED**

Intuitively, the ISP has sufficient tools (*i.e.*, it can separately charge both content providers and consumers) to engage in a perfect price squeeze even though CP1 can commit to the price of its complementary product before the ISP.[5]

4 Insufficient Instruments

The focus of the net neutrality has been on the prices charged by ISPs to content providers for access to consumers. We have shown, thus far, that regulating the ISP's ability to price discriminate between content providers has limited potential to ensure that rents flow to content providers to reward them for their investments. This arises because the ISP has sufficient instruments to extract rents even if the structure of pricing is regulated. In this section, we consider two situations in which the ISP no longer has sufficient instruments to extract all of the surplus: (a) when there is a regulatory cap on the ISP's access charges, and (b) when consumers have heterogeneous preferences with respect to content providers.

4.1 Capped Access Charges

Some proponents of net neutrality have advocated directly regulating access charges or even banning them all together. We now turn to consider this form of regulation which is a more restrictive form of weak content provider net neutrality. Specifically, suppose that the access charges, a_i, are regulated to be equal to one another and less than or equal to $\bar{a}$, where $\bar{a}$ is set by a regulator. In addition, suppose that the ISP can freely set its charges to consumers, t_i. Compared with weak content provider net neutrality, this regulation removes an instrument from the ISP's control. As we will now demonstrate, the effects of such a policy depend, in part, on whether negative prices are possible and whether quality can be used to dissipate content provider profits.

Return to the sequential pricing assumption that the ISP sets its prices before the CPs set theirs. If the ISP sets $t_i = v_1 + r_1 - \bar{a}$, then the maximum price CP1 can set is $p_1 = \bar{a} - r_1$, and CP2

[5] Proposition 2 identifies a specific family of equilibria but there are also others. For instance, there exists a family of equilibria in which CP1 drops out in stage one and the ISP extracts surplus generated by consumption of only CP2's services. The existence of this family of equilibria is a consequence of the timing, which makes CP1 indifferent between staying in and opting out at the first stage. We suspect that an equilibrium refinement (such as trembling hand perfection) may select the equilibrium in the proposition but it is beyond the scope of the paper to examine this technical extension.

cannot profitably attract any sales. Thus, the content providers earn 0 while the ISP earns the same amount as it would if there were no regulation. Observe that, under this regulatory regime, the ISP could also choose to set a different price to consumers depending on the identity of the content provider (*i.e.*, set $t_i = v_i + r_i - \bar{a}$), which would allow it to extract all of the surplus while having both content providers active in the case where consumers have heterogeneous preferences.

Note that, even if—in addition capping the access price charged to content providers—the regulator also imposed consumer net neutrality (*i.e.*, $t_1 = t_2 = t$), then the ISP would (if negative prices were possible) still be able to engage in the type of price squeeze examined in Section 3.1 above. Thus, the outcomes would be equivalent to strong net neutrality. If, on the other hand, negative prices were not possible, then some of the rents from advertising would be captured by CP1 as per the above discussion.

Setting $t_i = v_i + r_i - \bar{a}$ relies on it's being possible for content providers to charge negative prices when $r_i > \bar{a}$. If negative content prices are infeasible, then content providers can make positive sales only if $t_i \leq v_i$. Thus, in equilibrium, the ISP will earn $v_1 + \min\{\bar{a}, r_1\} - c$, while CP1 earns $r_1 - \bar{a}$. This finding reinforces the point made Gans (2015) regarding the role of payments between content providers and consumers in undermining network neutrality. Specifically, with regulation of this kind, for sufficiently low $\bar{a}$, the impact of regulation is to bias content providers towards models that emphasize advertising revenues rather than direct pricing to consumers.

When negative content prices are infeasible but each content can provider to choose the (costly) quality of its service, increased expenditures on quality can play a role similar to that of negative content prices—a content provider can lower its quality-adjusted price by raising its quality. To explore this possibility, we introduce additional notation:

- $v_i(e_i)$ is the value the consumer derives consuming content from content provider *i* when the provider expends e_i per customer on quality.
- e_i^* is content provider *i*'s efficient quality expenditure.

We assume that: $v_i(\cdot)$ is an increasing, concave function; $v_1(e) > v_2(e)$ and $v_1(e) + r_1 > v_2(e) + r_2$ for all $e > 0$; and $v_1(e_1^*) - e_1^* > 0$. We now establish the following result:

Proposition 3: *Imposing $\bar{a} = 0$ leads to socially excessive investment in content if $r_1 > e_1^*$ and efficient investment otherwise. Neither consumers nor content providers earn any surplus, while the ISP earns profits equal to* $v_1(\max\{r_1, e_1^*\}) - c$.

Proof: Suppose $r_1 > e_1^*$ and the ISP sets $t = v_1(r_1)$ and $a = 0$. The only way that CP1 can make positive sales without suffering a loss is to set $e_1 = r_1$ and $p_1 = 0$. CP2 cannot profitably make sales at any price or quality combination. The ISP would extract all of the surplus, $v_1(r_1) - c$. Inducing CPI to set a lower quality would reduce the ISP's profits, and the ISP cannot induce CP1 to choose a higher quality without lowering a, which would be unprofitable for the ISP given that $r_1 > e_1^*$.

Next, suppose $r_1 \le e_1^*$ and the ISP sets $t = v_1(e_1^*) - e_1^* + r_1$ and $a = 0$. Conditional on its service quality, e_1, the highest price CP1 can charge and still make positive sales is $\hat{p}_1(e_1) = v_1(e_1) - t = v_1(e_1) - r_1 - (v_1(e_1^*) - e_1^*)$. CP1's profits as a function of e_1 are equal to $\hat{p}_1(e_1) - e_1 + r_1 = v_1(e_1) - e_1 - (v_1(e_1^*) - e_1^*)$. CP1's profits are maximized by setting $e_1 = e_1^*$ and earning 0 profits. CP2 cannot profitably make sales at any price. The ISP extracts all of the maximum possible social surplus, $v_1(e_1^*) - e_1^* + r_1 - c$, which clearly is the most it could possibly earn given the rationality of other parties. **QED**

When $r_1 > e_1^*$, social surplus is dissipated by inefficient rent extraction. Intuitively, the equilibrium outcome can entail excessive content quality because there is little linkage between the marginal benefit of increased quality to consumers and the marginal value to the ISP, which forces the content provider's choice. By contrast, suppose that the ISP were required to set $t_i = c$ and $a_i = 0$. Then it is readily shown that the content providers would compete on quality and price, and the equilibrium expenditure levels and prices would be $e_1 = e_1^*$, $e_2 = e_2^*$, $p_1 = v(e_1^*) - v(e_2^*) + e_2^*$, and $p_2 = e_2^*$.

4.2 *Consumer Heterogeneity*

In an important sense, the original model is too simple. Because of the lack of consumer heterogeneity, the ISP needs very few pricing tools to fully extract content provider profits and consumer surplus. This section examines a model with consumer heterogeneity and shows that such heterogeneity creates the possibility that strong net neutrality affects the equilibrium outcome, but may do so in ways that harm welfare.

To explore the effects of consumer heterogeneity, (temporarily) return to the assumption that the CPs' quality levels are fixed and suppose there are two types of consumers: (i) a unit mass with values $v_1 = v_H$ and $v_2 = 0$; and (ii) a λ mass with values $v_1 = 0$ and $v_2 = v_L$, where $0 < v_L < v_H$. Thus, one set of consumers prefers CP1, while the other set prefers CP2, all else equal. λ parameterizes the degree of heterogeneity, with our baseline case corresponding to λ = 0.

Absent net neutrality regulation, the ISP fully squeezes the two content providers by setting $t_1 = v_H$, $t_2 = v_L$, and $a_i = r_i$. In each case, this squeeze will cause content providers to set p_i = 0. Consumers will make the optimal choices between the content providers, and both providers will have positive sales in equilibrium.

Similarly, under both forms of weak net neutrality, the ISP sets discriminatory prices such that consumers choose their preferred content providers but each content provider earns zero profits. The maximum price that can be charged by a content provider is $p_i = v_i - t_i$, which yields profits of $v_i + r_i - (t_i + a_i)$ per consumer. Under weak content provider net neutrality, the ISP will set $t_1 + a = v_H + r_1$ and $t_2 + a = v_L + r_2$, while, under weak consumer net neutrality, the ISP will set $t + a_1 = v_H + r_1$ and $t + a_2 = v_L + r_2$. Regardless, the ISP earns profits of $v_H + r_1 - c + \lambda(v_L + r_2 - c)$, which is what the ISP would earn in the absence of regulation.

With strong net neutrality regulation, the ISP is left with only two pricing instruments: the uniform t and a. In this case, the ISP will sometimes choose to squeeze CP1 at the expense of excluding CP2 from the market. In other cases, the ISP will allow CP2 to make positive sales but will squeeze it to extract all of its profits. In these latter cases, CP1 enjoys strictly positive profits. Observe that consumers earn no surplus in either case, as CP1 sets $p_1 = v_H - t$ and CP2 sets $p_2 = v_L - t$ in any equilibrium in which it makes positive sales. Summarizing these findings,

Proposition 4: *Under no regulation or either form of weak net neutrality, consumers make the optimal choices between the content providers, both providers have positive sales, and the ISP extracts all of the surplus,* $v_H + r_1 - c + \lambda(v_L + r_2 - c)$. *Under strong net neutrality: (i) if* $\frac{v_H + r_1 - (v_L + r_2)}{v_L + r_2} > \lambda$, then *only CP1 and type-1 consumers are active, and the ISP extracts all of the surplus,* $v_H - r_1 - c$; *(ii) if* $\frac{v_H + r_1 - (v_L + r_2)}{v_L + r_2 - c} < \lambda$, *then both content providers and all consumers are* *active, and CP1 earns profits of* $v_H + r_1 - (v_L + r_2)$ *while the ISP earns* $(1+\lambda)(v_L + r_2 - c)$ *; and if* $\frac{v_H + r_1 - (v_L + r_2)}{v_L + r_2} = \lambda$*, then equilibria of the types identified in (i) and (ii) both exist.*

This result is a variant of the standard trade-off inherent with heterogeneous consumers and an inability to engage in price discrimination: the ISP compares setting a high price, which excludes low-value transactions (*i.e.*, setting $t + a = v_H + r_1$ and earning $v_H + r_1 - c$) with setting a low price, which results in greater sales (*i.e.*, setting $t + a = v_L + r_2$ and earning $(1+\lambda)(v_L + r_2 - c)$).

When the proportion of low-value consumers is small (*i.e.*, λ is small), Proposition 4 demonstrates our earlier claim in the discussion of Proposition 1 above that strong net neutrality is a potentially blunt instrument in that it forces the ISP to foreclose on one content provider and its customers completely in order to supply the 'high end' of the market.[6] Thus, niche content providers become under-served under strong net neutrality—something that does not arise in the absence of regulation or under either form of weak net neutrality. In this case, strong net neutrality reduces total surplus—the ISP's profits fall, while consumer surplus and content provider profits remain equal to zero.

Proposition 4 also demonstrates that, when there are sufficiently many low-value consumers relative to high-value ones (*i.e.*, when λ is sufficiently high), strong net neutrality results in CP1's being able to capture some surplus along the lines initially claimed in Gans (2015). Of course, CP2 never captures any surplus despite being preferred by a large proportion of consumers. Thus, the hold-up problem is not eliminated by strong net neutrality, but the ability of the ISP to extract rents from the content providers is diminished.

We next consider dynamic effects. Gans (2015) found that strong net neutrality can enhance content providers' incentives to invest in quality, but will have no effect on ISP investment. We will now demonstrate that only the conclusion with respect to content provider investment incentives extends to the case of consumers with heterogeneous preferences.

We assume that investments are long lived relative to the rate at which the ISP can change it prices. Thus, we model investment by adding a new stage to the beginning of the game: content providers simultaneously make investments, with k_i denoting the investment made by content provider *i*. The game then proceeds as above: the ISP sets its prices; the content

[6] Hermalin and Katz (2007) analyze a model with a continuum of types and derive a similar result in which imposing net neutrality leads to the exclusion of the low end of the market.

providers then set their prices; and, lastly, consumers make purchase decisions. By Proposition 4, neither content provider earns positive quasi-profits (*i.e.*, profits gross of investment expenditures) in the absence of regulation or under either form of weak net neutrality. Hence, in these circumstances, the content providers have no incentives to investment regardless of the regulatory regime. Even under strong net neutrality, CP2 cannot earns positive quasi-profits, and $k_2 = 0$. By contrast, if $\frac{v_H(0)+r_1-(v_L(0)+r_2)}{v_L(0)+r_2} < \lambda$, then there is a range of investment levels such that CP1 earns $v_H(k_1)+r_1-(v_L(0)+r_2)$ under strong net neutrality. Over this range, the increase in CP1's revenues from an increase in its investment level is equal to its marginal contribution to total surplus. Thus, strong net neutrality can partially or—if $\frac{v_H(k_1^*)+r_1-(v_L(0)+r_2)}{v_L(0)+r_2} < \lambda$, where k_1^* the socially optimal investment level—even fully restore CP1's investment incentives if its customers constitute a sufficiently small proportion of all consumers (*i.e.*, λ is sufficiently large).[7]

Now consider net neutrality's effects on investment by the ISP. If one models the ISP's investment as an all-or-nothing entry or exit decision, then net neutrality reduces ISP investment incentives whenever it lowers the ISP's overall profits. However, a model in which the ISP has already entered and is considering incremental investments is also worth examining.

Let K denote the ISP's investment in network quality, and let $v_L(K)$ and $v_H(K)$ denote the dependence of the values of v_L and v_H on that investment (for convenience, the notation suppresses the influence of content-provider investment). We assume that both $v_L(\cdot)$ and $v_H(\cdot)$ are increasing, concave functions, with $v_H(K) > v_L(K)$ for all K. Note that the ISP's investment is not targeted at either group of consumers, although a given change in the common investment level may affect the two types of consumers differently.

By Proposition 4, the ISP appropriates the entire social surplus when there is no regulation, under either form of weak net neutrality, or under strong net neutrality when $\frac{v_H(K)+r_1-(v_L(K)+r_2)}{v_L(K)+r_2} > \lambda$. Hence, in all of these circumstances, the ISP has socially optimal investment

[7] This model, like many in the literature, takes a static view of ISP's incentives with respect to the hold up of content providers. In reality, especially a monopoly ISP would not want to appropriate all of the surplus from content providers when doing so eliminated their incentives to invest in content. Instead, the ISP would like to commit to leaving some surplus with content providers to induce investment that would increase the amount of total surplus, much of which would be appropriated by the ISP. Modeling this type of commitment would require an extension to repeated games which is beyond the scope of the present paper.

incentives regardless of the regulatory regime. Observe, however, that this does not mean that investment is unaffected by regulation. In particular, when strong net neutrality is imposed and $\frac{v_H(K)+r_1-(v_L(K)+r_2)}{v_L(K)+r_2} > \lambda$, the ISP's investment is efficient *conditional on its output decisions,* but is less than the investment that would have occurred absent network neutrality. Absent network neutrality, the ISP's marginal incentive to invest is $v_H'(K)+\lambda v_L'(K)-1$, which is also the social marginal value. By contrast, when strong net neutrality is imposed and $\frac{v_H(K)+r_1-(v_L(K)+r_2)}{v_L(K)+r_2} > \lambda$, the ISP does not derive any benefits from increasing v_L (type-2 consumers and CP2 are inefficiently excluded from the market) and the ISP's marginal investment incentive is only $v_H'(K)-1$.

When strong net neutrality is imposed and $\frac{v_H(K)+r_1-(v_L(K)+r_2)}{v_L(K)+r_2} < \lambda$, both types of consumer are present in the market and consumers that value CP2 service at $v_L(K)$ are the marginal consumers who drive the ISP's pricing. Spence (1975) demonstrated that a monopolist's incentive to invest in product quality is distorted when the marginal customer values incremental quality differently than does the average consumer. Hence, in this case, the Spence distortion arises when $v_H'(K) \neq v_L'$, and strong net neutrality may strengthen or weaken investment incentives depending on whether CP1's or CP2's customers are more sensitive to network quality. By contrast, in the absence of strong net neutrality, the ISP can screen consumers and effectively charge different prices to different consumer types, eliminating the Spence distortion.

In summary, we have shown:

Proposition 5: *In the absence of regulation or under either form of weak net neutrality, the ISP has socially optimal investment incentives. Under strong net neutrality:(i) if* $\frac{v_H(K)+r_1-(v_L(K)+r_2)}{v_L(K)+r_2} > \lambda$, *then the ISP has socially optimal investment incentives conditional on its sales; and (ii) if* $\frac{v_H(K)+r_1-(v_L(K)+r_2)}{v_L(K)+r_2} < \lambda$, *then the ISP's incentives are socially insufficient (excessive) if* $v_L'(K)$ *is less (greater) than* $\Delta_1 \geq C_1$.

5 Conclusion

Using a deliberately simple model, Gans (2015) identified powerful neutrality results (*i.e.,* that net neutrality regulations are often ineffective in terms of influencing market outcomes). In this paper we have corrected and extended this insight. The extensions illustrate the power the pricing-neutrality principle of Gans and King (2003). However, we have also shown that other effects of net neutrality quickly arise as one generalizes the model, making it very difficult to reach general theoretical conclusions about the effects of net neutrality. In particular, in addition to the Gans and King (2003) neutrality principle, there is another general force at work here: limiting the mechanisms through which a monopolist (here, the ISP) can extract rents will weakly lower the monopolist's profits and may also lower equilibrium total surplus as the monopolist is driven to use more inefficient means of rent extraction.

6 References

Choi, J.P., & Kim, B.C. (2010). Net neutrality and investment incentives. *RAND Journal of Economics*, *41*, 446–471.

Gans, J.S. (2015) Weak versus strong net neutrality. *Journal of Regulatory Economics, 47,* 183-200.

Gans, J.S., King, S.P. (2003). The neutrality of interchange fees in payment systems. *Topics in Economic Analysis and Policy, 3* (1), Article 1.

Greenstein, S., M. Peitz and T. Valletti (2016), Net Neutrality: A fast lane to understanding the trade-offs. *Working Paper* No.21950, NBER.

Hermalin, B.E., & Katz, M.L. (2007). The economics of product-line restrictions with an application to the network neutrality debate. *Information Economics and Policy*, *19*, 215–248.

Spence, A.M. (1975) Monopoly, quality, and regulation. *Bell Journal of Economics*, *2,* 417-429.