# Examining Net Neutrality Regulations for Subscription-Based Content

Evan Gee
Department of Economics
Boston University

November 8, 2012

**Abstract**

Many previous investigations of the economic issues surrounding net neutrality have assumed that content providers on the Internet obtain their revenue solely via advertising. While some businesses need not obtain revenue from consumers directly, an enormous number rely on either direct sales or a subscription model. I use a hotelling model of competition to explore how a direct connection between content providers and consumers might change the incentives for all involved. I show that while priority service can increase efficiency, if the Internet Service Provider can charge for priority it has a strong incentive to distort the content providers' market and little incentive to increase investment.

## 1 Introduction

The question of how to route information on the Internet sounds like a computer science problem but the answer has enormous economic implications because it will have a profound effect on what content people consume and how much they are willing to pay for it. This in turn will determine investment not only in new content but also in our infrastructure. The incentive to develop new infrastructure seems particularly important in the US which has fallen from 1st to 12th in the world in terms of average connection speed.[1] While the FCC weighed in on the issue in December 2010 (See FCC *10-201*), the issue is far from settled due to ongoing court battles, legislation, and even an apparent evolution in the FCC's interpretation of its' own rules.

[1] See Akamai Technologies [2012].

Virtually all previous work on the topic, focused on content providers who profited solely from advertising. While this simplifying assumption still covers a great number of areas, a number of the most interesting cases (e.g. Amazon, Netflix, Skype) charge consumers directly for goods and services and this has a significant impact on the incentives for all involved. This actually makes the market more flexible: because consumers directly pay content providers for their services, it is now possible for content providers to change their prices to better compete.

I intend to explore two major issues in this paper. In the short run what kind of prioritization makes the most efficient use of our infrastructure and what kind of incentives does the ISP have to actually implement that priority? In the long run how much infrastructure should we invest in and how does that change with prioritization?

## 1.1 Background

Net Neutrality is, in essence, a broad principle which says that Internet Service Providers (ISPs) should not alter the quality of service based on the origin or type of traffic. For example, under net neutrality an ISP couldn't slow down (or prioritize) traffic to video sites in general, nor to a specific site like Netflix. It does not, however, state that the service should not be altered based on the person requesting said information, so even under Net Neutrality ISPs are free to offer different levels of service to consumers. Schuett [2010] summarizes the issue as follows:

> One reason why the debate about net neutrality is often confusing to outsiders is that the term is used in connection with several distinct concepts. In essence, net neutrality is a theoretical benchmark with which various practices that depart from it are contrasted. The fact that network operators cannot distinguish between packets means, on the one hand, that they cannot determine their origin. It follows that they cannot charge the originator of a packet a fee for transmitting it to users. Thus, net neutrality implies a *zero-price rule*. On the other hand, the fact that network operators cannot distinguish between packets means they cannot discriminate in terms of price or quality of transmission depending on the type, the origin, or the destination of a data packet. Operators cannot engage in traffic management by, e.g., prioritizing traffic, favoring certain packets over others. Thus, net neutrality implies a *non-discrimination rule*.

The arguments for and against net neutrality regulations have been reasonably well established at this point. Those in favor of such regulations argue that ISPs are likely to discriminate in favor of their own services and that so-called 'termination fees' will stifle innovation and growth among content providers. Such fees might be particularly troublesome to new businesses because of the sheer complexity of arranging contracts with dozens of ISPs. Lee and Wu [2009] worry that even if service providers have a free 'slow lane' that they will have an incentive to degrade its quality in order to force content providers to pay.

Those against regulation argue that termination fees would help encourage future investment in bandwidth, bandwidth that might be badly needed in the near future if demand continues to increase. AT&T, for instance, claims that wireless traffic grew 5000% in three years (2006-2009) and in 2010 alone it invested over 23 billion dollars in wireless and wired networks. Furthermore it is not entirely clear that the current structure of the Internet really is a level playing field: large companies can already procure faster access through the backbone and the first-come-first-serve nature of the network may promote congestion-tolerant services like file transfer over real-time services like voice chat.

From the perspective of a governing body there is an entire spectrum of possible regulation ranging from a strict first-come-first-serve rule (net neutrality) to no regulation whatsoever. Somewhat more moderate voices have proposed a system of "equal treatment among similar applications" (see Wu [2003]) where discrimination is allowed as long as it is applied evenly to data of a given type. Finally the FCC's current system that requires transparency and forbids foreclosure or "unreasonable discrimination."(See FCC *10-201*) though more recent public statements by the chairman suggest the FCC is leaning toward a very restrictive interpretation of those rules.

## 1.2 Related Literature

There are several areas of related literature. The idea of charging for priority service has been extensively studied in the nonlinear pricing literature and broadly this paper comes to similar conclusions about the possibility of gains in efficiency.[2] In addition, the structure of the Internet resembles that of both the telephone and telegraph networks, and the concept of net neutrality there dates back to at least the 1800s.[3] Despite the superficial similarities, however, the technology of the Internet raises a number of issues not seen previously.

The literature directly on topic is relatively thin and many of the conclusions either ambiguous or contradictory. Some of the earliest papers on the topic were extensions of the two-sided market literature (see Rochet and Tirole [2003] and Armstrong [2006]). This makes sense to the extent that we see the ISP as a

[2]See Wilson [1993], among others.

[3]See the *Pacific Telegraph Act of 1860*, 36th Congress June, 1860

platform which facilitates interaction between consumers and content providers. In particular, both consumers and content providers generate positive externalities for the other, so in the optimal outcome the ISP may wish to subsidize one of the two sides. Economides and Tag [2009] find that in order for Net Neutrality to be the first-best outcome, content providers need to value additional customers more highly than customers value additional content providers as otherwise the ISP would not want to charge content providers at all. Caves [2010] points out that the parameter restrictions required for this result may not be reasonable, and that under other circumstances consumers may end up providing a large subsidy to content providers.

Musacchio et al. [2009] examine the possible negative externality effects ISPs might have on each other if each is a monopoly only in one region. Since each ISP captures the full benefit of a fee but the cost of a fragmented Internet is borne across all ISPs, each may overcharge content providers compared to the social optimum. The authors show that this may indeed be a serious problem depending on the elasticity of consumer demand and the level of advertising revenue.

More recent papers have grappled directly with the idea that content providers have varying sensitivities to congestion. Papers by Cheng et al. [2011] and Choi and Kim [2010] are mathematically most similar to this paper. Both use a Hotelling model of competition between content providers who subsist entirely on advertising. Cheng et al. [2011] find that doing away with Net Neutrality can be welfare improving if and only if only one content provider ends up with priority. Choi and Kim [2010] do not find a clear relationship between Net Neutrality and social welfare. They both also assume that consumer's utility functions are additive and that traffic to a given content provider is exogenous, restrictions that severely limit the kinds of decisions consumers might make.

Another paper dealing with congestion by Krämer and Wiewiorra [2012] uses a continuum of content providers with different sensitivities where, in equilibrium, only some content providers actually choose to enter the market. They find that doing away with Net Neutrality unambiguously improves welfare in the short run but may give very poor long run investment incentives. Their assumption that content providers don't compete in any sense and that consumers blindly visit every possible website a set number of times may cause their model to badly understate the negative effects of prioritization, however.

A new paper by Economides and Hermalin [2012] uses a screening model to examine how an ISP might want to prioritize a number of different types of content providers. This paper makes the much more reasonable assumption that consumers adjust their behavior based on actual transmission speeds. They conclude that in the short run Net Neutrality is often welfare-maximizing within the restricted set of outcomes that content providers and ISPs can be incentivized to pursue. In the long run (when firms can make long-term investments) the results are less clear and deviations from Net Neutrality may indeed be in the best interest of all parties.

## 2 A Model of Net Neutrality

Consider a model with three kinds of actors: consumers, a monopoly ISP, and two competing content providers. The timing in the model is as follows:

1. Content Providers each set a price for accessing their service.
2. The ISP sets a price for Internet access and decides what kind of prioritization to offer (if any).
3. Consumers decide whether or not to buy Internet access and, if so, which content provider to subscribe to.

Content providers compete with each other for customers based upon both the price for their good and the consumer's preference. $CP_1$ is situated at point 0 on a Hotelling line and sets a price for consumers $p_1$ while $CP_2$ is situated at point 1 on the line and sets a price of $p_2$. I assume that both content providers have an identical cost $mc$ for each additional consumer they serve.[4] Since the distance between the content providers has been arbitrarily set to 1 I assume that the actual transport cost for a distance of $x$ is $\eta x$. Thus the smaller $\eta$ the more directly the two firms compete.[5] Finally, depending on the kind of content that these two providers create their service may be more or less sensitive to congestion. Since these two content providers are assumed to be in direct competition, however, I assume that they have identical sensitivity and model it with a single parameter $\theta$. The timing of the model reflects the fact that content providers are typically national/international in their scope and thus unlikely to adjust their prices based upon the actions of a regional ISP.

The ISP is the sole provider of Internet access in its area. This assumption seems reasonable because even in markets with multiple providers the going price for access is typically far above marginal cost suggesting that competition isn't very intense. This assumption does rule out the possibility of different providers giving priority to different content providers as a way of differentiating their products. In the short run the ISP sets the price $a$ for access, and under Net Neutrality no

[4] With the creation of services like Amazon's S3 it's now possible for even very small content providers to get servers, bandwidth, etc at prices similar to what very big content providers have.

[5] For example, two firms like Amazon and Barnes and Noble selling exactly the same book might have an $\eta$ very close to zero while two firms like Netflix and Hulu have a substantially higher $\eta$ because they offer different TV shows and movies.

contract between the ISP and the content providers is allowed, so this is its only role. Without a strict Net Neutrality regulation the ISP can decide whether or not to offer some kind of priority service and what price to charge for it. In the long run the ISP also makes an investment in infrastructure which determines the average wait time for its services $w$.

Finally, consumers have a number of choices. They may choose to pay the ISPs price $a$ for Internet access or forgo it entirely. If they purchase Internet access they receive utility of $\frac{V}{1+w}$ where V represents their valuation for the Internet and 1+w represents the value lost due to waiting. This multiplicative structure is unusual compared to previous papers but I think more accurately describes the cost to consumers: rather than simply costing time, waiting actively reduces the amount of content someone consumes so the more they value the Internet the more they have to lose.

In addition, each consumer has an ideal location for content $x_i$ along the Hotelling line and consumers are equally distributed. A customer served by his/her ideal content would get a surplus of $\frac{1}{1+\theta w}$ where $\theta$ represents this service's sensitivity to congestion. The transportation cost to available content is $\eta$ for each unit of distance away from ideal content. Thus if a consumer has Internet access, given the Internet conditions (wait time $w$, sensitivity to congestion $\theta$), the prices set by the content providers ($p_1$, $p_2$), and the relative strength of the transport cost ($\eta$), consumers may choose to purchase from $\text{CP}_1$, $\text{CP}_2$, or nothing at all. I assume the services are sufficient substitutes that no one would not want to purchase both. Consumer i's willingness to pay for $CP_1$'s and $CP_2$'s services is respectively

$$WTP_1(x_i) = \frac{1-\eta x_i}{1+\theta w} \quad WTP_2(x_i) = \frac{1-\eta(1-x_i)}{1+\theta w}$$

Again, the multiplicative nature of this structure means that waiting changes not only the maximum possible utility from content but also the utility consumers actually get when transportation costs are taken into account.

Consumer $x_i$'s overall utility is

$$U(x_i) = max \begin{cases} 0 \\ \frac{V}{1+w} - a \\ \frac{V}{1+w} - a + \frac{1-\eta x_i}{1+\theta w} - p_1 \\ \frac{V}{1+w} - a + \frac{1-\eta(1-x_i)}{1+\theta w} - p_2 \end{cases}$$

## 2.1 Nash Equilibrium Under Net Neutrality

In order for content providers to actually be competing for customers in the middle, transportation costs need to be low enough; otherwise, content providers are best off simply charging monopoly prices and ignoring their competitor.

$$\eta \leq \frac{2}{3}(1 - mc(1 + \theta w)) \tag{1}$$

Similarly $V$ needs to be large enough that the ISP won't want to drive customers away just to capture more surplus in this market.

$$V \geq \frac{1 + w}{3}\left(\frac{1}{1 + \theta w} - mc\right) \tag{2}$$

In practical terms this is less strict than $V \geq 1$. Under (1) and (2) it is possible to show that the only Nash Equilibrium is for content providers to charge

$$p_1 = p_2 = \frac{\eta}{1 + \theta w} + mc$$

for the ISP to charge

$$a = \frac{V}{1 + w} + \frac{1 - (3/2)\eta}{1 + \theta w} - mc$$

and for all consumers to purchase Internet access and service from the closest content provider.[6]

Figure 1 shows a typical Net Neutrality outcome. The lower area represents ISP surplus, the middle areas represent content provider surplus, and the upper triangles represent consumer surplus.

# 3 Without Net Neutrality

If net neutrality is not enforced, there are several possible ways for ISPs to consider charging content providers. They might simply set a flat price for priority access, they might charge a rate per gigabyte of traffic, or they might hold an auction and sell priority access only to the highest bidder. Regardless of the way in which the ISP chooses to charge the content providers there are really only two possible outcomes: either the ISP sets a price $g$ at which either content provider can purchase priority over other Internet traffic, or the ISP charges a higher price $g^M$ to just one of the content providers for sole priority access.

I assume that anyone with priority access has a wait time $w^p = 0$ which is a reasonable approximation if only a very small percentage of total traffic has priority. Further let us assume that the wait time for anything that does not have priority is $w^{np} = \beta w$ where $\beta$ is determined by the overall level of traffic used by the content providers and is always $> 1$.

[6] See Appendix A.

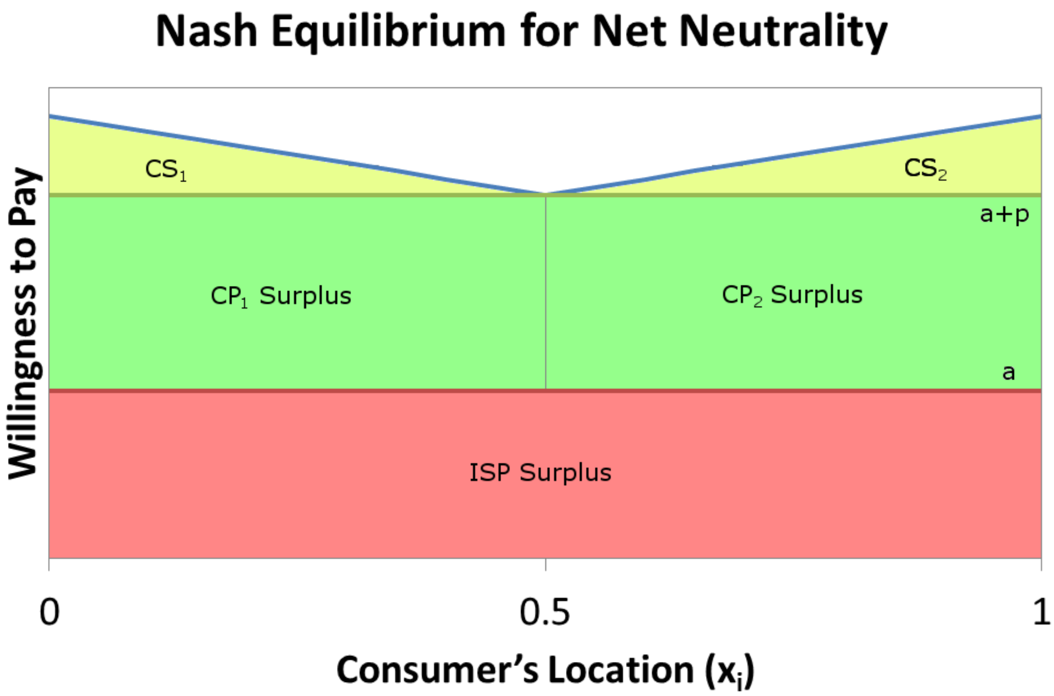


Figure 1: Nash Equilibrium when $V = 5$, $\eta = .55$, $w = .25$, $mc = .1$

### 3.1 Dual Priority

If both content providers pay for priority access the situation is remarkably similar to the one with net neutrality though consumer preferences shift as in Figure 2. With the same assumptions as before the unique Nash Equilibrium again has all consumers purchasing Internet and content from their closest provider with

$$p_1 = p_2 = \eta + mc$$

and

$$a = \frac{V}{1 + \beta w} + 1 - (3/2)\eta - mc$$

Finally the willingness of content providers to pay for priority is limited by the additional amount of profit they receive from consumers for having faster service. Since content providers have identical price, cost, and demand:

$$g = (1/2)\eta - (1/2)\frac{\eta}{1 + \theta w} = (1/2)\eta\frac{\theta w}{1 + \theta w}$$

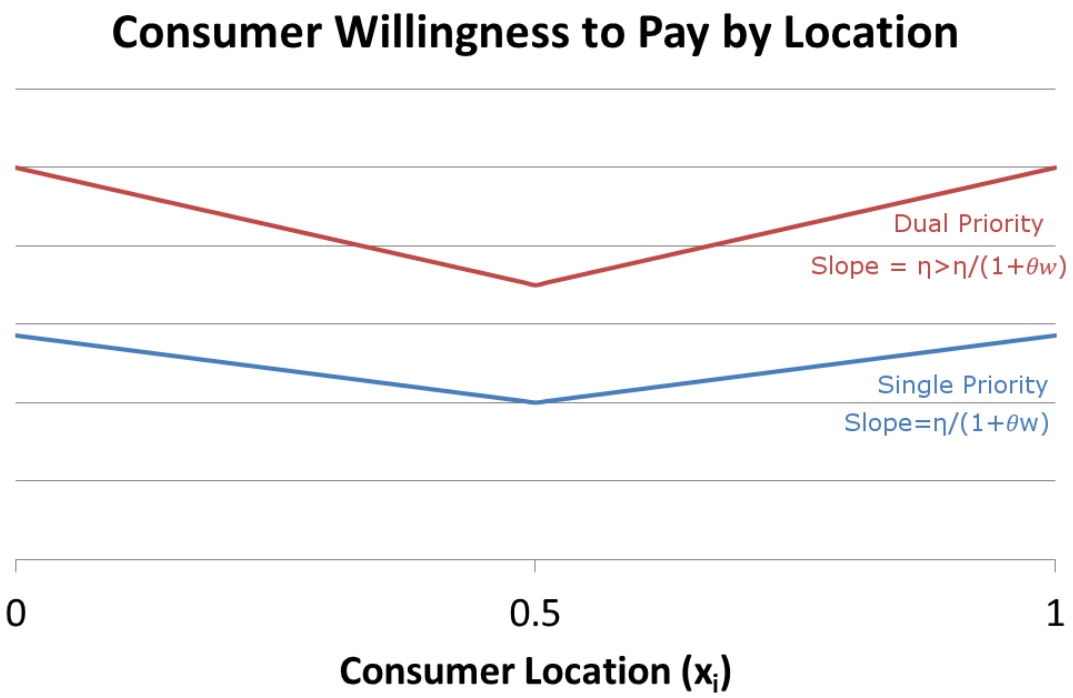


Figure 2: Consumer preferences under net neutrality and dual priority

## 3.2 Comparing Net Neutrality and Dual Priority

To recap:

**Under NN**

$$\bar{x} = 1/2$$
$$p_1 = p_2 = \frac{\eta}{1+\theta w} + mc$$
$$\pi_j = 1/2 \frac{\eta}{1+\theta w}$$
$$\pi_{ISP} = a = \frac{V}{1+w} + \frac{1-(3/2)\eta}{1+\theta w} - mc$$
$$CS = 1/4 \frac{\eta}{1+\theta w}$$
$$TS = \frac{V}{1+w} - mc + \frac{1-(1/4)\eta}{1+\theta w}$$

**Under DP**

$$\bar{x} = 1/2$$
$$p_1 = p_2 = \eta + mc$$
$$\pi_j = (1/2)\eta - g$$
$$g = (1/2)\eta \frac{\theta w}{1+\theta w}$$
$$\pi_{ISP} = \frac{V}{1+\beta w} + 1 - (3/2)\eta - mc + 2g$$
$$CS = (1/4)\eta$$
$$TS = \frac{V}{1+\beta w} - mc + 1 - (1/2)\eta$$

### 3.2.1 Consumers

Despite strictly higher prices for content in the dual priority case consumers as a whole are slightly better off because the ISP can't capture 100% of their new surplus. Recall that waiting reduces the consumer's actual surplus rather than just their maximum possible surplus, so reducing the waiting time actually increases the impact of transportation costs and thus the level of differentiation among consumers. Since the marginal consumer constrains the ISP from increasing the price of access any further, this increased differentiation results in more consumer surplus.

### 3.2.2 Content Providers

Content providers split the market equally in both cases. Despite getting to charge higher prices, the additional price paid for priority to the ISP ensures that they get exactly the profit they would under single priority.

### 3.2.3 ISP

$$\pi_{ISP}^{DP} - \pi_{ISP}^{NN} = V\left[\frac{1}{1+\beta w} - \frac{1}{1+w}\right] + (1 - (1/2)\eta)\frac{\theta w}{1+\theta w}$$

The ISP's change in profit is a bit more complicated than the others. By giving priority to the content providers it makes the rest of the Internet a bit slower reducing the benefit to consumers and thus the price the ISP can charge. Thus the first term of this equation is always negative and depends strongly on the value for the rest of the Internet $V$ and the amount of traffic the content providers require $\beta$. The second term represents the additional surplus that the ISP can extract from content providers for providing priority and is always positive. Note that the second term is decreasing in $\eta$, that is the more differentiated the firms the less interested the ISP will be in providing priority. This happens because again higher transportation costs mean a lower utility for the marginal consumer. The second term is also increasing in $\theta$ and $w$ so the higher the wait time or sensitivity the more the ISP will be interested in providing priority.

### 3.2.4 Social Welfare

$$TS^{DP} - TS^{NN} = V\left[\frac{1}{1+\beta w} - \frac{1}{1+w}\right] + (1 - (1/4)\eta)\frac{\theta w}{1+\theta w}$$

This equation looks a great deal like the one for the ISP. Again, the first part will always be negative and the second part will always be positive. This should make clear that in markets where $\theta$ is high (content providers are highly sensitive to congestion), $V$ is low (the service from content providers is valuable compared to other uses for the Internet), and/or $\beta$ is close to 1 (this service uses little traffic), moving away from Net Neutrality can be socially optimal. (See Figure 3.)

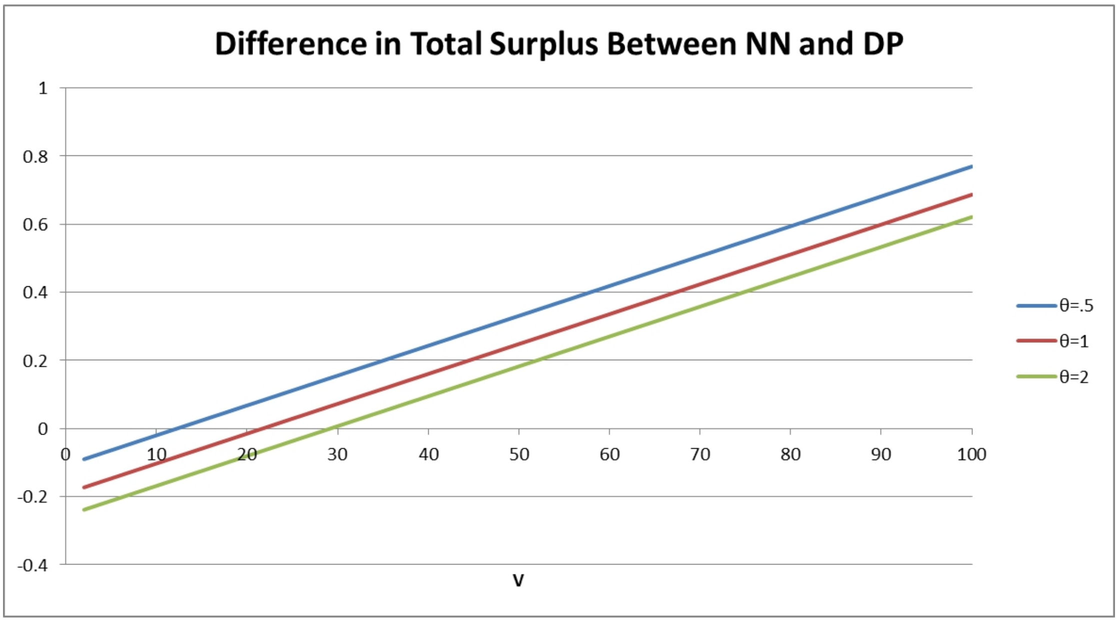


Figure 3: Difference in Total Surplus between net neutrality and dual priority as $V$ changes with $w = .3$, $\beta = 1.05$, $\eta = .7$ and $mc = .1$

In addition note that the benefit to the ISP for moving away from Net Neutrality is similar but will always be lower than for society as a whole. This is true because the ISP pays the full social cost in the form of lower payments from consumers but can not capture 100% of the social benefit from the change because the marginal consumer limits the consumer surplus that can be extracted. Figure 4 shows the $\theta$ required to make dual priority the preferred outcome for a given $V$ for both the ISP and society as a whole. Above the top line both the ISP and society want a dual priority scheme, below the bottom line both want net neutrality, but for values of $V, \theta$ that fall in between society would prefer dual priority but the ISP does not. This leads to Theorem 1.

**Theorem 1** *It may be both in the ISP's best interest and socially optimal to move away from Net Neutrality if the content involved is both sufficiently valuable and sensitive to congestion. Because the ISP bears the full social cost of priority but receives only some of the benefit, however, the ISP's incentive to switch is always smaller than the social incentive. Thus the ISP never implements a dual priority scheme when it is not socially optimal but may fail to implement it when the social benefit is low.*

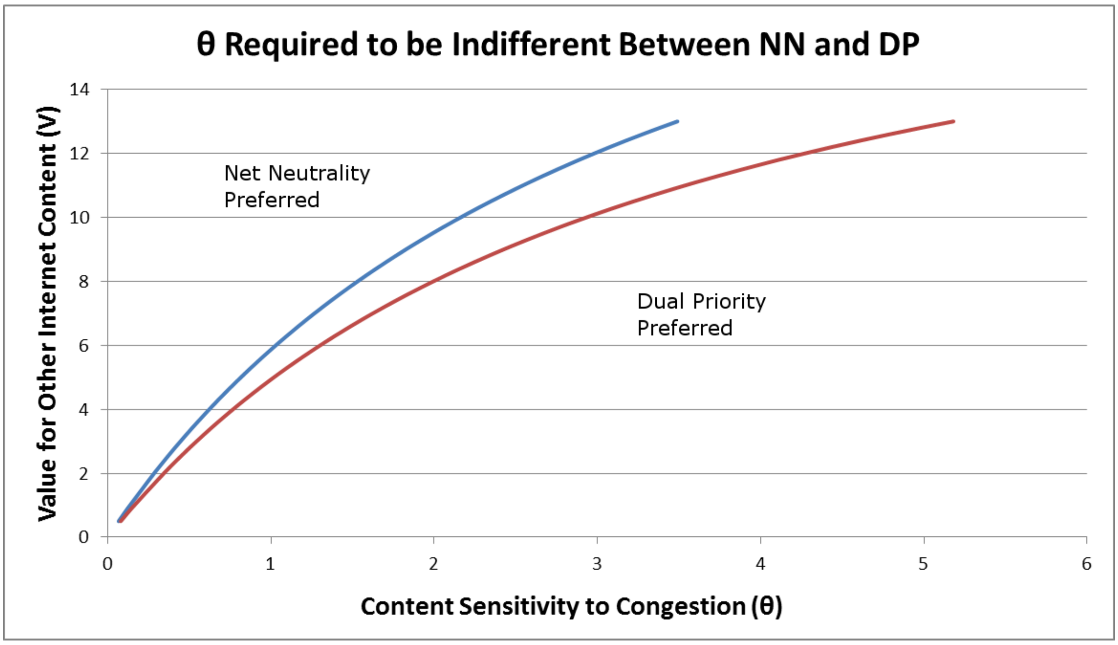


Figure 4: $\theta$ required to make dual priority optimal for a given $V$ when $w = .5$, $\beta = 1.2$, $\eta = 0.5$

### 3.3 Single Priority

If only one content provider gets access (without loss of generality let us assume content provider 1) firms no longer face symmetrical demand. Figure 5 shows the content provider with priority has much higher demand while the content provider without is actively worse off.

The Nash Equilibrium is no longer symmetric. Given $p_1$ and $p_2$ the indifferent consumer has $x_i = \bar{x}$ such that

$$1 - \eta\bar{x} - p_1 = \frac{1 - \eta(1 - \bar{x})}{1 + \beta^S\theta w} - p_2$$

$$\bar{x} = \frac{1 + \beta^S\theta w}{2 + \beta^S\theta w}\frac{1}{\eta}\left[1 + p_2 - p_1 + \frac{\eta - 1}{1 + \beta^S\theta w}\right]$$

Content providers have profits

$$\pi_1 = \bar{x}(p_1 - mc) \quad \pi_2 = (1 - \bar{x})(p_2 - mc)$$

and the first order conditions simplify to

$$2p_1 = 1 + p_2 + mc + \frac{\eta - 1}{1 + \beta^S\theta w}$$

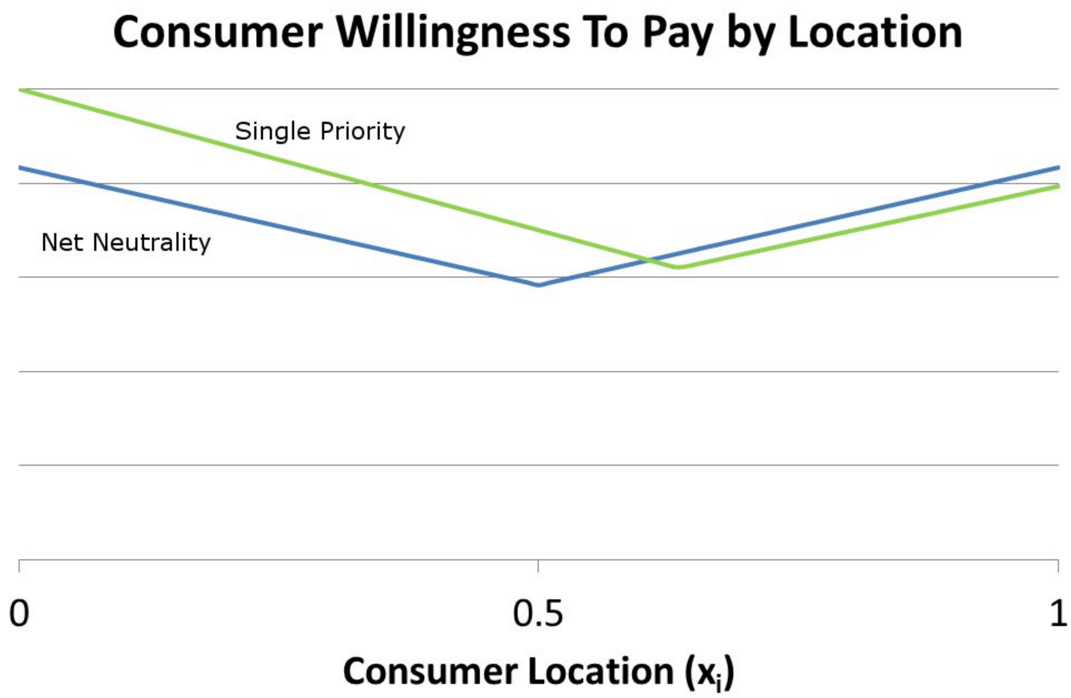


Figure 5: Consumer preferences under net neutrality and single priority

$$2p_2 = mc + p_1 - 1 + \eta + \frac{1}{1 + \beta^S \theta w}$$

which together give

$$p_1 = mc + 1/3 + \frac{\eta + (1/3)\eta\beta^S\theta w - 1/3}{1 + \beta^S \theta w}$$

$$p_2 = mc - 1/3 + \frac{\eta + (2/3)\eta\beta^S\theta w + 1/3}{1 + \beta^S \theta w}$$

The indifferent consumer is thus

$$\bar{x} = \frac{1 + \frac{1+1/\eta}{3}\beta^S\theta w}{2 + \beta^S \theta w}$$

which is always strictly greater than $1/2$ since $\eta < 1$. Figure 6 shows the Nash Equilibrium with consumer surplus (yellow), content provider surplus (green) and ISP surplus (red). The difference between the two green areas will be the price of prioritization $g^S$.

If the ISP is committed to only selling priority access to one CP, then not buying priority means allowing your competitor to buy it. Therefore content providers

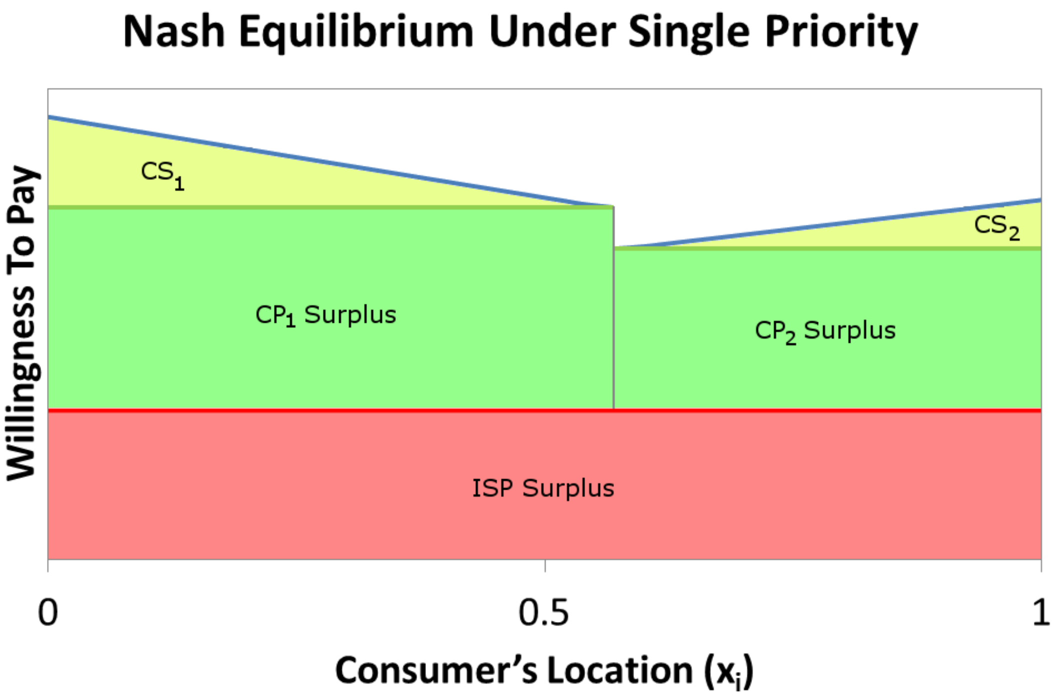


Figure 6: Nash Equilibrium under single priority.

should be willing to pay up to the difference in profit between having access and not having it. The price that makes them exactly indifferent is

$$g^S = \bar{x}^S(p_1 - mc) - (1 - \bar{x}^S)(p_2 - mc)$$

$$\pi_{ISP} = \frac{V}{1 + \beta^S w} + U_1(\bar{x}) - p_1 + g^S$$

Consumers Surplus is

$$CS = {}^1\!/\!{}_2(1 - U_1(\bar{x}))(\bar{x}) + {}^1\!/\!{}_2(\frac{1}{1 + \beta^S \theta w} - U_2(\bar{x}))(1 - \bar{x})$$

Finally Social Welfare is

$$TS = \frac{V}{1 + \beta^S w} + ({}^1\!/\!{}_2)\bar{x}(1 + U_1(\bar{x})) + {}^1\!/\!{}_2(1 - \bar{x})(\frac{1}{1 + \beta^S \theta w} + U_2(\bar{x})) - mc$$

## 3.4 Comparing Net Neutrality with Prioritization

Consumers strictly prefer some kind of priority system to net neutrality but only because it helps to differentiate them and thus limits the amount of surplus that the ISP can extract. Content providers are either no better off or actively worse off under prioritization but again the distinction is somewhat artificial because the ISP is extracting so much surplus. The most pressing comparisons, therefore, are ISP profit and social welfare.

#### 3.4.1 ISP profit

$$\pi^{NN} = \frac{V}{1+w} + \frac{1-(1/2)\eta}{1+\theta w} - mc$$
$$\pi^{DP} = \frac{V}{1+\beta w} + 1 - (1/2)\eta + \eta\frac{\theta w}{1+\theta w} - mc$$
$$\pi^{SP} = \frac{V}{1+\beta^S w} + 1 - \bar{x}\eta - (1-\bar{x})(p_1+p_2) - mc$$

By implementing a prioritization scheme the ISP reduces consumer welfare for other parts of the Internet and thus needs to reduce the price it charges to consumers. Because single priority involves giving priority to less total traffic this effect is less pronounced than under dual priority. The second part of the second and third equation represent the additional profit to the ISP from selling prioritization. When content providers are strongly differentiated ($\eta$ is large) they both charge high prices and as a result the ISP often makes more money from selling priority to both. When content providers are not strongly differentiated and competition is fierce, however, the ISP can often generate a great deal more revenue by offering prioritization to only one firm.

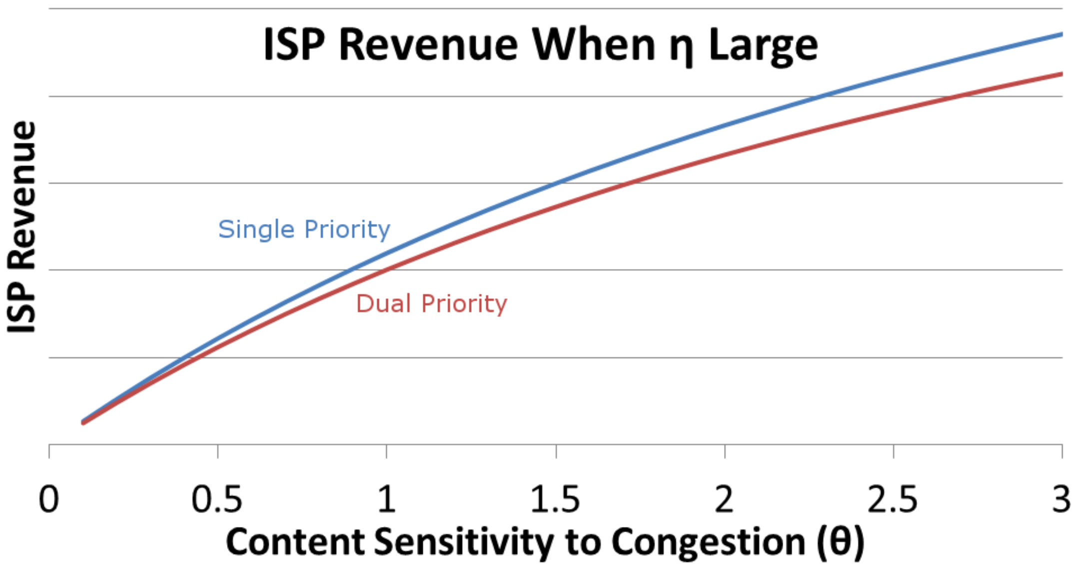


Figure 7: Total revenue from priority when $\eta$ is relatively large.

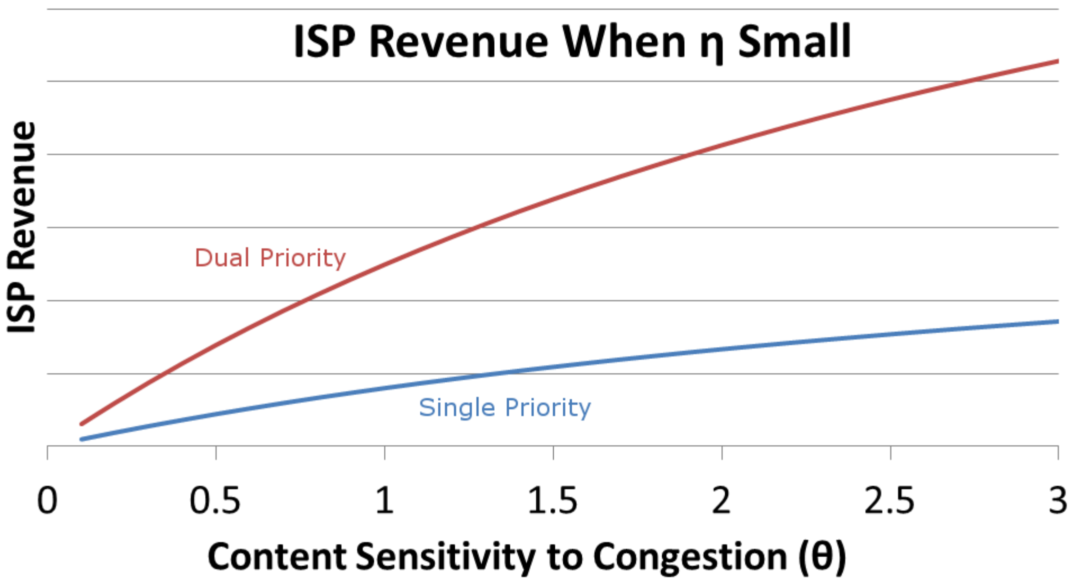


Figure 8: Total revenue from priority when $\eta$ is relatively small.

Depending on the exact parameters any one of the outcomes can be profit-maximizing for the ISP. For given values of $w$, $\beta$, $\beta^S$, $\eta$, and $mc$ the ISP's typical preferences are reflected in Figure 9. Dual priority is preferred for very high values

of $\theta$ (content providers are very sensitive), net neutrality is preferred when $V$ is very large (the current service from content providers is not worth very much), and single priority is preferred for moderate values of both variables.
Larger values of $\eta$ (more differentiation between firms) shift the curves toward each other, reducing the area where the ISP prefers single priority on both sides. A higher wait time increases the ISP's interest in both kinds of priority. Finally the value of $\beta$ and $\beta^S$ have an enormous impact on the overall outcome because they directly impact the cost of giving priority.

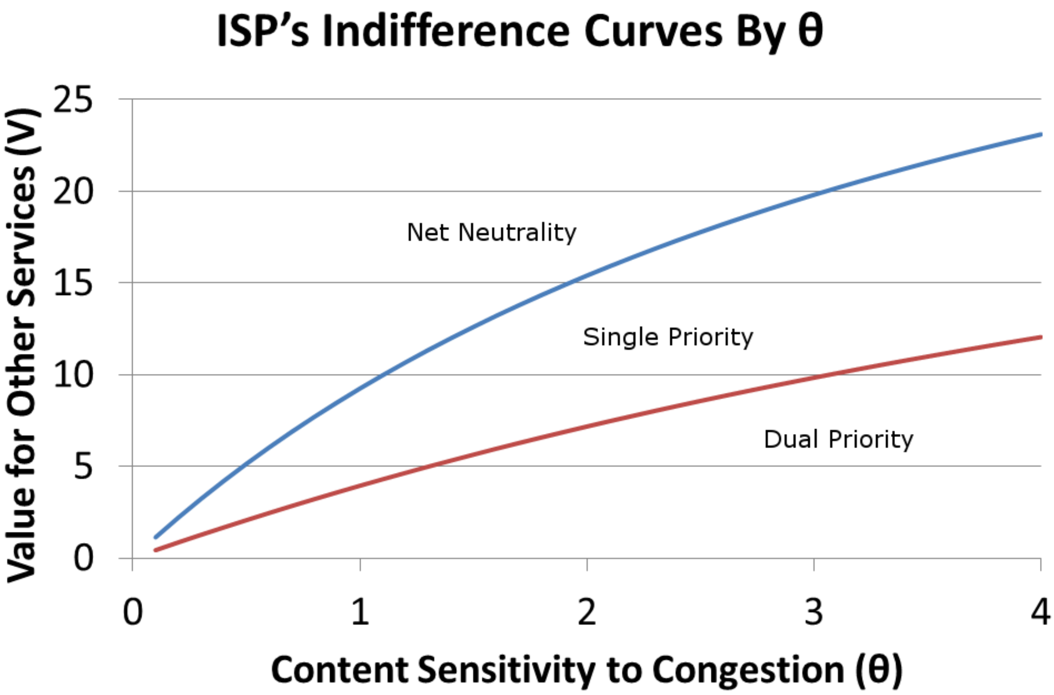


Figure 9: The ISP's preferred outcome by $\theta$ and $V$

### 3.5 Social Welfare

In every case the ISP bears the full cost of giving priority but receives something other than the social benefit of it so it is inevitably true that the ISP's best option is not what society would prefer. Using the same values as above Figure 10 shows the optimal social outcome.

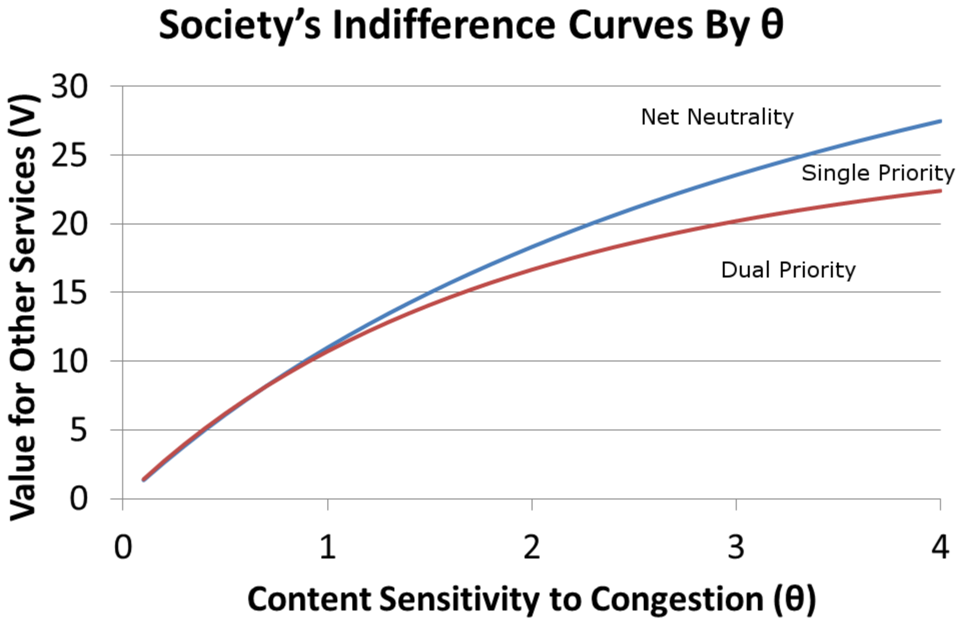


Figure 10: Society's preferred outcome by $\theta$ and $V$

Figure 11 shows a comparison between the socially optimal and ISP-preferred outcomes for a wide variety of $\theta$ and $V$ where green represents net neutrality, yellow represents dual priority, and red represents single priority. For every value of $\eta$ the ISP's preference for single priority is always too strong compared to the socially efficient outcome and as competition between content providers increases ($\eta$ falls) the gap becomes larger and larger.

Finally, Figure 12 represents the optimal choices for both the ISP and society if, as suggested above, the ISP is allowed to offer dual priority service but not allowed to charge for it. The ISP's incentive to give priority traffic is now lower than what society would like, but all parties agree except in borderline cases. As a result this suggests that a policy allowing prioritization could lead to efficiency gains.

**Theorem 2** *Any of the three prioritization schemes (net neutrality, dual priority, or single priority) can be both socially optimal and in the ISP's best interests, but*

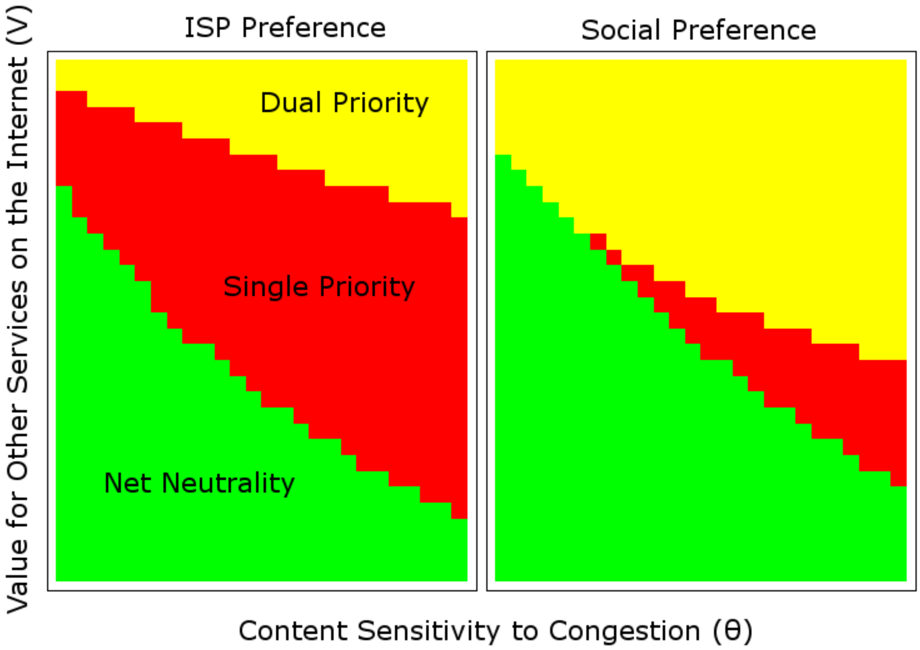


Figure 11: Preferred outcomes for $\theta \in [0.5, 3]$ and $V \in [1, 30]$ where red indicates single priority, yellow indicates dual priority, and green indicates net neutrality.

*the ISP need not prefer the socially optimal decision. In particular if it is allowed to charge for priority the ISP always has too strong an incentive to prioritize traffic. Content providers are never better off under either priority scheme due to the payments that the ISP manages to extract.*

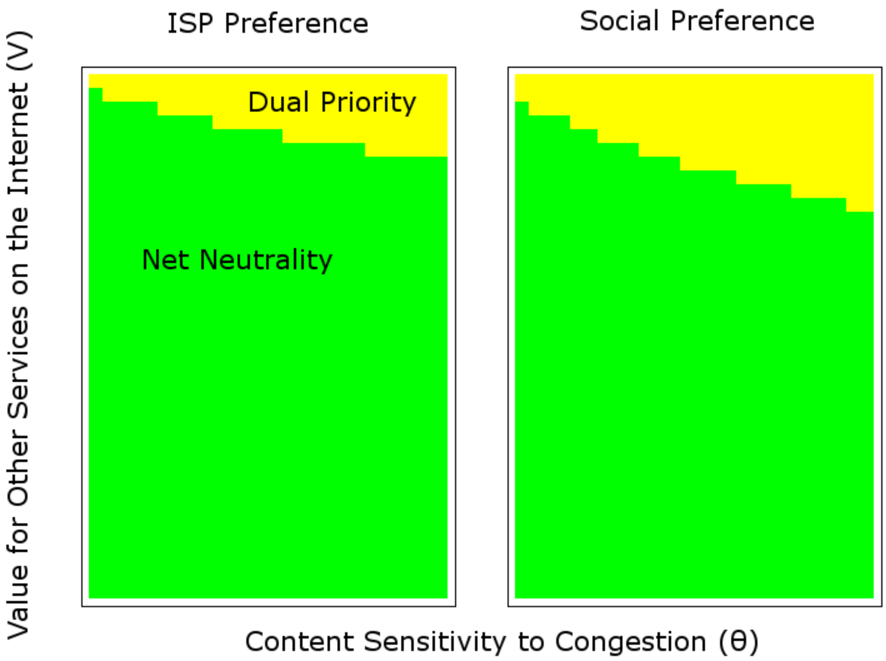


Figure 12: Preferred outcomes for $\theta \in [0.5, 3]$ and $V \in [1, 30]$ where yellow indicates dual priority and green indicates net neutrality.

# 4 Long Run Investment Incentives for the ISP

Until now I have focused on the short run where the question was about how to make the most efficient use of a limited resource (bandwidth). In the long run, though, the ISP has a choice to make about how much to invest and as a result what the average wait time will be for services using its network. Suppose then that the ISP can make a long run investment which sets $w$ by paying $c(w)$ where:

- $c(w) > 0$ and is everywhere differentiable
- $c'(w) < 0 \; \forall \; w$ since a lower w represents more investment
- $\lim_{w \to \infty} c'(w) = 0$
- $\lim_{w \to 0} c'(w) = -\infty$

## 4.1 Net Neutrality

$$\pi_{ISP}^{NN} = \frac{V}{1+w} + \frac{1 - 3/2\eta}{1 + \theta w} - c(w)$$

$$\frac{d\pi_{ISP}^{NN}}{dw} = \frac{-V}{(1+w)^2} - \theta\frac{1-3/2\eta}{(1+\theta w)^2} - c'(w)$$

$$-c'(w) = \frac{V}{(1+w)^2} + \theta\frac{1-3/2\eta}{(1+\theta w)^2}$$

Because a larger w results in a smaller cost, written this way the larger the right hand side the larger the ISP's incentive to invest. Note that social surplus is

$$TS^{NN} = \frac{V}{1+w} + \frac{1-.25\eta}{1+\theta w} - mc - c(w)$$

So the socially optimal level of investment is

$$-c'(w) = \frac{V}{(1+w)^2} + \theta\frac{1-.25\eta}{(1+\theta w)^2}$$

Since $\theta > 0$ and $\eta < 1$ the second term is always positive and thus the ISP's incentive to invest is too low compared to the social optimum.

## 4.2 Prioritization

Under dual priority:

$$\frac{d\pi_{ISP}^{DP}}{dw} = \frac{-\beta V}{(1+\beta w)^2} + (1/2)\eta\left[\frac{\theta}{1+\theta w} - \frac{\theta^2 w}{(1+\theta w)^2}\right] - c'(w)$$

which simplifies a bit to

$$-c'(w) = \frac{V\beta}{(1+\beta w)^2} - 1/2\frac{\eta\theta}{(1+\theta w)^2}$$

Social Surplus is

$$TS^{DP} = \frac{V}{1+\beta w} + 1 - (1/4)\eta - mc - c(w)$$

So the socially optimal level of investment is

$$-c'(w) = \frac{V\beta}{(1+\beta w)^2}$$

The first terms are identical so the difference hinges on the second term. Since $\theta$ and $\eta$ are always positive the second term will again always be negative, so even under dual priority the ISP's incentive to invest is too low.

Under single priority the ISP's investment will be:

$$-c'(w) = \frac{V\beta^S}{(1+\beta^S w)^2} - \frac{d\bar{x}}{dw}(p_1 + p_2 - \eta - 2mc) + (1-\bar{x})\left(\frac{dp_1}{dw} + \frac{dp_2}{dw}\right)$$

The first term represents the full social cost of giving priority while the second and third terms represent the change in ISP revenue from altering the wait time. Because $\frac{d\bar{x}}{dw} > 0$ the second term will always be negative and represents the loss of revenue from content providers: as w decreases priority service becomes less valuable to content providers. Finally the third term is positive for reasonable parameter values and represents the increased prices content providers can charge for better service. The second and third terms are relevant to the ISP's investment decision because the charge for single priority is based on the difference in profit for the two content providers.

## 4.3 Overall Long Run Investment Incentives

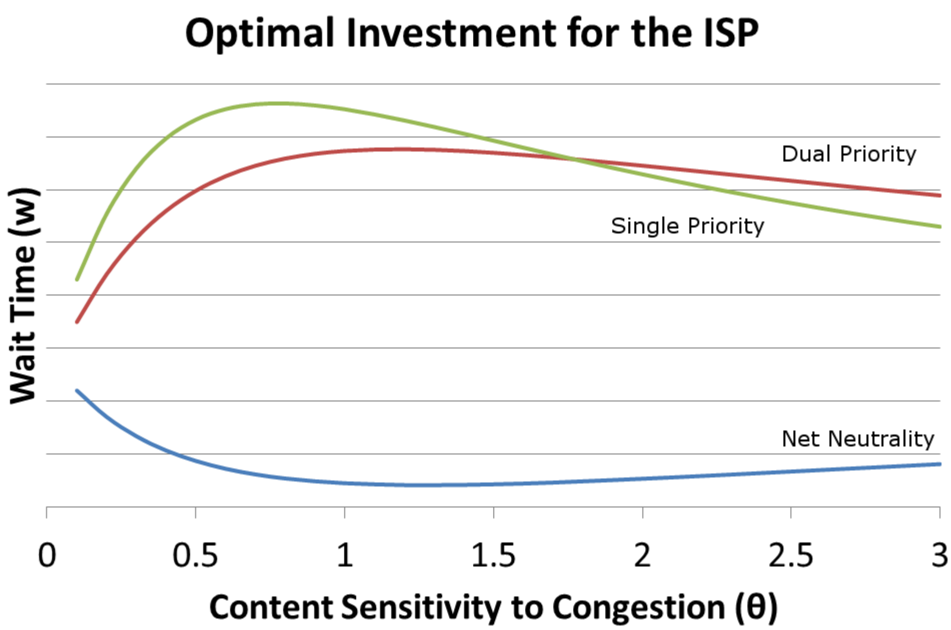


Figure 13: ISP's optimal investment in wait times ($w$) as a function of $\theta$. Higher wait times indicate lower investment.

Figure 13 shows the ISP's optimal choice of the wait time $w$ for each of the different priority options as $\theta$ changes. Since a lower wait time indicates more investment it is clear that both prioritization options result in substantially less investment than net neutrality. This may seem surprising given that a core argument against regulation is that ISPs would use the additional funds to invest more in their networks. Part of the reason for the reduced investment is that higher wait times make priority service more attractive so by reducing investment the ISP can increase fees from content providers. More importantly, though, because the ISP is making more efficient use of its' infrastructure less is required.

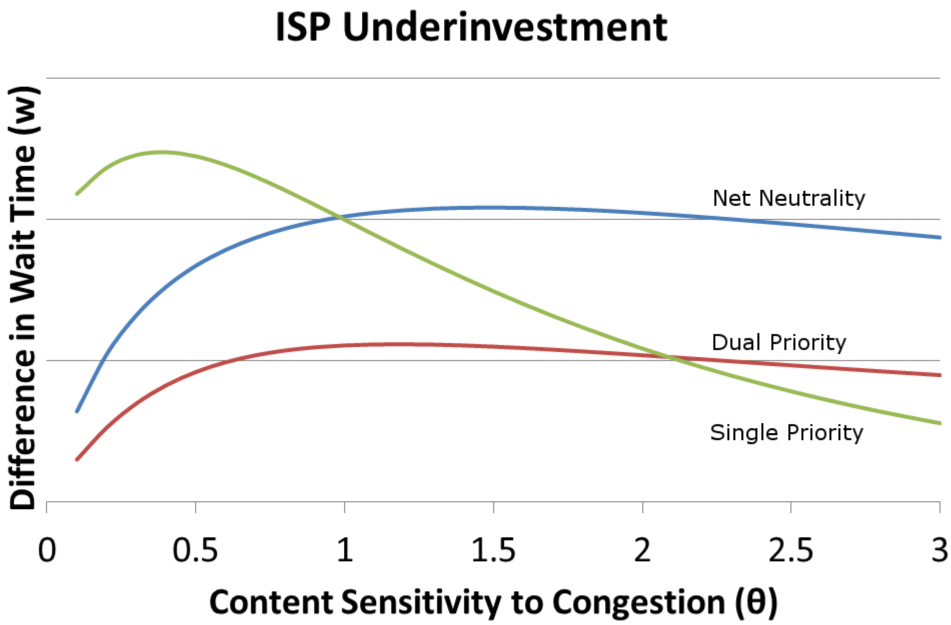


Figure 14: The difference between the ISP's chosen $w$ and society's optimal $w$. Positive values represent too little investment compared to the social optimum.

Figure 14 shows the ISP's optimal level of investment relative to the social optimum for each of the possible prioritization schemes. As I noted above, net neutrality does not give ISPs enough of an incentive to invest in their networks. While neither dual priority nor single priority appear to solve the issue entirely, both prioritization schemes can lead to substantially better results, especially for congestion-sensitive traffic where they are likely to be employed. Thus while the absolute level of investment may be lower, prioritization can result in an outcome much closer to the social optimum.

# 5 Conclusion

Allowing Internet service providers to move away from Net Neutrality has the potential to be socially beneficial by shielding high-value/low-bandwidth traffic from congestion. There is no question that the strict net neutrality approach precludes us from making the most efficient use of our infrastructure, so it would only seem justified if it were the least bad option. It is also clear that the no regulation approach is untenable because while it does give the ISPs an incentive to make more efficient use of bandwidth, it also gives them enormous incentives to distort the upstream market for content. Because the ISP has so much influence the fees it can plausibly extract from content providers are quite large and would likely substantially reduce the incentive for new kinds of content on the Internet. Furthermore I can find no support for the argument that ISPs will make up for it by investing more in their infrastructure.

The two more moderate approaches to regulation are much more interesting. Allowing Internet service providers to charge for prioritization but requiring them to sell it at a consistent price to any interested content provider might reduce the incentive to distort content markets but is likely to create a substantial enforcement challenge because so many service providers are vertically integrated. In addition, while the cost of prioritization is likely to be lower than under no regulation, it will still be substantial and thus decrease the incentive to create new content. Thus the most promising approach is allowing prioritization for different kinds of traffic but prohibiting service providers from charging for it. Despite the zero-price regulation for priority, ISPs would still have a substantial incentive to prioritize high-value traffic because the beneficial effects for consumers allow a higher price for Internet access. While the result may give service providers too little an incentive to implement priority, it captures a substantial portion of the benefits without the serious market distortions or reductions in content investment. Vertical integration still presents some serious regulatory challenges, but unlike trying to monitor the internal payments of a firm, the FCC and a number of consumer groups have already created a number of tests to ensure that service providers are technologically compliant.

# A Nash Equilibrium Under Net Neutrality

**Claim**: Under Net Neutrality the only Nash Equilibrium is for content providers to charge

$$p_1 = p_2 = \frac{\eta}{1+\theta w} + mc$$

for the ISP to charge

$$a = \frac{V}{1+w} + \frac{1-(3/2)\eta}{1+\theta w} - mc$$

and for all consumers to purchase Internet access and service from the closest content provider.
**Proof**:

## A.1 Consumers

Given the prices above all consumers get at least utility 0 from purchasing Internet access and a content subscription. The worst off consumer is $x_i = 0.5$ who has

$$U(0.5) = \frac{V}{1+w} + \frac{1-\eta x_i}{1+\theta w} - a - p_1 = 0$$

By our competition assumption above $p_1 \leq WTP_1(0.5)$ so every consumer with Internet access is strictly better off purchasing content than not.

## A.2 ISP

Given the behavior of consumers and the price set by content providers the ISP's optimal price is

$$a = \frac{V}{1+w} + \frac{1-(3/2)\eta}{1+\theta w} - mc$$

Charging less is clearly less profitable since there are no new customers able to enter the market. If the ISP charges more it instead has profit $\pi = aD(a)$ where demand

$$D(a) = 2\left(\frac{1-(1+\theta w)(a+mc-\frac{V}{1+w})}{\eta} - 1\right)$$

The first order condition on profit gives

$$0 = D(a) + aD'(a) = 2\left(\frac{1-(1+\theta w)(a+mc-\frac{V}{1+w})}{\eta} - 1\right) - a2\left(\frac{1+\theta w}{\eta}\right)$$

It is sufficient to show that this derivative is negative for every $a$ greater than or equal to the equilibrium price. Rearranging a bit gives

$$2a \geq \frac{1-\eta}{1+\theta w} - mc + \frac{V}{1+w}$$

Obviously if this is true for $a = \frac{V}{1+w} + \frac{1-(3/2)\eta}{1+\theta w} - mc$ it will be true for every larger $a$ as well. Plugging that in gives

$$\frac{V}{1+w} - mc + \frac{1-2\eta}{1+\theta w} \geq 0$$

By our competition assumption the upper bound for eta is $\eta \leq (2/3)(1 - mc(1+\theta w))$. Thus in the worst case we need

$$V \geq \frac{1+w}{3}\left(\frac{1}{1+\theta w} - mc\right)$$

## A.3 Content Providers

Given the equilibrium $a$ and $p_2$, if content provider 1 sets a price to $p_1$ it will have demand

$$D_1(p_1) = 1/2 - \frac{(p_1 - p_2)(1+\theta w)}{2\eta}$$

Since profit for content provider 1 is $\pi_1 = D_1(p_1)(p_1 - mc)$ the first order condition gives us

$$\frac{1+\theta w}{2\eta}(p_1 - mc) = 1/2 - \frac{(p_1 - p_2)(1+\theta w)}{2\eta}$$

Simplifying

$$2p_1 = \frac{\eta}{1+\theta w} + mc + p_2$$

Using a similar derivation, content provider 2's optimal response function is

$$2p_2 = \frac{\eta}{1+\theta w} + mc + p_1$$

Which implies that the optimal prices are

$$p_1 = p_2 = \frac{\eta}{1+\theta w} + mc$$

QED

# B Nash Equilibrium Under Dual Priority

**Claim**:Under dual priority it is optimal for content providers to charge

$$p_1 = p_2 = \eta + mc$$

for the ISP to charge an access fee

$$a = \frac{V}{1 + \beta w} + 1 - (3/2)\eta - mc$$

and a prioritization fee

$$g = \frac{\eta}{2} \frac{\theta w}{1 + \theta w}$$

and for all consumers to purchase Internet access and service from the closest content provider.
**Proof**: This situation is almost identical to the one before only the requirement on $V$ is

$$V \geq \frac{1 + \beta w}{3} (1 - mc)$$

which may be stronger or weaker than the previous one depending on the values of $\beta$, $\theta$, and $w$ but is again is typically weaker than $V \geq 1$

# C Nash Equilibrium Under Single Priority

**Claim**:
**Proof**:

# References

M. Armstrong. Competition in two-sided markets. *The RAND Journal of Economics*, 37(3):668–691, 2006.

K. Caves. Modeling the welfare effects of net neutrality regulation: A comment on economides and tag. 2010.

H. Kenneth Cheng, Subhajyoti Bandyopadhyay, and Hong Guo. The debate on net neutrality: A policy perspective. *Information Systems Research*, 22(1):60–82, March 2011.

Jay Pil Choi and Byung-Cheol Kim. Net neutrality and investment incentives. *RAND Journal of Economics*, 41(3):446–571, 2010.

Nicholas Economides and Benjamin E. Hermalin. The economics of network neutrality. *RAND Journal of Economics*, 2012. (Forthcoming).

Nicholas Economides and Joacim Tag. Net neutrality on the internet: A two-sided market analysis. *Working Paper, Stern School of Business*, 2009.

Jan Krämer and Lukas Wiewiorra. Network Neutrality and Congestion Sensitive Content Providers: Implications for Content Variety, Broadband Investment and Regulation. *Information Systems Research*, 2012. doi: 10.1287/isre.1120.0420. published online before print May 7, 2012.

R. Lee and T. Wu. Subsidizing creativity through network design: Zero pricing and net neutrality. 2009.

J. Musacchio, G. Schwartz, and J. Walrand. A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue. *Review of Network Economics*, 8(1):22–39, 2009.

J.C. Rochet and J. Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003.

F. Schuett. Network neutrality: A survey of the economic literature. *Review of Network Economics*, 9(2), 2010.

Robert B. Wilson. *Nonlinear Pricing*. Oxford University Press, 1993.

Tim Wu. Network neutrality broadband discrimination. *Journal of Telecommunications and High Technology*, 2003.