

# IPv4 address market and transition to the next protocol\*

Sarah George <sup>†</sup>

October 02, 2021

## Abstract

In this paper I study the market for internet addresses - a market created to mitigate short term scarcity of IPv4 addresses as firms transitioned to the next protocol. I analyze the inefficiencies that arise from the decentralized transition from IPv4 to IPv6. Inefficiencies arise due to the presence of adjustment costs, switching costs and network effects which are frictions in the market. Towards this end, I develop a dynamic model of firm behavior in the IPv4 market and IPv6 adoption. Then I collect a novel data to estimate the main parameters using simulated method of moments and use this to predict prices and IPv6 adoption. This model can explain the striking increase in prices observed and in the current environment near complete IPv6 adoption will take 33 years. I also find that enterprise firms (universities and companies that are mostly users) have 58% higher per-unit switching costs and content provider firms 24% higher compared to ISPs. In counterfactual simulations, I find that in the absence of switching and adjustment costs prices are lower and IPv6 adoption happens almost 20 years faster. I find that the producer surplus is 88% higher than in the baseline. In the second counterfactual I find that with inter-operable network effects across both protocols firms would adopt IPv6 almost immediately whereas prices reach a peak and fall within three years. Producer surplus in this scenario is lower than the baseline by 1%. In the third counterfactual I find the optimal adoption path that maximizes aggregate producer surplus and find that the surplus is higher by 25%.

---

\*University of Pennsylvania, Department of Economics, sgeorg@sas.upenn.edu

<sup>†</sup>I would especially like to thank Francesco Agostinelli, Juan Camilo Castillo, Amit Gandhi, Ken Moon and Aviv Nevo for their invaluable guidance and support. I thank Minji Bang, Yueyuan Ma, Amandeep Singh, Shunyi Zhao and all seminar participants at the University of Pennsylvania for their valuable comments and suggestions.

# 1 Introduction

A world without internet is unimaginable. Tasks in our daily life, sharing ideas, learning and entertainment have become increasingly coupled with the internet. An essential component for this exchange of data over the internet is an IP address. IP address or Internet Protocol address - is a 'number' assigned to every device connecting to the internet. It helps uniquely identify devices on the internet like a home address helps to uniquely identify each house. The first protocol - IPv4 was allocated from 1984 and was the de facto standard used by all firms. The initial allocation was given at zero price. Under IPv4 around 4 billion addresses were available and this could not be expanded. '172.16.254.1' is an example of an IPv4 address. Given the finite number of addresses and quick growth in the number of internet users it was clear since the late 1980's that the number of addresses available would run out soon.

This led to the development and deployment of the next version of the protocol: IPv6 in 1999. The number of addresses available under IPv6 was much larger with approximately  $3 \times 10^{38}$  addresses. The number of IPv6 addresses can be considered unlimited compared to the 7 billion devices in use today. The main difference between the two protocols is the number of addresses available under each protocol. IPv4 and IPv6 addresses are not directly compatible, and 'translation' is needed for devices on different protocols to be able to communicate with each other. Yet these 'translation' techniques are difficult to scale-up and lead to higher costs and increased latency. Hence these techniques are unlikely to be a permanent solution, thus ruling out the coexistence of IPv4 and IPv6 in the long run.

Firms don't use IP addresses in isolation rather they enable firms and users to connect to a network of other users. Thus, there is an increase in value from holding IP addresses as the majority i.e., there are network effects. The other two key features of an IP address are homogeneity and durability. Within a protocol every IP address is exactly the same and there is no product differentiation. IP addresses doesn't depreciate over time, hence can be considered a durable good.

In 2008 a market for IPv4 addresses was created motivated by the slow pace of IPv6 transition. Cloud service providers, datacenters are the biggest buyers in the market spending millions in each transaction, upto 0.5% their annual revenues. As the number of addresses are finite the constraint is binding and at the current price of \$30/address the market capitalization is over \$100 billion. Lack of inter-operability between IPv4 and IPv6 addresses along with the importance of network effects make this a very important market to study until transition to the next protocol is complete. Prolonged high prices in the market could have negative welfare consequences by preventing the entry of new network services into the internet ecosystem. Given their rapid growth, while cloud service providers are more likely to feel this scarcity in the short run,

this could eventually lead to price pass-through to their customers making this a broader concern for all internet users.

Similar markets are increasingly common in other areas as well. Carbon trading markets have been popular in many places - Europe, Canada and most recently China. These are used to encourage firms to transition to a low carbon economy through the price mechanism. Similarly, zero-emission vehicle credit transfer programs have been introduced in various states in the US to encourage firms to transition towards electric car production. Firms producing electric vehicles receive credit that they can sell to others in the market. There was a similar expectation with the creation of the market for IPv4 addresses - that it would allow firms with lower adoption cost to switch first. This was keeping with the spirit of the internet which grew largely outside the purview of strong regulation. Yet frictions in the market could lead to an excess delay in IPv6 adoption.

In this paper I analyze the inefficiencies that arise from the decentralized transition from IPv4 to IPv6 addresses. My research approach has been data-collection, development of a theoretical model, estimation of parameters and analyzing counterfactuals. This market has been largely understudied because of the lack of access to data. Data collection involved collecting micro data from the IPv4 market, IPv4 and IPv6 allocations over time and firm features for three regions between 1984-2020. I complement this with a novel dataset from a brokerage firm that has about 20% of the market transactions along with the transaction price. Since there is no official data collected about price in the market, this provides information into the evolution of price over time and across firms. Thus, I create novel datasets that will help in the continued study of this market and IPv6 adoption.

Three striking patterns emerge from the data. First, price in the market has been rising continuously since mid-2015. Price is measured as price/address as there is roughly linear pricing in the market - blocks of different sizes sell at the same price/address. Between 2015 and 2020 prices have more than doubled. Second IPv6 adoption rates vary across ISPs, content provider and enterprise firms. Third I provide suggestive evidence of sellers in the market facing adjustment costs. This is based on firms selling their IPv4 address holdings over time in multiple transactions and corroborated through interviews with sellers.

Motivated by the descriptive evidence I develop a dynamic model of firms making decisions in the market for IPv4 addresses and IPv6 adoption decision. Several industry features motivate the model. There are adjustment costs in freeing IPv4 blocks that are in currently use to be made available for sale. The current IPv4 holding of each firm affects their cost of IPv6 adoption which I call switching costs. Without a central marketplace finding a buyer/seller, completing the transfer proceedings is costly and is incurred for every single transaction and I call it the transaction fee. I use a dynamic model because IP addresses are

lasting goods and decisions made by a firm with respect to the IPv4 market and IPv6 adoption decision affect their future payoffs and choices. I abstract away from several features such as some firms also being users, the usage of NAT and CGNAT by ISPs that allow multiple devices to share one IPv4 address (only 17% of firms use it as measured by Richter (2017)) and translation techniques (this is a band-aid solution as explained earlier).

I estimate the parameters related to adjustment cost, switching cost, transaction fee, adoption fee and production parameters using a minimum distance estimator to match simulated moments from the model equilibrium to observed moments. The moments include number of buyer/sellers, fraction of firms adopting IPv6 (by type) and share of total addresses sold over time. The estimated model when simulated forward shows rising prices till \$60 (78% IPv6 adoption threshold) and it will take 33 years to reach almost complete IPv6 adoption. I find that enterprise firms have 58% higher per-unit switching costs and content provider firms 24% higher compared to ISPs.

Using the estimated model, I look at three counterfactual simulations. First, I evaluate the impact of adjustment and switching costs on prices, IPv6 adoption and producer surplus. Prices are lower than in the baseline. Price increases upto \$50 and then starts declining. Upon reaching a threshold of 52% IPv6 adoption prices start declining reaching a minimum in the last period. IPv6 adoption permeates 20 years quicker than in the baseline. Thus, compared to the baseline model prices in the market start declining at a lower threshold and reaches the peak 10 years earlier. The producer surplus is 88% higher than in the baseline model. Second, I evaluate the impact of network effects. With the same network effects across both the protocols, I find that the market adopts IPv6 almost immediately (within three years) with prices reaching a peak of \$60 and then dropping to below \$20. But due to the presence of adjustment costs firms sell their IPv4 holdings over time. The total producer surplus is 1% lower than in the baseline without network effects. Lastly I look at the optimal adoption path which leads to a much quicker adoption, 50% of firms adopting within 16 years. Producer surplus is 25% higher. This indicates that market frictions have a significant effect on the transition towards IPv6 as well as total producer surplus.

The many attempts by various countries to improve IPv6 adoption reflect the policy relevance of this transition. The Chinese government introduced additional regulations mandating 70% IPv6 adoption by end of 2025 which spurred IPv6 adoption in the country. The US military has established similar mandates to achieve 80% adoption. Thus, answers to these questions can help policy makers achieve transition to the next technology in the most efficient way.

## **Related Literature**

This paper relates to and builds on several strands of literature. The seminal work in this market was by Edelman and Schwarz (2015) that predicted declining price paths based on a durable asset over a finite horizon. This is based on the steady-state of the market where there is a last period  $T$  for IPv4 addresses. Building on this work my first contribution is developing a model to study the short-run transition from IPv4 to IPv6 that includes endogenous IPv6 adoption and network effects. With this addition I find that the model predicts prices to rise before reaching the threshold IPv6 adoption and drops from then on. There has been no empirical examination of the market yet. My second contribution is to bring together novel datasets that make empirical analysis possible.

The second strand of literature is the work studying technology adoption with network effects. Empirical work includes Caoui (2019) estimating the costs of standardization in the movie industry, Dubé, Hitsch, and Chintagunta (2010) studying tipping in the hardware industry with indirect software network effects and Ryan and Tucker (2012) studying network effects and heterogeneity in the adoption of a video conferencing technology within a firm. Theoretical work include Farrell and Saloner (1985) studying the inefficient delay in new standard adoption with network effects, Sen et al. (2010) analyzing the role of converters in technology adoption (IPv6 included as an example without IPv4 market), Ostrovsky and Schwarz (2005) studying the effect of uncertainty on standardization and Guérin and Hosanagar (2010) analyzing IPv6 adoption across types of firm based on quality of connection (without the market for IPv4 addresses).

I contribute to this literature in two ways. Unlike previous papers that have tried to study the market for IPv4 addresses or IPv6 adoption separately, I develop a model to study both aspects in a single framework. Second, I study the effect of technology adoption with network effects along with a secondary market. This work sets up the space for further work in areas such as adoption of electric vehicles and the regulatory credits program as well as the carbon trading markets.

The third strand of literature is the economics of internet infrastructure. This is a growing area answering a variety of economic questions within the internet realm. Bauer and Latzer (2016) in their handbook addresses questions related to technology adoption (e.g. fixed and wireless networks have been constantly upgraded), internet search as a two-sided market supported by advertising, online entertainment (related to bundling, pricing, content differentiation by OTT services) and many more. Greenstein (2020) provides a basic overview of the internet infrastructure.

My work draws on the techniques developed by Weintraub et al. (2010) on nonstationary oblivious equilibrium. Given the large number of firms that are present in this ecosystem and therefore the large state space I use non-stationary equilibrium which is line with a firm's knowledge: they know aggregate levels of IPv6 adoption rather than the state space of every firm. There is growing empirical application of

non-stationary oblivious equilibrium: Qi (2013), Buchholz (2016) and Saeedi (2019). In this paper I set the nonstationary oblivious equilibrium for the firm problem within in a dynamic general equilibrium. I develop the dynamic general equilibrium based on previous work in the literature including Humlum (2019), Lee (2005) and, Keane and Wolpin (1997).

The rest of the paper is organized as follows. Section 2 lays out the institutional setting details and hypothesis. Section 3 describes the data and presents descriptive evidence. Sections 4 and 5 details the model, empirical strategy and results. Section 6 analyzes the counterfactuals and section 7 concludes.

## 2 Institutional setting

This section will discuss the institutional settings related to IP addresses including the initial allocation of IP addresses, agents demanding IP addresses, IP address block sizes, IPv6 adoption and dual stacking in IPv6 adoption, market for IPv4 addresses and lastly a hypothesis regarding why the long and short run price diverge.

### 2.1 Initial IP address allocation

The IP address space has been managed by IANA (Internet Assigned Numbers Authority) - a standards organization that oversees global IP address allocation. They manage the allocation and maintenance of IP addresses through five regional RIRs (Regional Internet Registries) managing the task geographically.<sup>1</sup> ARIN (North America), APNIC (Asia) and RIPE (Europe) allow IPv4 address transfers between their organizations creating a single market for IPv4 addresses. In this study, I restrict to this common market for IPv4 addresses which accounts for the majority of the addresses.

### 2.2 Segments that demand IP addresses

Firms are the agents that demand IP addresses from the RIRs. They can choose between the two protocols: IPv4 and IPv6 as well as make decisions on how many addresses to hold on each protocol. They are heterogenous in their role and size. The three main types of firms by role are ISPs, content firms and enterprise firms. Examples of ISPs include Comcast, Verizon etc. ISP firms provide internet access to residential users and business firms by providing IP addresses to large and small users to allow them to

---

<sup>1</sup>The five RIRs in charge of the regions of North America, Asia Pacific, Europe, Latin America and Africa are ARIN, APNIC, RIPE NCC, LACNIC and AFRINIC respectively.

connect to internet while deriving revenue from them. Large users are assigned specific addresses whereas the smaller users are part of a dynamic pool (DHCP - dynamic host control protocol) where they are assigned addresses when they arrive and relieve the addresses once they are done. They also receive revenue from by routing packets of internet data. Content providers usually have their own IP addresses and rely on an ISP firm for internet connection. These firms provide content on the internet which is the main reason users and other firms access internet. Examples of content firms include Google, YouTube, Netflix etc. Enterprise firms include organizations, universities and companies that are mostly users, rather than ISPs, or content. They typically have a local network in addition to accessing internet. Examples include big firms such as HP, Xerox and universities

Firms further vary by their size within these three types. Size of the firm can be understood along two dimensions: cone size and number of users. Cone size is a metric used to measure the size and influence of a firm in the internet infrastructure measured as the number of firms that pay for data transit capturing the revenue the firm generates by routing data packet. Thus, firms with larger cone size can get larger revenues with the same number of IP addresses. The number of users captures the number of customers from whom the firm gets revenue. The demand for IP addresses increases with cone size and users.

Users are passive and care only about internet access (through their ISP) and not the IP version used. Hence, they are not agents in this model. Much of the internet technology developers (router/switch manufacturers, OS developers, internet applications) have already been IPv6 ready for many years thus allowing IPv4 and IPv6 addresses to be competing technologies, thus are not included in the analysis.

### **2.3 IP address block sizes**

IP addresses come in certain standard block sizes with contiguous addresses. The standard block sizes are identified using the slash notation, where a larger number represents smaller block sizes. Looking at IPv4 addresses for example, a /24 (read as slash 24) contains 256 addresses and is typically the smallest size given out. In the table below I include IPv4 block sizes, the number of IPv4 addresses in the block and classification of its size. IPv6 has a similar structure but with much more addresses in each block size.

Slash Notation	No: of IPv4 addresses	Size
/24	256	Small
/23	512	Small
/20	4,096	Small
/19	8192	Medium
/18	16,384	Medium
/16	65536	Medium
/15	131,072	Large
/12	1,048,576	Large
/8	16,777,216	Large

Table 1: IPv4 block sizes

## 2.4 IPv6 adoption - Dual stacking

Firms care about the number of other users and content that is available on either protocol. It has been noted that there are no direct network effects in this ecosystem - meaning firms don't care directly about how many other firms are on IPv6. Rather they care about indirect network effects that exists between firms and number of users. As the number of users increase in either protocol the value of holding address in that protocol increases. The ISPs make the decision about how many users receive IPv4 addresses and how many IPv6 addresses to connect to the internet based on their translation costs. To make the problem tractable without losing the objective of short run transition of firms, I assume that the number of users is an increasing function of number of firms with IPv6.

Due to the presence of this indirect network effects firms adopting IPv6 don't immediately decrease their IPv4 holdings. Rather they hold onto their IPv4 addresses and decrease dependence over time. Dual stacking refers to a firm supporting devices on both IPv4 and IPv6. Thus, by IPv6 adoption I mean that firms are dual stacking. This is similar to 'dual-homing' seen in other industries. Dual stacking is associated with deployment costs and continued higher operational costs as they need to support both the protocols in the interim. Higher operational costs include additional testing, troubleshooting etc. This makes it harder to be the first adopter as well as beneficial to be the last adopter.

There is no empirical measure of how actively firms are using IPv6. Instead, I look at two measures of IPv6 capability that I use to have a sense of IPv6 adoption. First, I look at the firms that have received at least one IPv6 block from an RIR. Second, I look at the firms that have routed at least one of these IPv6 block that they received. Naturally the number of firms with an IPv6 block is higher that the number of firms that have routed at least one of these blocks.



## 2.5 IPv6 adoption decision by a content provider

Here I look at an example of a content provider firm trying to adopt IPv6. In the initial stages of IPv6 adoption a firm looks at its own IPv4 address holdings. If it has enough it will choose to delay IPv6 adoption. A content provider firm facing a shortage of IPv4 addresses will compare the marginal benefit of an additional IPv4 address with its price (marginal cost) compared with marginal benefit of deploying IPv6 vs the marginal cost (which includes cost of deploying). The marginal benefit of an IP address depends on the number of users on the protocol. As the number of users on IPv6 increases the marginal benefit from an additional IPv6 address increases. As the price of IPv4 addresses rises in the market it raises the marginal cost of an additional IPv4 address. Thus, more IPv6 users and higher prices drive firms to adopt IPv6.

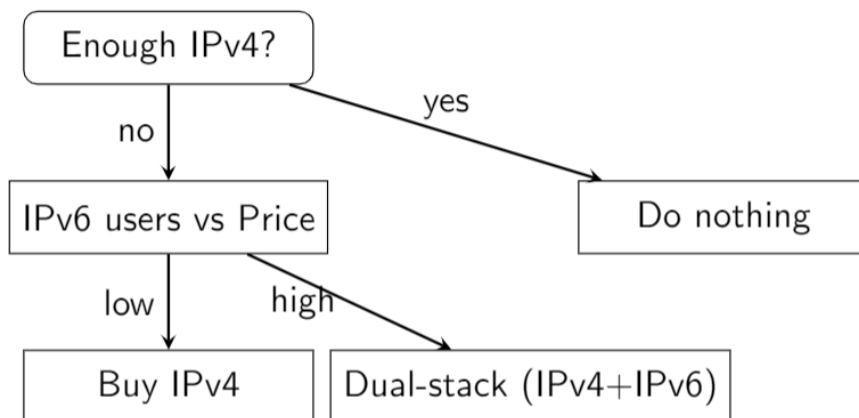


Figure 1: IPv6 adoption by a content provider

0

A similar decision process applies for ISPs and enterprise firms. This captures the main coordination challenge facing firms. It also brings to light the effect of initial allocation (through switching cost) on IPv6 adoption. As IPv6 adoption reaches a certain threshold the strength of IPv4 network effects starts to weaken and firms (even those not facing IPv4 scarcity) start switching towards IPv6. Thus, the industry would tip from the current standard to the next.

## 2.6 Market for IPv4 addresses

Though the market for IPv4 addresses was permitted from 2008, there is no central marketplace where buyers and sellers could turn to. Rather the onus is on firms to find a suitable buyer/seller and complete the transfers. This led to the springing up of organized exchanges and brokerages from 2011. Many of the

brokerage firms sell smaller blocks through online auctions and larger blocks through private negotiations. While firms buying addresses need to show projected demand to prevent speculative trade, there has been limited post-purchase verification. While the RIRs don't use the terminology of buying(buyer)/selling(seller) and instead use the definition of transfer of addresses between firms as they are deemed the final owner of IP addresses that are leased out to firms, for the purpose of clarity in this paper I stick to standard definitions of buyer/seller used in the economics literature.

Much of the buyers in the market have been data centers, web hosting and transit, organizations etc. These include Amazon Technologies, Charter Communications, Cluster Logic, Microsoft etc. Sellers have been firms that were allocated large blocks of addresses in the early years. These include MIT, Northeast Technology solutions, Dupont, Msen etc. There is no official rental market for leasing IPv4 addresses, though handful of firms have started as prices rose.

With finite number of IPv4 addresses (constant supply) prices in the market should be determined by demand. There are two factors affecting demand for IPv4 addresses: growth in number of users and user shift towards IPv6. An increase in the number of users can lead to both increase in demand as well a decrease in demand as some firms move towards IPv6. The net effect depends on the relative strength of the two factors. The price observed in the market is likely to be the net result of the interplay of these two factors.

### **2.6.1 How should prices and IPv6 adoption evolve over time?**

In the long run IPv6 is expected to become the de facto standard. At that time IPv4 addresses have zero fundamental value. In the short run the market starts with a positive price and keeps rising over time. I look into the factors that cause prices to rise in the short run.

Firstly, even though IPv4 addresses have a finite life, it is unclear ex-ante which is the last period as this depends on the evolution of IPv6 adoption. Thus, firms could have different information/expectation about the last period of IPv4 addresses and this creates a challenge of coordination. This prevents prices from falling over time. Rather the life of IPv4 addresses is tied to the pace of IPv6 adoption which is an endogenous choice of firms. In the model I look at an infinite period model where firms make IPv4 market and IPv6 adoption decisions over time. The last period for IPv4 addresses is endogenous and a simple backward induction cannot be used to determine price in any period.

Secondly, there are strong network effects in this ecosystem. Firms prefer using the same version as the majority as these IP addresses aren't used in isolation. Firms adopting IPv6 in the initial years still hold

onto their IPv4 addresses to maintain connectivity on both the protocols. Thus, IPv4 addresses become more valuable in the interim.

Thirdly, firms face switching costs and adjustment costs. Adjustment costs prevent firms from selling all the desired addresses in the market leading to a greater increase in demand. Similarly, the generous allocation of IPv4 addresses in the initial years causes firms to prefer IPv4 over IPv6 that delay IPv6 adoption. I call this as switching costs for firms trying to adopt IPv6. Switching costs are a function of the current IPv4 holdings.

Lastly, firms are heterogeneous in their role and size. This causes their incentives to vary in the market as well as for IPv6 adoption. These frictions have slowed down IPv6 adoption and has subsequently led to a rise in demand for IPv4 addresses. Thus, prices rise in the market in the short run diverging from its long run fundamental value.

### **3 Data sources and descriptive evidence**

This section describes the data used in this paper. There is no single data source that can be used to study this market and thus the source is equally important as this can be used for the continued study of this market. I explain more details about the dataset creation in the appendix. Finally, I highlight the descriptive statistics from the industry using these datasets.

#### **3.1 Data sources**

The first dataset is a confidential dataset from a brokerage firm from 2011 to 2021. It has 20% of the market transfers with price information. It includes data from both the daily online auctions for various sizes as well as data from privately negotiated transactions for larger sizes. It is a comprehensive dataset with information about buyer, seller, address block, date, price (starting, closing) and number of bids. They have transactions from 2011 which includes some of the earliest transactions in the market. While there is some general sense about price trends in this market, the price information in this dataset is unique in that I can see price at the transaction level and connect it with other public data sources.

The second dataset has all IPv4 market transfers from 2008-2021 (ARIN, RIPE and APNIC) excluding price. I restrict attention to these three internet registries as they allow transfers between all three registries creating a single transfer market. The other two agencies don't allow transfers outside their registry. This is

a public dataset with information about buyer, seller, address block and date of transfer at the level of the IP block that was transferred.

The third dataset has all IPv4 and IPv6 allocations (ARIN, RIPE and APNIC) over time (1984-2021). This is a public dataset available across the three different internet registries (ARIN, RIPE and APNIC). It includes information about firm, address block, date of allocation, usage status of block (IPv6). This information was downloaded, cleaned to create a dataset. This is a snapshot of the allocation across firms at a given date. The specific dates I use are included in the appendix.

The fourth dataset is about firm features. These are available publicly but in separate sources. I get information related to the role of firms (ISP/content/enterprise) and their cone size from CAIDA (Center for Applied Internet Data analysis). While the role of firms is available as a dataset to download, I scrape the cone size data using Python. I then get the number of users across each firm from APNIC labs. Together these features give a good picture of the firm characteristics.

A big challenge with these separate datasets has been that the unique identifier differs across the datasets, thus merging them is non-trivial. To get the total IPv4 addresses held by a seller in the initial allocation I look up the transferred IP addresses within the initial allocated block to connect the original owner and then total the addresses. I use an online database to connect the firm level features with the transfer dataset. While majority of the names matched exactly, the rest were merged manually.<sup>2</sup>I similarly connect the other pieces of the dataset. I explain in detail how I overcome this difficulty in the appendix. This allows me to study the whole picture rather than just look at individual slices of the data. Next, I show some key statistics from the IPv4 market over time.

Year	Avg no: of addresses	No: of buyers	No: of sellers	Total addresses transferred
2009	903	4	3	17,152
2010	2,624	3	3	10,496
2011	26,504	43	50	3,233,540
2012	39,030	82	78	6,518,020
2013	23,986	168	149	7,315,710
2014	15,840	504	424	17,044,200
2015	17,820	918	731	51,641,900
2016	10,047	1,526	1,037	31,528,200
2017	17,523	1,759	1,255	60,350,200
2018	12,482	1,836	1,452	46,933,200
2019	7,378	1,918	1,526	27,247,900

Table 2: Summary statistics over time from the market

In table 2 I look at the average number of addresses transferred in the market, number of buyers, number of sellers and total number of addresses transferred in the market over time. In general, I find that

<sup>2</sup>[www.BGPlookingglass.com](http://www.BGPlookingglass.com)

over time there are more buyers, sellers and more blocks transferred over time. I exclude 2020 as it stands out compared to the previous years which could also be due to the Covid-19 that disrupted the market as firms were busy taking care of other priorities. Around the start of the market the market price was low and only a few firms participated in the market. The number of buyers is larger than the number of sellers each period. As scarcity intensified the number of firms buying and selling in the market increases. Though From 2018 the number and volume of transactions started decreasing this could be because of the rapid rise in price.

Next, I look at some features across buyers and sellers in the market. In the table below I look at the size of IPv4 address blocks being transferred, total blocks of IPv4 addresses held by firms selling addresses in the market and the total years the addresses were held by the seller before being sold in the market. I calculate the total blocks of IPv4 blocks as the total blocks owned by firms as of early 2016 and years the addresses were held as of early 2016. Due to data limitations, I choose this period and this is further explained in the appendix.

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
IPv4 address size	20,287	12,871	115,229	256	512	2,048	4,194,304
IPv4 addresses held by sellers	19,510	695,932	4,436,316	256	4,096	95,232	58,615,552
Years held by seller	8,459	14	10	0	5	24	36

Table 3: Summary statistics of firms in the market

Comparing the size of the IPv4 address size in the market with the total block owned it is clear that firms are selling a smaller fraction of their address blocks in the market. The large standard deviation for the total IPv4 blocks held by the firms indicate the difference in the total blocks held among the sellers. There are firms holding the smallest blocks as well as the firm with the largest holding of 58 million blocks which is close to 3.6 times the 16 million blocks allocated that have sold addresses in the market. Looking at the years of holding addresses before selling it in the market, at the 25th percentile firms have held these blocks for 5 years and at the maximum 36 years (which would be one of the earliest firms that received address selling only in 2020).

In the table below I look at three features across all firms holding IPv4 blocks and not just firms selling in the market. This is a cross-section comparison of all firms as of early 2016. I look at the distribution of total blocks held, cone size and users.

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Total IPv4 addresses	67,638	1,206,383	6,782,475	256	1,024	65,792	182,560,256
Cone size	77,468	5	170	1	1	2	26,012
Users	77,468	39,695	1,439,894	1	1	1	292,944,972

Table 4: Summary statistics

Comparing the total blocks of firms selling in the market and all firms across the mean and percentiles it is clear that firms selling addresses are the firms that have received larger address blocks in the initial allocation. The large standard deviation for the total IPv4 blocks held by the firms indicate the difference in the total blocks held by firms. This is larger than among just the sellers. Similarly, the number of users has a lot of variation - it varies from 1 to close to 300 million users. Similarly cone size which is a measure of the transit revenue of the firm varies between 1 and 26,012 with a mean of 5.

I further categorize firms into three groups of ownership based on the total blocks held. Even within these groups there is great inequality in the total IPv4 blocks owned. While majority of the firms belong to the small group of owners, they hold less than 2% of total IPv4 addresses, majority of the address space is owned by firms in the medium category. A small number of firms own about 37% of the total address space.

Statistic	Small	Medium	Large
count	55,364.0	15,470.0	44.0
mean	1,029.6	130,513.1	28,239,057.5
std	1,048.4	649,116.5	33,346,620.0
min	256.0	4,352.0	15,008,000.0
25%	256.0	8,448.0	16,777,664.0
50%	512.0	32,768.0	18,403,840.0
75%	1,024.0	65,536.0	25,660,992.0
max	4,096.0	14,680,064.0	220,997,600.0
Total firms	78.1%	21.8%	0.1%
Total addresses	1.6%	61%	37.4%

Table 5: Distribution of blocks across sizes

In the next section I look at the key descriptive evidence from this ecosystem that will motivate my subsequent modeling choices.

## 3.2 Descriptive evidence

### Price/address

First, I look at the trend in price/address across all sizes over time in figure 2. Prices/address have been consistently rising since late 2015. This also corresponds to the time when the North American internet

registry (ARIN) exhausted allocating IPv4 addresses. From 2015 to 2020 price/address increased from \$7.5 to over \$21, which is more than a 100% increase.




Figure 2: Price/address over time

### Price/address across different block sizes

A similar trend in price/address is observed across different block sizes in figure 3. While there was more dispersion in the initial years across sizes, as the market picked up there is roughly linear pricing in the market. I use this information in modeling price in the market. There is a single price/address and the total purchase price across different blocks is given by: Purchase price = price/address  $\times$  number of addresses in the block




Figure 3: Price/address across different block sizes

*Notes:* I look at the the price/address across four different sizes in the market. These are the most commonly transferred addresses in the market and include small (/24), medium (/16) and large (/15) address blocks.

### IPv6 adoption

In figure 4 on the left panel, I look at the aggregate pace of IPv6 adoption. I look at the two measures of IPv6 adoption - the black line represents the percentage of firms that have received IPv6 address allocations, and the green line represents the percentage of firms that have routed at least one of these IPv6 addresses. Even though IPv6 has been around since 1999, the pace of firms adopting has been slow. In figure 4 on the right panel, I look at the pace of IPv6 adoption across firm roles. Even with firm roles IPv6 adoption varies with content providers leading and enterprise firms lagging. This is line with the initial hypothesis that ISP and content firms are likely to feel the effects from the shift in users to IPv6 stronger than the enterprise firms.



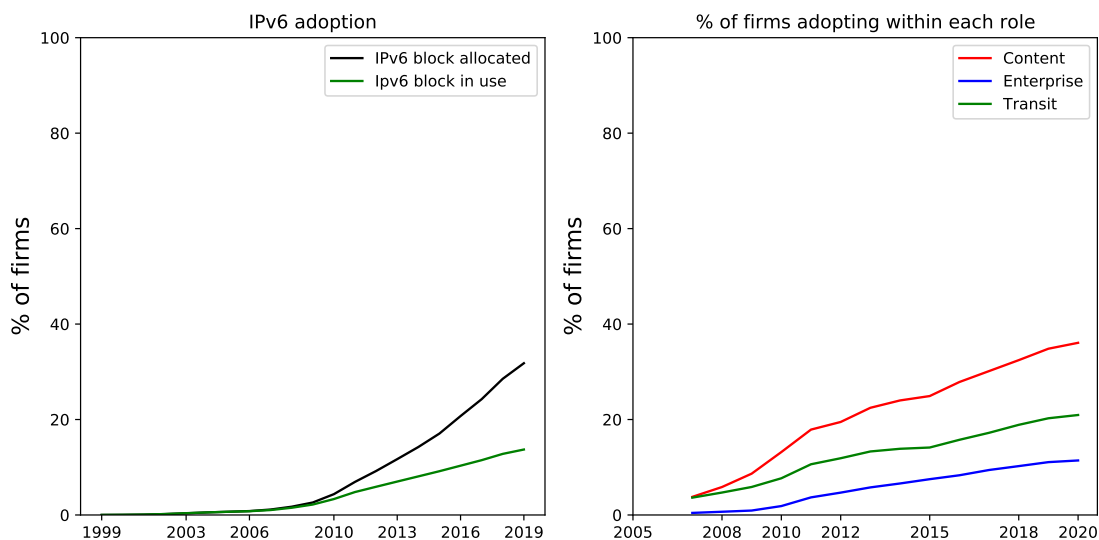


Figure 4: IPv6 adoption over time and across roles

*Notes:* In the left panel I show the rate of IPv6 adoption based on both the measures of IPv6 adoption - firms with IPv6 allocation and firms with routed IPv6 blocks. I use the second measure of IPv6 adoption for the right panel (which is based on firms with routed IPv6 blocks).

### Heterogeneity across firm roles and size

Even within roles firms vary in size across two dimensions: cone size and users. This heterogeneity affects their IPv6 adoption decisions and IPv4 market decisions. In the figure below I look at the histogram of users (logs) and cone size (logs) for firms across different roles. Firms vary in both cone size and number of users. For both cone size and users, ISP firms show the greatest variation, representing the difference in size across rural ISPs versus some of the biggest ISPs serving as the internet backbone. Enterprise firms have the lowest variation along both cone size and users.




Figure 5: Firm heterogeneity across roles

*Notes:* In the left panel I look at the distribution of cone size (in logs) across firms. In the right panel I look at the distribution of users (in logs) across firms. I restrict to the sample with cone size (logs) larger than 1 and number of users larger than 1. 16% of the firms have cone size (equal to 1). 40% of the firms have number of users equal to 1.

### Larger firms adopt IPv6 earlier

In figure 6 I look at a boxplot of cone size in logs (upper panel) and users in logs (lower panel) of firms adopting IPv6 over time. Larger firms adopt IPv6 before smaller firms. By 2014 the interquartile range almost collapses. Based on users while the median falls to the minimum by 2012, the minimum and maximum collapse by 2018-2020. This pattern holds true within firm roles as well.




Figure 6: Firm size of IPv6 adopters over time

### Early firms received larger addresses/user

Next, I provide evidence that early firms received larger IPv4 addresses/user which I define as total addresses/number of users. This is important to look at because firms with higher initial allocation could have higher costs of IPv6 adoption through switching costs. I look at the cross-section across all firms as of 2016. I regress the addresses/user on the internet registry, firm role age of allocation and cone size. I convert users and cone size to log. These variables explain upto 10% of the observed differences in addresses/user. Firms with older allocations have received significantly more addresses/user, whereas cone size has a significant negative effect. Firms in ARIN received significantly more addresses compared to RIPE and APNIC (which is line with the anecdotal experience and corroborates with the large number of sellers from ARIN in the market). Enterprise firms received more addresses/user compared to content and ISP firms, thus reinforcing their delay in IPv6 adoption.

$$\text{Addresses/user}_j = \beta_0 + \beta_1 \text{RIR}_j + \beta_2 \text{Role}_j + \beta_3 \text{Age}_j + \beta_4 \text{Cone size}_j + \epsilon_j$$

### Switching costs

Table 6: Results

<i>Dependent variable:</i>	
Addresses per user	
Age	0.1*** (0.003)
Cone size	-0.3*** (0.01)
ARIN	1.1*** (0.1)
RIPE	0.5*** (0.1)
Content	-1.1*** (0.1)
Enterprise	0.2 (0.2)
Transit	-2.6*** (0.1)
Constant	5.8*** (0.1)
Observations	46,547
R <sup>2</sup>	0.1
Adjusted R <sup>2</sup>	0.1
Residual Std. Error	3.9 (df = 46539)
F Statistic	822.7*** (df = 7; 46539)
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01

Next, I look at the correlation between of IPv4 holdings on IPv6 adoption status of a firm. IPv6 adoption status takes the value 1 if a firm has adopted IPv6 and 0 otherwise. I include the firm role, number of users and cone size in logs and an interaction between firm role and total IPv4 addresses held. I include the interaction term to capture any differences in IPv6 adoption across roles based on their IPv4 holdings. Age refers to the years of holding the IPv4 allocation. I run the following Probit regression:

$$Pr(\text{IPv6 adoption status}_j = 1) = \Phi(\beta_0 + \beta_1 \text{Role}_j + \beta_2 \text{Role}_j \times \text{IPv4 addresses}_j + \beta_3 \text{Users}_j + \beta_4 \text{Cone size}_j + \beta_5 \text{Age}_j)$$

I find that this explains upto 20% of the variation in IPv6 adoption status and most covariates are significant. Age has a negative coefficient, size has a positive coefficient and content firms have larger coefficients than ISP or enterprise firms. Enterprise firms having an additional IPv4 address has a lower probability of IPv6 adoption, whereas for a ISP firm this effect is positive but not significant. This there seems to be differential effect of total IPv4 addresses held and IPv6 adoption across firm roles. This is indicative of the presence of switching costs and that it varies across firm roles.

### IPv4 adjustment costs in the market

The last descriptive evidence I present is the presence of adjustment costs for firms selling in the market for IPv4 addresses. For this I look at the histogram of the number of transactions made by sellers in the market in figure 8. I look at firms in three different categories of IPv4 blocks owned - Small, medium and large.

Table 7: IPv6 adoption

	<i>Dependent variable:</i>
	Adopted IPv6
Enterprise	-0.2*** (0.02)
ISP	-0.2*** (0.02)
IPv4 addresses	0.04*** (0.002)
Users	0.004*** (0.000)
Cone size	0.1*** (0.001)
Age	-0.004*** (0.000)
Enterprise x IPv4 addresses	-0.005** (0.002)
ISP x IPv4 addresses	0.001 (0.002)
(Intercept)	0.1*** (0.02)
Observations	107,204
R <sup>2</sup>	0.2
Adjusted R <sup>2</sup>	0.2
Residual Std. Error	0.4 (df = 107195)
F Statistic	4,457.2*** (df = 8; 107195)

*Note:* \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

Without adjustment firms can sell as many addresses as they want in a single transaction. Adjustment costs limits the number of addresses sold in a single transaction. As a result, firms sell their addresses through multiple transactions. These adjustment costs arise as it takes time and effort for network engineers to identify IPv4 addresses that are currently in use but could be freed-up for sale. This would require the devices currently using those IP addresses to be moved to other addresses. This also corroborates with the interviews with large sellers in the market indicating the presence of adjustment costs.




Figure 7: Histogram of the number of transactions by sellers

The rest of the paper is concerned with developing tools that allows us to study the impact of the different frictions on prices and IPv6 adoption. This requires a model of firm decision making, which is the topic of the next section.

## 4 Model

This section introduces a dynamic model of heterogenous firms making IPv4 holding and IPv6 adoption decisions that explicitly incorporates the above-mentioned factors and defines the equilibrium. Using this model, I estimate the parameters and then use that for counterfactual simulations. Although some model details relate specifically to IP addresses it can be generalized to technology transition of any durable good with a secondary market. In the model, firms decide how many IPv4 and IPv6 addresses to hold each period, firms with only IPv4 decides whether to adopt IPv6 or not in that period. I use a dynamic model as IP addresses are durable goods and their decisions in one period affects their decisions and payoffs in later periods.

Time is discrete and infinite indexed by  $t = 1, 2, \dots, \infty$  where  $t$  stands for each year. The model environment is not stationary because it tracks the short-run evolution of the industry. Agents are heterogenous firms. The number of firms is fixed and indexed by  $j$ . There are three main sources of heterogeneity among firms which are fixed throughout the game. This is captured by  $\tau_j = (\eta_j, \lambda_j, \theta_j)$  representing role, cone size and users respectively. Firms based on their role can be of three types: ISP/content and enterprise,  $\eta_j \in \{\text{ISP}, \text{content}, \text{enterprise}\}$ . I categorize the cone size of firms into three types - small/medium/large,  $\lambda_j \in \{\text{Small}, \text{Medium}, \text{Large}\}$ . Lastly I categorize the number of users as low/high,  $\theta_j \in \{L, H\}$ .

In 2008 when the market started there were already firms that had received IPv6 addresses. So, I assume that there are two types of firms - firms that have already adopted IPv6 and firms that only have IPv4. I model users as passive in this model - they don't make any decisions. Rather a summary measure of total users on IPv4 and IPv6 is used. Users on IPv4 are denoted by  $N_t^4$  and IPv6 by  $N_t^6$  each period.  $N_t$  is assumed exogenous since the increase in number of internet users has nothing to do with the transition and increases both due to increase in the number of users going online for the first time and number of devices per user. For the observed periods I look at the total number of users in the world and then project a linear increase for the future. In the initial period:  $N_0^4 > N_0^6$ . I include the calculation of  $N_t^4, N_t^6$  and the relationship with the state of IPv6 adoption once I define the industry state.

IP addresses held by firm  $j$  in period is represented by  $x_{jt}$  for IPv4 addresses and  $y_{jt}$  for IPv6 addresses. Being durable goods a firm's IP holdings are available for use every period in the future unless it's changed

by the firm in the market. No firm can sell more addresses than they hold in that period.

The sequence of actions available to a firm in period  $t$  is:

- First IPv4-only firms choose between {adopt IPv6, wait}
  - Firms adopting IPv6 receive  $y$  IPv6 addresses in that period
- Next all firms can buy/sell IPv4 addresses in the market
  - Through the market firms can increase/decrease their IPv4 holdings. At the same time firms can change their IPv6 holdings for zero price

I choose this sequence of actions to reflect observed firm behavior. Firms adopting IPv6 progressively increase their IPv6 usage. To capture this I assume that all firms adopting start off with the same level of IPv6 addresses and can increase their IPv6 addresses over time. Similarly, all firms can adjust their IPv4 holdings each period in the market. This captures that IPv6 adoption doesn't automatically lead to decrease of all IPv4 holdings.

IPv6 adoption is assumed irreversible and dual stacking. IPv6 adoption cost is given by:  $A + \gamma_\eta x_{jt}$ . This includes a fixed adoption cost and a linearly increasing cost based on the IPv4 addresses held. The fixed cost captures costs associated with installation of new hardware, software and upgrading legacy applications (Paltridge (2014) and Tassely (2005)). The variable cost with IPv4 holdings captures the labor costs which includes training the IT staff to be proficient in IPv6 and troubleshooting. They would apply IPv6 configuration to routers, firewalls, computers and many other systems. This effort is assumed proportional to the effort in managing IPv4 addresses. Firms with larger IPv4 holdings typically have larger network teams to manage their holdings. Additionally, this also captures the switching costs as firms with larger initial IPv4 addresses have higher inertia to move towards IPv6. I allow  $\gamma_\eta$  to vary across firm roles - ISP/content provider/enterprise. This is to capture differences in switching costs across firm roles which leads to differences in IPv6 adoption across different roles.

There are three components to firm  $j$ 's per period profit. This includes the value a firm gets in every period from holding IP addresses, the cost of holding IP addresses and the revenue change that results from buying/selling addresses in the IPv4 market. By holding  $x_{jt}, y_{jt}$  IPv4 and IPv6 addresses firm  $j$  earn a value each period - which could represent revenue from customers, other firms or even implicit benefit:

$$R(x_t, y_t, \tau, N_t^4, N_t^6, P_t) = \lambda_j \theta_j (N_t^4 x_{jt}^\alpha + N_t^6 y_{jt}^\alpha)$$

Following Edelman and Schwarz (2015) IP addresses are considered inputs into production, with production function of the form:  $f(x) = x^\alpha$ , where  $0 < \alpha < 1$  capturing non-negative and diminishing marginal value from holding IP addresses. To introduce network effects, this is scaled by the number of users available on either protocol. Thus, given the difference in the number of users on both the protocols, the marginal value of an IPv4 and an IPv6 address is no longer the same. Lastly the value is scaled by both cone size and number of users to allow value to vary with firm size (cone size and number of users). This captures that firms with larger size get more value than smaller firms with the same number of addresses.

Firms face a cost of holding addresses each period. This represents the annual fee/leasing costs paid by the firms to the registries and network team expenses for overseeing use of addresses. This cost is assumed linear in the number of addresses given by  $c(x_{jt}, y_{jt}) = c(x_{jt} + y_{jt})$ . I set c at a low value of \$0.01 in line with the actual leasing costs.

Firms can buy and sell IPv4 addresses in the market each period. The net benefit from changing addresses in the market in a given period is:

$$b(x_t, x_{t-1}) = -P_t(x_{jt} - x_{jt-1}) - \phi(x_{jt} - x_{jt-1})^2 1(x_{jt} \leq x_{jt-1}) - F 1(x_{jt} \neq x_{jt-1})$$

With a positive price/address firms selling addresses in the market get extra revenue based on the market price and firms buying in the market have to pay extra.  $P_t$  stands for the IPv4 price/address in period t. IPv6 addresses have zero price as there is no scarcity. F stands for the fixed transaction cost in the market which captures the costs involved in arranging a suitable buyer/seller/broking firm and a clean address block (with no issues such as being blacklist/hijacked address).  $\phi$  stands for the adjustment cost as network engineers must re-optimize their usage before selling. This is a per-unit cost assumed to be convex to capture the increasing costs associated with re-arranging more addresses in a single transaction. This cost is incurred only by sellers. Combining all three parts the per-period profit can be summarized as:

$$\Pi(x_t, x_{t-1}, y_t, \tau, N_t^4, N_t^6) = R(x_{jt}, y_{jt}, \tau, N_t^4, N_t^6) - c(x_{jt}, y_{jt}) + b(x_{jt}, x_{jt-1}) + \epsilon_{jt}(x_{jt}, y_{jt})$$

The state space of firm j is  $k_{jt} = (x_{jt-1}, y_{jt-1}, \tau_j)$ . The industry state can be summarized as  $k_t = (k_{1t}, k_{2t}, \dots, k_{Jt}, N_t^4, N_t^6, P_t)$ . There is no uncertainty in state transition. The initial distribution of firms across IPv4 and IPv4+ IPv6 is obtained from the data. I assume that the number of users on either protocol is a function of the total number of users and the industry state.  $N_{t+1}^4 = g_4(N_t, k_t)$  and  $N_{t+1}^6 = g_6(N_t, k_t)$  where  $k_t$  stands for industry state. Here  $g_4, g_6$  are assumed to be increasing in both k and N. In the empirical



section I assume  $g_4 = N_t \times \text{Fraction of firms with IPv6}$  and  $g_6 = N_t - N_t^4$ . This captures the dependence between IPv6 adoption by firms and the fraction of users that are on IPv6 the next period. As more firms adopt IPv6 the number of users on IPv6 increases, thus increasing the benefit from adopting IPv6 and holding more IPv6 addresses. This is to capture the indirect network effects observed. Firms observe the number of users each period before taking actions.

### Bellman equations

Each firm maximizes its expected discounted profit function yielding value functions that satisfy the Bellman's equation. I drop the  $j$  subscript for ease of reading.

The dynamic programming problem of firms with both IPv4 and IPv6 in the market is given by:

$$V^6(k_t) = \max_{x_{jt}, y_{jt}} \{ \Pi(x_{jt}, y_{jt}, \tau_j, N_t) + \epsilon_{jt}(x, y) + \beta EV^6(k_{t+1}) \}$$

The dynamic programming problem for IP address holding adjustment for IPv4-only firms is given by:

$$V^4(k_t) = \max_{x_{jt}} \{ \Pi(x_{jt}, 0, \tau_j, N_t) + \epsilon_{jt}(x) + \beta EV^4(k_{t+1}) \}$$

The dynamic programming problem for firms deciding whether to adopt IPv6 is given by:

$$\begin{aligned} V^{4a}(k_t) = \max \{ & \Pi(x_{jt}, 0, \tau_j, N_t) + \epsilon_{jt}(0) + \beta EV^4(k_{t+1}), \\ & \Pi(x_{jt}, \bar{y}, \tau_j, N_t) - A + \epsilon_{jt}(1) + \beta EV^6(k_{t+1}) \} \end{aligned} \quad (1)$$

The error terms  $\epsilon_{jt}(x, y)$ ,  $\epsilon_{jt}(1)$  and  $\epsilon_{jt}(0)$  are idiosyncratic shocks that each firm has toward each action in that period. It is assumed to be independent and identically distributed (i.i.d) and follows extreme valued distribution. While each firm takes the price as given in the partial equilibrium, I embed the firm model within a general equilibrium model. This is because the equilibrium price is endogenously determined in the market such that it clears the market each period.

## 4.1 Equilibrium

Given the large state space and interest in the short-run dynamics of firms, I use the nonstationary oblivious equilibrium (NOE) concept when looking at the firm problem. This reduces the dimensionality of the state

space. Weintraub et al. (2010) first introduced the oblivious equilibrium concept which is based on the idea that simultaneous changes in an individual agent's state can be averaged out when there are many firms. This equilibrium can be used to approximate short run dynamics of an industry starting from a given initial state. With a large number of firms and no aggregate shocks the industry state follows a deterministic path, with each firm choosing near-optimal decisions based on its own type, state and the deterministic industry state. Weintraub et al. (2010) show existence results for NOE that converge to stationary strategy as time goes to infinity.

In this ecosystem there are close to 70,000 firms and thus each firm's individual decision in the IPv4 market and IPv6 adoption is unlikely to have a big effect on the industry state and the fraction of IPv6 users the next period. It is also unlikely that each firm knows the exact state space of all the other firms in the industry. I assume firms track the industry state through the fraction of firms with IPv6:  $\bar{h}_t$ . Each firm makes near-optimal decisions based on own state,  $\bar{h}_t$  and  $P_t$ . It makes sense to look at outcomes generated by strategies that become stationary as  $t$  goes to infinity, where the industry converges to a stationary state. This assumption allows me to use a finite horizon for estimation.

Non-stationary oblivious strategies for IPv4-only firms include optimal adoption decision ( $a^{no}$ ) and optimal IP holding decisions in the market ( $x^{4no}$ ). Optimal adoption decision is a sequence  $a^{no} = \{a_1^{no}, a_2^{no}, \dots\}$  where for each period  $t$ , firm  $j$  takes action  $a_t^{no}(\tau_j, k_{jt}) \in \{0, 1\}$ . Similarly optimal IPv4 holdings is a sequence  $x^{4no} = \{x_1^{4no}, x_2^{4no}, \dots\}$  where for each period  $t$ , firm  $j$  takes action  $x_t^{4no}(\tau_j, k_{jt})$ . IPv4+IPv6 firms make optimal IP holding decisions in the market  $\mu^{6no} = (x^{6no}, y^{6no})$ . Optimal holding decision is a sequence  $\mu^{6no} = \{\mu_1^{no}, \mu_2^{no}, \dots\}$  where for each period  $t$ , firm  $j$  takes action  $\mu_t^{no}(\tau_j, k_{jt})$ . Here unlike the Markov perfect equilibrium strategies the NOE strategies are not a function of the whole industry state but are dependent on time.

Firms assume the industry state to depend on time i.e., the expected industry state after  $t$  periods starting from an initial state and NOE strategies played by other firms. Given the NOE strategies of IPv4 and IPv4+IPv6 firms and an initial industry state, the deterministic fraction of firms with IPv6 adoption is defined as  $\bar{h}_t = E_{k_t}[h_t|k_0]$ .

I define ex-ante non-stationary oblivious value functions in period  $t$  when firm  $j$  follows  $a^{\tilde{no}}, x^{\tilde{4no}}, \mu^{\tilde{6no}}$  and all other firms use strategy  $a^{no}, x^{4no}, \mu^{6no}$  as

- For IPv4 firms adopting IPv6:

$$V_{\tau_j,t}^{\tilde{4}a}(k_{jt}|a^{\tilde{n}o}, a^{no}) = E_{\epsilon_{jt}}\{\Pi(x_{jt}, 0, \tau_j, N_t) + \epsilon_{jt}(0) + \beta V^{\tilde{4}a}(k_{t+1})\} \quad (2)$$

- For IPv4 firms in the market:

$$V_{\tau_j,t}^{\tilde{6}}(k_{jt}|a^{\tilde{n}o}, a^{no}) = E_{\epsilon_{jt}}\{\Pi(x_{jt}, y_{jt}, \tau_j, N_t) + \epsilon_{jt}(x, y) + \beta V^{\tilde{6}}(k_{t+1})\} \quad (3)$$

- For IPv4+IPv6 firms in the market:

$$V_{\tau_j,t}^{\tilde{4}}(k_{jt}|a^{\tilde{n}o}, a^{no}) = E_{\epsilon_{jt}}\{\Pi(x_{jt}, 0, \tau_j, N_t) + \epsilon_{jt}(x, y) + \beta V^{\tilde{4}}(k_{t+1})\} \quad (4)$$

Formally an equilibrium in this model is a sequence of price  $\{P\}_t$ , state  $\{\bar{h}\}_t$  NOE strategies  $a^{no}, x^{4no}$  for IPv4 firms and NOE strategies  $\mu^{6no}$  for IPv4+IPv6 firms such that taking  $\{P\}_t, \{\bar{h}\}_t$  as given

1. for IPv4 firms :

- $a^{no}$  maximizes  $V^{\tilde{4}a}$  for all states and time periods
- $x^{4no}$  maximizes  $V^{\tilde{4}}$  for all states and time periods

2. for IPv4 + IPv6 firms:

- $\mu^{6no}$  maximizes  $V^{\tilde{6}}$  for all states and time periods
- $\mu^{6no} = \{(x_1^{6no}, y_1^{6no}), (x_2^{6no}, y_2^{6no}), \dots\}$  represent nonstationary oblivious strategy for firms with IPv4 + IPv6

3.  $\{P\}_{t=1}^T$  clears the market each period

4. The evolution of  $\{\bar{h}_t\}, \{P_t\}$  is consistent with the strategy profile

## 5 Empirical strategy

In this section I present the assumptions and the model primitives used in solving the model, estimation and describe the estimation procedure. I assume that the industry converges to a stationary equilibrium by period  $T$  with probability arbitrarily close to one. In the stationary equilibrium there is industry-wide adoption of IPv6 - all firms have adopted IPv6, all users are on IPv6:  $N_T^4 = 0, N_T^6 = N_T$ , firms don't hold IPv4 addresses  $x_T^* = 0$ , they hold a constant amount of IPv6 addresses -  $y_T^*$  is constant - and price in the market  $P_T = 0$ . This allows me to numerically simulate industry evolution until  $T$  after estimating the parameters.

I use  $T$  equal to 33 years for the main analysis- which means that it would take 33 years from the start of the market in 2008 for the industry to be fully on IPv6. Since the market converges in 33 years the results wouldn't change for  $T$  larger than 33 years. I assume a time period to be one year. The annual discount factor is set to be  $\beta = 0.95$ . I assume that  $T$  equals 33 and then check and find that the industry converges to the stationary equilibrium by period  $T$ . I assume that the number of IPv4 and IPv6 addresses that a firm can hold lies on a grid uniformly between 1 and 10. For prices I use a grid from \$1 to \$70 evenly spaced by \$1. There are 70,000 firms and in the first period I observe that 5000 of these firms are already dual-stacked. The list of parameters to be estimated are market transaction fee  $F$ , adjustment cost  $\phi$ , switching costs for content  $\gamma_c$ , enterprise  $\gamma_e$  and ISP  $\gamma_i$ , fixed IPv6 adoption cost  $A$ , and revenue parameter  $\alpha$ . I summarize all the parameters as  $\theta = (F, \phi, \gamma_c, \gamma_e, \gamma_i, A, \alpha)$ .

To estimate the parameters I solve the model and simulate the moments. Solving the model can be understood in two steps. In the first step I solve the individual firm problem through backward induction for a given price and state. From period  $T$  the stationary equilibrium starts and so the optimal IPv6 holdings and value functions remain constant for a given state and so  $T$  can be considered the last period of the non-stationary equilibrium reducing the problem to a finite horizon. For period  $T$  using the assumptions from above I calculate the optimal IPv6 holdings for each state space (using FOC) and then calculate value function for the state space. Then I solve for the continuation values in  $T-1, T-2, \dots$  back to the first period using backward solution. In the second step I solve the equilibrium price that clears the market each period using a multiple shooting algorithm. I start by guessing a path of prices, then solving the dynamic problem of firms, simulating forward using the firm decisions and then look for the price that clears the market period-by-period. This algorithm is iterated until it converges in the price path and industry state. This algorithm is adapted from Lee (2005) and Humlum (2019). I relegate the algorithmic details of computing and simulating the model to the appendix.

## Moments used

Identification of the model parameters relies heavily on the moments chosen. Looking at the moments I find that they are sensitive to the structural parameters. The market transaction fee  $F$  is identified primarily through the number of buyers and sellers in the market each period. Higher values of  $F$  lead to smaller participation and lower values lead to higher participation. I also include the growth rate of prices over time to match the observed growth rate in prices. To identify the differential per-unit switching cost parameter across different firm roles I include the fraction of firms of different roles adopting IPv6 each time period. As the switching cost increases IPv6 adoption for that firm role would decrease and using the fraction of firms adopting IPv6 each period. As the fixed adoption cost increases the fraction of firms adopting IPv6 decreases. Thus, the fraction of firms adopting IPv6 along with the fraction of firms adopting IPv6 based on firms role identifies the fixed cost of adoption. The adjustment cost is primarily identified from the share of total addresses sold by a seller in the market. If the per unit adjustment cost increases, it becomes costlier to sell addresses and firms sell a smaller share of their holdings.

Overall, I estimate the 7 parameters using a combination of simple moments shown in the table below. I observe 10 values for all the moments (as its over time - 10 periods) except for the growth rate of prices for which I have 9 moments. Even though I simulate and solve the model for all 33 periods I construct moments from the first 10 periods that are observed in the data. In total I use the following 79 moments:

Moments	Number of moments
No: of buyers	$1 \times 10$
No: of sellers	$1 \times 10$
Share of total addresses sold	$1 \times 10$
Growth rate in prices	$1 \times 9$
Fraction of content w/ ipv6	$1 \times 10$
Fraction of enterprise w/ ipv6	$1 \times 10$
Fraction of isp w/ ipv6	$1 \times 10$
Fraction of firms adopting ipv6	$1 \times 10$

Table 8: Moments used for estimation

To estimate the parameters, I use an extensive grid search to initialize parameters. This is due to the possibility of being stuck in a local minima. In the second step I use simulated method of moments to estimate the parameters in the model. I use this estimation strategy and not a two-step estimation method due to data limitations. For each value of  $\theta$  I solve the equilibrium and simulate the model 1000 times for 33 periods each from the initial state. The SMM estimator is the vector of parameters that minimizes the distance between the simulated moments from the model and the empirical moments from the data,

$$\hat{\theta} = \arg \min_{\theta \in \Theta} [m(\theta) - m]'W[m(\theta) - m]$$

where  $m$  is the observed moments,  $m(\theta)$  is the simulated moments and  $W$  is the weight matrix. In the first step I set  $W$  equal to the identity matrix and calculate  $W$  again based on the first stage estimates and calculate the estimates again using the new  $W$  matrix. Following Bloom (2009) I use numerical derivatives to calculate standard errors for  $\hat{\theta}$ . I use a simulated annealing algorithm to account for potential discontinuities in the model and its moments.

This estimator is like the one referred to Hall and Rust (2003) as a simulated minimum distance (SMD) which minimizes a weighted distance between actual and simulated moments. Further this estimator belongs to the group of generalized method of moments (GMM) estimators introduced by Hansen (1982) and Pakes and Pollard (1989). Once I estimate the parameters, I calculate the producer surplus as the sum of the variable profits for each firm minus the fixed adoption cost and fixed transaction fee over every time period.

## 5.1 Estimation Results

The parameter estimates and their standard errors are presented in table 2. All the parameters (other than  $\alpha$ ) are statistically significant. Per-unit adjustment and switching costs are large compared to the market transaction fee and fixed IPv6 adoption cost. ISP firms have the lowest per-unit switching cost, enterprise firms have almost 58% higher per-unit switching costs compared to ISP firms and content firms have 24% higher per-unit switching costs compared to ISP firms, thus able to quantify the observed differential IPv6 adoption across firm roles.

Parameter	Estimated value	Standard error
IPv4 market transaction fee, $F$	10.0240	1.3901
IPv4 adjustment cost, $\phi$	8.0313	1.2094
Switching costs for content $\gamma_c$	4.0365	1.9308
Switching costs for enterprise $\gamma_e$	5.1326	1.1204
Switching costs for ISP $\gamma_i$	3.2416	2.8976
Fixed IPv6 adoption cost, $A$	34.0313	13.5386
Revenue parameter, $\alpha$	0.4335	12.6906

Table 9: Parameter estimates

There is no benchmark to compare these estimates with yet. Next, I simulate the model till  $T$  using the estimated parameters and look at the model's fit in figure 9.

## 5.2 Simulation results using estimated parameters

Here I plot the evolution of prices and IPv6 adoption using the estimated parameters compared with that observed in the data. The red line represents the actual data, and the blue line shows the model prediction.

In general, the model fits the basic trend in price and lag in IPv6 adoption well. The price pattern in the initial years follows the actual prices observed so far. Based on this price is expected to increase reaching a peak of \$60, but only slowly come down even with rising IPv6 adoption, reaching its minimum only by the last period. It takes almost 33 years for almost complete IPv6 adoption. The model is unable to capture the gradual increase in IPv6 adoption. Rather in the model IPv6 adoption stays relatively flat before starting to increase sharply. I calculate the total producer surplus to be \$35.25 million.




Figure 8: Model fit

## 6 Counterfactual Analyses

In this section I contrast the estimated model with the following three exercises. The first is to study the evolution of price and IPv6 adoption without any adjustment cost and switching cost. The second is to study the evolution of price and IPV6 adoption in the absence of network effects.

### 6.1 How would prices and IPv6 adoption evolve without adjustment and switching costs?

With the estimated parameters I approach the main question of the paper: What is the effect of market frictions on prices and IPv6 adoption? To answer this, I set the two main sources of friction - adjustment costs and switching costs - to be zero in the market. This means that firms can sell excess IPv4 blocks without any adjustment cost. As a result, more firms sell in the market leading to lower prices overall. At the same time without switching costs more firms can adopt IPv6 earlier, leading to overall faster adoption. Thus prices increase reaching a peak of \$50 and then fall gradually. Prices reach a peak when IPv6 adoption rate reaches 60%. Since network effects are still present it takes about 13 years for IPv6 adoption to permeate. Thus, IPv6 adoption happens 20 years faster than that predicted in the baseline. Calculating the producer surplus, I find it to be 65.40 million dollars, which is 85% higher than the initial prediction.



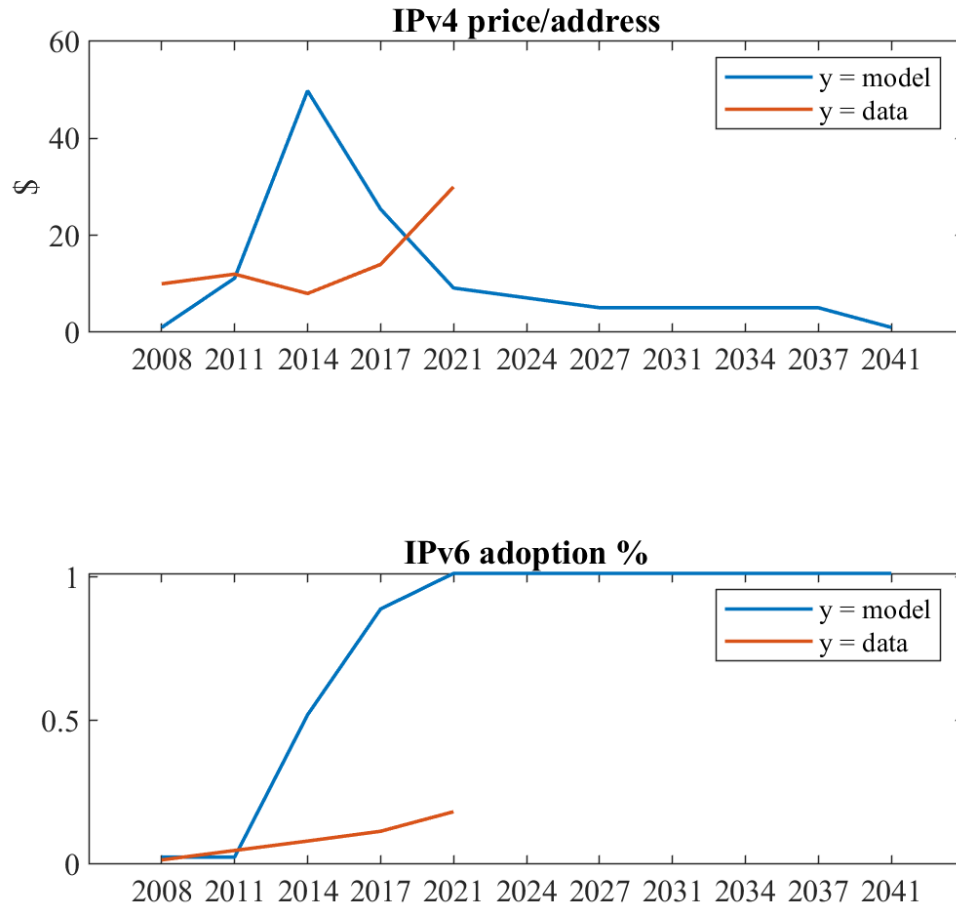


Figure 9: Counterfactual: Without adjustment and switching costs

## 6.2 How would prices and IPv6 adoption evolve with inter-operable network effects?

In the second counterfactual I let the network benefits that firms receive from both IPv4 and IPv6 addresses to be the same, while still allowing there to be all other frictions such as switching and adjustment costs. I do this by replacing the number of users on each protocol in profit function with the number of users

each period. Unlike in the baseline model where the network effects from that protocol was available only within that protocol based on the number of users in that protocol, now I assume that firms can get the same network effects (based on the total number of users) from having either address. In this case I find that there is no initial lag in IPv6 adoption by firms. IPv6 adoption happens almost immediately. This implies that network effects play a strong role in preventing the take-up of IPv6. Thus, compared to the baseline prediction IPv6 adoption happens almost 30 years earlier. Demand for IPv4 addresses keep prices high as adjustment costs are still present and prices increase due to the rising number of users. Calculating the producer surplus I find it to be \$35.17 million dollars. This is 1% lower than the one in the prediction.




Figure 10: Counterfactual: Inter-operable network effects

### 6.3 How does the decentralized adoption compare to the optimal adoption path?

To find the optimal adoption path I look at a social planner maximizes aggregate producer surplus taking as given the evolution of number of users. By converting to IPv6, firms increase the number of users on IPv6. This increased number of users on IPv6 is a positive externality on other firms. The increased number

of IPv6 users raises the option value from IPv6. This increased profits and option values are not internalized by firms adopting in a given period and thus industry profits are not maximized under the decentralized outcome.

In this benchmark, the social planner chooses a sequence of adoption decisions and IPv4, IPv6 holdings for each firm to maximize the producer surplus. This is a computationally hard problem to solve because of the large number of players in the industry. Since the objective of interest is aggregate producer surplus and the aggregate adoption path, I search for the highest aggregate producer surplus over the space of various adoption paths. For this, I search over the space of random adoption paths and select the path that maximizes the aggregate producer surplus. The paths are simulated by adding firm specific perturbations to their IPv4 adoption decisions. Each path lengths is 33 years. I draw 15,000 perturbations and average over the perturbations. Figure 11 shows the set of simulated paths of adoption and figure 12 shows a comparison of optimal adoption path by the social planner and the pace of IPv6 adoption observed in the data. The aggregate producer surplus is 25% higher than in the baseline. The optimal adoption path is almost linear, but faster than that observed in the data. In the optimal adoption path - within 16 years, more than 50% of the firms have adopted IPv6.




Figure 11: Feasible adoption paths

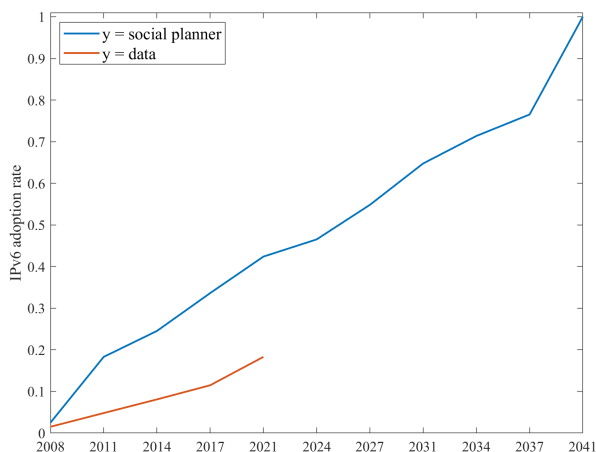


Figure 12: Social planner IPv6 adoption path

The main takeaways are the adjustment and switching costs, along with network effects can significantly delay IPv6 adoption leading to lower total producer surplus overall. This opens the discussion regarding the importance of subsidies that could lead to quicker IPv6 adoption by firms.

## 7 Conclusion

While transition markets seem to be gaining importance for transition through a price mechanism, the market is unable to solve the transition problem necessarily on its own in all cases. The aim of this paper was to study such a market - the market for IPv4 addresses by collecting a novel dataset, developing a dynamic model to study technology adoption in the presence of a secondary market. Using the model estimates I simulate the model forward and predict the evolution of price in the market of IPv4 addresses and IPv6 adoption.

I simulate two counterfactual simulations to study the effect of market frictions. I find that adjustment and switching costs lead to higher prices and keep prices positive until complete IPv6 adoption. IPv6 adoption happens 20 years faster compared to the baseline. Without network effects IPv6 adoption happens almost immediately. These findings indicate that in the presence of frictions the market mechanism could prove to be insufficient in technology transition. Thus, these findings have implications for regulators in the market for IPv4 addresses as well as more broadly other similar markets. Further, this paper introduces novel data and a framework that can help in the continued study of this market and other markets.

There are two caveats related to this analysis. The first caveat is that I ignore NAT and CGNAT -

services that allow multiple devices to connect using the same IP address (thereby extending the life of IPv4 addresses). While many firms are using this technology, I ignore this aspect as this is a band-aid solution and cannot be scaled-up without negatively affecting the quality of the internet connection. A similar argument follows for ignoring ‘translation’ techniques that allow IPv4 and IPv6 to inter-operate.

The second caveat is that I have not examined the possibility of a price bubble in this market. Experimental studies have found evidence of price bubble in a finitely lived durable asset. Theoretical work has laid the necessary conditions for a price bubble as: inefficient initial allocation, short-sales constraint and asymmetric information. All these conditions hold true for IPv4 addresses, and hence would be interesting to see if there is a price bubble in this market. Yet confidently testing for a price bubble remains challenging and thus remains outside the scope of this paper.

## References

- Bauer, Johannes M, and Michael Latzer. 2016. “The Economics of the Internet: An Overview.” *Handbook on the Economics of the Internet*.
- Buchholz, Nicholas. 2016. “Spatial Equilibrium, Search Frictions and Efficient Regulation in the Taxi Industry.(2016).” URL [https://Scholar.Princeton.Edu/Sites/Default/Files/Nbuchholz/Files/Taxi\\_draft.Pdf](https://Scholar.Princeton.Edu/Sites/Default/Files/Nbuchholz/Files/Taxi_draft.Pdf).
- Caoui, El Hadi. 2019. “Estimating the Costs of Standardization: Evidence from the Movie Industry.”
- Dubé, Jean-Pierre H, Günter J Hitsch, and Pradeep K Chintagunta. 2010. “Tipping and Concentration in Markets with Indirect Network Effects.” *Marketing Science* 29 (2): 216–49.
- Edelman, Benjamin, and Michael Schwarz. 2015. “Pricing and Efficiency in the Market for IP Addresses.” *American Economic Journal: Microeconomics* 7 (3): 1–23.
- Farrell, Joseph, and Garth Saloner. 1985. “Standardization, Compatibility, and Innovation.” *The RAND Journal of Economics*, 70–83.
- Greenstein, Shane. 2020. “The Basic Economics of Internet Infrastructure.” *Journal of Economic Perspectives* 34 (2): 192–214.
- Guérin, Roch, and Kartik Hosanagar. 2010. “Fostering IPv6 Migration Through Network Quality Differentials.” *ACM SIGCOMM Computer Communication Review* 40 (3): 17–25.
- Humlum, Anders. 2019. “Robot Adoption and Labor Market Dynamics.” *Princeton University*.
- Keane, Michael P, and Kenneth I Wolpin. 1997. “The Career Decisions of Young Men.” *Journal of Political Economy* 105 (3): 473–522.
- Lee, Donghoon. 2005. “An Estimable Dynamic General Equilibrium Model of Work, Schooling, and Occupational Choice.” *International Economic Review* 46 (1): 1–34.

- Ostrovsky, Michael, and Michael Schwarz. 2005. "Adoption of Standards Under Uncertainty." *RAND Journal of Economics*, 816–32.
- Paltridge, S. 2014. "Unclassified DSTI/ICCP/CISP (2014) 3/FINAL."
- Qi, Shi. 2013. "The Impact of Advertising Regulation on Industry: The Cigarette Advertising Ban of 1971." *The RAND Journal of Economics* 44 (2): 215–48.
- Richter, Philipp. 2017. "Empirical Analysis of the Effects and the Mitigation of IPv4 Address Exhaustion."
- Ryan, Stephen P, and Catherine Tucker. 2012. "Heterogeneity and the Dynamics of Technology Adoption." *Quantitative Marketing and Economics* 10 (1): 63–109.
- Saeedi, Maryam. 2019. "Reputation and Adverse Selection: Theory and Evidence from eBay." *The RAND Journal of Economics* 50 (4): 822–53.
- Sen, Soumya, Youngmi Jin, Roch Guérin, and Kartik Hosanagar. 2010. "Modeling the Dynamics of Network Technology Adoption and the Role of Converters." *IEEE/ACM Transactions on Networking* 18 (6): 1793–1805.
- Tassey, Gregory C. 2005. "Planning Report 05-2: IPv6 Economic Impact Assessment."
- Weintraub, Gabriel Y, C Lanier Benkard, Przemysław Jeziorski, and Benjamin Van Roy. 2010. "Nonstationary Oblivious Equilibrium." working paper.