



Pricing and Efficiency in the Market for IP Addresses

**Benjamin Edelman
Michael Schwarz**

Working Paper

12-020

June 9, 2013

Copyright © 2011, 2012, 2013 by Benjamin Edelman and Michael Schwarz

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Pricing and Efficiency in the Market for IP Addresses

Benjamin Edelman*
Harvard Business School

Michael Schwarz
Google

June 9, 2013

Abstract

We consider market rules for the transfer of IP addresses, numeric identifiers required by all computers connected to the Internet. Excessive fragmentation of IP address blocks causes growth in the Internet's routing table, which is socially costly, so an IP address market should discourage subdividing IP address blocks more than necessary. Yet IP address transfer rules also need to facilitate purchase by the networks that need the addresses most, from the networks that value them least. We propose a market rule that avoids excessive fragmentation while almost achieving social efficiency, and we argue that implementation of this rule is feasible despite the limited powers of central authorities. We also offer a framework for the price trajectory of IP addresses. In a world without uncertainty, the unit price of IP addresses is constant until all addresses are in use and begins to decrease at that time. With uncertainty, the price before that time is a martingale, and the price trajectory afterwards is a supermartingale. Finally, we explore the role of rental markets in sharing information about address value and assuring allocative efficiency.

1 Introduction

Every device connected to the Internet—from PCs to tablets, printers to cash registers—needs an Internet Protocol (IP) address. The current addressing standard, IPv4, uses addresses with 32 binary digits, allowing approximately 4 billion IP addresses. The world's central supply of IP addresses has reached exhaustion, and networks in most countries subsequently found that they could not easily obtain additional IPv4 addresses. While addresses may now be bought and sold, the institutions and rules of these transfers are not yet well-developed. Nor have economic models examined the unusual characteristics of this market.

*Edelman advises ARIN's counsel on matters pertaining to IPv4 exhaustion, IPv6 transition, and associated ARIN policy. This paper expresses only his personal view.

We thank Preston McAfee, Michael Ostrovsky, and two anonymous referees.

In this paper, we seek to speak to these gaps. Specifically, we use economic theory to examine the special structure of the market for IPv4 addresses, including the relevant externalities, applicable restrictions, and the conditions that could support a near-efficient outcome.

An important downside of IP address transfers is that when a buyer and seller trade, their transaction imposes a negative externality on everyone else—a problem resulting from the routing table entry records almost always created as a result of a trade. Indeed, for small trades the externality can be larger than the value of the trade. By performing a few large trades rather than many small trades, buyers and sellers can reduce the size of this externality, but they ordinarily have no incentive to do so. Moreover, as we discuss in Section 3.4, failure to address these externalities risks disproportionate burden on the router infrastructure that underlies Internet communications. The obvious solution is a Pigovian tax, but this is not feasible because the central authority lacks authority to impose or collect such a charge. However, the authority can impose certain other market rules.

We see several reasons why a market rule may be preferable to a complete market mechanism. First, some market authorities may lack the power to impose a specific mechanism. Second, an authority may lack information as to what mechanism is preferable. If market participants have better information, it may be preferable to let them negotiate a mechanism themselves.

We evaluate market rules on two distinct dimensions: First, among all equilibrium outcomes that are consistent with a market rule, how good is the worst outcome at achieving efficiency? Second, how much do agents seek to escape the rule, in the sense of being willing to pay to avoid complying? The “spartan” rule that we propose (Definition 4) is essentially non-invasive (Proposition 1) yet guarantees a near-optimal outcome (Corollary 2).

Although the rule we propose for IPv4 addresses is tailored to that market and neither the rule nor its theoretical guarantees are directly transferable to other markets, our approach to evaluating market rules might be useful in other markets where authorities have limited power. While the market design literature examines market mechanisms that fully specify market behavior, in practice many markets rely on central authorities that leave agents free to interact via mechanisms of their choosing so long as they respect certain limited market rules. Of the market rules used in practice, some might turn out to be splendid, while evaluation of others might reveal room for improvement. For example, consider the process by which students match with graduate degree programs. Since the 1960’s, the Council of Graduate Schools has required schools to grant students until April 15 to accept an offer. It is not obvious that a single round of acceptance is optimal—it eases a student’s decision conditional on acceptance, though it complicates a school’s offers. One can imagine other market rules such as multiple rounds of decision dates or rolling decisions, and alternatives would be attractive if the rule guarantees higher social welfare than the status quo and if the relevant authority has power to enforce it. Second, consider trading hours in financial markets. As markets grow, trading hours typically increase. But there is no guarantee that market forces will create a socially optimal number

of trading hours. In principle, reducing trading hours could increase liquidity, a valuable benefit for thinly-traded assets. Our approach to evaluating rules, including explicit consideration of the ability of an authority to enforce a rule, is particularly apt in light of some authorities' limited power.

We proceed as follows: In Section 2, we review the relevant literature, then in Section 3, we present the technical and institutional details that inform understanding of this market. In Section 4, we consider a one-period model of the market for IPv4 addresses and a mechanism that yields a (near) efficient outcome (notwithstanding the negative externalities we identify in Section 3.4).

In Section 5, we offer an equilibrium framework for modeling the price trajectory of IP addresses over time. In a world without uncertainty, the unit price of IP addresses is constant until all addresses are in use and begins to decrease at that time. With uncertainty, the price before that time is a martingale, and the price trajectory afterwards is a supermartingale. These results give testable predictions which can be evaluated as the market develops. These results also lay the groundwork for testing the external validity of important experimental findings including the widely-cited Smith et al. (1988) finding of price bubbles. Finally, we explore the role of rental markets in sharing information about address value and assuring allocative efficiency, and in Section 6 we conclude.

2 Relationship to the Literature

Low-cost electronic communications have sparked the transformation of many markets—both in creating the need to accommodate large communities of market participants, and in providing tools to make centralized markets mechanisms logistically viable. For example, keyword auctions let tens of thousands of buyers bid for millions of distinct objects—a context ill-suited to the negotiations traditionally used to sell advertising (Edelman et al. (2007)). At the same time, the Internet facilitates market mechanisms that would have been too cumbersome without modern information technology. For example, the spectrum auctions considered in McMillan et al. (1998) require quick transmission of other bidders' activities. The market for IPv4 addresses follows in part in this vein—another example of the Internet expanding feasible market structure. Of course the market for IP addresses is also notable in that its existence and operation are important for the continued growth of the Internet.

A line of market design papers lets the unusual requirements of a particular context guide the design of the corresponding allocation rules and institutions. Consider school choice (allocating placements based on rankings as well as proximity and sibling enrollment in Abdulkadiroglu et al. (2005)) and kidney allocations (planning transfer sequence and compatibility in Ashlagi et al. (2011)). In that spirit, we evaluate novel rules in service of the special concerns that arise in IP address markets.

Because the total supply of IPv4 addresses is fixed, our work has a natural connection to the literature on environmental resource economics. The classic Hotelling (1931) offers an optimal consumption path for natural resources.

In the context of IP addresses, one important difference is that environmental resources tend to be *consumed* (e.g. each gallon of gasoline can only be burned once), whereas IP addresses are durable. The durability of IP addresses serves to attenuate scarcity, although IPv4 addresses nonetheless run low due to rapid growth in demand (including more users seeking Internet access as well as multiple devices per user) and inefficient initial allocations.

In the language of the matching literature, IPv4 transfers are a many to many market with externalities because each buyer and seller can trade with multiple partners and because their trades impose an externality on others. Roth (1984) and Ronn (1990) study matching markets with externalities and obtain a number of negative results, for example that in a market with couples a stable match may not exist, and even determining whether a stable matching exists may be computationally hard. The positive results that we obtain are somewhat unusual in many to many markets with externalities. Our results are possible thanks to the special structure of the IPv4 environment: the externality depends on the number of trading partners but the agents only care about the total number of addresses they buy or sell, but not the identity of trading partners. Proposition 1 implies that this property is sufficient to manage externalities.

The impediments to migration to a new IP address standard are largely outside of the scope of this paper. Some of the challenges of migration from IPv4 to IPv6 are discussed in Edelman (2009). The general economic issues are largely captured in the literature on network goods and standard-setting. For example, consider the market for alternatives to gasoline: Drivers might be willing to switch from gasoline to electric vehicles if there were a network of charging stations; but without such vehicles, there is little impetus to build chargers, and vice versa. If the context of IP networks, it would be useful to coordinate on simultaneous transition to an alternative address technology in order to avoid the problems of Section 3.2. Farrell and Saloner (1986) present a similar switching problem in which agents may switch from an old standard to a new standard which they prefer unanimously. With switching opportunities arriving as independent Poisson draws, there exist parameter values in which agents never switch because each prefers that others switch first. More recently, Ostrovsky and Schwarz (2005) point out that coordinating transitions can be exceptionally difficult due to the risk that others will not make the transition as scheduled. This concern applies in full to IPv4 transition: To avoid the transition/compatibility costs in Section 3.2, a network would need other networks to cease using IPv4 at the same time. But it would be unrealistic to expect the entire Internet to make the switch simultaneously, and a tall order even for a network's closest communication partners to upgrade at the exact same time. The technical and economic impediments to IPv6 transition are largely beyond the scope of our paper, though we sketch the relevant constraints in Section 3.2 and discuss the implications in the conclusion.

3 The technologies and institutions of IP addressing

3.1 The institutions of IP addressing

Evaluating feasible rules for IPv4 markets requires understanding responsible institutions. We summarize the most important facts here. Edelman (2009) offers greater detail.

Like most of the Internet’s infrastructure, the IP addressing system is largely private. IP communication systems were initially developed pursuant to US government contract, but private parties have always administered address assignments. Under current allocation procedures, the Internet Assigned Numbers Authority (IANA), a California non-profit, assigns addresses to five regional internet registries (RIRs), which in turn assign addresses to ISPs and large networks. (In some regions, RIRs first assign addresses to country-specific internet registries, which in turn assign addresses to ISPs and networks.) Large networks tend to get their addresses directly from RIRs, while smaller networks and residential end-users obtain addresses from ISPs (often on a temporary basis).

Because RIRs are private organizations, they have little ability to penalize networks that flout their rules. Indeed, networks seem to value the limited powers of their RIRs, treating this as a philosophical matter (analogous to “limited government”). That said, RIRs do have some important powers. Notably, most ISPs check an RIR’s records before accepting a new customer purporting to use a given block of IP addresses. In particular, each RIR provides “WHOIS” listings which report the network authorized to use a given block of IP addresses, including the network’s name, mailing address, administrative and technical contacts, and email and telephone contact information. If a network sought to use a given block of addresses, but the network did not match the details listed in WHOIS for those addresses, most ISPs would not allow the customer to proceed. Historically, networks have found accurate WHOIS data sufficiently desirable—and RIR rules sufficiently unobjectionable—that networks have been willing to comply with RIR rules in order to obtain official address assignments, WHOIS listings, and associated benefits.

To date, RIRs have examined each network’s request to confirm the need for the requested addresses; addresses are issued only to networks that can justify their requests, e.g. through business plans, equipment purchases, and customer lists. Meanwhile, RIRs charge low fees. For example, the North American RIR, the American Registry of Internet Numbers (ARIN), charges just \$18,000 per year to its largest ISPs.

With modest fees and rapid Internet growth, IPv4 resources are in short supply. IANA announced in February 2011 that it had exhausted its reserves. APNIC, the RIR for the Asia/Pacific region, finished its supply in April 2011; RIPE NCC (Europe) ran out in September 2012; and ARIN (North America) is expected to run out by fall 2013.

3.2 Technical responses to IPv4 scarcity

Facing limited availability of IPv4 addresses, growing or newly-created networks have several options:

- Networks can move devices to an alternative addressing system called *IPv6*, which offers 128 binary bits of numbering capacity and therefore is in abundant supply. However, a device with only an IPv6 address cannot directly communicate with IPv4-only devices; for example, it cannot directly access a web site that has only an IPv4 address. Translation systems are not yet widespread, and some protocols may be difficult to translate between IPv4 and IPv6, making IPv6 a more appealing solution when others have adopted it also. Furthermore, until IPv6 is in widespread use, IPv6 communications will take longer and slower routes (limited to passing among IPv6-enabled routers). In addition, some have found that IPv6 systems less reliable than IPv4 (typically for lack of around-the-clock monitoring and troubleshooting).
- Alternatively, networks can seek to *make do with fewer* IPv4 addresses, typically by installing network address translation (NAT) devices which allow multiple devices to share a single public-facing IP address. By rewriting packet headers, NAT maintains the illusion that all local devices have a single public address. NAT devices have become common in the home gateways that many homes install in order to connect multiple computers to a single DSL or cablemodem connection. While most protocols function as expected through NAT, others do not or require significant customization, and NAT impedes certain kinds of innovation (Blumenthal and Clark (2001)). Furthermore, large-scale “carrier-grade” NAT would raise weightier questions of reliability and scalability, as well as reducing the likelihood of transition away from NAT in the coming years.
- Finally, networks can *buy* IPv4 addresses from others. Some networks may have more than they need—whether due to overly generous initial allocations, reduced requirements, or migration towards address-conserving alternatives (such as IPv6 or NAT). Meanwhile, others need more than they have—for example, new or growing networks.

In the long run, it is widely expected that networks will move to IPv6. But in the short run, IPv6 transition has been slow. For one, network effects encourage use of IPv4: with most users and sites on IPv4 only, it is natural to want IPv4 connectivity and unusual to want IPv6. Limited and untested translation systems—untested in part due to lack of customer demand—have further hindered transition. Meanwhile, NAT has not been deployed on the required scale and, even in the best case, adds complexity to network structure. For these reasons, it is currently anticipated that many new and growing networks will, in the short run, turn to IPv4 purchases to meet their IPv4 needs—prompting questions of the design of markets and institutions to facilitate such transfers.

3.3 The prospect of paid transfers of IPv4 addresses

Historically, paid transfers of IP addresses were impermissible. After all, with sufficient addresses available directly from RIRs, there was no need for sales between networks, and any buyer was either a scofflaw (e.g. a spammer needing limitless new addresses to send unsolicited email—a group networks and RIRs did not seek to assist) or a fool (who would be better served by a strong norm of obtaining addresses directly from RIRs). But as scarcity loomed, RIRs revised their policies to allow sales. For example, the ARIN transfer policy is codified in Number Resource Policy Manual (NRPM) at section 8.3.

Today the market for IPv4 is quite small—less than 100 transactions to date, covering a fraction of a percent of the Internet’s address resources. (Mueller et al. (2012)) That said, there is reason to expect IPv4 transfers to become large and important both in dollar volume and in enabling continued growth of the Internet. Indeed, IPv4 transfers will probably prove necessary to facilitate entry both by firms providing Internet infrastructure services (ISPs, hosting, cloud computing) and by ordinary Internet users (companies, organizations, and, via ISPs, even residential users). A well-functioning IPv4 transfer market is particularly valuable in light of the importance of entry in achieving price and quality improvements (as in Goolsbee and Syverson (2008)).

3.4 A negative externality from IPv4 sales

While trade in IPv4 addresses promises various benefits, including those sketched in preceding sections, transfers also prompt concerns. Most common is the fear that certain address transfers might threaten the Internet’s routing system—the infrastructure that transfers data from one network to another. The Internet’s routing architecture requires that each “default-free zone” (DFZ) router on the Internet keep a record of each block of addresses used anywhere else on the Internet. These records are stored in each router’s routing table—high-speed memory limited in size and limited further by the need to process and search the routing table exceptionally quickly. Importantly, a small block of addresses requires just as much routing table capacity as a large block. In particular, any block of addresses, of any size, requires one routing table row to indicate how to send traffic destined for that block of addresses. If large networks begin to acquire many small address blocks, rather than a few large blocks, the routing table could grow sharply. For example, a network might elect to buy 16 blocks of 256 addresses, rather than one block of 4,096, if the former is slightly cheaper. In that case, growth would be an order of magnitude more rapid.

Routing costs appear to be significant relative to address prices. Without centralized tracking of routing system participants or costs, it is difficult to measure the expenses, but network engineers have estimated the cost of adding routes in light of the number of routers in use, router lifetime, and costs of routers with varying capabilities. Herrin (2008) estimates a total social cost between \$4,000 and \$12,000 for each route announcement each year, proposing a point estimate of \$8,000. Meanwhile, Mueller et al. (2012) reports that more

than one third of IPv4 sales to date are of blocks of just 256 addresses—likely to have purchase prices less than \$3,000, well below the single-year externality imposed on others, suggesting that routing externalities are first-order important. Even the median IPv4 sale to date, for 1024 addresses, has a transaction price approximately equal to the single-year externality imposed on others. (At an estimated \$9 to \$12 per address, these sales garner revenue of \$9,000 to \$12,000.) Because externalities are reincurred each year, as the routing system is forced to operate at expanded size, externalities are actually even larger in the long run.

At best, growth in the routing table would require that networks upgrade their routers more frequently, with attendant costs as estimated by Herrin (2008). But particularly rapid growth could exceed router manufacturers’ ability to improve their offerings—destabilizing the Internet’s routing system.

The Internet’s routing system is decentralized by design, and no single entity—neither public nor private—directly controls the routing system. As a result, no one can easily impose rules on what routes may be added to the routing table. In principle a Pigovian tax could appropriately discourage creation of new routes, but with no one to collect that fee or penalize those who fail to pay, this approach is infeasible.

Meanwhile, there is an asymmetry in creating versus removing routing table additions: Once routes are created, there is ordinarily no practical way to remove them. To remove a route requires finding a network that can consolidate two nonadjacent blocks (each with a separate route announcement) into a single block—requiring that network to renumber some or all of the resources being consolidated. Because there is no fee to add a route and no reward for removing a route, no network would have an incentive to incur these costs. Indeed, unless a network designed its systems in anticipation of future renumbering, the network’s costs of vacating an address block are likely to exceed the total social cost of the block’s route announcement. Thus, efficient management of routing costs requires avoiding creating unnecessary routes in the first place; once these routes have been created, it is unduly costly to remove them.

Of course the world would not stand still while the Internet “breaks” under the weight of ill-functioning IP markets. If the problem were sufficiently acute, various drastic measures might look attractive—for example, a complete ban on address transfers. This would substantially end growth in the routing table, but it would be extremely socially costly due to allocative inefficiencies as well as adverse consequences for competition and innovation.

In principle, RIR policy can help avoid unnecessary growth of the routing table. In North America, ARIN initially established rules requiring each buyer to satisfy the entirety of its short-run need in a single transfer—disallowing multiple small transfers, and thereby avoiding unnecessary routing table additions. See ARIN NRPM rule 8.3 as first enacted, allowing transfers only to buyers which “can demonstrate the need for such resources, as a single aggregate, *in the exact amount* which they can justify under current ARIN policies” (emphasis added). However, other RIRs took a different approach. For example, APNIC’s transfer rules included no restrictions to prevent many small purchases. In a

2011 change, ARIN proposition 144 removed the “exact amount” requirement, allowing buyers to satisfy their needs via multiple small purchases.¹

In Section 4.2, we examine the welfare and efficiency properties of a rule that generalizes ARIN’s “exact amount” requirement.

4 One period model and market rule

In most asset markets, it is typical for large buyers and large sellers to split their purchases and sales into small pieces to be bought or sold separately. An unusual feature of market for IPv4 addresses is that such splits are socially costly for the reasons presented in Section 3.4. We therefore seek to devise market rules to discourage unnecessary splits, and in this section we develop a formal model to evaluate the performance of such rules.²

4.1 Notation and definitions

Consider a set of networks, each with an endowment of IPv4 addresses. We will use the size of the smallest tradable block of IPv4 addresses as the unit of measurement, so when we say that network k bought x_k addresses, we mean that k bought x_k times the smallest block. $x_k < 0$ means that k was a seller.

Denote by $f_k(x)$ the value that k derives from increasing (or the cost k incurs in decreasing) its IPv4 address holdings by x . Normalize $f(0) = 0$. We assume that $f()$ is increasing. We also assume decreasing differences—that the marginal benefit of an address decreases in the number of addresses the network holds (the discrete analogue of a negative second derivative). Formally, we require $f(x) - f(x - 1) \leq f(x - 1) - f(x - 2)$ for all x . This condition is needed to ensure the existence of the competitive equilibrium price.

Definition 1 *An outcome is a vector specifying how much each buyer purchases and how much each seller sells.*

Definition 2 *An outcome is feasible if $\sum_k x_k = 0$ and $x_k \in I$ for every k .*

Definition 3 *An outcome is efficient if $\sum_k f(x_k)$ is maximized subject to the feasibility constraints.³*

¹There are sound reasons to remove the requirement. First, unlike the spartan rule we propose, the “exact match” rule may cause an allocative efficiency. Second, ARIN’s efforts to prevent growth of the routing table create an obvious free-rider problem—North American networks paying the full cost of a public good that benefits networks worldwide. Some ARIN members recognized and discussed both these problems. At least as prominent in ARIN discussions was the sense that initial transactions had not caused undue growth in the routing table, and a suitable rule could be restored or created when needed.

²In a paper entitled “Minimizing Setup and Beam-On Times in Radiation Therapy”, Bansal et al. (2006) considers a mathematically similar model, albeit in an entirely unrelated substantive context.

³This definition of efficiency does not consider the negative externality presented in Section 3.4.

Due to the negative externality presented in Section 3.4, we will consider restrictions on trading to prevent excessive fragmentation of IPv4 blocks without inhibiting efficiency. The following definitions will assist in developing and analyzing such restrictions:

Definition 4 *Under the spartan rule, the two agents engaging in a bilateral trade must designate either the buyer or seller as “extinguished.” The spartan rule prohibits an extinguished agent from trades with other extinguished agents. An agent can be extinguished at most once.*

Notice the relationship between our concept of an “extinguished” agent and the “exact amount” requirement from Section 3.4: Whereas “exact amount” required that every buyer’s first trade be his last, our spartan rule allows either buyer or seller to be “extinguished” by a transaction, such that either the buyer or the seller ceases to trade further after that transaction. In Section 4.2 we examine the welfare and efficiency implications of this modified rule.

Because we offer a one-period model, an agent is extinguished forever. However, in practice IP allocations typically operate on a rolling basis, i.e. providing a network with the addresses it can justify for a three month period. As a practical matter, it would be logical for extinguished status to last for that same duration.

A sequence of trades satisfies the spartan rule if and only if each agent is extinguished at most one time, and if each agent has at most one trade with an agent extinguished by someone other than himself. Consequently, if a sequence of trades satisfies the spartan rule, then the same trades in a different sequence also must satisfy the spartan rule. Thus, we can define and evaluate spartan allocations without considering the order of trades.

An allocation is an outcome and the set of bilateral trades that lead to that outcome. More formally:

Definition 5 *An allocation specifies which seller(s) each buyer was matched with, the size of each transaction, the price of each transaction, and which party was designated as extinguished.*

An allocation is spartan if each agent is extinguished at most one time. In that case, regardless of the order of trades, the spartan rule is satisfied.

In preparation for addressing the externality presented in Section 3.4, we offer a notion of subdivision of a seller’s resources:

Definition 6 *A transaction between a buyer and a seller leads to a cut if, after the sale, the seller retains a positive quantity of resource.*

We clarify the definition with several examples. If a seller sells a positive quantity of resource to n different buyers and keeps some of the resource for itself, the number of cuts required is n . If a seller sells a positive quantity of resource to n different buyers and keeps none, the number of cuts is $n - 1$. No cuts occur if a seller sells all of its resources to a single buyer. This is consistent

with the externality identified in Section 3.4: If a seller sells all of its resources to a single buyer, it is likely that the seller's single entry in the routing table will be replaced by a single entry for the buyer, yielding zero net change.

4.2 Welfare and efficiency under spartan allocations

Let N denote the number of buyers and M the number of sellers.

Lemma 1 *A spartan allocation involving N buyers and M sellers never entails more than $N + M$ cuts.*

Proof. With N buyers and M sellers, the maximum total number of transactions is $N + M$. (Either a buyer or a seller must be extinguished in each transaction, and each entity is extinguished at most once, so the number of trades cannot be more than the number of participating entities, $N + M$.) Suppose seller m was matched to n_m buyers. Since the total number of transactions is no more than $N + M$, it must be the case that $\sum_{m=1}^M n_m \leq N + M$. ■

Proposition 1 *For any feasible outcome, there exists a spartan allocation that induces that outcome.*

Proof. Consider an outcome involving N buyers and M sellers. Denote by V_i^0 the quantity obtained in that outcome by buyer i , and denote by W_j^0 the quantity sold by seller j . Denote by V_i^t and W_j^t the number of units that buyer and seller respectively will acquire (or provide) in transactions that take place after period $t - 1$. (If $V_i^t = 0$, then by period t buyer i already made all his trades.) Consider a sequential mechanism where there is exactly one transaction in each period, with buyers randomly matched with sellers in arbitrary order. In period t , match some buyer k with some seller m . If $V_k^t > W_m^t$ (buyer k wants more than seller m can offer), they transact amount W_m^t —exhausting m 's supply, leaving $W_m^{t+1} = 0$, while k still seeks additional resources $V_k^{t+1} = V_k^t - W_m^t$. Alternatively, if $V_k^t \leq W_m^t$ (seller m wants to sell more than buyer k requires) then in period t the transacted amount is V_k^t . In other words, in each period the random buyer and seller transact with each other the largest possible amount, thereby extinguishing either the buyer or the seller in each transaction. The sequence of trades constructed through this process can implement any outcome (including the efficient outcome) with at most $N + M - 1$ cuts. ■

Corollary 1 *Suppose an allocation includes N buyers and M sellers. There exists a spartan allocation with no more than $N + M - 1$ cuts that induces an efficient outcome.*

Proof. Follows from the proof of Proposition 1. ■

Proposition 1 shows that even with the constraint provided by the spartan rule, any allocation remains feasible. Since we allow transfers between players, an inefficient allocation cannot be coalition-proof because the coalition containing all the players can move to an efficient allocation and arrange for transfers

to make every agent better off, according to the standard argument. This yields the following corollary:

Corollary 2 *Any coalition-proof spartan allocation is efficient, and entails no more than $N + M$ cuts. Furthermore, for any coalition-proof spartan allocation, there exists a competitive equilibrium price that supports it.*

The proposed market rule can be implemented easily in practice. Suppose agents know the competitive equilibrium price and hence agent knows whether it wishes to be a buyer or seller and how much it wants to buy or sell. Then the proof of Proposition 1 offers a constructive algorithm for obtaining a coalition-proof spartan allocation. Of course such an allocation can also be obtained via myriad other mechanisms—a centralized clearinghouse, multiple independent brokers, or a web site (such as the listing service ARIN now operates).

In contrast to the spartan rule, ARIN’s “exact amount” rule impedes efficiency in certain circumstances. Consider a buyer who wants to buy a larger quantity of addresses than any seller is willing to sell. Under an exact match rule, that buyer cannot make such a purchase at any price because there exists no seller whose resource exactly matches what the buyer wants to buy. But under the spartan rule, the buyer can buy all of the addresses offered by arbitrarily many sellers, extinguishing each seller in turn and thereby retaining the right to buy more from other sellers.

4.3 Properties of minimal allocations

The preceding section characterizes the maximum number of cuts required for spartan allocations and efficient outcomes. In this section, we examine the *minimal* number of cuts in order to minimize the externality flagged in Section 3.4.

Definition 7 *A minimal allocation is an efficient allocation with weakly fewer cuts than any other efficient allocation.*

Proposition 2 *There exists a spartan allocation that is also minimal.*

Proof.

Consider a connected component in a minimum allocation. Let M and N denote the number of buyers and sellers inside the connected component. Let h denote the number of sellers inside the connected component who are left with no resource after the last sale. Since we are considering a connected component, it must have at least $M + N - 1$ trades and thus the total number of cuts in a minimal allocation is at least $M + N - 1 - h$.

Now let us show that the same outcome can be implemented inside that connected component by a spartan allocation with the same number of cuts as in a minimum allocation. Since there are $M + N$ agents, there exists a spartan allocation with $M + N - 1$ trades. Note that the outcome implemented by the spartan allocation is the same as the outcome in the minimal allocation, so the

number of sellers who are left without resource is the same for both spartan and minimal allocation. Thus there exists a spartan allocation with $M + N - 1 - h$ cuts, the same as in the minimal allocation. ■

Corollary 3 *The number of extinguished agents in the minimal allocation equals the number of buyers and sellers minus the number of (disjoint) components in the graph representing the minimal allocation.*

Proof. It follows from Corollary 1 that the k -th component has at most $N_k + M_k - 1$ extinguished agents. Summing over all components, we obtain the upper bound on the number of extinguished agents. Now let us show that a component of a graph induced by minimal allocation cannot have less than $N + M - 1$ extinguished agents. From Proposition 2, it follows that without loss of generality we can consider a minimal allocation that is spartan and hence all agents are extinguished no more than once. Suppose a component has fewer than $N + M - 1$ extinguished agents. In that case the total number of links is at most $N + M - 2$ and hence at least two agents remain unextinguished. The trading partners of the unextinguished agents are extinguished (by transactions with unextinguished agents). Denote by Ω_{1h} and Ω_{2h} the sets of agents who are h links away from unextinguished agents 1 and 2. (Ω_{10} is agent 1 himself, and Ω_{11} is the set of agents who are connected to agent 1 and hence extinguished by that trade.) Note that agents in a set Ω_{1k} are extinguished by agents in a set Ω_{1k-1} . Consequently, any agent connected via a chain with agent 1 is extinguished by agents closer to agent 1, and the same is true for those surrounding other unextinguished agents. Thus, the clusters around unextinguished agents can never be connected to each other, which contradicts the assumption that they belong to the same component of the graph. ■

Note the total number of cuts must be less than or equal to the number of extinguished agents. Indeed, the number of extinguished agents equals the number of trades, and each trade leads to at most one cut. As a result, Corollary 3 gives an upper bound on the total number of cuts.

The preceding results establish the favorable characteristics of the spartan rule: it achieves efficiency and is consistent with achieving the fewest possible number of cuts. At the same time, the spartan rule generalizes the ARIN rule presented in Section 3.4. Relative to the mechanism ARIN adopted, the spartan rule is more flexible in that it allows large buyers to more readily find the addresses they seek, yet it does so without inviting undue growth in the routing table.

5 Price trajectory

The previous section considers a static world: from the outset, all networks know their values for IPv4 addresses, and these values do not change over time. The preceding section ignored dynamic aspects of the market for IPv4 addresses. But in fact IP addresses are long-lived assets, and networks' needs change over time.

In this section, we consider the dynamics of the IPv4 market, characterizing the price trajectory of IPv4 addresses.

In the previous section, we showed that even with integer constraints and even with the trading constraint of the spartan rule, the same efficient allocation obtains as in an unrestricted commodity market. Hence, for technical convenience this section assumes that IPv4 addresses are homogeneous and perfectly divisible.

Our examination of pricing builds on an understanding of long-term prices: Although IPv4 addresses are currently valuable, it is expected that scarcity will eventually be relieved as networks migrate to IPv6. In particular, once much of the Internet supports IPv6, there will be less need for IPv4 addresses, though supply will remain unchanged. As a result, it is reasonable to assume that at some period, T , the value of IPv4 address will become zero. But what happens to prices en route to that point?

To explicitly model the pricing process, we treat IPv4 addresses as a production input. Let $f_{it}(x_{it})$ denote the value that network i derives from having x_{it} more IPv4 addresses in period t than the network was endowed with before the first period. (If x_{it} is negative, the network is a seller.)

5.1 Price Trajectory in the Complete Information Case

We begin by considering a world without uncertainty or private information. (In section 5.2 we consider a world where information about demand for IPv4 arrives over time.)

We assume that the marginal value derived from an additional IP address is non-negative and diminishing, that is $f'_{it}(x) \geq 0$, $f''_{it}(x) \leq 0$, for all $x \in R$, $i \in \{1, \dots, N\}$ and $t \in \{1, \dots, T\}$. The supply of IPv4 addresses is exogenously fixed, that is $\sum_{i=1}^N x_{it} = 0$ for any $t \in \{1, \dots, T\}$. At the start of each period, networks can buy and sell IPv4 addresses, and we will show that there is a unique competitive equilibrium price p_t in each period. The payoff of network i is the sum of period payoffs $\sum_{t=1}^{T-1} f(x_{it})$ plus the proceeds from buying and selling IPv4 addresses.⁴

The competitive equilibrium price trajectory of IPv4 is a vector of prices p_1, \dots, p_T such that the market clears in each period if all networks trade in each period in order to maximize the sum of payoffs across periods. More formally, assuming that each network solves the maximization problem $\max[\sum_{t=1}^T f_{it}(x_{it}) - p_1 x_{i1} - \sum_{t=2}^T (p_t(x_{i,t} - x_{i,t-1}))]$, we will say that the price vector clears the market if there exists a solution $x_{it}^*(p_1 \dots p_T)$ to the individual maximization problem such that $\sum_{i=1}^N x_{it}^* = 0$ for every t .

Proposition 3 *In a world without uncertainty, there exists a unique competitive equilibrium (CE) price vector $p_1 \dots p_T$ with the following properties:*

⁴By adding dollar payoffs from different periods, we implicitly assume that the interest rate is zero.

1. The price is non-increasing over time. That is, for any t and τ , if $t < \tau$ then $p_t \geq p_\tau$.
2. Prices are constant when addresses are not scarce. More formally, if during period t there exists at least one network that does not utilize some of the addresses that it owns (i.e. because $f'_{it}(x_{it}) = 0$), then $p_t = p_{t+1}$.

Proof. We begin by showing that a competitive equilibrium exists. Rearranging and using the assumption that $p_T = 0$, a network's maximization problem can be rewritten as $\max[\sum_{t=1}^{T-1} f(x_{it}) - x_{it}(p_t - p_{t+1})]$. Thus, the maximization problem is separable, with the network maximizing $f(x_{it}) - x_{it}(p_t - p_{t+1})$ in each period. In particular, if the competitive equilibrium price vector exists, then the optimal course of action in period t is independent of actions taken in other periods. We can interpret $g_t = p_t - p_{t+1}$ as the rental rate of IPv4 addresses in period t .

Note that in a competitive equilibrium, $g_t = 0$ if and only if there exists a vector of $x_{1t}^* \dots x_{Nt}^*$ such that $f'_{it}(x_{it}^*) = 0$ for all i and $\sum_{i=1}^N x_{it}^* \leq 0$. In this case, there is no scarcity of IPv4 in period t . In a world without scarcity, the existence of a competitive equilibrium is obvious.

Now consider the case of scarcity and $g_t > 0$. We showed above that in a deterministic competitive equilibrium, a network's profit maximization problem is reduced to $\max(f_{it}(x) - xg_t)$ in each period. Note that the assumptions for $f_{it}(\cdot)$ guarantee that the profit maximizing quantity x exists and is unique for every value of $g > 0$. Thus we can consider a demand function $x_{it}(\dot{g})$ that assigns to each positive rental price the corresponding demand by network i in period t . Our assumptions on f_{it} guarantee that demand is continuous and non-increasing in rental price \dot{g} .

We now demonstrate that the price trajectory is non-increasing. Consider the equilibrium rental price. If there is no scarcity in period t , then the competitive equilibrium rental price is zero. For the case of positive rental price, we previously established that $x_{it}(g)$ is continuous and non-increasing. If there exists a unique competitive equilibrium price in period $t+1$, there exists unique equilibrium rent in period t and hence there also exists a unique competitive equilibrium price in period t . By assumption there is no scarcity in period T and all subsequent periods, which means that the competitive equilibrium price in period T must be $p_T = 0$. We showed that if there is a unique competitive equilibrium price in period T , there must exist a unique competitive equilibrium price in the previous period, and recursing across previous periods implies that the price in period t equals the sum of rental rates in the remaining periods $p_t = \sum_{\tau=t}^{T-1} g_\tau$. Since rents are non-negative, the price is non-increasing in time.

Finally, consider periods in which there is no scarcity of IPv4 addresses. Because rent equals zero in those periods, the price of IPv4 addresses does not change during such periods. ■

We recognize that there is not currently a formal rental market for IPv4; such a market may or may not be desirable, and may or may not emerge in the future. Regardless, the rental market is a useful framework for thinking about

pricing of IP addressees. For example, if such purchases and sales were fully permitted by applicable rules, the act of buying an IP address in period t and selling it in period $t + 1$ can be interpreted as renting the address for period t at a price g_t .

We pause to offer intuition on the preceding results. Our analysis began with the insight that in any period without scarcity, the rental price of IPv4 must be zero. In other words, as long as there are blocks of IPv4 addresses that are not in use, equilibrium behavior of market participants requires that the rental price of IPv4 must be zero. (Of course, a zero *rental* price does not mean the prevailing market price of IPv4 should be zero. Rather, the market price of an IPv4 address incorporates expectations about future rental prices of that address.) In a world with a risk-free interest rate of zero and with no uncertainty, the price of IPv4 addresses must remain constant as long as there are some unused IPv4 blocks.

In contrast, consider the first period when there are no unused IPv4 blocks available. In this period, true scarcity begins, and the implied rent of an IPv4 address becomes positive. Notably, the market price of an IPv4 address declines each period by an amount equal to the implied rent of that address during the prior period. Using the preceding analyses, Figure 1 presents plausible price profiles for the complete information case (without uncertainty).

With the right framing, some might consider it obvious that IPv4 addresses drop in value over time—after all, each day these resources draw a bit closer to obsolescence. That said, we believe our analysis of IPv4 price profiles is particularly likely to be counterintuitive to non-economists. Many would find it counterintuitive that prices are at their highest before scarcity actually sets in, and that prices only decrease as networks search for sources of the scarce resources at issue.

5.2 Price Trajectory in a World with Uncertainty

In a world with uncertainty, the price of IPv4 addresses may increase or decrease as new information arrives. For example, prices would likely be influenced by new information about growth in the number of Internet users or changes in the speed of implementation of IPv6.

We model the arrival of new information by assuming that in each period there is a public signal s_t that influences demand for IPv4 addresses. Networks receive the signal s_t in the beginning of period t , then trade addresses as desired. The signals are jointly distributed according to a distribution $\rho(s_1, s_2, \dots, s_T)$ that is common knowledge. We assume that ρ is well-behaved in the sense that all conditional distributions exist and are integrable. We denote by $f_{it}(s_1 \dots s_t, x_{it})$ the payoff of a network i in period t , which depends on the history of signals and on the number of IPv4 addresses that the network already holds. Other assumptions remain as in the prior section.

The definition of competitive equilibrium price needs to be modified to include uncertainty. We now define a competitive equilibrium price as a set of functions $p_t(s_1, s_2, \dots, s_t)$ that clear the market for every realization of signals in

each period, provided all networks maximize their expected payoff. We use a similar argument as in the previous section to establish properties of competitive equilibrium.

Proposition 4 *In a world where publicly observable signals about the value of IPv4 arrive over time, there exists a unique competitive equilibrium (CE) price vector $p_t(s_1, s_2, \dots, s_t)$ for $t = 1 \dots T$ with the following properties:*

1. *Expected prices are non-increasing over time. That is, $E[p_t] \geq E[p_{t+1}]$.*
2. *Expected prices are constant when addresses are not scarce. If during period t there exists at least one network that does not utilize some of the addresses that it owns (i.e. because $f'_{it}(x_{it}) = 0$), then $E[p_t] = E[p_{t+1}]$. As new information arrives during a period when addresses are not scarce, the price trajectory is martingale.*

Proof. We follow the proof of Proposition 3 but generalize to a world with uncertainty.

Consider a network's optimization problem in period $T - 1$. By assumption, migration to IPv6 eliminates scarcity of IPv4 beginning in period T , so in period $T - 1$ there is no remaining uncertainty: Everyone knows the price in the next period will be zero. With no uncertainty, Proposition 3 applies from then forward, guaranteeing that a competitive equilibrium price, $p_{T-1}(s_1, s_2, \dots, s_{T-1})$, exists and is unique.

Now consider an agent deciding how many IPv4 addresses to acquire in period $T - 2$. The agent already knows signals $s_1 \dots s_{T-2}$ but does not know the realization of the signal in period $T - 1$. Consequently, from the agent's point of view, p_{T-1} is a random variable: for any realization of s_{T-1} , the agent knows the value of p_{T-1} , and the agent also knows the joint distribution of all signals and hence knows the distribution of the signal in period $T - 1$ conditional on previous signals. Thus, given signals s_1, \dots, s_{t-2} , the expected value of the competitive equilibrium price in period $T - 1$ is well defined, and we denote it by $E[p_{T-1}|s_1 \dots s_{T-2}]$.

Let us define $g_{T-2} = p_{T-2} - E[p_{T-1}|s_1 \dots s_{T-2}]$, which we interpret as the (expected) rental price for a period. Then an agent in period $T - 2$ considers the maximization problem $\max[f_{it-2}(s_1, s_2, \dots, s_{t-2}, x_{it-2}) - x_{iT-2}g_{T-2}]$. Thus period $T - 2$ demand is a function of the rental price, and the problem matches the world of certainty, save only for dependence on the expected rental rate in period $T - 2$ (rather than the actual rental rate in the case without uncertainty).

We have shown that if there exists a competitive equilibrium price in period t , then there must exist a competitive equilibrium price in the previous period $t - 1$. By assumption, the competitive equilibrium price in period T is zero, so the competitive equilibrium price exists and is unique, and the expected decline in price in period t is equal to the rental price in period t . ■

The environment that we consider is reminiscent of the environment in the influential Smith et al. (1988), a laboratory experiment which found a tendency

towards a price bubble when subjects traded a security that pays dividends for T periods followed by zero value. Smith et al. chose this structure because the finite horizon eases subjects' efforts to anticipate future asset values. To our knowledge, the Smith et al. result was never tested in the field, as it is unusual to find a real-world asset that has sufficiently simple pricing structure to enable a meaningful test for price bubbles. However, if IPv4 prices will trend downwards as a supermartingale, this provides a possible context for testing Smith et al., albeit requiring efforts to isolate the effects of uncertainty.

5.3 The Benefits of a Rental Market and Constraints on IP Address Rentals

The preceding sections characterize the price trajectory in a world with public availability of all information influencing the price of IPv4 addresses. But in practice, some information will typically be private to individual participants. Depending on the information structure available to market participants, a rental market can assist in spreading information and putting IPv4 resources to use.

Finance, general equilibrium and industrial organization models often assume that prices correctly aggregate private information, thus leading to allocative efficiency. In many contexts, there are good reasons to expect that markets aggregate information and achieve allocative efficiency. Indeed, work dating to Adam Smith recognizes that markets aggregate information, and a series of papers have formalized the mechanisms of information aggregation. It is well known that in a private value setting, an English auction leads to allocative efficiency. In a common value setting, Ostrovsky (2012) shows that, if trading is frictionless, prices correctly aggregate all private information (even if agents obtain multidimensional private signals). However, in environments with both common and private value components, markets may fail to aggregate information, a possibility first explored by Maskin (1992).

There are good reasons to suspect that IPv4 address transfers may not reach an efficient allocation. We can think of the value of IPv4 address to a network as a sum of the network's private value of using it during the present period plus the resale value in a future period. The value from renting an address is a private value component and the resale value in future periods is a common value component. Efficient information aggregation results do not apply in this context due to the presence of both common value and private value components, with participants having private information about both. In particular, each network will have private information about its private value for IPv4 addresses in the current period as about its future demand for IPv4. Some networks may also have greater insight into factors affecting the common value of IPv4 addresses. For example, thanks to their special roles in providing or deploying transition technologies, some networks may enjoy superior information about IPv6 deployment and about IPv4 address sharing. More generally, networks have private signals containing information about both their individual private value for IPv4 in the present period as well as information about future demand

for IPv4.

The following example illustrates why IPv4 transfers may not lead to an efficient outcome. Consider a network that has unused IPv4 addresses. If the network has unused addresses after scarcity occurs, declining market prices provide an incentive for the network to sell the addresses quickly to avoid the price drop that results from the positive real rental price. However, suppose the network also has a private signal about the common value. For example, the network might believe that future demand will be higher than currently anticipated by most other market participants. With this private information, the network seeks to hold IPv4 resources in anticipation of future price increases not expected by others. Note that the market will fail to aggregate this network's information because others cannot tell whether the network is accumulating IPv4 because it has high private value for using them in the current period or because it has private information that the price will increase in future periods. As a result, a network may elect to speculate in IP addresses by continuing to hold addresses that it does not use.

In principle, a rental market can help spread information and put IPv4 resources to their highest and best use. Returning to the problem posed in the preceding paragraph: If a rental market existed, the network could rent its unneeded addresses—collecting the positive per-period rental price and achieving allocative efficiency by putting scarce resources to use, yet continuing to enjoy the benefit of possible increase in the market price. In contrast, without a rental market, the network's private information is not passed to the market, and prices do not include the network's information. The essence of the problem is that address value includes both private value and common value components. The availability of a rental market would absorb the common value component, allowing market participants to trade on their private signals.

In evaluating IPv4 rental markets, a first question is whether networks would accept renting addresses, rather than using addresses substantially permanently. As noted in Section 3.4, changing addresses is ordinarily burdensome—identifying myriad systems and dependencies, a chore not unlike inventorying possessions when moving to a new home. But with suitable planning, changing addresses can become more routine—more like moving from one furnished apartment to another. For example, ISPs that connect standard home users can send automated instructions to change users' addresses, whereas ISPs that host servers typically face greater costs due to variation in configurations and dependencies. If IPv4 addresses were available for rent on favorable terms, some networks would surely put the addresses to use.

Furthermore, some aspects of IP addresses seem to lend themselves to a rental market. For example, there is little prospect of latent or concealed damage to an IP address. To the extent that IP addresses can be damaged, via spamming or other behaviors that harm the reputation of an address block, most such reputation information is quickly and publicly available. Compare concealed damage to a home or automobile, markets where rentals have proven problematic for some participants. With address condition fully observed to market participants, providers and renters can write contracts to disallow dam-

age and require payment in case of damage.

That said, other aspects of IP addresses make rentals difficult. For example, there is no strong central authority or other obvious process to “repossess” rented IP addresses at the conclusion of a rental. Indeed, if a renter continued to use the rented addresses and continued to announce their use via the Internet’s routing system, the provider would have little ability to reclaim the addresses from the renter. Conflicting address announcements test the stability of the routing system, but in any event most networks would resolve conflicts in favor of whichever network provides more useful and better-known content on a given address block—a rule that tends to favor recent use over contract rights. This problem could be solved by cryptographic verification of routing rights, as in Resource Public Key Infrastructure (Lepinski and Kent (2012)), though such systems have not yet been deployed by network operators and have prompted concerns about reliability, security, and overly-centralized authority. Finally, community norms may also stand in the way: many ARIN members expressed concern at the prospect of allowing IP addresses to be bought and sold, and ARIN rules only allow sale of addresses to networks, further excluding speculators. In the short run, it seems rentals will probably remain difficult and limited.

Policy discussions at ARIN suggest a commitment to putting IPv4 resources to use, avoiding waste, and making transactions easy and safe for participating networks. It seems rentals could advance these objectives: For example, a vibrant rental market would let a network provide its unneeded addresses for use by others, even if the network anticipates a future increase in market prices (and therefore does not want to sell the addresses). Furthermore, rentals could discourage efforts to try to “time the market” in buying or selling addresses. Policy changes could help facilitate these benefits. For example, a RIR could offer time-limited WHOIS updates, letting a network provide addresses for temporary use by others, with WHOIS guaranteed to revert after some predefined period.

The importance of rental markets extends beyond the market for IPv4 addresses. In general a durable good cannot be allocated efficiently if the only way to trade on information about future value is to hold the good in the current period; in that case, someone without a productive use for the good may nonetheless hold it in anticipation of price increases, to the exclusion of others who could immediately put it to productive use. Consequently, even if the market for the good is frictionless, a rental market is needed to achieve efficient allocation. Some markets have already developed mechanisms to facilitate rentals—for example, in the context of real estate, there are well-established rules for evicting a renter whose rights have expired. The lack of similar institutions for IP address rentals is likely to stymie such rentals. Meanwhile, real estate rentals entail genuine and inevitable frictions (moving costs, possibility of concealed damage, etc.), which pose important limits on short-term rentals in that context. In contrast, for some networks (most notably, providers of residential internet access) it is straightforward to move from one IP address range to another. Thus the potential benefits of IP address rentals should be

particularly large, if institutions develop to support such rentals.

Rental markets are likely to offer larger benefits when there is significant uncertainty about future prices. Technological change tends to create such uncertainty. In the IP address context, the long-run move to IPv6 leaves uncertainty as to the duration of continued use of IPv4. We see similar transitions in myriad other contexts: Transition from offset printing to digital printing yields legacy printing presses of uncertain value. Ever-more advanced microprocessor fabrication processes yield drops in prior generation plants. The prospect of self-driving cars raises questions of value of traditional cars and buses. Market participants may have differing views and information about transition dynamics, but rental markets can help aggregate this information and achieve allocative efficiency.

5.4 The Current State of the IPv4 Market

Whatever the price trends and trajectories, networks' most immediate concern is typically the *level* of prices—an urgent question for networks seeking to plan future expansion. RIRs report the resources that have been transferred through their paid transfer systems, facilitating the tabulation and analysis in Mueller et al. (2012)—reporting a total of 83 transactions totaling 6 million IPv4 blocks through June 2012. That said, RIR records lack information about prices: RIRs have no operational reason to collect prices, and to date no RIR has required networks to disclose prices as a condition of transfer.

To date, price data is available only for a subset of transactions: sales from networks in bankruptcy proceedings. Consistent with standard practice of litigation records open to public review, bankruptcy filings reveal the amounts received for debtors' assets, including IPv4 addresses. In one widely-publicized 2011 transaction, Microsoft paid \$7.5 million for approximately 666,000 IPv4 addresses—\$11.25 per address. Other transactions from sellers in bankruptcy have brought prices that are broadly similar.

It might seem natural to use initial transactions to predict future prices. But we question whether sales from bankrupt estates are useful in predicting prices. By all indications, early sales reflect buyers and sellers in unusual conditions. As of May 2013, networks could still obtain IPv4 addresses from ARIN to satisfy up to three months of documented need, at prices that are effectively zero. This explains why early sales of IPv4 addresses are so few and why early sales are unrepresentative. From a buyer's perspective, a paid transfer is most attractive if the buyer has a special need for advance allocation or predictability. (While a network can claim only three months of addresses directly from ARIN, rules allow a network to buy up to 12 to 24 months of addresses for a fee.) Early sellers also appear atypical, e.g. companies in bankruptcy that are unable to await a thicker future market. In short, we doubt that early transactions fit a simple model of IPv4 price trajectory, and early sales probably do not predict the future price of IPv4 addresses.

6 Concluding remarks

The impending scarcity of IPv4 addresses calls for economic research to facilitate suitable institutions and market rules. In many markets, participants had the benefits of a period of years to design such systems, and often geographically-isolated markets served as parallel laboratories to invent and test alternative approaches. In contrast, IPv4 addresses have been issued at de minimis price and will continue to be issued that way until RIRs have fully exhausted their supply—delaying the development of market institutions. Then, IPv4 scarcity will arrive in one fell swoop, raising the stakes for adjusting rules and offering less opportunity to refine rules over time. Thus, even though RFC 1883 formalized the successor technology IPv6 in 1996, to date networks have had little to no economic incentive to move to IPv6, nor any market-based economic incentive to use or transfer IPv4 resources efficiently. Meanwhile, because IPv4 resources are inherently portable, experiments in one region tend to affect behavior elsewhere—limiting the potential for regional innovation. Here too, we are reminded of experience in other markets. For example, the spectrum market quickly grew from non-existent to huge (due to a regulatory action creating the market). For lack of practitioners bringing decades of experience tinkering and fine-tuning, economists were distinctively important in designing that market. The same may prove true for IPv4 addresses.

Our proposed spartan rule builds on ARIN’s initial restriction—seeking to address the same negative externality, but offering somewhat more flexibility in order to avoid unnecessary deviations from efficient allocations. We believe ARIN could adopt our rule within its existing transfer framework and with minimal administrative burden.

Economic science might benefit from IPv4 markets almost as much as IPv4 markets benefit from economic science. As IPv4 markets develop, there will be an opportunity for empirical economic research on these and other questions. For example, will prices in fact be linear in block size, or will large blocks carry disproportionately higher (or lower) prices? Will prices over time be supermartingale, as our results suggest? Will the market experience a price bubble that confirms the findings of the experimental literature?

Figures

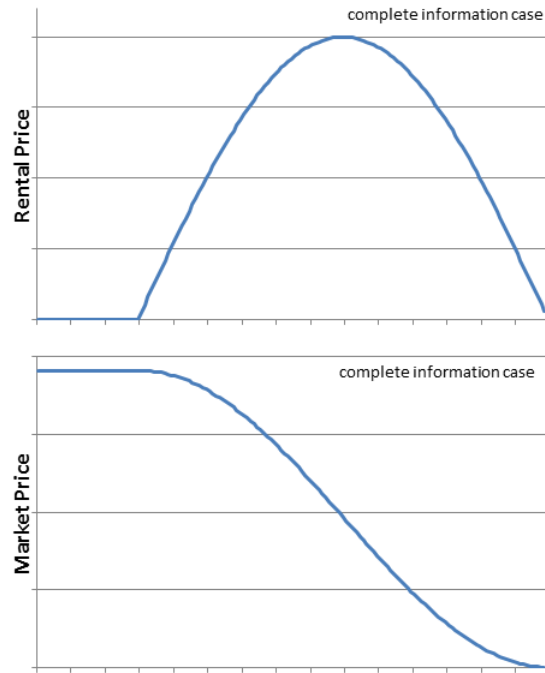


Figure 1: Price Profile of IPv4 Addresses

7 References

- Atila Abdulkadiroglu, Parag A. Pathak, Alvin E. Roth, and Tayfun Sonmez. The Boston public school match. In *American Economic Review, Papers and Proceedings*, volume 95, pages 368–371, 2005.
- Itai Ashlagi, Duncan S. Gilchrist, Alvin E. Roth, and Michael A. Rees. Nonsimultaneous chains and dominos in kidney paired donation – revisited. *American Journal of Transplantation*, 11:984–994, 2011.
- Nikhil Bansal, Don Coppersmith, and Baruch Schieber. Minimizing setup and beam-on times in radiation therapy. In Josep Díaz, Klaus Jansen, José D. P. Rolim, and Uri Zwick, editors, *APPROX-RANDOM*, volume 4110 of *Lecture Notes in Computer Science*, pages 27–38. Springer, 2006. ISBN 3-540-38044-2.
- Marjory Blumenthal and David Clark. Rethinking the design of the Internet: the end-to-end arguments vs. the brave new world. *ACM Transactions on Internet Technology*, 1:70–109, August 2001. ISSN 1533-5399.
- Benjamin Edelman. Running out of numbers: Scarcity of IP addresses and what to do about it. In Sanmay Das, Michael Ostrovsky, David Pennock, and Boleslaw K. Szymanski, editors, *Auctions, Market Mechanisms and Their Applications*, volume 14 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 95–106. Springer, 2009. ISBN 978-3-642-03820-4.
- Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97:242–259, 2007.
- Joseph Farrell and Garth Saloner. Installed base and compatibility: Innovation, product preannouncements, and predation. *American Economic Review*, 76: 940–955, 1986.
- Austan Goolsbee and Chad Syverson. How do incumbents respond to the threat of entry? evidence from the major airlines. *The Quarterly Journal of Economics*, 123(4):1611–1633, 2008.
- Bill Herrin. What does a BGP route cost? <http://bill.herrin.us/network/bgpcost.html>, 2008.
- Harold Hotelling. The economics of exhaustible resources. *The Journal of Political Economy*, 39:137–175, 1931.
- Matt Lepinski and Stephen Kent. RPKI architecture, February 2012. RFC 6480. <http://tools.ietf.org/html/rfc6480>.
- Eric S. Maskin. Auctions and privatization. In H. Siebert, editor, *Privatization*, pages 115–136. Institut für Weltwirtschaften der Universität Kiel, 1992.

- John McMillan, Paul Milgrom, Bradley Miller, Bridger Mitchell, Daniel Vincent, and Robert Wilson. Simultaneous ascending auctions with package bidding. 1998.
- Milton Mueller, Brenden Kuerbis, and Hadi Asghari. Dimensioning the elephant: An empirical analysis of the IPv4 number market. *Proceedings of the Telecommunications Policy Research Conference*, 2012.
- Michael Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80:2595–2648, 2012.
- Michael Ostrovsky and Michael Schwarz. Adoption of standards under uncertainty. *RAND Journal of Economics*, 36:817–833, 2005.
- E. Ronn. NP-complete stable matching problems. *Journal of Algorithms*, 11: 285–304, 1990.
- Alvin Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, pages 991–1016, 1984.
- Vernon L. Smith, Gerry L. Suchanek, and Arlington W. Williams. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, 56:1119–1151, September 1988.