Maintaining the Progress of IPv6 Adoption^{☆,☆☆}

Mehdi Nikkhah

*Multimedia and Networking Lab, Electrical and Systems Engineering Department,
University of Pennsylvania, 200 S 33rd St. Philadelphia, PA 19104, U.S.A.*

Abstract

Even though IPv6 adoption has accelerated in recent years, the complete migration of the Internet still faces many challenges. There are multiple factors that can potentially affect, negatively or positively, the future adoption of IPv6 by various Internet stakeholders. This situation begs the question of “what could be done to avoid derailing the IPv6 adoption progress?” and “how different factors can help maintain its acceleration?” There has been significant interest in those questions, and the paper proposes a series of models to investigate and shed light on them. The results confirm the effect of a number of known factors, while also providing new insight. Particularly, they highlight the destabilizing impact of disagreement across Internet Service Providers (ISPs) on immediate migration to IPv6, and show the benefits of minimal coordination among them in offering IPv6 as an option. They also show what affects technology transition in a large network with complex dependencies such as the Internet. Using robustness analyses, the findings are shown to hold in the presence of different assumptions and scenarios.

Keywords: IPv6, adoption, Modeling, Two-sided Market, Game Theory

Guide to the color codes:

- Text added in this revision.
- Text removed in this revision.

1. Introduction

IANA announced in February 2011 that the free pool of IPv4 addresses is depleted, and even if IPv4 addresses scarcity has not yet materialized every-

[☆]A preliminary version of this work was presented at the IFIP Networking 2014 Conference [1].

^{☆☆}This work was supported by NSF grant CNS-1116039.

Email address: mnikkhah@seas.upenn.edu (Mehdi Nikkhah)

where, we are slowly but surely headed in that direction. IPv6 was designed to address this issue, and even though recent studies show its adoption is accelerating, there are hurdles that can impede or slow down its progress in the future. Although these hurdles are not (anymore) of a technical nature, years of technology disparity between IPv4 and IPv6 caused a marginal adoption of IPv6 across major Internet stakeholders [2], which in addition to incompatibility of the two technologies forced the use of translation mechanisms to allow IPv6-only users access to the IPv4-only Internet. These translation mechanisms are widely used today by ISPs such as CERNET2 in China, and Verizon Wireless and T-Mobile in the U.S. CERNET2 [3] (an academic network), already had over 400k *IPv6-only* users in 2009, is expected to reach 3 million by the end of 2015 (see [4, 5]), and uses “IVI”, which translates IPv4 traffic to IPv6 and vice versa. Similarly, Verizon Wireless and T-Mobile are now primarily relying on IPv6 addresses for new cell-phone subscribers [6, 7], and use “NAT444” and “464XLAT” as their translation mechanisms, respectively. While necessary for a transition, the quality degradation those mechanisms introduce [8, 9, 10, 11] reduces motivation for the new users to adopt IPv6. This is an instance of hurdles in front of the progression of IPv6 adoption in the future. Our initial intuition was that besides the above instance, the distributed structure of the Internet can also affect the progression of IPv6 adoption. Specifically, the benefit of migrating to IPv6 depends to a large extent on what others in the Internet do. This is not an uncommon situation (*e.g.*, see [12] for a related discussion in the context of Internet security protocols), but uncertainty in the decisions of others can significantly delay the adoption of a new technology.

A goal of this paper is, therefore, to explore and explain strategies that can derail or speed up the current progress of IPv6 adoption. These strategies require careful assessments as we are dealing with a highly decentralized system (the Internet). To better understand the extent to which these strategies can affect IPv6 adoption, several simple yet representative scenarios and models were developed. The focus of these models is on the decision making process of independent and decentralized stakeholders across the Internet, and how those decisions can affect IPv6 adoption. We acknowledge up-front the many simplifying assumptions these models rely on (a necessity in most modeling efforts), and their lack of completeness. However, they incorporate major aspects of IPv6 adoption decisions, namely, (i) heterogeneity in the Internet stakeholders making decisions; (ii) a representative sample of available technology options; and (iii) the dependencies that exist across decisions.

Our findings from these models indicate that independent decision making process of ISPs can negatively affect IPv6 adoption. In other words, disagreement between ISPs on connectivity option offerings, adds uncertainty to the factors that affect IPv6 adoption decisions of the Internet stakeholders, and makes it hard to identify winning strategies. As a result of this uncertainty, migration to IPv6 slows down, or at the very least becomes haphazard. Another finding of the models is that even minimal coordination among ISPs in offering connectivity options, *e.g.*, an Internet-wide consensus on offering IPv6 as one of the connectivity options, can significantly improve our abilities to identify

strategies that hasten the IPv6 migration process. Although consensus alone is not sufficient, it makes it easier for the Internet stakeholders to identify winning strategies that can, at the same time, speed up the migration of the Internet to IPv6.

The paper’s contributions are, therefore, two-fold:

- (i) It shows how distributed decision making of the Internet stakeholders, in the presence of competing solutions to the problem of IPv4 address scarcity, can negatively affect identifying winning strategies, and therefore, linger the (current) uncertainty in IPv6 adoption; and
- (ii) It illustrates how the introduction of limited coordination among ISPs, which is not in itself enough for IPv6 success, can help determine the impact of different parameters on IPv6 adoption, and hence, facilitate a smoother migration process.

The rest of the paper is structured as follows. Section 2 discusses the framework of the problem, including the Internet stakeholders, connectivity options, and scenarios. Sections 3 introduces the models in two categories of disagreement and consensus. Section 4 and 5 explore the outcome of the models with a certain set of assumptions, and provide the key findings. Section 6 investigates the robustness of our findings to different modeling assumptions and extensions. Section 7 briefly reviews related works, with Section 8 summarizing the paper’s findings and recommendations.

There are too many differences between this paper and its preliminary version [1], and therefore, we only list the major changes, namely, generalized models, extended range of numerical analyses, and robustness analyses section.

2. Problem Framework

There are many factors that arguably affect the adoption of IPv6, and any (tractable) model is unlikely to account for all of them and their variations across stakeholders. Our models operate within a certain framework, and this section specifies the outline of that framework by introducing the Internet stakeholders, their connectivity options, the inter-dependencies between their decisions, and the scenarios in which they interact.

2.1. Internet stakeholders

We distinguish between *three* types of Internet stakeholders: Internet Service Providers (ISPs), Internet Content Providers (ICPs), and Internet Content Consumers (users). ISPs derive revenues from providing Internet connectivity to both ICPs and users, and are, therefore, concerned with the choices and costs of the technologies used to implement this connectivity. They make the ultimate decision to offer IPv6 connectivity to the other two stakeholders, hence, they play the most significant role in IPv6 adoption across the Internet. ICPs obtain the bulk of their revenues from users that connect to them through ISPs. Hence, their focus is on the quality of their connectivity to users and how it may affect their revenues, as well as any cost they may incur to upgrade their

existing infrastructure to support a new connectivity option, *e.g.*, IPv6. Finally, users purchase Internet connectivity from ISPs, and use it primarily to connect to ICPs (and to a lesser extent to each others). Hence, they are affected by the cost of Internet connectivity and by its quality.

2.2. ISP's connectivity options

ISPs are the providers of Internet connectivity, and therefore control technology choices. Although IPv6 adoption is on the verge of happening, implicit to our modeling effort lies the fact that IPv6 still faces competing solutions. Among those available technology choices ISPs may choose from to accommodate customer growth, we consider three representatives.

The first choice an ISP can make is to simply continue using public IPv4 addresses. This has the advantage of full compatibility with the current Internet, but given the growing scarcity of public IPv4 addresses is likely to quickly involve added costs, *e.g.*, to purchase public IPv4 addresses from an address market such as Hilco Streambank IPv4 Address Marketplace.

The second option an ISP can rely on is to use private IPv4 addresses together with Carrier-Grade NATs (CGNs). Unlike public IPv4 addresses, private IPv4 addresses can be reused and so are not scarce. CGNs are required to allow connectivity to the public Internet, but the technology behind CGNs is mature. Private IPv4 addresses also have the benefit of letting ISPs defer a potentially expensive upgrade of their network to IPv6. The main disadvantage (to the ISP) is the cost of CGNs, which grows as more users are assigned private IPv4 addresses.

IPv6 is the third option. IPv6 addresses are not scarce, but like private IPv4 addresses will require some form of “translation,” *e.g.*, NAT64 [13] or DS Lite [14], to allow IPv6 users to communicate with the IPv4 Internet. IPv6 $\leftrightarrow$ IPv4 translation is less mature than that for private IPv4 addresses, and may therefore be initially more expensive. On the flip side, even if the exact time-frame remains unclear, the need for translation, and therefore its cost, should disappear as the Internet eventually migrates to IPv6.

2.3. Decision dependencies

As alluded to, although ISPs choose Internet technologies, their decisions, including *pricing*, depend heavily on users and ICPs. For example, an ISP offering both IPv6 and (public) IPv4 connectivity might discount the IPv6 service, thereby attracting users to that option and lowering the need for (expensive) public IPv4 addresses. However, more IPv6 users also means higher translation costs, unless this entices more ICPs to become IPv6 accessible thereby lessening the need for translation. This creates a complex web of dependencies, whose impact is amplified by the distributed decision process that prevails in the Internet. As we shall see, this can make devising sound (profit maximizing) strategies difficult if not impossible. [To understand the dependencies better, we present the system dynamics diagram of IPv6 adoption by ICPs and users in Fig. 1.](#) It can be seen that there are externalities between users and ICPs as

well as externalities due to competition among ICPs (captured by thin links in the diagram). However, in our models, we ignore the latter, since it is hard to quantify them. Note how the overall access quality for IPv6 users improves when more ICPs adopt IPv6. This is due to the fact that the traffic from IPv6 users does not go through translation devices when an ICP is IPv6-enabled. This is a positive feedback (shown by R for Reinforcement), because larger IPv6 user-base also means higher incentives for ICPs to adopt IPv6. As alluded to earlier, translation quality degradation is shown to negatively impact the overall access quality for IPv6 users, and the number of IPv6 users makes this quality worse due to congestion in translation devices. We show in the next sections that these dependencies indeed play a critical role in IPv6 adoption, and by breaking only one of the links in the web of dependencies, the outcomes change drastically.

Figure 1: System dynamics diagram of IPv6 adoption by ICPs and users

2.4. Scenarios

In many technology adoption instances presence of multiple entrants, and lack of consensus on a single choice among stakeholders can prevent a full market penetration by any of those choices. While competition of alternative solutions can be helpful in keeping the evolution of a technology on the right track, consensus on one choice makes a full market penetration faster and easier. In the case of IPv6, a full market penetration is required, if the Internet is to avoid permanent traffic translation, therefore, the Internet Engineering Task Force (IETF) standardized IPv6 as the replacement for IPv4. However, due to the hurdles in front of IPv6 adoption, other alternative solutions have become popular among some ISPs.

As different ISPs manage separate Autonomous Systems (ASes), their decisions are to some extent independent of each other. This heterogeneity among ISPs can lead them to offer (at least temporarily) different connectivity solutions. Since ISPs provide Internet connectivity, their heterogeneous decision making has a more significant impact on IPv6 adoption compared to other Internet stakeholders. Therefore and in order to investigate this impact, we consider two major scenarios: (i) a scenario in which ISPs disagree on immediately offering IPv6 connectivity to their users; and (ii) a scenario in which all ISPs offer IPv6 along with other connectivity options to their users. Next, we describe these two scenarios in more details.

2.4.1. Disagreement on offering IPv6

In this scenario, one ISP is always assumed to offer IPv6, as otherwise the outcome is trivial, *i.e.*, stagnation in IPv6 adoption, while the other ISP offers either public or private IPv4 addresses.

Given that the main competition IPv6 faces is the incumbent IPv4 Internet, we consider the case of two ISPs, one having embraced IPv6 as the technology of choice for its new customers¹, or an ISP that plans to join the market but does not have the option of acquiring IPv4 addresses, while the other ISP has decided to defer any migration and to simply acquire additional public IPv4 addresses to accommodate new customers. The first ISP needs to deploy address translation devices to allow its new (IPv6) customers to connect to the legacy IPv4 Internet. This cost grows with the number of users that choose IPv6, and decreases as more ICPs become IPv6 accessible². Conversely, while the second ISP does not incur translation costs, it needs to purchase public IPv4 addresses for its new customers. Those costs are expected to rise as public IPv4 addresses become scarcer. Note that in this scenario, we do not assume that an ISP with a current IPv4 user-base switches instantly to an IPv6-only infrastructure. The main focus of this scenario is however, on ISPs that are going to join or have recently joined the market, but do not have the option of acquiring IPv4 addresses, which is the case in many Asian countries with already exhausted IPv4 address pools.

Another variation of this scenario is when no ISP wants to incur the cost of purchasing more public IPv4 addresses (or those addresses are unavailable for purchase). ISPs that defer upgrading to IPv6 would then rely on private IPv4 addresses. Offerings based on either IPv6 or private IPv4 addresses both require translation (CGNs) to connect to the public IPv4 Internet. Translation costs for private IPv4 are likely to be lower than for IPv6, if only because of more mature technology and/or greater operational familiarity and compatibility with the current Internet. On the flip side, translation costs for private IPv4 keep increasing as more users join, independent of how many ICPs become IPv6 accessible. We describe this scenario in Appendix A.

2.4.2. Consensus on Offering IPv6

In this scenario, there exists a global consensus on offering IPv6 (along with other service types), as a technology of choice to users, hence, all ISPs offer IPv6 and another service, *e.g.*, public IPv4.

On the technology choice front, this scenario is identical to the first one, namely, both IPv6 and public IPv4 are available as connectivity options. The main difference is that the two options are now systematically offered by all ISPs, and therefore priced internally, as opposed to competitively, to maximize their

¹T-Mobile has recently started to only assign IPv6 addresses to its Android 4.4 users (see [7]).

²Translation costs are assumed proportional to the volume of traffic that needs to be translated, *i.e.*, higher capacity devices are needed.

own profit. The price difference is a means of modeling, and can be interpreted as the cost of extra services that ISPs offer (for free) along with their IPv6 services, but charge users for those same services in IPv4, *e.g.*, static addresses (<http://www.vo.lu>) etc. This scenario is equivalent to having a monopolistic ISP that sets the price of both connectivity choices.

3. Models

Based on the scenarios of the last section, we developed models that capture the interactions and decision dependencies of ISPs, ICPs and users. As alluded to in section 2.3, the decisions of users depends on the decisions of ICPs and ISPs, and vice versa. ISPs are the selectors of the technology and affect the interactions of the other two stakeholders through their decisions. This framework is common to other environments, *e.g.*, gaming platforms, where the number of game developers and the number of gamers are affected by the decisions of the console provider. Analyzing these frameworks is typically through a two-sided market setting [15]. The ISP is the market maker through its offering of connectivity options, while users and ICPs are the two sides of the market that derive value from each other through the ISP.

We assume that at each step, new and existing users evaluate the Internet connectivity choices available to them through their local ISP(s)³ and select the one yielding to the highest *utility*. Even though many users are oblivious to the underlying connectivity technology, here, the option of having IPv4 or IPv6 is simply a choice between a high or low quality connection, since a user that chooses IPv6 goes through translation to access more than 90% of the websites [16], and this makes the overall connection quality lower than that of an IPv4 user. Therefore, assuming users choose between IPv4 or IPv6 is equivalent to the assumption that users are sensitive to quality degradation and, when given the choice, choose between different quality tiers (with varying costs). This sensitivity to quality is shown in a study by Dobrian *et al.* [17], which focuses on user engagement when the video quality changes, and another study by Zander *et al.* [18] that focuses on sensitivity of game players to connection quality. A similar model which excludes users as decision makers is presented in [2].

One obvious shortcoming of this model is the lack of inertia in decision making of users, *i.e.*, every user decides at each time epoch, therefore, in section 6 we investigate the robustness of our results in scenarios where the users face some form of inertia, *e.g.*, contractual agreements. Moreover, Users are assumed heterogeneous, but primarily in their sensitivity to connectivity quality⁴.

³According to <http://www.broadbandmap.gov/summarize/nationwide>, over 99% of the U.S. population can choose from two or more ISPs, while this figure is 90% in Europe (see <http://goo.gl/MjTPJ6>).

⁴Coarser grain heterogeneity is also possible, *e.g.*, between, say, residential and enterprise users, but adds significant complexity to the model. Similarly, heterogeneity in price sensitivity can also be included, but with again a cost in terms of complexity.

We further assume (see [10] for a related discussion) that address translation devices, if used, are the main contributors to degradation in connectivity quality/functionality. It is true that other factors can affect connectivity quality, but they are present whether in IPv4 or IPv6, and today, the main difference between native IPv4 connectivity and IPv6 (or private IPv4) is the presence of translation devices (see [2] for a comparison between the maturity of IPv6 compared to IPv4 in practice).

ICPs are part of the current Internet, therefore, they already have a public IPv4 address, and their only decision is whether or not to become IPv6 accessible. They incur a cost when doing so (upgrading their existing IPv4 infrastructure and/or update of operational processes), but unlike users that can revert their decisions, an ICP’s decision to become IPv6 accessible is irreversible (once incurring the upgrade cost). Next, we present the utility functions of the Internet stakeholders.

3.1. Users utility

Users derive a *unit* value from Internet connectivity, with price and quality affecting their overall utility. An alternative model assumes heterogeneous values for different connectivity options, however, since we use pricing as the control knob of the ISPs, the former presentation is chosen (the outcomes are nevertheless similar). Recall that quality is assumed to be primarily affected by (the presence of) translation devices. A user’s utility is then captured through the following expression:

$$U_{\text{user}}(\sigma) = 1 - \underbrace{p_R}_{V_R} - \sigma a_R \gamma_R, \quad (1)$$

where R indexes connectivity options, p_R is the price of type R connectivity ($p_{\text{pub. IPv4}} > p_{\text{IPv6}} > p_{\text{priv. IPv4}}$) (alternatively V_R is the value of option R), $a_R \in [0, 1]$ quantifies (translation) quality impairments for connectivity option R , if any (a_R is 0 for public IPv4 and positive for both private IPv4 and IPv6), γ_R is the fraction of the Internet (ICPs) affected by those impairments, and σ denotes a user sensitivity to quality impairments.

3.2. ICPs utility

ICPs derive revenue from users, and that is affected by connectivity quality [19]. A major factor in an ICP’s decision to become IPv6 accessible⁵ is, therefore, the impact this decision can have on the revenue it generates from IPv6 users, and how this compares to the cost of upgrading to IPv6 (or convincing its hosting provider to upgrade). Revenue improvements depend on the

⁵As participation in events such “World IPv6 Launch Day” demonstrates, there are obviously many other possible reasons for an ICP to become IPv6 accessible. However, even when those other motivations prevail, the importance of preserving connectivity quality remains, *e.g.*, through the enforcement of some form of “white-listing.”

number of IPv6 users and how they are affected by the ICP’s adoption of IPv6. In particular, and as shown in [2], IPv6 connectivity quality is now mostly on par with IPv4, hence, the main benefit of native IPv6 access is to eliminate the need for translation.

The cost of upgrading to IPv6 is largely a function of the “size” of the ICP’s infrastructure. For simplicity, this size is assumed proportional to the Internet user-base (the traffic volume an ICP sees grows with the Internet). The net utility in(de)crease an ICP derives from becoming IPv6 accessible can, therefore, be captured as follows:

$$\Delta_6(\text{ICP}) = \beta n_6 a_6 - S_{\text{infra}} \theta c_6 \quad (2)$$

where βn_6 is the fraction of IPv6 users that an ICP can benefit from, a_6 is the per-user revenue gain from eliminating translation, and θc_6 is the per-user upgrade cost of the ICP’s infrastructure (of size S_{infra}). β and θ capture heterogeneity in revenue and cost, respectively, across ICPs. In addition to the cost component of upgrading the infrastructure of an ICP, there are costs associated with “breakage”, misconfigurations, etc., however, other studies (such as [2]) have shown that native IPv6 is now on par with IPv4 in operational environments, and therefore, these costs are not excessively higher in IPv6 compared to IPv4. Moreover, even if such costs still exist, they are dwarfed by the infrastructure upgrade costs, and therefore, we ignore them in our models. *Note that ICPs gain the same revenue from IPv4 and IPv6 users when those users connect through a native (non-translated) connection. However, when a user goes through translation devices to access an ICP, the quality of access drops. Therefore, when an ICP adopts IPv6, those users who used to access it through translation, would experience the higher quality of native connection, hence, increasing the revenue of the ICP.*

3.3. ISP utility

An ISP’s utility (profit) depends on revenues derived from users⁶ (*i.e.*, home users, or small businesses), and costs associated with providing services. In our models, we focus on added cost components that an ISP incurs to provide services to its new users in the future, and ignore the costs associated with maintenance or infrastructure upgrade. As costs differ across connectivity options, we introduce the ISP’s utility function separately for each.

3.3.1. Public IPv4 only

An ISP that only offers public IPv4 connectivity has a utility function of the form:

$$\Pi_{\text{pub. 4}} = n_4 p_4 - C(n_4 - 1)_+^2 \quad (3)$$

⁶We ignore revenues from ICPs or enterprise customers, as they are mostly independent from an ISP’s connectivity choices.

n_4 is the number of users willing to pay p_4 for public IPv4 connectivity, while $C(n_4 - 1)_+^2 = C \max(0, n_4 - 1)^2$ is the acquisition cost of the $(n_4 - 1)$ additional public IPv4 addresses the ISP needs beyond the “unit” block it already owns (to accommodate its existing users). The quadratic function used for address acquisition costs seeks to capture the growth in the price of public IPv4 addresses due to increasing scarcity. Section 6 changes this assumption, and investigates the impact of other functions on the model’s outcome. Note that the ongoing maintenance costs are ignored, as they are incurred by all ISPs regardless of connectivity options.

3.3.2. IPv6 only (and IPv6↔IPv4 translation)

An ISP offering IPv6 connectivity has a utility of the form:

$$\Pi_6 = n_6 p_6 - D_6 n_6 \gamma_6, \quad (4)$$

with n_6 the number of users choosing IPv6 connectivity at a price of p_6 , and $D_6 n_6 \gamma_6$ the translation cost for those users. This expression assumes each user generates 1 unit of traffic distributed uniformly across ICPs, so that if γ_6 ICPs are not IPv6 accessible, $n_6 \gamma_6$ units of traffic must be translated at a unit cost of D_6 . Note that there are multiple cost components an IPv6 ISP incurs, such as cost associated with infrastructure upgrade, translation, breakage, customer support, etc. In our model, we ignore the first one, namely, the infrastructure upgrade cost, since it is mostly incurred during the life-cycle upgrades of the ISP’s routers, switches, etc., or can be seen as an initial investment to enter the market. [This strategy is aligned with what many entities suggested, including but not limited to the U.S. Department of Commerce IPv6 Task Force suggests \[20\], Cisco \[21\], etc. Also, this is based on the fact that many ISPs around the world have already adopted IPv6, but not many of them actually offer it to their users \[22\].](#) Our focus is on those costs that an ISP “with an IPv6 infrastructure in place” incurs, exclusively in the future, to provide IPv6 services to its user-base. Therefore, the cost component of the above utility function, *i.e.*, $D_6 n_6 \gamma_6$, includes costs associated with providing translation services (*e.g.*, CGNs), solving issues related to CGN instability, and added customer support due to service breakages (for an overview of these costs see [10, 23]). This is a per user cost, and it depends on the volume of translated traffic, *i.e.*, when all ICPs adopt IPv6, this cost component shrink to zero.

3.3.3. Public IPv4 and IPv6

An ISP offering both public IPv4 and IPv6 has a utility that is simply the sum of Eqs. (3) and (4) and is of the form:

$$\Pi_{46} = n_4 p_4 - C(n_4 - 1)_+^2 + n_6 p_6 - D_6 n_6 \gamma_6. \quad (5)$$

The next subsection explains the decision mechanism of the Internet stakeholders, and the timing of those decisions.

3.4. Decision Mechanisms and Timing

In all scenarios, ISPs first announce a price for each of the connectivity options, with users then choosing one in a best response fashion, *i.e.*, they select the option that maximizes their utility. ICPs decide whether or not to become IPv6 accessible in the third and last stage of the game, again in a best response manner and based on the number of users that have chosen IPv6. ISPs are assumed aware of the rationale and economic incentives guiding users and ICPs decisions, *e.g.*, based on surveys of users and ICPs. Hence, they set prices that maximize their own profit, *i.e.*, by solving the above sequential decision process in reverse order. In the disagreement scenario, where not all ISPs agree to offer IPv6 immediately, we assume the two ISPs compete for the users by playing a best response game between themselves, and their strategies are the prices of the services. In the consensus scenario, however, the problem reduces to an internal optimization of a single ISP that offers both services, namely, IPv4 and IPv6 connectivity.

An alternate game would have users and ICPs aware of each others decisions, deciding simultaneously rather than sequentially. This assumes that users are able to predict how ICPs will respond to their decisions and vice versa, and makes for a more complex and possibly less realistic game (neither users nor ICPs may have access to the necessary information). More importantly, the outcomes are similar to those of the simpler sequential game. As a result, we focus on the latter. Another alternative scenario is when the decision making process of ICPs is in a different time scale compared to ISPs and users, *i.e.*, ICPs re-evaluate their decisions less often than ISPs and users. We relegate the analysis of this scenario to Appendix C.

Note that a model incorporating external incentives, *e.g.*, government subsidies, etc., would have the same framework and outcomes; the only difference would be the scaling of the cost factors. In fact, in Section 5, where we analyze the outcome numerically, a decrease in the costs can be interpreted as external incentives. Additionally, the models and their assumptions hold even if in the real world content consumption is local, *i.e.*, users in a geographical region consume the content generated and mostly hosted in that same region. This is because the models do not have any restrictions on whether all ISPs and ICPs in the world should participate in this game.

4. Model Solution

4.1. Disagreement Scenario

This section considers scenario 2.4.1, which involves (two) ISPs competing for users and offering different connectivity options. One ISP relies on IPv6, but the other has deferred upgrading to IPv6. Instead, it chooses to either incur the (growing) cost of acquiring public IPv4 addresses, or to assign private IPv4 address to new users and rely on translation (CGNs) to connect them to the public Internet. Here we present the solution to the former scenario (public

IPv4 vs. IPv6), and relegate the latter (private IPv4 vs. IPv6) to Appendix A.

Specifically, we assume *rational* and *myopic* ISPs that engage in a repeated multi-stage game played each time the Internet user population increases by $\delta < 1$ new users. Again, in this scenario, one ISP offers IPv6 and the other stays with public IPv4 connectivity. Public IPv4 has an edge when it comes to connectivity quality ($a_6 > 0$), but that edge is present only for the fraction γ_6 of ICPs that require translation. Conversely, the disadvantage of public IPv4 is the likely cost of acquiring additional public IPv4 addresses.

As per Eq. (1), users' utility depends on price (p_R), quality of connectivity (a_R), and the fraction γ_R of ICPs affected by quality impairment associated with connectivity option R . γ_R is assumed known to users, and in the case of IPv6 depends on the outcome of the previous round of the game, *i.e.*, how many ICPs have become natively accessible. For tractability purposes, we assume that σ , *i.e.*, the sensitivity of users to quality impairments, is uniformly distributed in $[0, 1]$. Section 6 relaxes this assumption, and investigates the outcome with a more general distribution.

Hence, in round i of the game and assuming IPv6 and public IPv4 ISPs announced prices of p_6 and p_4 , users and ICPs decisions proceed as follows. Based on Eq. (1), a user with quality sensitivity σ chooses IPv6 if, $1 - p_6 - \sigma a_6 \gamma_6^{(i-1)} \geq 1 - p_4$, where $\gamma_6^{(i-1)}$ denotes the fraction of ICPs not yet IPv6 accessible after round $(i - 1)$ (this information is available after each round, with $\gamma_6^{(0)} = 0$ for completeness). Hence, the fraction $\sigma_6^{(i)}$ of (new and existing⁷) users choosing IPv6 in round i satisfies

$$\sigma_6^{(i)} = \begin{cases} 0 & \text{if } p_4 - p_6 < 0 \\ \frac{p_4 - p_6}{a_6 \gamma_6^{(i-1)}} & \text{if } 0 \leq p_4 - p_6 \leq a_6 \gamma_6^{(i-1)} \\ 1 & \text{if } p_4 - p_6 > a_6 \gamma_6^{(i-1)} \end{cases} . \quad (6)$$

The dependency on the price differential $p_4 - p_6$ is intuitive. For example, when IPv6 is priced higher than IPv4, IPv6 adoption is zero, while when the discount for IPv6 is larger than the quality penalty perceived by the most quality sensitive user ($\sigma = 1$), then all users select IPv6.

An ICP reevaluates its IPv6 adoption decision once knowing the outcome of users' decisions. Again, for tractability purposes, ICPs are assumed to derive homogenous revenues from users (*i.e.*, $\beta = 1$), but we relax this assumption in Section 6.

From Eq. (2), ICPs adopt IPv6 if the difference between the added revenue, $n_6 a_6$ (remember $\beta = 1$), this generates and the upgrade cost $S_{\text{infra}} \theta c_6$ is positive. The latter depends on the current size of the ICP's infrastructure, S_{infra} , which is proportional to the Internet user-base in round $(i - 1)$, *i.e.*, $1 + (i - 1)\delta$

⁷In section 6 we show that our results remain qualitatively similar even if users have inertia in decision making, *e.g.*, contractual agreement, etc.

(where 1 is the size of the existing Internet user population). Conversely, the revenue increase created by becoming IPv6 accessible is proportional to the number of users choosing IPv6 in round i , *i.e.*, $n_6^{(i)} = (1 + i\delta)\sigma_6^{(i)}$. Assuming θ is uniformly distributed in $[0, 1]$ (which we again relax in Section 6), ICPs for which becoming IPv6 accessible yields a positive profit in round i are those with $\theta \leq \theta_6^{(i)}$ (conversely, the fraction of IPv6 accessible ICPs after round i is $\gamma_6^{(i)} = 1 - \theta_6^{(i)}$), where

$$\theta_6^{(i)} = \begin{cases} \frac{ka_6}{c_6} & \text{if } p_4 - p_6 > a_6\gamma_6^{(i-1)} \\ \frac{ka_6\sigma_6}{c_6} & \text{if } \frac{\gamma_6^{(i-1)}(1-\gamma_6^{(i-1)})c_6}{k} \leq p_4 - p_6 \leq a_6\gamma_6^{(i-1)}, \\ 1 - \gamma_6^{(i-1)} & \text{Otherwise} \end{cases}, \quad (7)$$

where for notation simplicity $k = \frac{1+i\delta}{1+(i-1)\delta}$ is the relative growth in user population between rounds $(i-1)$ and i .

The first expression of Eq. (7) corresponds to all users selecting IPv6, *i.e.*, $\sigma_6^{(i)} = 1$, which yields the maximum possible adoption of IPv6 among ICPs. IPv6 adoption progressively decreases as fewer users select IPv6 (second expression), down to no less than $1 - \gamma_6^{(i-1)}$, which reflects the fact that ICPs that upgraded to IPv6 in an earlier round do not revert their decisions.

Eqs. (6) and (7) are known to the two competing ISPs, which use them to optimize their own utility functions, as expressed in Eqs. (3) and (4). This yields the following expressions for optimal prices, where for simplicity we omit the index i and use $\gamma_6 = 1 - \theta_6$.

$$p_4^* = \underset{p_4}{\operatorname{argmax}} \{ (1 + i\delta)(1 - \sigma_6)p_4 - \quad (8)$$

$$C(((1 + i\delta)(1 - \sigma_6)) - 1)_+^2 \} \\ p_6^* = \underset{p_6}{\operatorname{argmax}} \{ (1 + i\delta)\sigma_6 p_6 - D_6\sigma_6(1 + i\delta)\gamma_6 \}. \quad (9)$$

The two equations are coupled through Eqs. (6) and (7).

Explicitly solving this joint optimization is difficult⁸. It can be formulated as the solution of a best response game between the ISPs, each successively announcing and reacting to the other's price. In general, the game does not have a Nash Equilibrium to which prices would converge. In particular and as illustrated in section 5.1, instances of "cycles" in the ISPs' search for optimal prices arise in many cases. In other words, competition (or disagreement) between ISPs on the basis of connectivity makes identifying rational operating (pricing) points difficult.

⁸Analytical solutions can be obtained, but are mostly negative results, *e.g.*, the absence of a Nash Equilibrium, which do not shed insight into the problem. Hence, we resort to numerical investigations to explore the solution space.

Interestingly but not surprisingly, dependencies between Internet stakeholders' decisions are largely responsible for this. In particular, if ICPs' decisions were independent of those of users (or proceeded at a much slower pace), the game would typically admit a unique Nash Equilibrium (see Appendix C).

4.2. Consensus Scenario

A scenario where all ISPs offer IPv6 and another alternative, *e.g.*, public IPv4, is equivalent to a monopolistic ISP that serves all of the users. Choices need to be preserved, as users (and ICPs) are likely to remain heterogeneous in their willingness to accept a migration to IPv6. However, connectivity options should not be the basis on which ISPs compete. In other words, this scenario is equivalent to a pricing problem including a single provider with two types of products.

Consider an ISP offering its users (new and existing) the choice between traditional public IPv4 connectivity and IPv6 connectivity at prices of p_4 and p_6 , respectively. As in the previous section, users that opt for IPv6 must undergo translation when connecting to the γ_6 fraction of ICPs that are not yet IPv6 accessible. As before, translation introduces impairments of relative magnitude a_6 . Similarly, the ISP incurs a cost of D_6 per unit of traffic that needs translation. The ISP has an existing user-base of unit size, and therefore owns a unit-size block of public IPv4 addresses. If it needs additional public IPv4 addresses, it acquires them at a cost that, again as before, grows quadratically, *i.e.*, based on Eq. (3). ICPs decide to become IPv6 accessible following the same process as that of last section. We describe next how the ISP selects the prices p_4 and p_6 that maximize its profit.

Growth in the Internet user population again proceeds in steps of size δ that coincide with epochs where the ISP adjusts its prices p_4 and p_6 . Choosing optimal prices involves solving the following optimization problem

$$\begin{aligned}
 (p_4, p_6) = \operatorname{argmax}_{(p_4, p_6)} \bigg\{ & \\
 (1 + i\delta) \left(1 - \frac{p_4 - p_6}{a_6 \gamma_6^{(i-1)}} \right) p_4 - & \\
 C \left(\left((1 + i\delta) \left(1 - \frac{p_4 - p_6}{a_6 \gamma_6^{(i-1)}} \right) \right) - 1 \right)^2_+ & \\
 + (1 + i\delta) \left(\frac{p_4 - p_6}{a_6 \gamma_6^{(i-1)}} \right) p_6 & \\
 - D_6 (1 + i\delta) \left(\frac{p_4 - p_6}{a_6 \gamma_6^{(i-1)}} \right) (1 - \theta_6^{(i)}) \bigg\}, & \quad (10)
 \end{aligned}$$

where $\gamma_6^{(i-1)} = 1 - \theta_6^{(i-1)}$ is known, while $\theta_6^{(i)}$ needs to be replaced by its expression from Eq. (7). Note that different expressions must be used for $\theta_6^{(i)}$

depending on the value of $p_4 - p_6$. It is the need to consider those different cases that makes solving Eq. (10) cumbersome though not impossible.

Except for the fact that the optimal price for public IPv4 always satisfies $p_4 = 1$ (actually just below 1 to ensure positive utility), the expression for an explicit solution for Eq. (10) sheds little light on the role of different parameters, the reader is referred to [24] for details, and we instead rely on numerical examples to explore the range of outcomes. The next section discusses the results of numerical analyses of the models.

5. Results

In this section, we first present the impact of disagreement on IPv6 adoption among ISPs. Then we elaborate on how consensus among ISPs changes the adoption prospect; finally we discuss the real-world findings these models afford. For better readability and tractability of the discussions, we put forward in this section Table 1, which lists the variables introduced in previous sections in addition to their roles. The first row presents a parameter, and the second row briefly defines its corresponding role. For instance, the first entry points to the translation quality impairments that a user or an ICP experiences, and it is denoted by a_6 or a (which we use interchangeably).

Variable	Role
a_6 or a	Translation quality impairments
$\sigma_6^{(i)}$	Fraction of IPv6 users
$\gamma_6^{(i-1)}$	Fraction of non-IPv6 ICPs
p_4	Cost of IPv4 service
p_6	Cost of IPv6 service
β	ICP revenue heterogeneity
θ	ICP cost heterogeneity
n_4	Number of IPv4 users
n_6	Number of IPv6 users
C	Cost of IPv4 Addresses
D_6	Translation Cost

Table 1: List of variables

5.1. The Impact of Disagreement

ISPs' inability to converge to jointly optimal prices is primarily because the coupling between users and ICPs' decisions introduces two distinct strategies for the IPv6 ISP, and correspondingly a discontinuity in its utility function. When the price of public IPv4 connectivity is high enough, it is best for the IPv6 ISP

to heavily discount⁹ its connectivity to attract many users and in turn convince many ICPs to become IPv6 accessible, which lowers translation costs. This, however, triggers a price decrease from the public IPv4 ISP to recoup part of its lost user-base, and then forces the IPv6 ISP to itself lower its price to maintain a sufficiently attractive discount. This eventually results in a public IPv4 price that is too low to allow the IPv6 ISP to give a large enough discount. The better strategy for the IPv6 ISP is then to reduce its discount and attract fewer users. Each user generates a higher revenue, and because there are few of them, translation costs are low. This occurs while guaranteeing a non-negative utility function for the ISP, *i.e.*, the IPv6 ISP finds an alternative strategy that yields to a “higher” revenue and chooses that over the current strategy. This behavior is basically capturing the outcome of a real world scenario in which ISPs compete with each other, and are not able to converge to a Nash Equilibrium. This cyclical pattern is shown in Fig. 2 that plots each ISPs’ best-responses as a function of the other’s price, and includes an instance of a pricing cycle.

Figure 2: Cycle in ISPs best response game

Cycles occur when the utility gap between the two strategies of the IPv6 ISP is large enough to ensure that its best-response function and that of the IPv4 ISP do not intersect. Fig. 3 explores how often this arises across a reasonable range of configurations. The price of IPv4 addresses is chosen to have a normalization constant $C = 1$, so that the quadratic cost function for IPv4 addresses yields a value of 1 when the number of IPv4 Internet users reaches $n_4 = 2$, *i.e.*, doubles. In other words, doubling the size of the current IPv4 Internet yields a public IPv4 address price equal to the value of Internet connectivity itself. This choice reflects the fact that according to current statistics there were about 2 billion Internet users by the end of 2012, and given the $\approx 50 - 75\%$ utilization of the address space, a doubling of IPv4 users is then still possible. The per-user IPv6 conversion cost, c_6 , is assumed to be ten percent of the base value of Internet connectivity.

Fig. 3 shows the outcome of the game played by the two ISPs as a function of unit translation costs, D_6 , and the non-native connectivity quality impairment a

⁹This is happening in some places even today [25, 26].

Figure 3: IPv6 vs. public IPv4 competition

user incurs, a_6 . D_6 is varied from zero to ten percent of the per user IPv4 address acquisition cost, while a_6 is varied from zero to the per user IPv6 conversion cost an ICP incurs. The figure illustrates the presence of cycles in a wide range of configurations, and in particular as soon as both a_6 and D_6 slightly grow from zero. Similar results emerge from a scenario with an IPv6 ISP and a private IPv4 ISP, but we relegate them to Appendix A. These results show that disagreement between ISPs on offering IPv6 as a connectivity option can potentially derail the current progress of IPv6 adoption, by making it hard for the Internet stakeholders to identify winning strategies.

5.2. The Benefit of Consensus

The previous section illustrated the difficulty of devising effective strategies, when ISPs tackle public IPv4 address shortage with competing connectivity options. The intent of this section is **not** to argue that to migrate to an IPv6 Internet, we need to forfeit competition among ISPs. This would be neither realistic nor meaningful. Instead, we want to argue for shifting competition away from connectivity choices, *i.e.*, have a consistent offering of connectivity choices among ISPs.

Unsurprisingly, IPv6 adoption and the ISP's pricing strategy are directly affected by C , the normalization constant for the cost of acquiring additional public IPv4 addresses, and D_6 , the translation cost of one unit of traffic. In addition, two other parameters indirectly affect the ISP's strategy because of how they influence users and ICPs decisions, namely, c_6 , the per-user cost of upgrading an ICP's infrastructure to IPv6, and a_6 , the relative magnitude of the impairment translation causes, *i.e.*, the loss in quality-of-experience for users and the related revenue loss for the ICPs.

It is possible to scope the ranges some of those parameters can span, *e.g.*, $C \leq 1$, but a complete sampling of this four-dimensional space is impractical. We rely instead on several figures to report how the outcome changes as some parameters vary, while others remain fixed. The figures help identify parameters that have a significant effect on IPv6 adoption by both users and ICPs; hence suggesting possible strategies.

Figure 4: ICPs adoption levels for small C

Figure 4 illustrates an intuitive outcome, and in the process offers some level of “sanity checking” of the model. In particular, it confirms the expected negative impact on ICPs’ adoption of decreasing their adoption costs, c_6 , while a_6 remains constant. However, this figure alone ignores the effect of a_6 , which as we discuss next, can have an ambivalent effect. a_6 is the only factor that couples ISPs, ICPs and users, hence, it is important to investigate its impact on the future of IPv6 adoption.

Figs. 5 and 6 plot four quantities as a function of a_6 , while c_6 is assumed to stay put¹⁰. The left hand-sides of the figures give the cumulative per-user discount the ISP offers to its users, and the ISP’s total profit when the size of the Internet user population grows by 100%. The right hand-sides of the figures give, after doubling the size of the Internet user population (*i.e.*, growth by 100%), the (final) fractions of users that have opted for IPv6, and ICPs that have become IPv6 accessible. The figures report the results for two different configurations, namely, small and large values of C , and for each configuration consider different ratios between translation costs and IPv4 address acquisition costs, *i.e.*, D_6/C takes values 0.1, 1 and 10. These figures identify two parameters that have an impact on IPv6 adoption, namely, a_6 and D_6 .

5.2.1. Effect of a_6

Consider first the effect of a decrease in the level of impairment, a_6 , that translation imposes. Such a decrease can (initially) make IPv6 more attractive to users by lowering the penalty they incur when accessing ICPs that are not yet IPv6 accessible. This can increase the number of users that choose IPv6, which can in turn entice more ICPs to become IPv6 accessible; possibly starting a positive feedback loop in IPv6 adoption. On the flip side, a lower a_6 value also decreases the potential per-user revenue gain ICPs derive from becoming IPv6 accessible. This makes it more likely that revenue increases won’t offset adoption costs; hence reducing ICPs’ adoption of IPv6. This would in turn make IPv6

¹⁰Although c_6 remains constant throughout the analysis, in the above ratio it serves as a normalization factor.

Figure 5: Total profit, discount & adoption levels for small C

Figure 6: Total profit, discount & adoption levels for large C

less attractive to users, and having fewer users opting for IPv6 would further reduce its attractiveness to ICPs. As we can see, the role of changes in a_6 on IPv6 adoption is unclear, and the figures help elucidate under which conditions changes in a_6 improve IPv6 adoption.

First, the figures illustrate that an increase in a_6 (*i.e.*, in the ratio a_6/c_6 since c_6 is not changing) systematically results in higher IPv6 adoption by ICPs and to a lesser extent users. In the case of ICPs, one can think of the ratio a_6/c_6 as the ICPs' per user return from IPv6 adoption. An increase in this return motivates more ICPs to make such an adoption choice. When this increase is through an increase in a_6 (rather than a decrease in c_6 that is trivially beneficial to both ICPs and users), the greater number of ICPs that opt to become IPv6 accessible, offsets the larger penalty that IPv6 users suffer when accessing IPv4-only ICPs. In other words, users experience greater impairments when accessing ICPs that still require translation, but because there are fewer such ICPs, the impact is mitigated.

The number of user that choose IPv6 is not however, overly affected by an increase in a_6 , at least not until it reaches a certain threshold. This is because while there may be more ICPs that can be accessed natively over IPv6, this benefit is offset by the greater impairments users experience when accessing the remaining ICPs. However, after a_6/c_6 reaches a threshold, the IPv6 adoption level shifts from a steady state to 100% (*i.e.*, “bang-bang” solution). When and how this shift happens is, however, affected by the relative magnitude of IPv4 address acquisition costs, C , and IPv6 address translation costs, D_6 . Consequently, the number of users that choose IPv6 is not overly affected even if differences exist. In other words, while both Figs. 5 and 6 establish that a larger ratio a_6/c_6 benefits IPv6 adoption by both ICPs and users, ensuring a complete migration to an IPv6 Internet requires a large enough value. How large this value needs to be depends on a number of factors, and in particular, on the IPv4 address acquisition cost and the translation cost (*i.e.*, C and D_6), which as we discuss next introduce some interesting behaviors in their own right.

5.2.2. Effect of C and D_6

To understand the role of these two parameters, consider the scenarios where the return on IPv6 adoption is less than 1 for ICPs in Figs. 5 and 6. In these instances, ICPs have limited incentives to become IPv6 accessible. Low translation costs¹¹, D_6 , afford the ISP enough leeway to price IPv6 competitively and convince a number of users, and consequently, a number of ICPs to adopt IPv6. This is reflected in the higher adoption levels of both users and ICPs as D_6 decreases. However, when translation cost increases, the ISP is unable to provide enough incentives to users and ICPs, therefore, the adoption levels are low.

When IPv4 address acquisition cost, C is high, and $a_6/c_6 > 1$ (Fig. 6), the

¹¹Note that changes in translation costs are relative to the value of C , which explains in part the similarity of Figs. 5 and 6.

shift in the level of IPv6 adoption by users is abrupt, because the high cost of IPv4 addresses entices the ISP to aggressively discount IPv6 early on to quickly convince ICPs and users alike to adopt IPv6. This is also confirmed by the left hand-side curves in Figs. 5 and 6 that report the total discount the ISP offers to entice users to adopt IPv6, and the ISP's total profit, which shows this abrupt transition is beneficial to the ISP by saving large discounts given to users.

In contrast, Fig. 5 shows that when IPv4 address acquisition cost, C , is low and $a_6/c_6 > 1$, the ISP's behavior depends on the translation cost D_6 . In particular, with low translation costs (D_6), an ISP may initially offer only a limited discount for IPv6, which can prevent the full adoption of IPv6, and prolong the coexistence of IPv4-only and IPv6-only Internets. In other words, if IPv4 addresses remain cheap for an extended period of time, it not only prolongs the transition to an IPv6 Internet, it may also make the transition significantly more expensive by deterring many ICPs from migrating early on, since the later an ICP adopts IPv6, the bigger infrastructure it has to migrate.

The next section summarizes our findings, while suggesting guidelines to avoid strategies that can derail IPv6 adoption in the future.

5.3. Findings

This section summarizes five major findings models of the previous section posit:

(i) Disagreement between ISPs on what connectivity options they offer to their users has a deteriorating effect on IPv6 adoption. Since the future growth of the Internet depends on IPv6 adoption, this delay can eventually affect all types of ISPs, regardless of their connectivity offerings. Therefore, in order to make the migration of the whole Internet to IPv6 smoother, ISPs should avoid offering *only* alternative connectivity options to IPv6.

(ii) Consensus among ISPs on offering IPv6 adoption along with other connectivity options, helps identify strategies that lead to a smooth IPv6 adoption progress. In other words, when ISPs all offer IPv6 as one of their connectivity options, devising successful migration strategies becomes easier for other major Internet stakeholders.

(iii) ICPs' Infrastructure upgrade cost (c_6) plays a key role in adoption of IPv6 across not only ICPs, but also users, by facilitating better connectivity quality compared to translated traffic. Therefore, technology developers can play a role in adoption of IPv6 by lowering the price of their products. Equivalently, governments can subsidize these products and accelerate IPv6 adoption.

(iv) Although the existence of translation devices (IPv4 $\leftrightarrow$ IPv6) is necessary for an uninterrupted migration to an IPv6 Internet, the presence of high quality translation devices (*i.e.*, small a_6) has a negative impact on IPv6 adoption by ICPs, since these devices facilitate a loss-less access to IPv4-only content for IPv6-only users, hence, eliminating incentives for IPv6 adoption by ICPs. Even though this was known in other technology adoption instances [27, 28, 29], the models of this paper confirm its impact in the case of IPv6 adoption.

(v) Translation costs incurred by ISPs D_6 can have an ambivalent effect on IPv6 adoption: when translation impairment level (a_6) is small, smaller

translation costs (D_6) accelerates IPv6 adoption by ICPs and users, and also contributes to higher profits for ISPs; however, when a_6 is large, larger D_6 and small C values benefit IPv6 adoption.

These five findings are the outcome of the models of previous sections, which were solved, for tractability and better presentation, with some simplifying assumptions, *e.g.*, homogenous revenue for ICPs, uniform distribution of users' sensitivity to quality, etc. In order to investigate the potential impact of these assumptions on our findings, the next section relaxes/alters them and provides the analyses for each case.

6. Robustness Tests

In order to capture the real world phenomena in a tractable way, every model makes some simplifying assumptions. Our modeling effort is not an exception, however, to ensure that our findings are not simply artifacts of those assumptions, we perform a series of robustness analyses. Therefore, we consider a number of variations in those assumptions, and accordingly solve the models with such variations, and compare their results and findings with those of Sections 5 and 5.3, respectively. Next, we present a summary of these variations.

(i) Heterogeneity in revenue and cost of ICPs: In the current solution, we assume that ICPs derive revenue from users homogeneously. In reality, however, the revenues are heterogeneous, and mostly depend on the volume of an ICP's traffic, *i.e.*, popularity. Furthermore, we assume the adoption costs are uniformly distributed across ICPs, while in the real world, larger ICPs possibly benefit from economies of scale. Hence, in this variation of the model, we assume the revenue and cost of an ICP depends on its popularity.

(ii) Users' sensitivity to quality impairments (σ): In our current model, this parameter is assumed to be uniformly distributed, while in reality, the distribution is most probably not uniform. It is hard, if not impossible, to determine the exact distribution in these scenarios, however, if similar (qualitative) outcomes emerge with different distributions, the findings are on a more solid footing. Therefore, in this variation we assume users' sensitivity to quality impairments follows a different distribution.

(iii) Users' decision making inertia: Currently, our model assumes that all users make decisions after each announcement of prices by ISPs. That might not be practical in all scenarios, *i.e.*, many ISPs impose contractual agreements on their users. Thus, in this variation of the model, we incorporate inertia in users' decisions through contractual agreements.

(iv) IPv4 address acquisition cost: In our current model, this cost is captured via a quadratic function, which reflects a scenario with growing scarcity of IPv4 addresses. However, another scenario involves ISPs with extra IPv4 addresses providing the market with enough resources to keep the cost constant. Therefore, in this variation of the model we assume the cost of IPv4 address acquisition grows linearly.

(v) The per-user cost of IPv6 adoption by ICPs (c_6): In our current model, this cost is assumed to stay put over time. However, it is more likely

that it decreases as the technology matures. Hence, this variation of the model incorporates a decreasing per-user cost of IPv6 adoption.

6.1. Heterogeneity of ICPs

ICPs are heterogeneous in many aspects. However, it is not practical, from a modeling standpoint, to capture all of those dimensions. Moreover, from the perspective of our model, it is only the net effect of those heterogeneities that matters. Therefore, here we only focus on one aspect of heterogeneity among ICPs, namely, popularity. We present next the changes required in the utility function of ICPs to incorporate this aspect.

The popularity is captured in Eq. (2) through parameter β , which, in previous sections, was assumed to be equal to 1 (*i.e.*, $\beta = 1$) yielding to a homogenous revenue across all ICPs. Here, we assume β is distributed in $[0, 2]$ with a certain probability distribution¹². We assume that most ICPs are of similar popularity, but some are either highly popular or unpopular. This leads to the choice of a unimodal bell-shaped distribution function, which is adequate to show how differences in popularity of ICPs can potentially change our findings¹³. Without loss of generality, and to extend the tractability of our model, we use a unimodal Kumaraswamy distribution function for β .

Moreover, in order to incorporate the economy of scale in IPv6 adoption decisions of ICPs, we choose θ , the parameter that captures adoption cost heterogeneity across ICPs (in Eq. (2)), to be a function of their popularity. In other words, we choose θ to be a decreasing function of β , to reflect the economy of scale, *i.e.*, a more popular ICP with a larger user-base incurs (on average) lower per-user IPv6 adoption costs.

Although changes in β and θ only affect the decision making process of ICPs, their incorporation in the model help capture a more general, and perhaps more realistic, picture of the IPv6 adoption problem. Next we show the impact of these changes on our earlier model formulation.

6.1.1. Formulation

Without loss of generality, we choose $\theta = \frac{2-\beta}{2+\beta}$, which bounds θ in $[0, 1]$, and is a decreasing function of β . It can be easily shown, through numerical analysis, that this choice is as good as any other decreasing function.

While the expressions for users and ISPs remain the same as Eqs. (6), (A.5) and (A.6), $\gamma_6^{(i)}$, the fraction of ICPs that do not choose IPv6 at the end of round i , changes as follows:

¹²Choosing $[0, 2]$ instead of $[0, 1]$ only makes it easier to compare the outcome with the current outcomes, without loss of generality.

¹³Similar results emerge when we use a bimodal distribution function.

$$\gamma_6^{(i)} = \begin{cases} \text{if } p_4 - p_6 > a_6 \gamma_6^{(i-1)} \\ \mathcal{K}_{A,B} \left(\frac{\sqrt{(2 + \frac{c_6}{ka_6})^2 + 8 \frac{c_6}{ka_6}} - (2 + \frac{c_6}{ka_6})}{2} \right) \\ \text{if } \frac{\gamma_6^{(i-1)}(1 - \gamma_6^{(i-1)})c_6}{k} \leq p_4 - p_6 \leq a_6 \gamma_6^{(i-1)} \\ \mathcal{K}_{A,B} \left(\frac{\sqrt{(2 + \frac{c_6}{ka_6 \sigma_6})^2 + 8 \frac{c_6}{ka_6 \sigma_6}} - (2 + \frac{c_6}{ka_6 \sigma_6})}{2} \right) \\ \text{Otherwise} \\ \gamma_6^{(i-1)} \end{cases} \quad (11)$$

where $\mathcal{K}_{A,B}(\cdot)$ denotes the Kumaraswamy function with parameters A and B (the rest of the parameters are introduced in Eq. (7)). With the proper values for A and B , one can achieve a uni-modal distribution function in $[0, 1]$. Next, we compare the findings from this variation of the model with our original findings.

6.1.2. Results and Discussion

Comparing Fig. 7 with Fig. 3, one can see that this scenario yields to similar outcomes as the original scenario, *i.e.*, in the majority of the cases the disagreement between ISPs ends in a cyclical behavior. Therefore, the first finding of Section 5.3, *i.e.*, detrimental effect of disagreement between ISPs on technology options, holds even in the presence of heterogeneity in revenue and cost of ICPs.

The results of the consensus scenario with popularity incorporated in ICP's decisions, are labeled "Uni-modal β " in Figs. 8, 9, 10, and 11. In this scenario we assume a small C value, and $D_6 = 0.1C$, hence, we can compare them with the corresponding curve in Fig. 5, which is also plotted (for more convenience) in the above figures under the label "Original Scenario". A similar figure to Fig. 5 is relegated to Appendix B.

Comparing "Uni-modal β " curves with the "Original Scenario" (in the above figures), one can see quantitative differences, *i.e.*, the Uni-modal β scenario yields to slightly lower discounts given by ISPs, higher total profits for ISPs, and different adoption rates for users and ICPs. The main reason behind these differences is that in this scenario compared to the original scenario, more popular ICPs derive higher benefits from IPv6 adoption, therefore, their adoption rate is higher. Higher adoption rate by popular ICPs causes lower translation costs by ISPs, and therefore lower discounts and higher profits. Users also have greater incentives to adopt IPv6 in the presence of higher adoption rate by ICPs. This is only true for small values of translation impairments (a), and for larger values of this variable most of the incentives fade away, and cause lower overall adoption rates by users and ICPs. These quantitative differences nonetheless, the outcomes are qualitatively similar, *i.e.*, the patterns are almost similar across all metrics, namely, discounts by ISPs, ISPs' profit, and adoption rates by users and ICPs. In other words, findings (ii) to (v) of Section 5.3 hold in the presence of heterogeneity in revenue and cost of ICPs.

As alluded to earlier, using a uni-modal or bi-modal distribution function for β results in similar outcomes. Also, we choose the parameters of the Kumaraswamy function (*i.e.*, A and B), randomly (as long as it results in a uni-modal or bi-modal distribution function), therefore, the robustness test is not capturing a single scenario. In other words, the above robustness test shows that ICPs' heterogeneity does not qualitatively affect the findings of Section 5.3, regardless of the underlying distribution.

Figure 7: IPv6 vs. public IPv4 competition — single-modal β

Figure 8: ISP's total per-user discount offered to users

6.2. Users Sensitivity

In the original scenario, we assumed that sensitivity to quality impairments is distributed uniformly across users. To investigate the dependency of our findings to this assumption, it is best to substitute the distribution with a general distribution function. However, this seems impractical given that we rely on numerical analysis. Therefore, we only focus on a uni-modal bell-shaped distribution function, which is commonly seen in real world phenomena.

6.2.1. Formulation

In our model, sensitivity to quality impairments is specific to users, hence, the model formulation remains the same for ICPs and ISPs. The changes to users is reflected via σ_6 as follows:

Figure 9: ISP's total profit

Figure 10: Final IPv6 adoption by users

Figure 11: Final IPv6 adoption by ICPs

$$\sigma_6^{(i)} = \mathcal{K}_{A,B}\left(\frac{p_4 - p_6}{a\gamma_6^{(i-1)}}\right) \quad (12)$$

which is similar to Eq. (6) except for $\mathcal{K}_{A,B}(\cdot)$ that as alluded to earlier denotes the Kumaraswamy distribution function with parameters A and B .

6.2.2. Results and Discussion

Similar to the previous section, comparing Fig. 12 and Fig. 3, one can see that the choice of distribution function for users' sensitivity to quality impairments (σ), does not change the outcomes qualitatively. Therefore, the first finding of Section 5.3 holds in the presence of a different distribution for σ .

We plot the outcomes of the consensus scenario in Figs. 8, 9, 10, and 11 under the label “Uni-modal σ ”. Again, in this scenario the value of C is small and $D_6 = 0.1C$ (a similar figure to Fig. 5 is relegated to Appendix B). Comparing these curves with the corresponding “Original Scenario” curves, one can see slight quantitative differences, which are mainly an artifact of a different distribution for σ . However, the patterns are qualitatively similar to the original scenario. In other words, findings (ii) to (v) of Section 5.3 hold in the presence of a different distribution for σ .

We choose the parameters of the Kumaraswamy function randomly, as long as it yields to a uni-modal distribution function for σ_6 . Therefore, the above test captures a wide range of functions that take values in $[0, 1]$, and shows that the findings are not qualitatively affected by the choice of this function.

Figure 12: IPv6 vs. public IPv4 competition — single-modal σ

6.3. Users Inertia

In our initial solution we assumed all users re-evaluate their decisions after each price announcement by ISPs. In reality, this might be violated due to the inertia in decision making process of users; the most common source of this inertia is involvement in a contractual agreement, which bars users from changing their services at any time. Hence, we need to examine the validity of our findings in the presence of such inertia. In the consensus scenario, this assumption does not play a significant role, because the monopolistic ISP can

assign services to its users at any given time, in order to maximize its profit (*i.e.*, contractual agreements are internal to the ISP). However, in the scenario with disagreement among ISPs on offering IPv6, we need to analyze the model with the assumption that not all users can change their services after each announcement.

6.3.1. Formulation

The original formulation of the problem does not change, except that in Eqs. A.5, A.6 and 10, we substitute $(1 + i\delta)$, with a fraction of users that can make decisions at epoch i . In other words, instead of the total user population, here only a fraction of them make a decision. Without loss of generality, we assume that users can change their services every year, therefore, at any given time epoch i , only $1/12$ of the existing users can change their services.

6.3.2. Results and Discussion

Fig. 13 demonstrates the outcome of the analysis for the above variation of the model. Comparing it with the results of our original scenario plotted in Fig. 3, one can see the similarity of results. In other words, this result confirms that inertia in decision making of users does not change our first finding of Section 5.3 regarding the disagreement scenario. The rest of the findings also remain intact, since the contractual agreements do not affect them.

Fig. 13 shows that the existence of cycles becomes even more widespread in the presence of contractual agreements. In other words, if the term of the contractual agreements become longer, the cycles become the norm. We confirmed this by looking into a scenario in which existing users are not allowed to change their services ever (*i.e.*, life-time contractual agreement). This impractical scenario also results in a qualitatively similar outcome as Fig. 13. This demonstrates that contractual agreements, regardless of their terms, do not affect the findings of our models qualitatively.

Figure 13: IPv6 & public IPv4 consensus — contractual agreement

6.4. IPv4 Address Acquisition Cost

In the original scenario, we assumed that the cost of IPv4 address acquisition grows quadratically with the number of addresses (the per address cost grows

linearly). We made this assumption to capture the increasing cost of IPv4 addresses as they become scarcer. However, an alternative scenario is when ISPs with extra IPv4 addresses sell their addresses in the corresponding markets, and therefore, the supply of addresses meets the demand. In this situation, the cost of IPv4 addresses grows linearly with the number of addresses (per address cost remains constant). Here, we examine the extensibility of our findings to such scenario.

6.4.1. Formulation

We only need to substitute the quadratic function in the ISP's utility with a linear function. In other words, in Eqs. A.5 and 10, we replace $C((1+i\delta)(1-\sigma_6)-1)_+^2$ with $C((1+i\delta)(1-\sigma_6)-1)_+$.

6.4.2. Results and Discussion

We plot the outcome of the disagreement scenario in Fig. 14, which is quite similar to the outcome of the original scenario in Fig. 3, *i.e.*, the cycles are present predominantly. Therefore, the first finding of Section 5.3 holds for this variation of the model.

We also plot the results of the consensus scenario in Figs. 8, 9, 10, and 11 under "Linear IPv4 Acquisition Cost". Once again, this scenario is for a small value of C and $D_6 = 0.1C$ (with a similar figure to Fig. 5 relegated to Appendix B). Comparing these curves with the corresponding "Original Scenario" curves, one can observe quantitative differences. However, these differences do not affect the overall patterns of adoption, *i.e.*, the patterns are qualitatively similar. Therefore, findings (ii) to (v) of Section 5.3 extend to this variation of the model.

The cost of IPv4 addresses is not going to decrease in the near future, at least until a near complete adoption of IPv6. Therefore, the linear cost function used in the above robustness test, which is equivalent to a constant per-address cost, is the best case scenario, *i.e.*, the lower bound of address cost is when it is non-decreasing. The above test shows that the findings of this scenario, are qualitatively similar to those of Section 5.3, and therefore, we can easily conclude that the findings of any scenario with a super-linear function for IPv4 address acquisition cost, is also qualitatively similar to those of Section 5.3.

6.5. Per-User Cost of IPv6 Adoption by ICPs

In the original scenario, we assumed that the per-user cost of IPv6 adoption for ICPs remains constant over time. However, in reality, as technology matures the cost of adoption also decreases. Here, we investigate the validity of our results in such scenario.

6.5.1. Formulation

The formulation remains the same for almost all equations, and the only difference is that we replace c_6 in Eq. (7) with $c_6(i)$, a decreasing function of i .

Figure 14: IPv6 vs. public IPv4 competition — linear IPv4 address acquisition cost

6.5.2. Results and Discussion

Comparing Figs. 15 and Fig. 3 shows that temporal decrease of c_6 does not change the overall pattern of our results, *i.e.*, the first finding of Section 5.3.

Comparing the curves labeled “Decreasing c_6 ” and “Original Scenario” in Figs. 8, 9, 10, and 11, it can be easily seen that there are significant quantitative differences (a similar figure to Fig. 5 is relegated to Appendix B). However, this is not surprising, as decreasing the cost of IPv6 should speed up the adoption process by ICPs, which consequently causes higher profit for ISPs and higher adoption rates for users. This exactly re-enforces what the third finding of Section 5.3 states. Also, the qualitative similarities between the outcomes of this scenario and the original scenario validates findings (iv) and (v) of Section 5.3.

Several decreasing functions were used for c_6 in the above analysis, and the results, even though quantitatively different, all agree qualitatively. In other words, the choice of the decreasing function for c_6 does not affect the findings of Section 5.3.

Figure 15: IPv6 vs. public IPv4 competition — decreasing c_6

7. Related Works

Explaining the slow progress of IPv6 adoption has been the focus of much prior work (see [30, 31] for a recent overview). Earlier works focused on identifying technical issues that created initial adoption hurdles [32, 33, 34, 35], but as those were eventually addressed, the attention shifted towards measuring IPv6 adoption progress [36, 37, 38, 39, 40, 41, 42, 16, 43, 44], as well as exploring the role that economic forces may be playing [45, 46, 47, 48, 49, 50, 51].

Those latter works bear the most direct relevance to the investigation presented in this paper, with [45] echoing many of the same themes we identify, including the importance of coordination, albeit without analytical support. Casting IPv6 adoption as a game was proposed in [51], but one with Autonomous Systems as the sole players, *i.e.*, it did not account for either users or ICPs. The use of two-sided markets to capture dependencies between the decisions of Internet stake-holders was suggested in [48], but used simply to assess the impact of changing certain parameters, *i.e.*, it did not explore the possibility or competition between ISPs nor how the presence of coordinated revenue maximization by ISPs would influence the outcome.

There is also a vast literature on two-sided markets, and the reader is referred to [52, 53] for recent surveys. The most relevant works deal with competing platforms [54, 15], *e.g.*, IPv4 and IPv6, but the absence of pricing for one side of the market (the ICPs) in our context makes for a very different (and simpler) focus. There is also a closely related work on the two-sided market analysis of Internet content delivery [55], however its focus is not on IPv6 adoption.

8. Conclusion

The paper’s motivation is to shed light on what can potentially affect, positively or negatively, the current accelerated progress of IPv6 adoption. It proposes a number of models that capture the dependencies between different Internet stakeholders and consider various connectivity options. The paper then explores how these dependencies and connectivity options affect the decisions of Internet stakeholders in migrating to IPv6. A first set of scenarios consider ISPs that respond to IPv4 address scarcity differently, namely, by using different connectivity options. The ISPs compete to attract users on the basis of connectivity options, and the models demonstrate how in these scenarios devising an effective IPv6 adoption strategy may be difficult, and that stakeholders can potentially derail the current progress of IPv6 adoption.

The paper then explores an alternative scenario in which ISPs still preserve the ability to offer multiple connectivity options, but all of them have a consensus on offering IPv6 as one of those options. While this is not by itself sufficient to maintain the current acceleration in IPv6 adoption, it affords a more predictable look at how different factors can affect the progress. In particular, it helped identify translation impairments a as an important factor that can negatively or positively affect the adoption of IPv6 in the future, *i.e.*, larger values can motivate ICPs to adopt IPv6 to avoid loss of revenue and at the same time

deter users from adopting the technology, while smaller values reduces incentives for ICPs and motivates users. The model also helped identify the role IPv4 address acquisition cost C and translation cost D_6 play. In particular, it showed that low IPv4 address acquisition costs can derail the progress of IPv6 adoption at least temporarily. Even if the latter is to some extent due to the myopic decision making process of ISPs in the models, the presence of a more strategic decision does not change the outcomes significantly.

The models on which this paper's investigation is based have numerous obvious limitations, and fail to capture the impact of many other factors. For instance, it is true that not all ICPs have the same impact on users, and the largest ones play differently. Such ICPs can make a difference in the direction of IPv6 adoption, as they host the most popular content consumed by the users, and their migration can incentivize ISPs to offer IPv6 to their users. However, in this paper it is assumed that every unit of content consumed by the users is distributed uniformly across all ICPs, regardless of their sizes, as it is hard if not impossible to capture the role of different ICPs in a tractable mathematical model. However, they capture the interdependencies between different Internet stakeholders and their decisions. As such the results offer some insight into the potential factors that can derail or accelerate the progress of IPv6 adoption in the future.

Acknowledgments

The author would like to acknowledge Dr. Guérin's contributions to the preliminary version of this paper (presented at IFIP Networking 2014 Conference [1]).

References

- [1] M. Nikkhah, R. Guérin, Migrating to IPv6 - The Role of Basic Coordination, in: Proc. IFIP Networking Conference, 2014.
- [2] M. Nikkhah, R. Guérin, Migrating the Internet to IPv6: An Exploration of the When and Why, IEEE/ACM Transactions on Networking (TON), To appear.
- [3] J. Wu, J. H. Wang, J. Yang, CNGI-CERNET2: an IPv6 deployment in China, ACM SIGCOMM Computer Communication Review 41 (2).
- [4] State of IPv6 in China, (Available at <http://goo.gl/sXG6g0>).
- [5] IPv6 Deployment Best Practice by China Telecom, (Available at <http://goo.gl/80ewsD>).
- [6] IPv6 at Verizon Wireless, (Available at <http://goo.gl/t3hd9i>).
- [7] T-Mobile Goes IPv6 Only on Android 4.4 Devices, (Available at <http://goo.gl/WZUBB6>).

- [8] S. Alcock, R. Nelson, D. Miles, Investigating the impact of service provider NAT on residential broadband users, INFOCOM, San Diego, CA, USA.
- [9] About Carrier Grade NAT (CGN), (Available at <http://goo.gl/S2gF5e>).
- [10] C. Donley, Ed., L. Howard, V. Kuarsingh, J. Berg, J. Doshi, Assessing the impact of carrier-grade NAT on network applications, Internet Draft (2013).
- [11] Y. Liu, J. Pan, The impact of NAT on BitTorrent-like P2P systems, in: Proc. IEEE P2P, 2009.
- [12] A. Ozment and S.E. Schechter, Bootstrapping the adoption of Internet security protocols, in: Proc. WEIS, 2006.
- [13] M. Bagnulo, P. Matthews, I. van Beijnum, Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers, RFC 6146 (Proposed Standard) (Apr. 2011).
URL <http://www.ietf.org/rfc/rfc6146.txt>
- [14] A. Durand, R. Droms, J. Woodyatt, Y. Lee, Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion, RFC 6333 (Proposed Standard) (Aug. 2011).
URL <http://www.ietf.org/rfc/rfc6333.txt>
- [15] J.-C. Rochet, J. Tirole, Platform competition in two-sided markets, J. European Econ. Assoc. 1 (4).
- [16] M. Nikkhah, R. Guérin, Y. Lee, R. Woundy, Assessing IPv6 through web access a measurement study and its findings, in: Proc. ACM CoNEXT, 2011.
- [17] F. Dobrian, V. Sekar, A. Awan, I. Stoica, D. Joseph, A. Ganjam, J. Zhan, H. Zhang, Understanding the impact of video quality on user engagement, ACM SIGCOMM Computer Communication Review 41 (4) (2011) 362–373.
- [18] S. Zander, G. Armitage, Empirically measuring the qos sensitivity of interactive online game players, in: Proc. ATNAC, 2004, pp. 511–518.
- [19] E. Schurman, J. Brutlag, Performance related changes and their user impact, available at <http://velocityconf.com/velocity2009/public/schedule/detail/8523> (2009).
- [20] I. T. Force, Technical and economic assessment of internet protocol version 6 (ipv6), White Paper, US Department of Commerce, National Institute of Standards and Technology, National Telecommunications and Information Administration.
- [21] Cisco IPv6 Q&A, (Available at <http://goo.gl/nwn4yE>).

- [22] List of IPv6 ISPs, (Available at <http://goo.gl/QSxep3>).
- [23] L. Howard, T. W. Cable, Internet access pricing in a post-ipv4 runout world, White Paper.
- [24] M. Nikkhah, R. Guérin, Migrating to IPv6 - The Role of Basic Coordination, Tech. report, University of Pennsylvania, (Available at <http://goo.gl/gWJd95>) (2013).
- [25] CERNET2 IPv6 experience, (Available at <http://goo.gl/fhesar>).
- [26] IPv6 Transition Revisit, (Available at <https://goo.gl/oW0DVX>).
- [27] S. Sen, Y. Jin, R. Guérin, K. Hosanagar, Modeling the dynamics of network technology adoption and the role of converters, *IEEE/ACM ToN.* 18 (6).
- [28] J. P. Choi, Do converters facilitate the transition to a new incompatible technology? a dynamic analysis of converters, *International Journal of Industrial Organization* 14 (6) (1996) 825–835.
- [29] D. Joseph, N. Shetty, J. Chuang, I. Stoica, Modeling the adoption of new network architectures, in: *Proc. ACM CoNEXT*, 2007.
- [30] J. Czyz, M. Allman, J. Zhang, S. Iekel-Johnson, E. Osterweil, M. Bailey, Measuring ipv6 adoption, in: *Proc. of ACM SIGCOMM*, 2014.
- [31] J. Czyz, M. Allman, S. Iekel-Johnson, E. Osterweil, M. Bailey, Assessing IPv6 adoption, Tec. Report TR-13-004, International Computer Science Institute (2013).
- [32] K. Cho, M. Luckie, B. Huffaker, Identifying IPv6 network problems in the dual-stack world, in: *Proc. ACM SIGCOMM NetT Workshop*, 2004.
- [33] S. Narayan, S. Kolahi, Y. Sunarto, D. Nguyen, P. Mani, Performance comparison of IPv4 and IPv6 on various Windows operating systems, in: *Proc. ICCIT*, 2008.
- [34] Y. Wang, S. Ye, X. Li, Understanding current IPv6 performance: A measurement study, in: *Proc. 10th IEEE Symp. Comp. Comm. (ISCC)*, 2005.
- [35] X. Zhou, P. Van Mieghem, Hopcount and E2E delay: IPv6 versus IPv4, in: *Proc. PAM*, 2005.
- [36] R. Beverly, M. Luckie, L. Mosley, K. Claffy, Measuring and Characterizing IPv6 Router Availability, in: *Passive and Active Measurement*, Springer, 2015.
- [37] V. Bajpai, J. Schönwälder, IPv4 versus IPv6-Who connects faster?, in: *Proc. IFIP Networking Conference*, Toulouse, France, 2015.
- [38] k. claffy, Tracking IPv6 evolution: Data we have and data we need, *ACM SIGCOMM Comput. Comm. Rev.* 41 (3).

- [39] L. Colitti, S. H. Gunderson, E. Kline, T. Refice, Evaluating IPv6 adoption in the Internet, in: Proc. PAM, 2010.
- [40] A. Dhamdhere, M. Luckie, B. Huffaker, k. claffy, A. Elmokashfi, E. Aben, Measuring the deployment of IPv6: topology, routing and performance, in: Proc. ACM IMC, 2012.
- [41] S. Iekel-Johnson, C. Labovitz, D. McPherson, H. Ringberg, Tracking the IPv6 migration, Tech. Report TR-2008-01, Arbor Networks (2008).
- [42] E. Karpilovsky, A. Gerber, D. Pei, J. Rexford, A. Shaikh, Quantifying the extent of IPv6 deployment, in: Proc. PAM, 2009.
- [43] N. Sarrar, G. Maier, B. Ager, R. Sommer, S. Uhlig, Investigating IPv6 traffic – what happened at the World IPv6 Day?, in: Proc. PAM, 2012.
- [44] A. Hovav, R. Patnayakuni, D. Schuff, A model of internet standards adoption: the case of ipv6, *Information Systems Journal* 14 (3) (2004) 265–294.
- [45] A. Dul, Economics of IPv4 transfer market on IPv6 deployment, (Available at www.quark.net/links.html) (2011).
- [46] B. Edelman, Running out of numbers: Scarcity of IP addresses and what to do about it, in: Proc. First Conference on Auctions, Market Mechanisms and Their Applications, 2009.
- [47] H. Elmore, L. J. Camp, B. Stephens, Diffusion and adoption of IPv6 in the ARIN region, in: Proc. WEIS, 2008.
- [48] R. Guérin, K. Hosanagar, Fostering IPv6 migration through network quality differentials, *ACM SIGCOMM Comput. Comm. Rev.* 40 (3).
- [49] G. Huston, Is the transition to IPv6 a 'market failure'?, available at <http://www.potaroo.net/ispcol/2009-09/v6trans.html> (2009).
- [50] Internet addressing: Measuring deployment of IPv6, Tech. report, OECD, (Available at www.oecd.org/sti/ict/ipv6) (2010).
- [51] T. Trinh, L. Gyarmati, G. Sallai, Migrating to IPv6: A game-theoretic perspective, in: Proc. IEEE LCN, 2010.
- [52] J.-C. Rochet, J. Tirole, Two-sided markets: A progress report, *The RAND J. Econ.* 37 (3).
- [53] R. Roson, Two-sided markets: A tentative survey, *Review of Network Economics* 4 (2).
- [54] M. Armstrong, Competition in two-sided markets, *The RAND J. Econ.* 37 (3).
- [55] N. Zhang, T. Levä, H. Hämmäinen, Value networks and two-sided markets of Internet content delivery, *Telecommunications Policy* 38 (5) (2014) 460–472.

Appendix A. Disagreement Scenario — IPv6 vs. Private IPv4

In this scenario, one ISP offers IPv6 addresses to new users, while the other relies on private IPv4 addresses. Both require translation (IPv6 $\leftrightarrow$ IPv4 and Private IPv4 $\leftrightarrow$ Public IPv4) to communicate with the public IPv4 Internet. Both types of translation equally affect connectivity quality, as measured by a common parameter, a . The greater maturity of Private to Public IPv4 translation benefits a Private IPv4 solution, since $D_4 \leq D_6$. On the flip side, IPv6 users incur translation penalties only for the fraction γ_6 of ICPs not yet IPv6 accessible.

As with the scenario of IPv6 vs. Public IPv4 (4.1), we describe next the decision process of users and ICPs, and how the two ISPs select their prices. For simplicity, and unlike the previous scenario where existing (public IPv4) users also had the option to adopt IPv6, we assume that only new users decide on which connectivity option to choose. Allowing existing (public IPv4) users to make such a choice requires pricing a third option (public IPv4), which adds significant complexity without qualitatively affecting the results.

Appendix A.1. Decision Mechanism & Solution

After the two ISPs announce prices of p_6 and p_{p4} , (new) users choose an ISP as per Eqs. (A.1) and (A.2), where $\sigma_6^{(i)}$ and $\sigma_4^{(i)}$ denote the fraction of users choosing IPv6 or private IPv4 addresses in round i , respectively. In particular, a user with quality sensitivity σ prefers IPv6 over private IPv4 if $1 - p_6 - \sigma a \gamma_6^{(i-1)} \geq 1 - p_{p4} - \sigma a$. Note that since $\gamma_6^{(i-1)} \leq 1$, this implies that prices verify $p_{p4} \leq p_6$. Note also, that it is possible that there exists a value σ_{NA} such that for $\sigma \geq \sigma_{NA}$, $1 - p_6 - \sigma a \gamma_6^{(i-1)} \leq 0$, *i.e.*, users that are very sensitive to connectivity impairment will altogether opt out of connecting to the Internet. This can also arise in the previous scenario, albeit much more rarely as the availability of the public IPv4 option typically ensures that high σ users have access to a suitable alternative¹⁴.

$$\sigma_{p4}^{(i)} = \begin{cases} \frac{(p_6 - p_{p4})}{a(1 - \gamma_6^{(i-1)})} & \text{if } p_{p4} > \frac{p_6 - 1 + \gamma_6^{(i-1)}}{\gamma_6^{(i-1)}} \\ \frac{(1 - p_{p4})}{a} & \text{if } p_{p4} < \frac{p_6 - 1 + \gamma_6^{(i-1)}}{\gamma_6^{(i-1)}} \end{cases} \quad (\text{A.1})$$

$$\sigma_6^{(i)} = \begin{cases} \frac{a(1 - \gamma_6^{(i-1)}) - (p_6 - p_{p4})}{a(1 - \gamma_6^{(i-1)})} & \text{if } p_{p4} > \frac{p_6 - 1 + \gamma_6^{(i-1)}}{\gamma_6^{(i-1)}} \\ & \& p_6 < 1 - a\gamma_6^{(i-1)} \\ \frac{(1 - p_6)}{a\gamma_6^{(i-1)}} - \frac{(p_6 - p_{p4})}{a(1 - \gamma_6^{(i-1)})} & \text{if } p_{p4} > \frac{p_6 - 1 + \gamma_6^{(i-1)}}{\gamma_6^{(i-1)}} \\ & \& p_6 > 1 - a\gamma_6^{(i-1)} \\ 0 & \text{if } p_{p4} < \frac{p_6 - 1 + \gamma_6^{(i-1)}}{\gamma_6^{(i-1)}} \end{cases} \quad (\text{A.2})$$

¹⁴It arises only for combinations of large C and D_6 values, *i.e.*, very high acquisition costs for public IPv4 addresses and very high translation costs.

Once users have selected their connectivity option, ICPs proceed with their decisions as in the previous section.

$$\theta_6 = \frac{ka\sigma_6}{c_6} \quad (\text{A.3})$$

$$\theta_6 = \begin{cases} \min(\max(\frac{ka(1-\gamma_6^{(i-1)})-k(p_6-p_{p4})}{c_6(1-\gamma_6^{(i-1)})}, 1-\gamma_6^{(i-1)}), 1) \\ \text{if } p_{p4} > \frac{p_6-1+\gamma_6^{(i-1)}}{\gamma_6^{(i-1)}} \ \& \ p_6 < 1-a\gamma_6^{(i-1)} \\ \min(\max(\frac{k(1-p_6)}{c_6\gamma_6^{(i-1)}} - \frac{k(p_6-p_{p4})}{c_6(1-\gamma_6^{(i-1)})}, 1-\gamma_6^{(i-1)}), 1) \\ \text{if } p_{p4} > \frac{p_6-1+\gamma_6^{(i-1)}}{\gamma_6^{(i-1)}} \ \& \ p_6 > 1-a\gamma_6^{(i-1)} \\ \min(\max(0, 1-\gamma_6^{(i-1)}), 1) \\ \text{if } p_{p4} < \frac{p_6-1+\gamma_6^{(i-1)}}{\gamma_6^{(i-1)}} \end{cases} \quad (\text{A.4})$$

$$p_{p4} = \underset{p_{p4}}{\operatorname{argmax}} \{ (1+i\delta)(\sigma_{p4}p_{p4} - D_4\sigma_{p4}(1+i\delta)) \} \quad (\text{A.5})$$

$$p_6 = \underset{p_6}{\operatorname{argmax}} \{ (1+i\delta)\sigma_6p_6 - D_6\sigma_6(1+i\delta)\gamma_6^{(i)} \} \quad (\text{A.6})$$

Solving Eq. (A.5) requires an optimization with respect to the conditions of Eq. (A.4), however, those conditions are cumbersome, hence, we resort to numerical analysis.

The Impact of Disagreement: Fig. A.16 offers a perspective similar to that of Fig. 3, and reports the outcome of the ISPs' price selection process for a range of configurations. There are some differences, *e.g.*, migration to IPv6 arises rarely, but there is nevertheless a broad range of parameters for which cycle are present. As argued earlier, this makes devising pricing strategies difficult and is likely to contribute continued uncertainty in deciding how to settle on a migration strategy.

Figure A.16: IPv6 vs. private IPv4 competition

Appendix B. Robustness Tests — Figures

Here, we present a set of figures similar to Fig. 5 for the robustness tests of Section 6. These figures show that despite slight quantitative differences introduced by various assumptions in the model, the outcomes are qualitatively similar to the original solution of Section 4.

Figure B.17: Total profit, discount & adoption levels for small C — single-modal β

Figure B.18: Total profit, discount & adoption levels for small C — single-modal σ

Figure B.19: Total profit, discount & adoption levels for small C — Linear IPv4 acquisition cost

Figure B.20: Total profit, discount & adoption levels for small C — Decreasing c_6

Appendix C. ICPs Lag Behind Users & ISPs

Here, we investigate a model in which the timing of ICPs' decisions is different from that of the ISPs and users. In other words, ICPs less frequently make decisions compared to the ISPs and users. From a modeling standpoint, we can remove the decision making of ICPs, and update the fraction of IPv6 accessible ICPs infrequently. In this setting, we can analytically find the Nash Equilibrium of the game played by ISPs (in a disagreement scenario).

This setting is analytically similar to the one where ICPs make decisions *after* each price announcement, but the timing of their decision making remains unknown to ISPs. The model formulation remains the same (as the original model of Section 3), except for the IPv6 ISP's utility function, which instead of $\gamma_6^{(i)}$ uses the previous fraction of IPv4-only ICPs ($\gamma_6^{(i-1)}$). In other words, instead of predicting what ICPs will do after announcement of prices, the IPv6 ISP uses the known number of IPv6-accessible ICPs. Next, we provide the analytical solution to this model.

$$\begin{aligned}
\sigma &= \frac{p_4 - p_6}{a\gamma_6^{(i-1)}} \\
\Rightarrow n_6 &= s\sigma, \quad n_4 = s(1 - \sigma) \\
f(p_4) &= n_4 p_4 - C(n_4 - 1)_+^2 \\
g(p_6) &= n_6 p_6 - D_6 n_6 \gamma_6^{(i-1)} \\
f'(p_4) &= 0 \\
\Rightarrow s - \frac{2s}{a\gamma_6^{(i-1)}} p_4 + \frac{s}{a\gamma_6^{(i-1)}} p_6 + \\
\frac{2Cs^2}{a\gamma_6^{(i-1)}} - \frac{2Cs}{a\gamma_6^{(i-1)}} - \frac{2Cs^2}{(a\gamma_6^{(i-1)})^2} p_4 + \frac{2Cs^2}{(a\gamma_6^{(i-1)})^2} p_6 &= 0 \\
\Rightarrow p_4^* &= \frac{1}{2 + \frac{2Cs}{a\gamma_6^{(i-1)}}} \left(\left(1 + \frac{2Cs}{a\gamma_6^{(i-1)}}\right) p_6 + a\gamma_6^{(i-1)} + 2Cs - 2C \right)
\end{aligned}$$

$$\begin{aligned}
g'(p_6) &= 0 \\
\Rightarrow \frac{s}{a\gamma_6^{(i-1)}}p_4 - \frac{2s}{a\gamma_6^{(i-1)}}p_6 + \frac{D_6s}{a} &= 0 \\
\Rightarrow p_6^* &= \frac{p_4}{2} + \frac{D_6\gamma_6^{(i-1)}}{2}
\end{aligned}$$

where s is the total user population at epoch i , p_4^* and p_6^* are the best response prices of the IPv4 and IPv6 ISPs, respectively. Depending on the values of these prices, the solution space is divided into three regions:

Region 1:

$$\begin{aligned}
0 \leq p_4^* - p_6^* &\leq a\gamma_6^{(i-1)} \quad \& \quad n_4 > 1 \\
a > D_6 \quad \& \quad s &\geq \frac{3a}{2a + D_6} \\
\text{OR} \\
a < D_6 \quad \& \quad s &\geq 1 + \frac{\gamma_6^{(i-1)}(D_6 - a)}{2C}
\end{aligned}$$

In this region, the Nash Equilibrium can be found as follows:

$$\begin{aligned}
p_4^* &= \left(aD_6(\gamma_6^{(i-1)})^2 + 2D_6\gamma_6^{(i-1)}Cs + \right. \\
&\quad \left. 2(a\gamma_6^{(i-1)})^2 + 4a\gamma_6^{(i-1)}Cs - 4a\gamma_6^{(i-1)}C \right) / (a\gamma_6^{(i-1)} + 2Cs) \\
p_6^* &= \left(aD_6(\gamma_6^{(i-1)})^2 + 2D_6\gamma_6^{(i-1)}Cs + \right. \\
&\quad \left. 2(a\gamma_6^{(i-1)})^2 + 4a\gamma_6^{(i-1)}Cs - 4a\gamma_6^{(i-1)}C \right) / (6a\gamma_6^{(i-1)} + 4Cs) \\
&\quad + \frac{D_6\gamma_6^{(i-1)}}{2}
\end{aligned}$$

Region 2:

$$\begin{aligned}
0 \leq p_4^* - p_6^* &\leq a\gamma_6^{(i-1)} \quad \& \quad n_4 \leq 1 \\
a > D_6 \quad \& \quad 1 < s &< \frac{3a}{2a + D_6}
\end{aligned}$$

The Nash Equilibrium in this region is obtained as follows:

$$\begin{aligned}
p_4^* &= \frac{p_6}{2} + \frac{a\gamma_6^{(i-1)}}{2} \\
p_6^* &= \frac{p_4}{2} + \frac{D_6\gamma_6^{(i-1)}}{2} \\
\Rightarrow p_4^* &= \frac{D_6\gamma_6^{(i-1)}}{3} + \frac{2a\gamma_6^{(i-1)}}{3}
\end{aligned}$$

Region 3:

$$a < D_6 \quad \& \quad 1 < s < 1 + \frac{\gamma_6^{(i-1)}(D_6 - a)}{2C}$$

In this region, there can be two Nash Equilibria:

$$\begin{aligned}
&p_4^* = p_6^* \quad \text{No IPv6 Users} \\
&\text{OR} \\
&p_4^* - p_6^* \geq a\gamma_6^{(i-1)} \quad \text{No IPv4 Users}
\end{aligned}$$