Journal of Economic Perspectives—Volume 23, Number 3—Summer 2009—Pages 21–36

Priced and Unpriced Online Markets

Benjamin Edelman

Some online resources are free and others are not—but it can be hard to predict which resources are in which category. Do users pay for web-based e-mail? Sometimes they do, as in the case of lifetime e-mail service from Pobox, but often they do not, as in the case of Hotmail from Microsoft or Gmail from Google. Do users pay for wireless Internet access at "hotspots"? Historically such access carries a fee, as it does at the T-Mobile HotSpots at many airports and hotels, but now a number of retail outlets are providing such access without a fee, such as Panera Bread, Whole Foods, and recently Starbucks. Do users pay for software? Many users do pay for a wide variety of commercial software, yet other software like the Linux operating system and the FireFox web browser can be downloaded without charge.

Zero prices offer important benefits, even relative to small positive prices. For one, fee-free access reduces transaction costs—eliminating the need for billing systems as well as, in many cases, account setup, user names, and the like. Furthermore, zero prices seem to create an environment of experimentation and progress for products and consumers (Lessig, 2002; von Hippel, 2001). Finally, consumers overwhelmingly favor zero-price products, even beyond what might be predicted by their ordinary efforts to maximize consumer surplus (Shampanier, Mazar, and Ariely, 2007).

Yet experience in other contexts offers cause for concern. Although marginal costs may be near zero for many levels of use of online resources, costs generally eventually increase as usage nears a capacity constraint given by technological capability or system design. More generally, experience in other contexts repeatedly reveals overconsumption, scarcity, and even hoarding when resources are

■ *Benjamin Edelman is Assistant Professor of Business Administration, Harvard Business School, Boston, Massachusetts. His e-mail address is ⟨bedelman@hbs.edu⟩.*

provided without charge. For examples, see Parry (2002) on highway congestion, Starkie (1998) on airport landing slots, and Baumol and Oates (1988) on pollution.

With competing forces both supporting and opposing zero prices, typical Internet-related activities—like surfing the web, web searches, and e-mail, along with behind-the-scenes practices like domain names and the allocation of IP (Internet Protocol) addresses—present a natural context to reevaluate our sense of the tradeoffs that arise between free and a positive price.

Surfing the Unmetered Web

To a user unfamiliar with the Internet, it may seem odd that web browsing often carries the colloquial term "surfing." Yet web browsing and literal surfing have a number of similarities. Both activities typically proceed without a detailed plan: surfers follow the waves, while web browsers follow the links. Just as a surfer pays no per-minute fee, web browsing typically proceeds without a cost for each additional page.

For current U.S. Internet users, paying per downloaded web page seems almost inconceivable; modern U.S. Internet access is overwhelmingly offered through flat-rate pricing. However, online services previously used two-part tariffs, with both a monthly fee and an additional charge proportional to usage. Initially, no major commercial dial-up service offered flat-rate pricing. But in a 1996 change, AOL offered unlimited access for $19.95 per month (Levinson and Odlyzko, 2008). After AOL added flat-rate pricing, its average customer usage tripled (Odlyzko, 2001). At that time, almost all U.S. residential telephone service already allowed unlimited local calls of unlimited duration with no extra fees. So, once AOL ended its per-minute fees, customers could use dial-up AOL as much as they wanted over their local phone lines without additional out-of-pocket cost.

During the same time period, dial-up users outside the United States typically paid per-minute fees for local phone calls. Thus, even if a user's Internet service provider allowed unlimited usage, the user's total cost still increased in proportion to connection duration. For example, even without usage caps at French Internet service providers, French dial-up usage remained at roughly the low rate AOL achieved before its switch to flat-rate pricing. When New Zealand phone service in 1999 switched from per-minute billing to a flat rate per month, Internet usage promptly jumped to roughly the level seen at AOL (Odlyzko, 2001). Figure 1 illustrates some of these patterns.

Users' strong preference for flat-rate pricing of Internet access is well-established. For example, in experiments that repeatedly adjusted access prices presented to California ISDN users, Chu (1999) finds that users are willing to pay a premium for flat pricing above and beyond the usage-based pricing that would otherwise apply to their actual usage. (ISDN was an early technology for increased Internet access speed and phone companies typically offered ISDN with per-minute fees.) Lambrecht and Skiera (2006) show similar patterns in 2003 data on German users' choice of DSL service package. (DSL provides high-speed Internet access

Figure 1
Usage Growth when Marginal Prices Drop

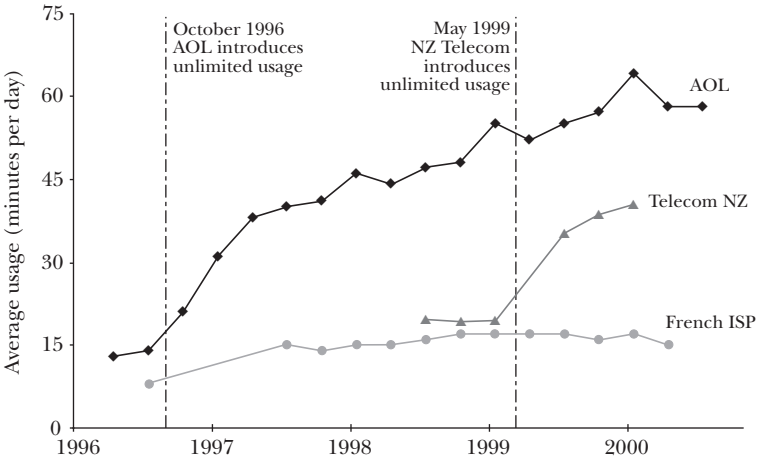

Source: Adapted from Odlyzko (2001) with permission.

over telephone lines.) Levinson and Odlyzko (2008) catalogue various reasons for user sensitivity to access charges: insuring against future usage spikes, avoiding the hassle or mental transaction costs associated with considering non-flat-fee pricing, and reducing the salience of fees. Users' preference for flat pricing predates the Internet: a 1979 Bell study showed that as many as 80 percent of users chose flat-rate telephone service even though measured service would have reduced their cost (Cosgrove and Linhart, 1979). The Internet's variety of content makes flat pricing all the more compelling: users would be hard-pressed to assess the quality of a site before viewing the site's materials, so paying per download would discourage the exploration of new or unfamiliar sites.

The cost structure of broadband Internet technology makes flat-rate pricing possible. Initial network construction costs are high—involving rights of way, cabling, and equipment. But once a network is in place, transferring more data entails minimal additional cost: Burnstein (2007) reports that medium-sized Internet service providers pay $0.10 per gigabyte (including the cost of links to other carriers, as well as transit costs charged by those carriers), while Clark (2008) estimates a range of prices per gigabyte from $0.06 to $0.18.

For household use, a gigabyte can be viewed as either a relatively large or a relatively small amount. Most households use the Internet for shopping, checking e-mail, finding information, and sometimes downloading music or watching brief YouTube videos. For these purposes, a few gigabytes a month of data will suffice. However, households that make greater use of online video will need considerably more capacity. Downloading a single movie can take three gigabytes or more. Increased use of web-based video, video-conferencing, certain kinds of online gaming, and online backup services would sharply increase data transfer requirements. Table 1 presents evidence on monthly network usage by broadband users

Table 1
ComScore Users' Bandwidth Usage as of July 2008

Monthly bandwidth usage (in gigabytes)	*Proportion of users*
< 1GB	44.3%
1GB–5GB	36.6%
5GB–10GB	11.0%
10GB–50GB	7.34%
50–100GB	0.53%
100GB–250GB	0.19%
> 250GB	0.003%

Source: ComScore, 2008.
Note: This analysis understates true household bandwidth consumption because comScore tracks usage at the level of the individual computer. When multiple household computers share a single Internet connection, comScore does not aggregate their bandwidth usage.

tracked by comScore: 44.3 percent of users transferred less than 1 gigabyte (GB) of data in July 2008, although 8 percent of users transferred 10GB or more.

As households begin to make greater use of online video, some Internet service providers have recently started to impose usage caps. In July 2008, Frontier Communications announced plans to limit customers to five gigabytes of usage per month—intended, a company spokesperson explained, "to make heavy users pay their fair share" (Grace, 2008). Other U.S. Internet service providers are following suit, albeit with somewhat higher limits—20 gigabytes in certain AT&T service areas, 40 gigabytes for certain Time-Warner service areas, 75 gigabytes at Cox, and 250 gigabytes at Comcast. Outside the United States, limits often remain considerably lower. For example, the basic high-speed Internet access plan offered by Telstra, the largest Australian Internet service provider, includes just 600 megabytes per month for $69.95 in Australian dollars (roughly $55 in U.S. dollars), while competitor Optus offers two gigabytes for $39.99 Australian (U.S. $32). In the United Kingdom, British Telecom offers ten gigabytes of monthly bandwidth in its basic £15.99 (U.S. $28) service.

Ordinarily, competition might be expected to protect consumers from pricing schemes they dislike, such as usage caps. But consumers typically have few choices for high-speed Internet access: most U.S. consumers can choose between at most two high-speed providers, typically a phone company (providing DSL service over telephone lines) and a cable television company (providing high-speed data service over coaxial television wiring).

Usage caps could hinder the spread of new services that tend to be particularly bandwidth-intensive. Some of the fastest-growing web features over the past few years have been web-based video, videoconferencing, and online file backup. For example, Slingbox lets a user watch home television on that person's computer screen anywhere an Internet connection is available. However, had Slingbox's developers seen that bandwidth caps would constrain usage of Slingbox, they might

not have bothered to invent Slingbox. More generally, even if usage caps exceed users' current requirements, the caps could nonetheless impede development of future services. This concern is more than speculative: In 2005, British Telecom's basic package offered users just 1 gigabyte of monthly usage. If usage had been equally limited at all providers worldwide, video services like Slingbox or YouTube might never have reached the market.

Lee and Wu (in this issue) address the related questions of network neutrality—whether Internet service providers may intentionally favor some traffic over others. Because video is among the most bandwidth-intensive uses of broadband Internet access, usage caps disproportionately limit video services. A complicating factor here is that some Internet service providers also own cable television franchises. These providers' investments give them a special incentive to disrupt online video: impeding online video protects demand for their cable television service. Thus, even if bandwidth caps seem to be content-neutral—affecting all content equally—the caps often disproportionately interfere with online video services.

Web Search and Advertising-Supported Services

How much would a user pay to use Google web search if it was not available for free? In other contexts, users pay substantial fees for information services. For example, even when universities commit to long-term campus-wide subscriptions to LexisNexis Academic, fees typically exceed $20 per user per year. For individual users, LexisNexis charges $1 to $4 for a single document from copies of public records, $3 for a news article, and $4 to $12 for financial information. Thus, for individuals seeking a series of documents, LexisNexis fees can quickly reach the triple digits.

Web-based information providers generally reverse the LexisNexis business model—offering service at no charge to the end-users who browse their sites but charging advertisers for the right to be included. By relying on advertising revenues, web-based information providers can avoid fees to consumers, which in turn encourages use and appeals to the widespread preference for free goods. Advertising support also eliminates the need for micropayments: if users were to pay for each search, many small payments would be required—a task poorly handled by existing payment systems. Evans (in this issue) presents the online advertising industry in greater detail. In short, advertising seems to make "free" an equilibrium in certain parts of the online content economy.

One key limitation of the advertising-supported model is that some services facilitate advertising more than others. For example, search engines have proven valuable for advertising because users' search terms reveal their intentions, including (in many instances) intentions about purchases. If a user searches for "laptop computer" or "hotel in LA," it is straightforward to identify advertisers particularly likely to want to reach that user. In contrast, other online services struggle to find suitable ads. What ads are best shown to a user reading news of a hurricane or an election? A site might be better able to select appropriate ads if it knew more about

the user—for example, what the user *previously* browsed or searched for. But tracking, storing, and analyzing long-term web usage creates considerable technical complexity as well as inevitable privacy concerns (Schwartz et al., 2007). Thus, advertising supports certain websites and web-based services fairly well, but others not so well. To wit, see extraordinary profits at Google, while online news sites consistently struggle.

Advertising is the most prominent example of online services funded by bundled offerings, but there are others. Some firms offer software for free—but charge for technical support. For example, Sun Microsystems offers the widely-used MySQL database at no charge—but consulting, training, and technical support all have fees. The Linux operating system is also available for free, but vendors like Red Hat and Ubuntu provide paid versions with technical support, warranties, and certain legal protections.

The Freedom to E-mail

Internet e-mail, the Simple Mail Transport Protocol (SMTP), has been free since its invention. There is no fee to send or receive e-mail, for no central power runs the Internet's mail servers. Instead, each participating mail server sends messages to other servers as needed. This decentralized design facilitated expansion by anyone wishing to join the network. In contrast, early proprietary competitors for electronic mail faced considerable additional complexity. Such competitors often charged per-message fees—at MCI Mail, $0.50 and up based on message length—so they had to carefully tabulate all messages sent and received. Furthermore, their proprietary systems required them to develop special gateways to transfer messages to and from other e-mail services—limiting their ability to expand quickly.

More recently, e-mail's openness has proven its major stumbling block. The same system that lets a person write to a long-lost friend equally lets a spammer write to millions of strangers. Symantec (2008) estimates that unsolicited commercial e-mail now constitutes 80 percent of e-mail volume, an estimated 100–120 billion bulk e-mails per year (SpamUnit, 2008; IronPort, 2008). The volume of unsolicited commercial e-mail has exploded to engulf legitimate mail: as users obtain increasingly effective filters, bulk mailers send more and more messages in order to reach a constant number of users. As of 2001, mail filter Barracuda Networks found that just 5 percent of e-mail was spam. Figure 2 depicts the rapid rise to present levels.

Extraordinary volumes of bulk e-mail are possible in large part because senders transmit mail through others' computers. Initially, spam senders enlisted "open relay" mail servers that allowed use by the general public. More recently, at least 85 percent of spam is transmitted through "botnets"—end-user computers that have fallen victim to security exploits that grant control to remote operators (Marshal, 2008). Use of these infected computers is available on a rental basis through numerous "botmaster" intermediaries. Competition pushes prices as low as $0.03

Figure 2
Percent of E-mail that is Spam

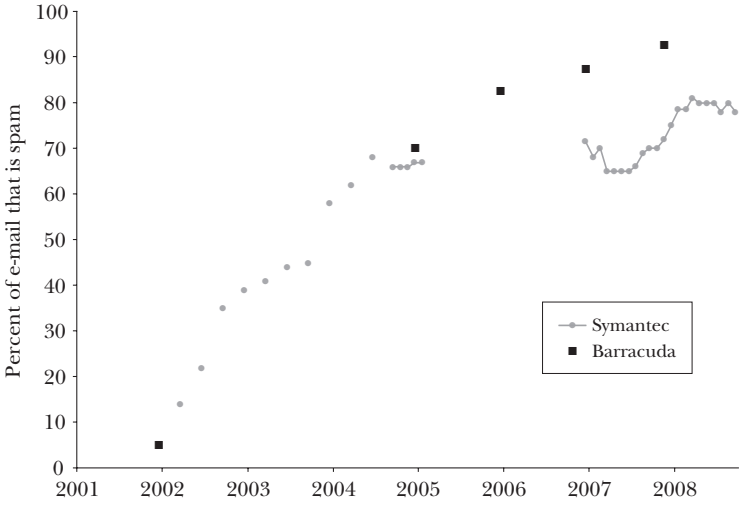

Source: Adapted from press releases from Symantec/Brightmail and Barracuda Networks.

per computer per week (LaMacchia, 2005), and a single computer can transmit (at least) thousands of messages per day. In principle, investigations can track spam back to its source, but in practice the many intermediaries make such analysis impractical.

Bulk e-mail imposes significant costs on recipients. Historically, industry sources placed high estimates on the value of employee time spent deleting junk e-mail—perhaps as much as $549 to $877 per user per year (Ipswitch, 2008; Nucleus Research, 2007; Barracuda Networks, 2008). More recently, improved filters have reduced the time required to delete spam, but users now face the problem of false positives—filtering errors moving legitimate messages to "Junk Mail" folders where they are effectively lost. Meanwhile, bulk mail also causes network costs estimated at more than $300 of storage per employee (typically high-cost server-based storage that holds spam pending deletion) and $10 to $25 per employee spent on mail filters (Fontana, 2003).

The decentralized design of e-mail makes bulk messages particularly hard to block. Many filters inspect message contents to try to identify unwanted mail. But such inspection is computationally burdensome when applied to millions of items. Furthermore, identifying spam through content analysis proves surprisingly difficult: Legitimate messages and spam both often come from strangers, and both often link to previously-unknown web sites. Distinguishing spam from mailing lists, automatic notifications, and requested commercial announcements is difficult, too. Of course, spammers continually adjust their messages to avoid detection.

A more robust defense against unwanted e-mail would authenticate each message to verify its sender, then confirm each sender's reputation in a database of individual or group experience. However, widely-used e-mail standards offer no

clear method for authentication: Despite encryption and digital signatures installed by leading-edge users, such systems are uncommon, and it is generally difficult to confirm that a message actually comes from its supposed sender (the address listed on the "FROM:" line). Several groups sought to add such confirmation, most prominently via the Sender Policy Framework (SPF), by cross-checking a mail server against the address listed as the message's source. While SPF can identify some counterfeit messages, its implementation has been slow: although SPF has been available since 2003, recent analyses report that just 8 to 13 percent of .com domains have adopted it (Usr/Local 2007; SPF-all, 2008; Measurement Factory, 2007). SPF's slow adoption reflects the underlying incentives: a site has little incentive to configure its servers to publish SPF data until others are validating SPF listings. But until SPF data is widespread, few servers will validate SPF listings. So e-mail's distributed implementation gives no single party sufficient incentive to take action.

Even a small fee would undermine the economics of spam. For example, convicted Denver spammer Min Kim in 2004 held a database of at least 200 million e-mail addresses and had made a profit of $250,000. If Kim sent each recipient a single message, then his profit per recipient was approximately one-eighth of a cent; if he sent multiple messages to each recipient, as most bulk mailers do, his profit per message was even less. Analyzing the effectiveness of nearly half a billion spams sent by a large botnet, Kanich et al. (2008) estimate that spam now offers a response rate of roughly one in twelve million—a sharp reduction thanks to improved filters and increasingly skeptical users. Thus, a fee of even a fraction of a penny per message would suffice to deter spam. It seems unlikely that such a fee would deter any reasonable user from sending ordinary personal messages.

Pundits have called for fee-based mailing systems for several years (Enyart, 2002). Bill Gates suggested "e-mail stamps" in a World Economic Forum speech in 2004, predicting that spam would be much reduced within two years because such fees would undermine the incentives behind bulk e-mail (Hansell, 2004). Yet there are several reasons why no widely-used fee-based mailing system exists. Processing e-mail payments would require robust authentication and tracking—a far cry from current openness of e-mail. Furthermore, most implementations of detailed message tracking entail centralized records of who sent mail to whom, but such records would invite both litigation and regulation. Additional complexity would come from inevitable pressure to exclude certain mailings from fees: for example, e-mails for announcements, notifications, mailing lists, and the like. Finally, the necessary institutions simply do not exist; no single e-mail provider is large enough to start the process. Thus, the price of e-mail remains zero, and spam remains widespread.

Why does "free" make sense for browsing and search but not for e-mail? Part of the difference seems to come from the decision-making sequence. In the context of browsing and search, free access encourages providers to develop materials users want—such as useful websites and search engines. In contrast, in the context of e-mail, zero-price access invites a business where a sender counts one answer in ten million as a "success." More generally, with materials pushed to unwitting recipients, "free" shifts from benefit to detriment.

Paying for Domain Names

The preceding examples—web browsing, search, and e-mail—each offer access with a zero price on the margin. Not so for domain names, where a contractor has managed to charge fees that total billions of dollars.

In 1983, computer scientists Paul Mockapetris and Jon Postel designed the Domain Name System (DNS)—servers and communication standards to locate the servers associated with the domain names users request. To request a new domain in a major top-level domain, like the now-familiar .com, an interested user sent a brief e-mail to Postel. Postel offered domains on a first-come-first-served basis, and an applicant could obtain any domain not already claimed by someone else. Postel's modest costs were paid by a series of grants from the Department of Defense, so he saw no need to charge domain recipients.

The Internet's rapid growth strained the early domain name system. Even with an assistant processing routine requests, interest in .com and other top-level domains came to exceed Postel's interests and capabilities. In 1993, the National Science Foundation put Postel's domain name allocation functions out to bid, and Virginia software developer Network Solutions won the contract. Initially, Network Solutions charged the NSF a fixed fee of $5.9 million per year, and Network Solutions continued Postel's policy of providing registrants with domain names at no charge. But in 2005, the NSF approved a Network Solutions request to begin charging domain registrants—ending NSF's payment to Network Solutions but requiring that each registrant pay Network Solutions $50 per year per domain name.

With Network Solutions' move to charge domain registrants, the domain name system was positioned for rapid growth. Crucially, per-domain fees resolved long-standing ambiguity as to registration of multiple domains. Previously, each entity was generally asked to limit itself to a single domain. For example, in 1993 Digital Equipment Corporation requested digital.com, but registration staff opposed the request since the company already held dec.com (Mueller, 2004). But with a price on each domain, Network Solutions was willing to provide each registrant with as many domains as it cared to purchase.

Meanwhile, demand had increased sharply—from fewer than 7,000 .com domains in 1993 to 382,000 in 1995 to four million in 1999. Network Solutions' revenue grew correspondingly, while many of the firm's costs were fixed thanks to automation. With fast-growing profits reaching more than $26 million in 1999, Network Solutions achieved a $21 billion acquisition by VeriSign in 2000.

Through 1999, Network Solutions was the sole provider of .com domains. Domain registrants often complained of Network Solutions errors or delays; in hindsight, management conceded that its service was "the worst in the industry" (DomainInformer, 2007). But under 1998 instructions from the Department of Commerce and subsequent oversight by the Internet Corporation for Assigned Names and Numbers (ICANN), more than 150 registrars now offer customer-facing functions relating to domain registration. VeriSign still provides the centralized

database "registry" services, but by agreement with ICANN, VeriSign divested its direct relationships with registrant customers.

In some respects, domain name pricing and operations are working smoothly. VeriSign's wholesale price is now $6.86 per domain per year, and competition among registrars limits markup of VeriSign's price. For example, the largest registrar, GoDaddy, posts a list price of just $10.69 per domain. These lower prices yield dramatic benefits for consumers. For example, consider the four million domains registered by 1999: Annual renewals for the subsequent ten years would have cost $350 to $500 per domain at 1999 prices. But with registrar competition, actual renewal expense fell to roughly $100 to $150 per domain—yielding a total benefit to consumers of more than a billion dollars. Newly-registered domains—the more than 70 million .com sites registered since 1999—make the true gains to consumers even larger.

Yet there is ample reason to question other aspects of the domain name market. Much of the registrars' costs come from renewal reminders, payment processing, and customer support for renewal and payment. That is, soliciting and receiving payment consumes most of the payment—which is reminiscent of a tollbooth where collection costs approach revenues. Domain registrants might be better off with a registration system that charged one-time fees or otherwise reduced ongoing administrative costs.

Furthermore, with ICANN approval, VeriSign has begun to increase its registry fee: 2007 and 2008 brought increases from $6 to $6.42 to $6.86 per .com per year. VeriSign claims that its price increases reflect growing costs such as increased security efforts. But critics note that VeriSign benefits from economies of scale and declining costs of information technology inputs (such as servers and bandwidth). Critics thus worry that VeriSign's charges reflect its market power—that .com domains have few close substitutes. Indeed, competing registries offer domains like .biz, .info, and .us. But these domains have less cachet, so typical .com registrants perceive a need to pay even increased VeriSign fees.

Registrants' incentives further impede registry competition. Having invested in a .com domain, whether through business cards, TV ads, or word-of-mouth, a company is ill-equipped to move to another web address. Furthermore, VeriSign's price increases are small for any individual company, so no single company has much incentive to seek fee reductions. Yet fees are large when totaled across tens of millions of domains and across many years.

Shortfalls in the domain name system remain widespread and widely discussed. For example, the "Whois" listings of domain ownership are often inaccurate (for discussion, see Edelman, 2002), and a number of domains seem to infringe on the rights of trademark holders (Edelman, 2008). But these deficiencies are not obviously attributable to pricing rules. Despite some degree of waste and inefficiency from the billing apparatus and from sole-sourced functions, the domain name system at least works reliably, without the uncontrolled attacks seen in e-mail.

Can lessons from domain name pricing help prevent unwanted e-mail? In contrast to decentralized e-mail servers burdened with spam, domain names at least offer a strong central authority that reliably operates its resource. A strong central

player could offer an alternative e-mail system—perhaps an improved version of the private messaging systems already included on sites like Facebook. Such a system would robustly authenticate participants, blocking most spam. But new problems would arise: just as VeriSign sees opportunity to increase its .com prices, a centralized private e-mail system operator could seek to claim much of the value its service creates—perhaps charging high (or discriminatory) per-message fees. When compared with this possibility, the current state of e-mail looks considerably more palatable.

Running Out of Numbers: *De Minimis* Pricing and the Allocation of IP Addresses

Every computer connected to the Internet needs an IP address (Internet Protocol address) to identify itself, to label its messages to other computers, and to receive incoming communications. The Internet's primary numbering system, called IPv4, features 32 binary digits, allowing 2^{32} distinct addresses (approximately four billion). Nonprofit administrative organizations, known as Regional Internet Registries (RIRs), receive applications from interested networks (such as Internet service providers, companies, and universities), confirm each applicant's bona fides, and determine how many addresses each applicant requires.

At present, IPv4 addresses are provided at minimal cost to any network that can make a legitimate claim to need them. Prices are intended only to cover costs of distributing the numbers, not to claim a share of the numbers' value. In fact, prices are fixed for broad swaths of allocation sizes, as shown in Table 2. Notice that regional Internet registries offer flat fees to "extra large" networks of 2^{18} or more addresses; specifically, the largest U.S. Internet service providers need pay the RIR for North America, the American Registry for Internet Numbers or ARIN, just $18,000 per year. The absolute price for requesting more numbers is low, and given the fixed pricing within broad swaths of allocation sizes, the marginal price of additional addresses is often effectively zero. With a finite supply of numbers and a trivial price, these resources are speeding towards exhaustion. At the current rate of consumption, IPv4 addresses will become unavailable in 2012 (Huston, 2008).

Anticipating a shortage of IPv4 addresses, engineers have created an alternative address space with greater capacity. The IPv6 standard extends IP addresses to 128 binary digits (approximately 3.4×10^{38} addresses, more than three billion billion billion). But a computer with only a v6 address cannot access the existing Internet directly. Instead, a v6 computer needs assistance from a translation server in order to, for example, browse the web. Furthermore, translation typically cannot accommodate unusual or nonstandard applications like custom corporate software, multiplayer games, or various video services. Right now, networks have enough v4 addresses, so few networks are currently willing to pay extra for v6-capable systems. As a result, hardware and software developers have hesitated to build such systems—preventing networks from moving to v6 even if they wanted to do so. In

Table 2
IP Address Fees, North America

Network size	*Annual fee*
Less than 2^{12} addresses	$1,250
2^{12} to 2^{13} addresses	$2,250
2^{13} to 2^{16} addresses	$4,500
2^{16} to 2^{18} addresses	$9,000
More than 2^{18} addresses	$18,000

Source: ARIN (American Registry for Internet Numbers), 2009a.

Edelman (forthcoming), I offer additional discussion of the technical and incentive impediments to v6 transition. The U.S. Department of Commerce (2005) estimates that v6 transition will take 25 years and will cost $25 billion.

The lack of price signals in the IP numbering system has invited inefficient use of existing numbers and also has impeded transition towards alternatives. Because regional Internet registries offer addresses at minimal cost, networks have had little incentive to economize on usage. To the contrary, networks may seek to hold extra addresses as a hedge against any future shortage. In addition, some networks previously received an abundance of addresses—an artifact of early technical constraints that initially limited allocations to powers of 2^8 addresses (that is, networks of size exactly 2^8, 2^{16}, or 2^{24} addresses) rather than allowing flexible network sizes that better matched networks' actual requirements.

With suitable incentives, large early recipients may be willing to make addresses available to others. Other addresses could come from networks that have ceased operations (but have so far failed to return addresses) or from networks where a change in business focus reduced address requirements. Superficially, allowing transfers of such numbers may seem trivial. Indeed, transfers are under discussion at regional Internet registries worldwide. APNIC, the RIR in the Asia/Pacific region, has moved to allow transfers between any willing provider and any willing recipient. But just as a person cannot easily transfer a credit card number or license plate number, so too are IP address transfers historically restricted. Exhibit 1 recites ARIN's current policy prohibiting such transfers.

While trading IP numbers might extend the life of IPv4, practical and logistical concerns limit the gains from such trading. First, address transfers affect the routing system that directs messages between networks across the Internet. The routing system supports hierarchical aggregation: If an Internet service provider serves many customers using a single contiguous block of IP addresses, those many customers all require just a single entry in the routing table. However, as a result of trading, many Internet service providers may end up with multiple discontinuous blocks of addresses. If many networks begin to use multiple sets of discontinuous addresses, more routing entries would be required and network operators would have to upgrade their routers more often than in the past. There are 150,000

Exhibit 1
ARIN Policy on IP Address Transfers

"Number resources are non-transferable and are not assignable to any other organization unless ARIN has expressly and in writing approved a request for transfer. [N]umber resources are not 'sold' under ARIN administration. Rather, number resources are assigned to an organization for its exclusive use for the purpose stated in the request, provided the terms of the Registration Services Agreement continue to be met and the stated purpose for the number resources remains the same. . . .

"ARIN will consider requests for the transfer of number resources only upon receipt of evidence that the new entity has acquired the assets which had, as of the date of the acquisition or proposed reorganization, justified the current entity's use of the number resource."

Source: ARIN (American Registry for Internet Numbers), 2009b.

routers around the world at an estimated cost averaging $30,000 per router, and with regular updates, the routing system is believed to cost on the order of $2 billion per year. The demands of routing data between multiple discontinuous blocks of addresses could enlarge this cost sharply, conceivably even exceeding router manufacturers' ability to increase performance sufficiently (Li, 2007).

Second, allowing trading of numerical IP addresses might allow wealthy networks to buy up large swaths of addresses to impede competitors' expansion. This worry has analogues in other markets. For example, European competition authorities have investigated whether paid transfers of airport landing slots might tend to reduce or exclude competition (U.K. Office of Fair Trading, 2005).

Of course, one can imagine writing rules to address these issues. For example, a transfer policy could disallow subdivision of large blocks. Alternatively, a transfer policy could require that each address recipient satisfy its entire need in a *single* transfer—not several transfers of smaller address blocks—which would still allow multiple smaller networks to purchase a large block jointly from a large network. This approach to v4 transfers is currently under consideration at ARIN (Leibrand, 2008).

But even if IPv4 addresses come to be used much more efficiently, networks will ultimately need to move to an expanded address space like v6. Setting a higher positive price for IP address numbers and allowing transfers within v4 can help facilitate this transition. For example, a positive transfer price encourages existing networks to vacate v4 addresses: Every address returned yields funds that a network can use for other purposes. Furthermore, a higher price on v4 space creates a cost to v4 expansion, encouraging networks to expand in other ways. As a result, v4 prices and transfers create a direct financial benefit to developing and installing v6-capable systems. Thus, v4 transfers could push networks towards v6 in a way that has proven infeasible with v4 space widely available at minimal cost.

Transfers can also mitigate some costs of v6 transition. Without transfers, transition to v6 would probably occur in a more-or-less arbitrary order. But some networks are likely to have higher switching costs than others: some technologies better lend themselves to early use of v6; some networks have newer equipment that

Table 3
Alternative Pricing Rules

	Centrally administered	*Distributed*
Zero price	Linux, free software	E-mail
Positive fixed price but zero marginal price	IP addresses	Web browsing
Positive marginal price	Domain names	Commercial software

can better accommodate v6; some network staff have (or can more easily obtain) the required expertise. Networks with high v6 switching costs can pay to receive space from networks that can vacate v4 space more easily, which would reduce the total costs of transition.

Looking Forward

Online services typically have low marginal costs, despite high fixed costs. With total cost at most loosely related to quantity, prices become unpredictable. As Table 3 shows, some online services are offered at zero price, some at a positive fixed price with zero marginal price, and some at a positive marginal price—filling the field of plausible pricing models. Setting a low or zero price increases a service's appeal and gains more users—but yields low or no revenue. Nonetheless, zero prices appear to be sustainable when there are adequate profits in complementary businesses like advertising or technical support. Indeed, once development or installation expenses have been incurred, social welfare may be maximized with a zero price that encourages usage. In other words, there are sensible reasons to oppose paying for each web page, each e-mail, each domain name, or each IP address.

Flat or zero prices serve to accelerate usage—reducing technical complexity, simplifying billing, and attracting users. But the same low prices that drive early adoption can also create future problems, like an excess of unwanted e-mail or a shortage of available IP addresses. In a few markets, such as domain names, allocations have switched from zero prices to positive prices. But revised pricing can also create large wealth transfers, as with the considerable fees paid to domain registration providers. Usage-based pricing of U.S. broadband providers raises similar questions, and it remains unclear whether users and policymakers will accept the metered usage future that network operators now seem to favor.

Experience reveals no necessary pattern in pricing of online resources. Rather, online service pricing reflects a process of discovery and innovation with a series of transitions between free and priced online services. Services may tend to be free when their marginal costs come particularly close to zero, but free services face pressure from various kinds of technological change (like the increased use of video that dramatically raises bandwidth requirements) as well as cost increases induced by outside attacks (like spam) and capacity constraints (like a shortage of

IPv4 addresses). Yet even these shocks may spur further changes in due course, again reducing marginal cost and restoring the feasibility of zero-price access.

Meanwhile, significant fortunes can be made both in offering material at zero price (as in the business of Internet search), and in instituting a positive price for something previously free (as with domain names). Reminiscent of the old adage about losing money on every unit but making it up in volume, online markets challenge norms about who should pay, when, and why.

■ *I thank K. C. Claffy, Peter Coles, Tom Eisenmann, Al Roth, and Tom Vest for their suggestions and feedback. Matthias Baeuml provided excellent research assistance.*

References

ARIN (American Registry for Internet Numbers). 2009a. "Fee Schedule." Available at https://www.arin.net/fees/fee_schedule.html.

ARIN (American Registry for Internet Numbers). 2009b. "ARIN Number Resource Policy Manual," Version 2009.2, April 1, 2009. Section 8.1. ⟨https://www.arin.net/policy/nrpm.html⟩.

Barracuda Networks. 2007. "Barracuda Networks Releases Annual Spam Report." http://www.barracudanetworks.com/ns/news_and_events/index.php?nid=232.

Barracuda Networks. 2008. "Spam Cost Calculator." Available at http://www.barracudanetworks.com/ns/resources/spam_cost_calculator.php.

Baumol, William J., and Wallace E. Oates. 1988. *The Theory of Environmental Policy.* 2nd ed. Cambridge: Cambridge University Press.

Burnstein, Dave. 2007. "Price of Bandwidth Continues Dropping." *Future of TV,* October 21. http://web.archive.org/web/20080202122723/http://www.futureoftv.net/.

Chu, Karyen. 1999. "User Reactions to Flat-Rate Options under Time Charges with Differentiated Quality of Access: Preliminary Results from INDEX." Presented at ISQE'99, Workshop on Internet Service Quality Economics.

Clark, David. 2008. "A Simple Cost Model for Broadband Access: What Will Video Cost?" Presented at the Telecommunications Policy Research Conference. http://cfp.mit.edu/publications/docs/DDC.Cost.analysis.TPRC.R1.pdf.

comScore. 2008. "Bandwidth Consumption: US Home Broadband Users." July 2008. Report available through comScore's paid "Media Metrix" service.

Cosgrove, James, and Peter Linhart. 1979. "Customer Choices under Local Measured Telephone Service." *Public Utilities Fortnightly,* August 30, 104(5): 27–31.

DomainInformer. 2007. "An Interview with Champ Mitchell, Chairman and CEO, Network Solutions." http://www.domaininformer.com/interviews/interview/070321NetSol.html.

Edelman, Benjamin. 2002. "Large-Scale Intentional Invalid WHOIS Data." http://cyber.law.harvard.edu/archived_content/people/edelman/invalid-whois/.

Edelman, Benjamin. 2008. "Typosquatting: Unintended Adventures in Browsing." *McAfee Security Journal,* Fall, pp. 34–7.

Edelman, Benjamin. Forthcoming. "Running Out of Numbers: Scarcity of IP Addresses and What to Do About It." In *Proceedings of the First Conference on Auctions, Market Mechanisms and Their Applications.* Springer-Verlag Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering series. Springer-Verlag.

Enyart, Robert. 2002. "Fee-Based Message Delivery System." United States Patent Application 20060041505.

Fontana, John. 2003. "Tallying the True Cost of Spam." *Network World,* November 17.

Grace, Tom. 2008. "Frontier to Limit Internet Usage." *The Daily Star,* August 16.

Hansell, Saul. 2004. "Speech by Gates Lends Visibility to E-Mail Stamp in War on Spam." *New York Times,* February 2.

Herrin, Bill. "What Does a BGP Route Cost?" http://bill.herrin.us/network/bgpcost.html.

Huston, Geoff. 2008. "IPv4 Address Report." Available at http://www.potaroo.net/tools/ipv4/index.html.

Ipswitch. 2008. "Ipswitch Messaging Division Spamometer Reports High Costs for Spam." September 3. http://www.marketwire.com/press-release/Ipswitch-895727.html.

IronPort. 2008. "2008 Internet Security Trends." Available at http://www.ironport.com/securitytrends/.

Kanich, Chris, Christian Kreibch, Kirill Levchenko, Bandon Enright, Geoffrey M. Voelker, Vern Paxson, and Stefan Savage. 2008. "Spamalytics: An Empirical Analysis of Spam Marketing Conversion." CCS'08 (Conference on Computer and Communications Security, October 27–31, 2008, Alexandria, Virginia.)

LaMacchia, Brian. 2005. "Security Attacks and Defenses." 47th Meeting of International Federation for Information Processing. Powerpoint available at http://www.laas.fr/IFIPWG/Workshops&Meetings/47/WS/08-LaMacchia.pdf.

Lambrecht, Anja, and Bernd Skiera. 2006. "Paying Too Much and Being Happy About It: Existence, Causes and Consequences of Tariff-Choice Biases." *Journal of Marketing Research,* May, 43(2): 212–23.

Leibrand, Scott. 2008. "Possible Revisions to 2008-6." ARIN-PPML Message (message in the official public policy discussion list of the ARIN.) http://lists.arin.net/pipermail/arin-ppml/2008-October/012335.html.

Lessig, Lawrence. 2002. "The Architecture of Innovation." *Duke Law Journal,* 51(6): 1783–1801.

Levinson, David, and Andrew Odlyzko. 2008. "Too Expensive to Meter: The Influence of Transaction Costs in Transportation and Communication." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences,* 366(1872): 2033–46.

LexisNexis. 2007. "LexisNexis Academic." Catalogue. Available at http://www.lexisnexis.com/Academic/catalog/2007pdfs/Academic2007Domestic.pdf.

LexisNexis. 2009. "LexisNexis by Credit Card." Webpage. http://web.lexis.com/xchange/ccsubs/cc_prods.asp.

Li, Tony. 2007. "Router Physics—IP Markets." ARIN XX (20th meeting of the ARIN), October 17, 2007.

MacDonald, Jake. 2007. "Attack of the Killer Bots." *Canadian Business Online,* June 4.

Marshal. 2008. "Srizbi Now Leads the Spam Pack." http://www.marshal.com/trace/traceitem.asp?article=567.

Measurement Factory. 2007. "DNS Survey: October 2007." Available at http://dns.measurement-factory.com/surveys/200710.html.

Mueller, Milton. 2004. *Ruling the Root: Internet Governance and the Taming of Cyberspace.* Cambridge, MA: MIT Press.

Nucleus Research. 2007. "Spam Costing US Businesses $712 per Employee Each Year." April 2. http://nucleusresearch.com/news/press-releases/nucleus-research-spam-costing-us-businesses-712-per-employee-each-year/.

Odlyzko, Andrew. 2001. "Internet Pricing and the History of Communications." *Computer Networks,* August, vol. 36, pp. 493–517.

Parry, Ian. 2002. "Comparing the Efficiency of Alternative Policies for Reducing Traffic Congestion." *Journal of Public Economics,* 85(3): 333–62.

Roth, Alvin. 2007. "Repugnance as a Constraint on Markets." *Journal of Economic Perspectives,* 21(3): 37–58.

Schwartz, Ari, et al. 2007. "Consumer Rights and Protections in the Behavioral Advertising Sector." Comment to the FTC. http://www.cdt.org/privacy/20071031consumerprotectionsbehavioral.pdf.

Shampan'er, Kristina, Nina Mazar, and Dan Ariely. 2007. "Zero as a Special Price: The True Value of Free Products." *Marketing Science,* 26(6): 742–57.

SpamUnit. 2008. "Spam Statistics." http://www.spamunit.com/spam-statistics.

SPF-all. 2008. "Stopping eMail Forgery." Webpage. http://spf-all.com/stats.html (as of October 20, 2008).

Symantec. 2008. "The State of Spam: A Monthly Report—July 2008." Doug Bowers, executive editor; Dermot Harnett, editor.

Starkie, David. 1998. "Allocating Airport Slots: A Role for the Market?" *Journal of Air Transport Management,* 4(2): 111–16.

U.S. Department of Commerce. 2005. *Planning Report 05-2: IPv6 Economic Impact Assessment.* Prepared by RTI International for National Institute of Standards and Technology. http://www.nist.gov/director/prog-ofc/report05-2.pdf.

U.K. Office of Fair Trading and Civilian Aviation Authority. 2005. "Competition Issues Associated with the Trading of Airport Slots." http://www.oft.gov.uk/shared_oft/reports/oft_response_to_consultations/oft832.pdf.

Usr/Local. 2007. "SPF/Sender ID Adoption Rates—April 2007." Available at http://www.usrlocal.com/stories.phtml?sid=1716 (last accessed November 3, 2008).

von Hippel, Eric. 2001. "Innovation by User Communities: Learning From Open-Source Software." *MIT Sloan Management Review,* 42(4): 82–86.