

Diffusion and Adoption of IPv6 in the United States

Hillary A. Elmore, L. Jean Camp, Brandon P. Stephens
School of Informatics
Indiana University Bloomington
{elmoreh, ljcamp, bpstehe}@indiana.edu

Abstract

What are the best case, current projection and reasonably optimistic cases for the adoption of the IPv6 protocol? This paper will analyze IPv6 routes from the ARIN Project to quantify the current adoption rate in order to offer three distinct views of three possible paths.

The data source consists of three hourly points from the last 60 months of IPv6 route records. This data will then be compiled and plotted on three distinct dissemination curves to determine three possible adoption timelines.

This paper will also discuss potential gains in security from adoption of the IPv6 protocol, along with a background discussion of analysis methods and the use of dissemination curves as a means of reporting findings. The paper will not include assertions about exhaustion of IPv4 space; however, it will inform research on the potential implications of this resource exhaustion.

Predictive S-curves will be used to demonstrate at what rate adoption will occur over the next 8, 10 and 12 years domestically. This will aid in showing whether or not domestic IPv6 adoption is on track with global adoption.

1. Introduction

In its first conception in 1977, it was believed that the address limits contained within IPv4 would never be an issue. As time progressed, the limits of IPv4 began to become apparent. The Internet became more popular and ubiquitous than ever imagined at the time of conception. Now, not only do desktop computers require IP addresses, but mobile devices are beginning to require connectivity through IP protocols. This increase in needed IPv4 addresses is beginning to exhaust available addresses. With the limit of unique IPv4 addresses (4,294,967,296) being approached—only 17% remains as of October 2007¹—a new protocol had to be developed in order to accommodate the continual expansion of the global network (Huston, 2007).

Enter IPv6.

IPv6 was heralded as “the” new wave protocol that would be responsible for securely connecting the world. It was designed to solve many possible problems faced by IPv4. After IPv6 was introduced, IPv4 adapted to include several features that IPv6 implemented.

This paper illustrates that adoption has been slow at best. We argue that this is due primarily to the misalignment of incentives for the adopters and possibly due to

¹ ICANN Factsheet on IPv6,
<http://www.icann.org/announcements/factsheet-ipv6-26oct07.pdf>, accessed Jan. 22, 2008.

lackluster domestic governmental policy. We provide three possible market responses to current US government policy. Since IPv6 made its groundbreaking introduction into the IP standards realm in 1998, it has made a mere ripple in the wake of Internet traffic.

2. Related Work In Economics Of Information Security

When considering the adoption of IPv6 in terms of its potential security ramifications, it is useful to look at previous literature regarding the economics of information security. In particular, the issues of patching behavior and privacy valuation can be related to the adoption (or lack thereof) of IPv6.

In addition to these specific issues, some of the largest issues in the study of the economics of information security also apply to the adoption of IPv6. For information security to be effective, incentives must be properly aligned and must not allow hidden actions (Anderson, 2006). In the case of IPv6 adoption, security incentives are not properly aligned. The party who suffers when private information is lost is an individual, not the adopter who has chosen not to implement IPv6. The adopter chooses not to apply IPv6 because the benefits are too low, the presence of uncertainty, the risk too high, and their own perceived risk too low to justify the cost.

Network externalities are also applicable to the adoption of IPv6. Especially in terms of security, an increase in the size of the network increases the value of the network to each user (Anderson, 2006). As a result, a small network of deployment results in small total benefit. With the small network, the cost of adopting a new technology is greater than the benefits gained by the added

protections of the network. In marketing terms, the technology may never reach critical mass. To address the problem of network externalities limiting adoption of a superior technology, Ozment and Schechter propose a combination of subsidies, mandates and bundled technologies (WEIS 2006).

Patching behavior literature is generally based on the fact that not everyone who can apply a patch does so. Cavusoglu et al. note that there are four main reasons individuals and firms fail to apply patches (2006). First, too many vulnerabilities exist to individually apply patches. Secondly, patches will not be applied until they are trusted, and they cannot be trusted until they have been tested on a system. Third, there is no standardization in the distribution of patches, and finally, patches require testing after installation.

These difficulties in patching apply, in many ways, to the difficulties in IPv6 adoption. The most obvious parallels are to the testing, trust, and installation aspects of patches. Though most routers are now sold with IPv6 capabilities, the move from IPv4 to IPv6 could still cause disruptions in service, as well as temporary increases in security vulnerabilities (Rowe, 2006). Regardless of the capabilities with which a device is sold, IPv6 must be enabled on devices for them to work. A product advertised with a full IPv6 stack may not have a fully functional stack. Because of the complexity of IPv6, two implementations may be compliant but suffer subtle failures in interoperability. The increases in switching costs to IPv6 coupled with the bootstrapping effect provide a plausible basis for the current slow adoption of IPv6 in the United States (Rowe, 2006).

3. Related Work In Diffusion

Technological diffusion has been studied in various disciplines for the past several decades, most notably in economics and sociology (Hall, 2004). Bass presents the epidemic, information-based model in *A New Product Growth for Model Consumer Durables*. This model assumes that, for all consumers except innovators, pressure to adopt a new technology increases as time and the number of other adopters increases (Bass, 1969). It also assumes that once a party is aware of the technology, it will be adopted almost immediately. This model naturally lends itself to the development of an s-shaped diffusion curve; where diffusion of the technology first is spurred by innovators, but as the number of adopters increases their influence will diminish, as will the number of consumers who have not yet adopted the new technology. Thus, the rate of adoption slows as diffusion reaches its peak, completing the s-curve.




Figure 1: The Technology S-Curve (Bass, 1969)

The epidemic model is most useful when studying the gradual impact of a new innovation. Knowledge of new technology takes longer to spread than, for example, knowledge of a world event because information on world events can be summarized, simplified, and broadcast from a common source (Geroski, 2000). New technologies also often have both hardware and software components. Hardware installation and factual descriptions of the

protocol can be broadcast from a single source, and the information itself changes slowly. In practice, however, adopting IPv6 for a specific network and utilizing it in dynamic network conditions requires experience. The tacit knowledge that comes only from experience cannot be broadcast, but must be learned first-hand or through mentoring. Therefore, delay in adoption in the epidemic model can be attributed to the time it takes for the base of knowledgeable parties to reach a critical mass. Intuitively, simpler software is likely to diffuse more quickly than complicated software.

Though the epidemic (or population) model of diffusion is the most popular, it may not be the most applicable in all situations and for all purposes. The probit model of diffusion assumes that technology will diffuse and examines not the overall diffusion, but the diffusion to individual firms (Geroski, 2000). To do this, the probit model assumes that every firm sets a threshold of profitability. When profitability of the new technology is below the threshold, the firm will not adopt the innovation, but when the profitability rises above this internally defined threshold, the firm will choose to adopt the new technology. The S-curve in this model is determined by the change in profitability of a particular technology and firms' changes in their threshold for adoption as more information about the technology becomes available. Therefore, the probit and epidemic models differ in that the epidemic model assumes that adoption of a new technology will occur when a firm becomes aware of it, while the probit model explains, to some degree, the apparent hesitation of firms to adopt new technologies even after they are aware of the technology. The probit model and the S-curve are not mutually exclusive. This is particularly true when profitability (benefit) is a function of

the number of previous adopters (e.g. network effects).

Though the literature on technological diffusion makes a strong case for the S-curve model, like all models it has flaws. In particular, though diffusion can be studied retrospectively, using the S-curve model to predict and prescribe new technologies is problematic, particularly when applied within individual firms (Christensen, 1992). Improvement in a technology is difficult to predict, and is sensitive to external factors. Simply forecasting that accepted technology is approaching its natural improvement limit can have a direct downward effect on the technology's growth trajectory (Christensen, 1992). This effect both causes and is furthered by a subsequent decrease in engineering resources devoted to the displaced technology, as well as an increase in resources toward the innovative technology. The limit of IPv4 has been announced several times, but new technologies have been found. However, a confirmation of that limit—such as the development of a market for IPv4 addresses—could increase IPv6 adoption.

Much of the diffusion literature is built upon two basic assumptions: first, that new technology—once it is released—and the old technology do not change during the diffusion process, and second, that the new technology is better than the old technology. This first assumption is often proven wrong in the real world, since the quality of the new product does, in fact, increase during diffusion, and that this improvement causes the equilibrium adoption point to rise continuously (Chow, 1967). This has certainly been observed with IPv4—for example, with DHCP. Hall argues that it cannot be assumed, either, that the old technology does not change during the diffusion of the new technology, as old

technologies can experience a “last gasp” improvement in an effort to remain dominant (Hall, 2004). Changes in either or both of the technologies delay adoption of the new technology because changes increase the uncertainty of the benefits of the new technology and can increase the risks, especially if the old technology is able to incorporate aspects of the new. Again, this applies to the case at hand. One of the major benefits of IPv6 is the increased security of the protocol and the difficulty it presents to a potential attacker. However, since the introduction of IPv6, many of the security enhancements have been applied to IPv4, thereby reducing the benefits of switching.

4. Data Experiment Setup

The major question that this analysis seeks to answer is: given current adoption rates, when might IPv6 have significant domestic market penetration? We extrapolated current route data into five models using three sets of assumptions. The first model, the most pessimistic, is a best fit for current data, assuming no exogenous influence on the demand for IPv6. Exogenous influences include resource exhaustion, policy action, a major security event, or other forced market tipping point.

The second analysis assumes the creation of a tipping point. For example, the US Department of Defense or governmental adoption could create a positive network externality. This can function as a tipping point, where all organizations connected to the Department of Defense, for example, move to more quickly adopt IPv6.

The third analysis is the most optimistic. It includes the best possible case that can reasonably be derived from current data. By selecting the last 3 months of data and fitting it to an exponential growth curve we found

the most optimistic but defensible adoption rate for IPv6, exogenous influences excluded.

This experiment originally focused on attempting to see which factors were important in determining why a firm was or was not adopting the IPv6 protocol. This was to be done by looking at the ARIN Project data set and filtering out firms in different sectors of the market. After this was performed, criteria would be set up to decide why a firm was adopting as compared to all other firms in that sector. Unfortunately after analyzing the data, we realized that there was not sufficient multi-sector adoption to prove any worthy prediction of the market factors. Domestic adopters consist almost entirely of network service providers and the US government. Probit analysis for the United States therefore seems premature.

We have therefore adopted a macro S-curve approach to IPv6 adoption. We based our analysis on publicized routes from ARIN. In calculating the IPv6 route growth, we can measure the gains over time and predict the rate of increase in the adoption of the IPv6 protocol. We have begun work on evaluating whether the domestic timeframe is concurrent with international adoption timeframes, or how far ahead of or behind global adoption the domestic market is.

In an attempt to gather points at which open IPv6 routes were at their average use, we chose three points each month. These points were every first and third Wednesday and every second Friday of the month. During the Wednesdays, a time period as close to 2 p.m. was chosen and during the Friday selection, a time period as close to 11 a.m. was chosen. By doing this, we have possibly inaccurate data for open routes that would

negatively affect this experiment. We based this selection in part on other communications technology—for example, avoiding 9am Monday as a possible peak.

5. Data Analysis

Predicting the future growth of the IPv6 adoption rate by the current data using best-fit results in the S-Curve shown in Figure 2 below. The data that goes along with Figure 2 shows that at the current rate of adoption, it will take approximately 18 years for a 20% implementation of the IPv6 protocol. As for an 80% implementation of IPv6, it will take nearly 22 years. Unfortunately with fewer and fewer IPv4 address remaining available each day, IPv6 will need to be adopted far before the above listed dates. This analysis does not take into account demand push (i.e. exhaustion) but only supply pull. Resource exhaustion is a significant component in IPv6 diffusion globally. According to ICANN², Mexico's DNS distributor will stop allocating addresses January 2011, and many other countries are following close behind in similar policies. Therefore, systems will need to be in place much more quickly than the current domestic adoption rate can manage.

5.1 Standard Model

Using past data we cannot perfectly predict future adoption. In our analysis and modeling we identified a crossover point at roughly 4% of routes, which occurs with a standard model at 11.1 years in the future. To what extent will resource exhaustion need to drive domestic IPv6 adoption before

² ICANN Factsheet "IPv6—The Internet's Vital Expansion", <http://www.icann.org/announcements/factsheet-ipv6-26oct07.pdf>, accessed Jan. 22, 2008.

the classic diffusion analysis predicts exponential growth? And what effect would this have on domestic US adoption?




Figure 2

5.2 Best Case

It is possible for exponential growth in the number of adopters to occur. Mapping an exponential trend line to the last three months of data available in January 2008 produces Figure 3. As shown in Figure 3, major adoption still does not start to occur until approximately 11 years from 2008. Similar to the predictive S-curve, the timeframe is simply too long for domestic adoption to be properly aligned with international adoption.

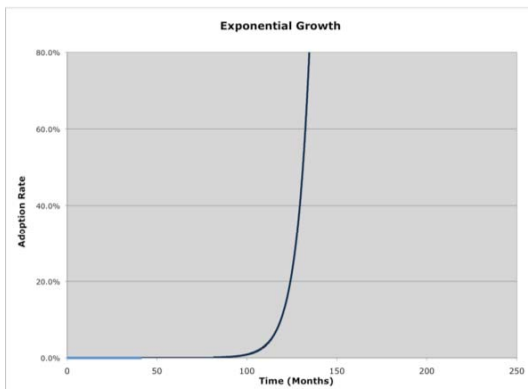


Figure 3

5.3 Forcing Function

In Figure 4, the curves that are represented are 8, 10 and 12 year predicted adoption timelines, respectively. What adoption timeframes are needed in the domestic market to align US and international adoption rates? What near-term goals are needed for a feasible chance for long-term success? An example of a possible exogenous event creating a tipping point is a Department of Defense adoption of IPv6. Variants of the success of this planned adoption can be seen in the different diffusion curves projected in this section. The possible adoption of all .mil and .gov sites could provide a projection point for best case and an optimistic case. These curves assume that full DoD adoption will result in greater overall adoption. The data point for full DoD adoption is roughly based on current IPv4 allocation (e.g. 3%). We implemented aggressive exponential curve fits, constrained by the theory and the last three data points. For an 8-year adoption plan the following goals are needed. In 2009, IPv6 must have .119% diffusion; in 2010, .595%; and 2.9% diffusion in 2011.

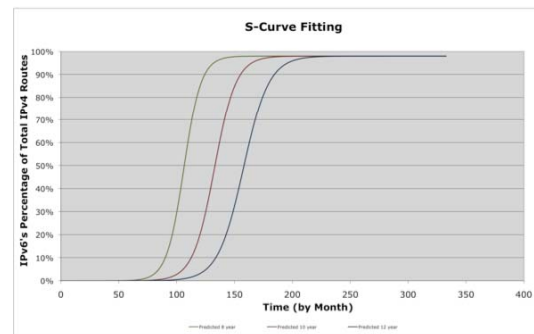


Figure 4

Even though the adoption timeline of 8 years is best case, this may be inadequate in terms of responding to IPv4 exhaustion, as the projected date of IANA's unallocated address pool exhaustion is November 28th, 2010, along with RIR's projection as being

December 5th, 2011 (Huston, 2007). At this point, two of the major distributors of IPv4 addresses will no longer be able to provide more addresses.

6. Economics

In order to understand why the creation and specification of IPv6 has yet to result in its widespread diffusion, we began by taking an economic approach to analyze possible sources of the current adoption rate. We focus on the misaligned incentive structures that are faced by IPv6 adopters. In doing this we believe that we can determine a partial cause to the slow adoption of the IPv6 protocol.

Adoption of new protocols in the technology world can be very costly, not only in monetary terms, but also in time spent understanding and deploying the technology (Rowe, 2006). In “Could IPv6 Improve Network Security?”, Rowe presents estimates for the costs of implementation and the benefits gained from the implementation of IPv6. Rowe begins by estimating the incremental cost of labor and training required for the IPv6 conversion at approximately 25 billion dollars. This amount would be spent over an estimated 25-year implementation period, which seems like a very large cost for the system, but it is negligible when compared to the hardware and software costs associated with the conversion, less than 1% estimated (Rowe, 2006). This large discrepancy in the cost of adoption versus benefits places a large burden on initial adopters of IPv6, one that most companies cannot bear in terms of maintaining comity with stockholders. Ironically, IPv6 may increase near-term security vulnerabilities because of the relative immaturity of the software (Ozment, 2006). In addition to the inherent problems of a younger code base, lack of employee

experience may increase misconfiguration. Misconfiguration is an extremely common event (Mahajan, 2002). The unintended interactions of software components are a major cause of real world vulnerabilities. Each component may be secure independently but create vulnerabilities through their interactions (Landwehr, 1994). All of these costs weigh heavily on the shoulders of early adopters. Unfortunately, as stated by Jaffe, “the initial benefits obtained by early adopters might fall significantly below the costs of adoption (2005).” This can be a large negative incentive for early adopters as they tend to incur most of the cost and tend to have to wait for long term gains for a return on their investment.

Market interventions to offset the costs incurred by early adopters promote adoption. In subsection 5.3 we modeled three adoption paths with the same type of forcing function. This can be done several ways. The first that we will discuss are subsidies, or a grant provided by the government. Subsidization of technology-based adoption is not uncommon in the global economy, and even though it is not currently being done domestically for IPv6, it is being done globally. Take, for instance, South Korea, China, Japan, and the European Union (Hovav, 2006). All of these countries are taking strides towards implementing IPv6 and are doing so by making adoption a key governmental policy. They are mainly supporting adoption via production subsidies associated to developing and implementing new Internet technologies, mainly IPv6. These subsidies provide a positive incentive for the companies to produce in that sector. Subsidies also reduce the cost of adoption by lowering long-term production costs. By reducing the cost to implement and develop, more adopters are able or willing to attempt

the IPv6 conversion. A similar effect can be produced with a tax credit, as it provides monetary compensation for those adopting IPv6.

Another form of cost displacement is a fine. Fines allow the adopter to “receive a conceptual benefit” by complying and not having to pay the fine (Ozment, WEIS 2006). Fines are commonly used as a negative incentive. This in turn causes them to not produce the full expected effect as people respond better to positive incentives than to negative incentives. Content-based adoption incentives have been implemented in a few markets. However, these have proven ineffective (Camp).

By providing incentives like those mentioned above, the United States may be able to induce more adopters to enter the market. This in turn will drive prices down as there will be a larger supply of products including software, hardware and support for IPv6. As the prices go down, the demand will continue to increase until it hits near market saturation, either by simple replacement of old hardware, or by actively adopting. This will of course take time, but with positive incentives to adopt, typically supported through strong governmental policy, the rate of adoption can cause the adoption timeline to become relatively short.

7. Implications And Conclusions

On a global scale, the benefits of IPv6 adoption far outweigh the costs for developing and late-adopting nations. Developing nations stand to see significant benefits from developing IPv6 infrastructure at this point. However, given the current expenditures on IPv4 in the United States and the investment cost necessary to switch from IPv4 to IPv6, this may not be the best

option for the U.S. and other developed countries with existing IPv4 infrastructure. According to the European Union (Force 2004)’s IPv6 Task Force Steering Committee’s 2004 report, “the actual level of IPv6 deployment [in Europe] is still imperceptible.”

Though IPv6 purports to address many of the security flaws of IPv4, the fact that many of these security upgrades have subsequently been applied over IPv4 limits the benefit of this aspect of IPv6 adoption. In addition, as Rowe suggests, the transition to IPv6 will inevitably result in unforeseeable new security vulnerabilities (2006). The marginal benefits of IPv6 over IPv4—especially since many of the security enhancements of IPv6 have been implemented over IPv4—and the high switching cost, it might be more beneficial for the U.S. to continue operating over IPv4 until a successor to IPv6 is developed. Though research toward a new communications technology is underway, no viable alternative to Internet Protocol (either version 4 or 6) has emerged.

This paper finds that in the US, IPv6 is not being adopted at a rate comparable to other countries, largely as a result of misaligned incentives and a lack of governmental policy support for the switch from IPv4 to IPv6. The current high switching costs and low perceived benefits of switching are discouraging the adoption of IPv6 among service providers and early adopters. Governmental support—in the form of subsidies, training, demand pull, information provision and policy changes along with possible tax cuts—could provide incentives for early adopters to switch to IPv6. However, as seen from the data analysis, even this many not adequately increase adoption rates and enable the US to implement IPv6 in a timely manner. As the

estimate for IPv4 exhaustion is uncertain, it is difficult to determine the exact time at which IPv4 will no longer support more addresses. The unanswered question is if there is a price point for IPv4 addresses that will drive IPv6 adoption. Alternatively, will there be another network addressing innovation that will serve to expand the life of IPv4 another decade (e.g. daughter of NAT)?

Should investment be made to force acceptance of IPv6, which is a tightly bundled product? Or should the U.S. and European governments immediately invest in IPv4 expansion alternatives? Without action on adoption or unbundled alternatives Europe and the US risks having chronically insecure, unnecessarily expensive, second generation IP networks.

IPv6 adoption may have to be encouraged. As Ozment and Schechter propose, some combination of government encouragement of IPv6 development in the form of subsidies, adoption mandates, and bundling of technologies could help increase adoption rates. Yet Camp suggested and time has

illustrated that information provision, subsidies, training and demand pull have proven inadequate for eliminating even those vulnerabilities based on logical errors. Arguably, because IPv6 technology is present in most routers sold today, simply making it available—as in the bundled technology strategy—is inadequate to truly hasten its adoption. Subsidies and mandates together might prove sufficient incentives for firms to switch to IPv6. Yet as history illustrates, no incentive is as great as the market compulsion of resource exhaustion.

As an alternative to IPv6 is it feasible to double the number of bits in IPv4 and require security adaptations for new IPv4 addresses? While there are strong technical arguments for IPv6 adoption, there are also unforeseeable potential complications. For example, highly mobile systems and the undertheorized nature of routing storms both present uncertainty. The market is adopting DNSSEC and to a lesser degree IPsec. Is IPv4 with security a pipedream or a possible solution to the coming Western address crisis?

Works Cited

Anderson, Ross, and Tyler Moore. "The Economics of Information Security." *Science*, October 27, 2007: 610-613.

Bass, Frank M. "A New Product Growth for Model Consumer Durables." *Management Science* 15, no. 15 (1969): 215-227.

Camp, L. Jean, and Catherine Wolfram. "Pricing Security." In *The Economics of Information Security*, edited by L. Jean Camp and Stephen Lewis, 17-35. Springer-Kluwer, 2004.

Cavusoglu, Huseyin, Hasan Cavusoglu, and Jun Zhang. "Economics of Security Patch Management." *The Fifth Workshop on the Economics of Information Security*. Cambridge: WEIS 2006, 2006.

Chow, Gregory C. "Technological Change and the Demand for Computers." 57, no. 5 (1967): 1117-1130.

Christensen, Clayton M. "Exploring the Limits of the Technology S-Curve. Part I: Component Technologies." *Production and Operations Management* 1, no. 4 (1992): 334-357.

EU IPv6 Task Force. "IPv6 Task Force Steering Committee Final Report." *IPv6 Task Force*. November 29, 2004. www.ipv6tf-sc.org/html/public/ipv6tf-sc_pu_d4v1_9.pdf (accessed February 29, 2008).

Geroski, P.A. "Models of Technology Diffusion." *Research Policy* 29 (2000): 603-625.

Hall, Bronwyn H. "Innovation and Diffusion." In *Handbook on Innovation*, by Jan Fagerberg, David C. Mowery and Richard R. Nelson. Oxford: Oxford University Press, 2004.

Hovav, Anat, and Jung Kim Yoo. *Determinants of IP Version 6 Adoption*. IEEE Computer Society, 2006.

Huston, Geoff. "IPv4 Address Report." <http://www.potaroo.net/tools/ipv4/> (accessed December 13, 2007).

Jaffe, Adam B., Richard G. Newell, and Robert N. Stavins. "A Tale of Two Market Failures: Technology and Environmental Policy." *Ecological Economics* 54 (2005): 164-174.

Landwehr, C. E., A. R. Bull, J. P. McDermott, and W. S. Choi. "A Taxonomy of Computer Program Security Flaws." *ACM Computing Surveys* 26, no. 3 (September 1994): 211-254.

Mahajan, R., D. Wetherall, and T. Anderson. "Understanding BGP Misconfiguration." *Proceedings of the 2002 Conference on Applications, Technologies, ARchitectures, and Protocols for Computer Communications*. Pittsburgh: ACM, 2002. 3-16.

Ozment, A., and Stuart E. Schechter. "Milk or Wine: Does Software Security Improve with Age." *15th USENIX Security Symposium*. USENIX, 2006.

Ozment, A., and Stuart E. Schechter. "Bootstrapping the Adoption of Internet Security Protocols." *The Fifth Workshop on the Economics of Information Security*. Cambridge: WEIS, 2006.

Rowe, Brent, and Michael Gallaher. "Could IPv6 Improve Network Security? And, If So, at What Cost?" *I/S: A Journal of Law and Policy* 2, no. 2 (2006): 231-266\.