

Economic and Technical Drivers of Technology

Choice: Browsers

Timothy F. Bresnahan and Pai-Ling Yin

First Draft: Nov 05, 2003

This Draft: August 12, 2005

Abstract

The diffusion of new technologies is their adoption by different economic agents at different times. A classical concern in the diffusion of technologies (Griliches 1957) is the importance of raw technical progress versus economic forces. We examine this classical issue in a modern market, web browsers. Using a new data source, we study the diffusion of new browser versions. In a second analysis, we study the determination of browser brand shares. Both analyses let us examine the importance of technical progress vs. economic forces. We find that the critical economic force was browser distribution with a complementary technology, personal computers (PCs). In both of our analyses, distribution had a larger effect on the rate and direction of technical change than technical browser improvements. This shows an important feedback. Widespread use of the Internet spurred rapid expansion of the PC market in the late 1990s. Our results show that the rapid expansion in PCs in turn served to increase the pace of diffusion of new browsers and thus move the economy toward new mass market commercial Internet use.

Timothy F. Bresnahan
Landau Professor in Technology and the Economy
Stanford University
Department of Economics
579 Serra St.
Stanford, CA, 94305 USA
tbres@stanford.edu

Pai-Ling Yin
Harvard Business School
Morgan Hall 241
Soldiers Field
Boston MA 02163
pyin@hbs.edu

1) Introduction

A new invention creates a technological opportunity. The diffusion of the new technology to the economic agents who will use it determines the rate and direction of realized technical change in the economy.

The classical model of diffusion of a new technology (Griliches 1957) emphasizes adopters' incentives. Valuable technical progress is an incentive to adopt; the greater the advance of the latest technology over earlier alternatives, the more rapidly it will tend to diffuse. Yet adopters of many new technologies wait rather than adopt immediately.¹ They are held back by the fixed costs of learning about new technologies, of installing them, or by the costs of adapting technologies to particular uses. Actions by sellers of a new technology, such as distribution and marketing, can reduce the influence of these fixed costs and speed up diffusion. Ultimately, the pace of diffusion, and thus the pace of realized technical change in the economy, is determined not only by the pace of technical progress but also by economic forces like adoption costs and the success of market distribution (Griliches 1957).

The direction of technical progress is often determined by the same forces. The relative rate of diffusion of competing technologies depends on adopters' incentives. All else equal, a superior technology or a technology that is more effectively marketed and distributed will be adopted more quickly, while a harder-to-install technology will be adopted more slowly. The pace of adoption of competing technologies is particularly important in industries with *de facto* standards, like computing. The competing technology that diffuses more rapidly will often set the standard.

In this paper, we study the diffusion of new and improved versions of World Wide Web (WWW) browsers from 1996 to 1999. We focus on commercial browsers from Microsoft and Netscape. The classic issues in the diffusion of technology appear in new forms.

The invention and commercialization of the browser triggered the widespread use of the Internet and created the opportunity for commercial applications such as Internet e-

¹ Hall (2003) provides a modern survey of the literature documenting that many technologies diffuse slowly and seeking to explain the wide variation in the rate of diffusion.

commerce, entertainment, and advertising-supported web pages. These commercial applications exploited improvements in each of several new browser versions released over the period we study. These applications could only reach a mass market if the improved browser versions were running on the personal computers (PCs) of many consumers. Thus, the pace of widespread diffusion of new browser versions determined the rate at which the Internet came to serve a mass market with valuable new commercial applications. For that reason, we study the diffusion of new and improved browser versions.

The interaction between Netscape and Microsoft led to another important market outcome. Widespread use of a browser of one brand would set a *de facto* standard for connection between the Internet and PCs. A Netscape standard seemed likely during the early stages, but the “browser war” ended with a *de facto* Microsoft standard in place.

Browser technical progress thus gives us two phenomena to study. We examine the pace of diffusion of new browser versions in order to study rate of technical progress in transformation of the Internet into a mass market commercial platform. We examine the reversal of the two main commercial browser brands’ market positions in order to study the setting of a *de facto* standard.

The diffusion of commercial improvements to the browser, while very important, has not yet been studied by economists. In particular, no explanation has been proposed for a phenomenon noted by market participants at the time: later versions of both brands of browser diffused more slowly than earlier versions.

In contrast, the browser market brand switch from Microsoft to Netscape has been widely studied², but the explanation remains controversial because of the antitrust trial. The controversy follows classical considerations in the diffusion of technology. One side argued that faster technical progress by Microsoft led to the reversal, while the other side emphasized the role played by distribution and marketing. Turned into a behavioral

² See Cusumano & Yoffie (1998), Shapiro & Varian (1999), and Windrum (2000) for analysis of the browser war events. See Manes & Andrews (1993) Wallace (1997), and Bank (2001) for Microsoft-centered histories. See Liebowitz and Margolis (1999), Fisher (1999) Schmalensee (1999a and 1999b), and Bresnahan (2002), as well as the interchange among Evans and Schmalensee (2000a, 2000b) and Fisher and Rubinfeld (2000a, 2000b) for analysis closely related to the antitrust case. Jenkins et al. (2004) studies the effect of browser brand tying and exclusion on brand choice, but by modeling dynamic investment decisions by Netscape and Microsoft rather than user decisions.

hypothesis about browser users, the two sides differed on whether users adopted the latest and greatest browser (technical progress explanation) or tended to use the browser that came with their computer (distribution convenience explanation).

In this paper, we create a new dataset and use it to study the determinants of both the pace of diffusion of new browser versions within brand and the shift in brand shares. We are able to measure the impact of technical progress and of distribution on both outcomes. This industry gives us the opportunity to study distribution in particular detail because of contracts which required or blocked the distribution of particular browsers. Distribution and technical progress vary separately because the contracts differentially affected the distribution of the same browser to users of different kinds of computers.

Our model of diffusion is entirely conventional, except for a particular feature of the PC industry in this era. The rapid expansion of demand for PCs meant that the browser was also diffusing into a rapidly expanding field of potential adopters instead of diffusing into a fixed field of potential adopters.

We reach the same conclusion, in separate analyses, about the pace of diffusion of new versions within brand and about brand choice. In each case, distribution played a larger role than did technical progress in determining the market outcomes. We quantify the effect of these forces in our analysis.

2) Browser Market Background

In this section, we review the market background with the goal of making clear what technical progress we study and why it is important.

a) Browser Invention

A web browser lets an individual computer user easily find, observe, and retrieve information on the Internet. With a browser, users can access information like news, entertainment, maps, etc. The browser is also the users' gateway to new mass market online applications, including electronic commerce, email, and so on.

It had been evident to market participants for many years that a technology to connect users to online content would be valuable. Many computer and telecommunications industry firms attempted to create mass markets in online

applications in the late 1980s and early 1990s, with little success. Demand for browsers grew quickly because the browser fulfilled this long-felt need.³

The Internet dates back to the 1970s, but a number of inventions in the early 1990s transformed it in a direction more suitable for mass market use. These included first the World Wide Web (WWW) and then the browser.⁴ The WWW is largely defined by a set of standard “protocols” for connecting computers together.⁵ One of the things a browser does is let a personal computer communicate through those protocols.

These inventions largely came out of academic science. Mosaic, the first browser, was invented at the University of Illinois. As a result, the key protocols for the Internet and the WWW were largely in the public domain and the technologies at the heart of the browser itself were not protected by any strong intellectual property rights.

In 1993 and 1994, an “Internet Mania” took off. Especially in universities, users demanded browsers in growing numbers, acting on their incentive get online. At this stage, these were simple freeware browsers like Mosaic. Users’ incentive to get online grew as more online information and applications appeared. There were more and more web sites – simple ones, at this stage -- for them to visit. “Webmasters” (those who built websites) also benefited from greater browser usage. More browser usage meant a larger audience for online information and more customers for online commerce. The “mania” was a positive feedback loop between the growing number of websites and the growing number of users.

Very few people had used the Internet until the advent of the browser as a connection technology. With the browser, many more people could use the Internet easily. The Internet Mania spread beyond universities creating a large-scale demand for connecting to the Internet.⁶ Internet service providers (ISPs) saw their market grow as a result. A retail ISP sector emerged and spread very rapidly (Downes & Greenstein 2002).

³ See Mowery & Simcoe (2002) for more on the history of the Internet’s development, and Wiseman (2000) and Vogelsang & Compaine (2000) for essays on the economic and political impact of a commercialized Internet.

⁴ Gottardi (2003) presents a diffusion model showing the impact of the browser on Internet use.

⁵ There is an alphabet soup of protocols and standards in the Internet and in the WWW. We will avoid discussion of these specifics and expansion of these acronyms, to the extent possible. There are also a set of semi-public standard-setting bodies for these protocols, like the W3C, to which we pay little attention, since the important standard-setting activity in the era we study was *de facto* and commercial.

⁶ Forman et al. (2005) discuss the Internet diffusion process.

Many existing online services, such as America OnLine (AOL), began to transform themselves, also offering ISP services (Swisher 1998). Others were entrepreneurial entrants.

Access to the Internet made the PC more valuable. Web browsing, email, and instant messaging were among the new and important applications that raised the value of PCs. The overall effect was an increased demand for PCs.⁷

These developments transformed two technologies. The PC had been used primarily for non-networked applications such as word processing and spreadsheets; it was transformed into a connected communications node. The Internet had largely been used by a small number of public or academic users; it was transformed into a ubiquitous commercial technology.

The role of interactions among complementary technologies as the center of these events is widely studied. Greenstein (2004), in a forthcoming book on Internet geography, describes the importance of complementary infrastructures in Internet diffusion. This useful literature has not yet addressed the end users' browser adoption decisions.

b) The Commercial Browser

Several of the inventors of the browser sought to commercialize it by founding Netscape Communications. In 1994, they released their first browser, which we abbreviate as NS1. Over the next several years, Netscape continued technical progress in the browser ranging from better display ("frames," and so on) to more secure interaction with servers (computers that hosted websites). This technical progress would make the browser and the WWW more commercially valuable by enabling more complex applications such as e-commerce and advertising-supported websites.

Netscape introduced five major versions of its browser between 1994 and the end of the "browser war" in 1999. Each included new technical progress of value to commercially oriented webmasters, and each provided new features valued by users. We

⁷ Goolsbee and Klenow (2002) examine computer demand in this era, finding that externalities across households, email, and Internet use are important drivers of demand.

list these major versions in Table 1.⁸ For ease of comparison, we rename Netscape version 4.5 as version 5.

Table 1 Major Netscape browser versions in our analysis

NS1	NS2	NS3	NS4	NS4.5 (“NS5”)
includes 1.x, 0.x	includes 2.x	includes 3.x	up to 4.4x	includes 4.6, 4.7

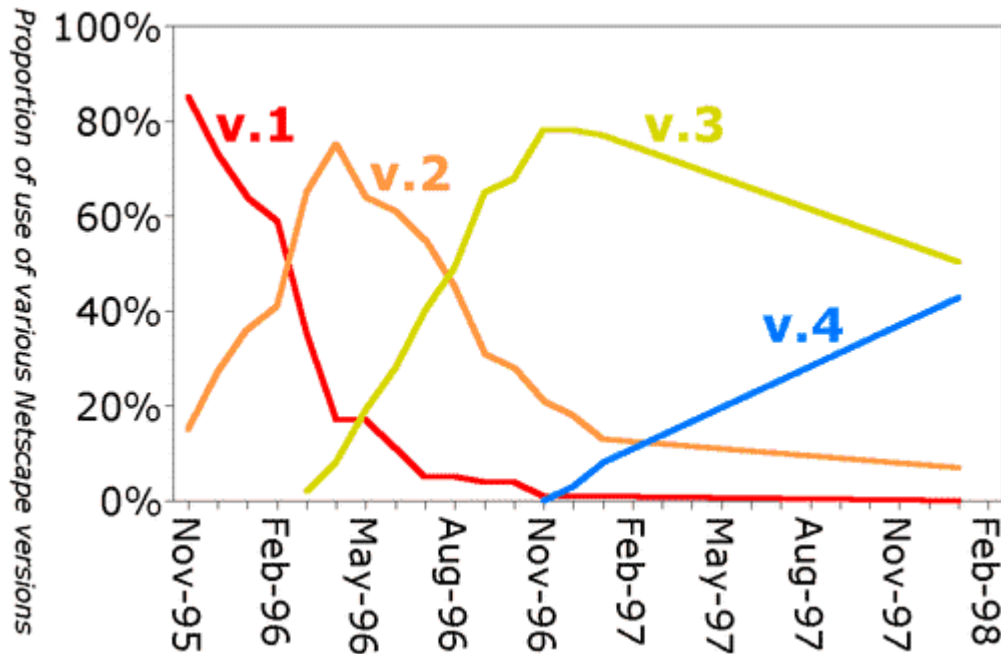
We study the diffusion of new browser versions to users. We do not study the parallel adoption of new and improved WWW technologies by webmasters. The distinction matters in thinking about technical progress in browsers. Some types of technical progress in browsers, such as “rendering” images more quickly, was directly useful to the user and thus provided an immediate incentive to adopt. However, many improvements in browsers were designed to permit webmasters to make more advanced web sites. This type of technical progress would only give users an incentive to adopt after websites took advantage of the improvements. Webmasters, in turn, could only get a wide audience for their more advanced websites if there was widespread usage of new and improved browsers. As a result, the rate of diffusion of the newest browser version depended on the time it took for webmasters to a) become convinced that the newest version would be adopted by a sufficient number of users and b) subsequently release advanced websites.

Webmasters cared about the rate of diffusion of newer browsers for a number of reasons, including the need to make websites work with both older and newer versions. Most importantly, the advanced websites that they designed often had commercial applications, and widespread diffusion was necessary to reach a mass market. Statistical tools appeared that would let a webmaster look at a particular website’s “server log” to see what versions of browsers its customers were using. A marketing research industry, with firms such as AdKnowledge and InterSe, emerged to measure the pace of diffusion of new browser versions and to report it to webmasters.

⁸ There are many minor browser releases, updates and so on. We aggregate numerous small versions into major browser versions in our empirical analysis in order to have each version represent a substantial advance over the earlier one. The minor versions we include in each major version are listed in the last row of the Table. For example, Netscape Navigator versions numbered less than 2 (1.45, 0.98, etc.), are aggregated to major version NS1 in our analysis.

Webmasters noted a slowing pace of diffusion of new browser versions.⁹ This presented them with a tradeoff between offering the best available web pages and gaining the highest volume of usage. In a website / magazine aimed at webmasters, Nielsen (1998) made a browser versions diffusion graph which we display as Figure 1. Nielsen used data from InterSe and AdKnowledge because those tracked the commercial websites of greatest importance to webmasters contemplating new features. In this figure, he shows the share of each major version of Netscape browsers out of total usage of all Netscape browsers. His conclusion was that Netscape users were adopting the newer versions more slowly than they adopted earlier versions.¹⁰

Figure 1 Nielsen's Netscape browser diffusion graph



To study the adoption of technologies which supported those mass market commercial applications, we focus on the diffusion of version 2 and later commercial browsers from Netscape and Microsoft. As our sample begins, some users are already using pre-commercial browsers (such as Mosaic) and version 1 browsers from Netscape

⁹ See Nielsen (1998). Also see, for example, the discussion thread “Browser usage stats” [sic] in the online discussion forum comp.infosystems.www.authoring.site-design of September 2nd to 4th, 1997.

¹⁰ Nielsen also finds slowing adoption for IE.

and Microsoft. We measure technical progress relative to the base of version 1 commercial browsers.¹¹

c) The “Browser War”

We now turn to a more controversial part of the history, the “browser war” between Netscape and Microsoft.¹²

Netscape introduced the browser as a modular component. That meant that the Netscape browser ran on all kinds of computers, including Macintoshes, UNIX machines, the latest version of Windows PCs, and earlier versions of Windows. (Introduction dates for Netscape browsers on a variety of computers can be seen in Table 3.) A user could use a Netscape browser regardless of the web content accessed or kind of computer she used. The open modular component strategy also meant the browser’s “protocols” for communicating with web sites were open and documented. Netscape encouraged the development of server software to communicate with the browser by other firms and also wrote such software itself.

Netscape’s goal was ubiquitous distribution of its browser in order to set a standard for web communications. It made distribution agreements with firms selling new computers, ISPs, online services, and others. Netscape also made its browser available for free download on its website and for sale in retail stores.

The second commercial browser firm of any importance was Microsoft, producer of Internet Explorer (IE). Through the period of the Internet Mania, Microsoft did not supply browsers. Once Microsoft realized the browser was a competitive threat to the Windows operating system (OS), it began a rapid program of imitating Netscape’s innovations. Released in August of 1995, the first version of Internet Explorer, IE1, was a poor imitation. After that, Microsoft worked to catch up with improved versions of its browser. The major versions of Microsoft browsers used in our analysis are listed in Table 2.

¹¹ Of course, nothing is ever an entirely new application, and one might follow Nordhaus (1997) by modeling the technical level of a broad “online services” category which could include the predecessors to the widely used Internet. Before the widespread use of the Internet, however, online services reached approximately 2.5% of their current total users and had a fundamentally different approach to technology that lacked the universality and openness of the modern Internet.

¹² See footnote 2 for studies of this period.

Table 2 Major Microsoft browser versions in our analysis

IE1	IE2	IE3	IE4	IE5
includes 1.x	includes 2.x	includes 3.x	includes 4.x	includes 5.x

Browser quality improved over time at both innovator Netscape and imitator Microsoft. It is clear that quality increase was more rapid at Microsoft, but whether Microsoft caught up in quality is debatable. Schmalensee (1999a and 1999b) and Liebowitz and Margolis (1999) use measures that show Microsoft catching up by version 3, while Fisher (1999) and Bresnahan (2002) argue that the catch-up occurred later, if at all.¹³

The impact of the differential rate of technical progress in the two brands of browsers is also controversial. Schmalensee (1999a and 1999b) and Liebowitz and Margolis (1999) argue that differential rates of technical progress explain browser brand market share changes over time. In contrast, Bresnahan (2002), Fisher (1999), and Fisher and Rubinfeld (2000a) argue that distribution played a substantial role in the browser brand shift.¹⁴ Examining these very different hypotheses is one of the empirical goals of this paper.

The marketing and distribution strategies of the two browser firms had similarities and differences. Both Netscape and Microsoft attempted to distribute their browsers widely and charged zero prices at the margin.¹⁵ Each hoped to have a ubiquitous browser in order to set standards for connecting personal computers to the Internet.¹⁶

¹³ Liebowitz and Margolis (1999) and Schmalensee (1999a and 1999b) used a relative browser quality index based on expert opinion. They proceeded by examining software reviews in personal computer magazines and counting the number of reviews which recommended Netscape, those which were tied or mixed, and those which recommended Microsoft. Liebowitz and Margolis used similar data for many other applications categories. Fisher (1999) based his analysis in market research conducted by Microsoft and in the opinions of Microsoft marketing managers, both from documents in the antitrust case. Bresnahan (2002) quotes extensively from the Microsoft marketing managers.

¹⁴ This and many other economic issues in the antitrust case were debated in a useful point/counterpoint format by Fisher and Rubinfeld (2000a and 2000b) and Schmalensee and Evans (2000a and 2000b).

¹⁵ This applies to marginal prices in our sample period. Earlier, Microsoft's plan had been to distribute Internet Explorer in a Windows "plus pack" and charge separately for it. Microsoft abandoned that plan after recognizing a competitive threat from the widespread use of the Internet and never charged separately for Internet Explorer. Netscape, at the beginning of our sample, used a plan of price discrimination in favor of individual users, who got prices of zero, while corporate users were charged a price. Netscape later abandoned this in favor of zero prices to all users.

¹⁶ For a review of the standards literature, see David & Greenstein (1990) and Stango (2004).

There were important differences between the two firms' strategies as well. Microsoft began by introducing its browser only for the latest version of Windows. IE1 was introduced at the same time as a new version of Windows, Windows 95, and only worked with that OS. At the time, the vast majority of PC users were running earlier versions of Windows, notably Windows 3.1, with a minority running Apple Macintosh. These existing PC users could not use IE1. Starting with IE2 in late 1995, Microsoft introduced new browser versions both for the newest version of Windows and for other kinds of computers (sometimes with a lag).

Because we are interested in widespread diffusion, we focus on four mass market OSs for PCs. Three are from the Windows family (Windows 3.1¹⁷, Windows 95, and Windows 98), and the fourth is Apple Macintosh. These OSs are the ones most likely to be used by individual end users.

In Table 3, we report the introduction date for each of the major browser versions on each OS, gleaned from the suppliers' websites. There are some gaps in the table because of supplier choice (IE1 was available only for Windows 95) or because the browser was obsolete at the time of the OS's introduction (the 1994 version of Netscape's browser, NS1, was not relevant to the 1998 Windows 98). The table clearly shows Netscape's strategy of writing for all existing OSs and Microsoft's changing strategy.

Another aspect of Microsoft's distribution and marketing strategy was to compel widespread distribution of IE browsers and to prevent widespread distribution of Netscape browsers. Microsoft imposed contractual restrictions on sellers of computers (which required them to distribute IE with new computers) and used threats to stop computer sellers from distributing Netscape with new computers. Microsoft signed contracts with ISPs and online services (such as AOL) which required them to distribute IE to subscribers and strictly limited distribution of Netscape. These contractual agreements grew more restrictive and more widespread over the course of our sample period.¹⁸ For example, the default browser distributed with Macintosh computers was

¹⁷ In our data, Windows 3.1 includes older versions recorded as "16 bit Windows."

¹⁸ For a general overview of this topic with many links to original antitrust case documents, see Bresnahan (2002) or US Department of Justice et al. (1999) Section V. B.

Netscape until August of 1997, when a contract between Microsoft and Apple came into effect, making IE the default browser on Macs.

Table 3 Basic browser timing facts:

Introduction date for each browser on each operating system

	Windows 3.1	Windows 95	Windows 98	Macintosh
NS1	Dec. 1994	Aug. 1995		Dec. 1994
NS2	March 1996	March 1996		March 1996
NS3	Aug. 1996	Aug. 1996		Aug. 1996
NS4	June 1997	June 1997	Aug. 1998	Sep. 1997
NS5		Oct. 1998	Oct. 1998	Oct. 1998
IE1		Aug. 1995		
IE2	April 1996	Dec. 1995		April 1996
IE3	Dec. 1996	Aug. 1996		Jan. 1997
IE4	Feb. 1998	Oct. 1997	Aug. 1998	Jan. 1998
IE5	March 1999	March 1999	March 1999	2003

The efficacy and purpose of the distribution restrictions was hotly disputed in the antitrust case. Davis & Murphy (2000) argued that the bundling of IE with Windows was itself user-friendly technical progress rather than a restriction on distribution, Microsoft's economic expert, Richard Schmalensee, argued that other means of distribution, such as downloading a browser over the Internet or buying it in a retail store, rendered the restrictions on computer manufacturers and on ISPs and online services irrelevant.

In contrast, the government emphasized the effect of distribution and restrictions on distribution.¹⁹ Adopting the analysis of distribution advanced by Fisher (1999), it

¹⁹ Two forces led to the government winning this part of the antitrust case. One was the large volume of Microsoft documents and testimony concerning distribution and restrictions on it. A Microsoft browser marketing executive said in court that the purpose of restrictions was this: "we did specifically ask that ISPs distribute Internet Explorer by itself when they distributed Internet Explorer, so that we would not lose all of those side-by-side user choices." Cameron Myhrvold (1999), 2/10/99, at 62:7-25. The other reason was the Court's rejection of Schmalensee's statistical analysis.

rejected Microsoft's defense contention that distribution was entirely unrelated to the success of IE.

Much of the government's argument was drawn from browser marketing documents inside Microsoft. The marketing managers selling browsers focused on two elements of the fixed costs as particularly important. One was simple distribution convenience. Users might conveniently get a browser with a new computer or when they signed up for Internet service. Less conveniently, they might download a new browser from the web and install it on their own computer. The second fixed cost identified by marketing managers was the complexity of installing a new piece of software on a computer.

Computer users are heterogeneous in the degree to which adopting new technologies involves large fixed costs. This heterogeneity in tastes is reflected in heterogeneous inertial behavior. Users vary in their valuation of convenience. They may have a faster or slower web connection.²⁰ Users can vary in their ability to install new software on their computer or in their confidence in their ability. For this reason, less-sophisticated users tend to avoid new software if it requires installation. Less sophisticated users also tend to be uninformed about new products and may have a fixed cost simply of learning of the existence of a new browser. The implication of all these heterogeneous fixed costs is that some users might delay adoption of a new browser version.

How many users might delay adoption is, of course, an empirical question. At one extreme, there are users who occasionally get new computers and use the software that came with their computer until they get the next one. At the other extreme are people who always download the latest version of software. In between are people who download the latest version or get it in a store when they come across a task they cannot complete without it.

The prevalence of these distinct behaviors depends not only on the size of the fixed costs but also on their distribution in the population. Distribution of fixed costs of adoption in the population is the determinant of whether inertia in diffusion is a

²⁰ Over the period we study, modem speeds grew faster while browsers also grew larger, so the time costs of a download remained roughly constant for the average user.

quantitatively important force. It is also the determinant of the role of mass distribution and marketing. Since these costs are distributed in a population containing many new users, introspection about the costs of downloading is an unpromising method. We rely instead on data.

We are able to examine the quantitative importance of mass distribution in two dimensions. Our largest contribution comes in examining the role of distribution vs. raw technical change in the diffusion of new browser versions.²¹ We also revisit the distribution vs. technical change controversy in our empirical work on browser brand choice. Our contribution there consists of using better data and of examining measures of both technical progress and distribution.

3) Model

The goal of our empirical work is to examine the roles of distribution and technical progress in both the diffusion of new versions of each brand of browser and in the brand shift.

Several different technical and economic forces can accelerate the diffusion process. A new technology may be a substantial improvement over the one it replaces. Agents may all learn about the technology at about the same time. Network effects may lead to bunching of adoption times. Agents may experience similar costs and benefits of the new technology, so they reach the decision to switch to it at around the same time. Switching costs and inertia, if they are present, may also be broadly similar across agents. These forces favor a rapid adoption process once diffusion has begun.

The diffusion process can be slowed by a number of forces as well. The costs of adopting a technology may be high, or the costs sunk to an older technology may be high, for some users. Information about the new technology may be poor at first. Agents may wait for better versions of a technology and will have an incentive to wait if there are fixed costs of adopting a new version. For valuable new technologies, economic institutions are likely to arise to partially redress these retarding forces. Reliable

²¹ Our closest predecessor is Goldfarb (2004), who looks at the role of universities on Internet adoption.

information sources and effective distribution channels, for example, may overcome the inertia that these retarding forces build around older technologies.

Our model of diffusion of new versions is closely related to the technology adoption and diffusion model of Griliches (1957). Like him, we emphasize the economic return to a technology adopter. Like him, we use aggregate data to study the diffusion of browsers into a similarly situated population. Like him, we study a diffusion process that takes place over time and use cross section differences in the environment of groups of adopters as exogenous predictors in the rate of diffusion. Indeed, much of our economic interpretation follows his, notably in seeing the adoption decision as limited by the frictions of distribution (possibly overcome by marketing and distribution) and advanced by the attractiveness of new technologies. Where we differ primarily is in context, in examining the direction as well as the rate of technical progress, and in our emphasis on the changing size of the population that might adopt the technology (the expanding field of adopters).

The browser diffused into a rapidly expanding field of potential adopters. The widespread use of the Internet caused by the browser meant that a large number of people were buying new PCs and opening new ISP subscriptions. We examine the possibility that this expansion of the diffusion field affected the rate and direction of diffusion of browsers.

a) Browser Version Analysis (within Brand): Pace of Diffusion

In our diffusion analysis, like any other, the dependent variable measures the tendency to choose the new technology. Conceptually, in our study, that means the tendency to choose the newest version of a brand of browser available to a particular customer. We condition on the kind of computer (OS) used and the brand of browser chosen; i.e. we focus on the choice of the newest browser version within the set of available versions of a particular brand of browser.

The start date of the diffusion process for each version of each brand of browser on each OS is the availability date given in Table 3. Call the newest version of each brand of browser at a given date, b^* . The dependent variable is a function of the share of all users of a particular brand of browser on operating system o at date t who are using b^* , $S_{b^*o|brand}$. Our model is

$$\ln\left(\frac{S_{b^*ot|brand}}{1-S_{b^*ot|brand}}\right) = X_{b^*ot}\beta + \varepsilon, \quad \text{Equation 1}$$

where X_{b^*ot} denote regressors associated with the newest browser of a particular brand on operating system o at time t , and ε is an iid error term.

The relevant concept of technical progress in this within-brand model is the improvement of each new version, from a user perspective, over the previous versions. We will measure technical progress by the rate at which the new version diffuses. Accordingly, the X s always include $INTR_{b^*ot}$, the number of months since the newest browser was introduced on the operating system. This allows the coefficient on $INTR_{b^*ot}$ to capture one potential cause of the slowing rate of diffusion seen in Figure 1: decreases over time in the rate of improvement of new browser versions.

Of course, the coefficient on $INTR_{b^*ot}$ measures all the forces that tend to make a technology diffuse that are not captured in other time-varying X s. For example, we cannot control for website improvements which led users to want the latest browser. We nonetheless interpret a browser that diffuses faster (has a higher coefficient on $INTR_{b^*ot}$) as one that has advanced technically relative to pre-existing browsers. This interpretation almost certainly overestimates the importance of technical progress as a cause of browser diffusion, particularly for later browser versions given that improvements in websites is an omitted variable.

The other X s in our model are measures of distribution. The relevant concept of distribution is whether the browser came with the user's computer (or with their ISP account). A user who got a new version of a browser with their computer will not need to bear any costs of learning that it exists, downloading it, or installing it. In the data section below, we define regressors which measure the probability that a particular browser was distributed with a users' computer or by their ISP.

Our first goal is to estimate the distinct effects of these technical progress and distribution measures. Our second analytical goal stems from the systems nature of the PC and Internet industries and from the importance of some large-scale feedback loops. The upturn in PC demand resulting from the invention of the browser occurred rapidly enough potentially to impact the diffusion of improvements to the browser. So, too, did the rapid rise in ISP subscriptions. We will push into novel analytical territory

empirically by measuring the impact of the rapid increase in demand for these complementary technologies on the diffusion of browsers.

b) Browser Brand Analysis: Direction of Diffusion

There is another aspect to users' browser choice other than version: brand. In our brand models, we estimate a linear probability model. The dependent variable, denoted $S_{IE,ot}$, is IE's share of commercial browsers in a particular month t on a particular OS o . The numerator is the number of users of all versions of IE. The denominator is users of all versions of IE plus all versions of Netscape. Our interpretation of this linear probability model is entirely standard. In the brand analysis, the relevant concept of technical progress is relative, i.e., how much faster is IE improving than Netscape? Similarly, the relevant concept of distribution is relative, i.e., how much better distribution is IE gaining on a particular OS than is Netscape? Except that they are relative, our measures of technical progress and distribution in the brand analysis follow our measures in the version diffusion analysis.

The zero marginal prices for browsers mean that it is not possible to address browser quality by hedonic pricing methods (Griliches 1988). As a result, we will use Schmalensee's (1999a) relative browser quality measure in part of our empirical work. The concerns of the hedonic literature about valuation based on quality assessments by experts limit its usefulness,²² so we will also use a much less restricted set of quality measures.

We do not attempt to control for network effects in browser brand choice, since our data are aggregate. We instead employ a reduced form model for the aggregate share of a brand. This allows us to estimate the equilibrium effects of changes in technical progress and distribution (some of which may flow through network effects) but not to separately estimate users' valuation of network effects vs. technical progress or distribution.

²² In the case of PC industry magazine reviews, there is a particular interpretational problem related to the distinction between a reviewer's assessment of best technology and his recommendation. In markets with network effects, and especially in PC software markets, it is commonplace to recommend that users choose the product which is going to be most popular even if it is not the best. See, e.g., "Nine Timeless Tips for Tech Buyers," Kevin McKean, *PC World*, June, 2002, p. 19. ">>**Tip Lean toward what's popular**: What a shame to have to offer this advice. But the best technology doesn't always win in the market." [emphasis in original].

Conceptually, the two analyses described in the last two subsections attempt to measure the relative importance of the technical progress and distribution theories explaining the pace and direction of diffusion and the events of the marketplace in the late 1990s.

4) Data Set

We use aggregate data for browser usage based on browsing at a web site at the University of Illinois Urbana-Champaign (UIUC).²³ In this section, we describe the data source (since it is novel) and our sample, dependent variable, and regressors.

The UIUC computer center keeps monthly logs made by its World Wide Web servers. We chose UIUC because it has maintained those logs consistently since early in the history of widespread use of the Internet. Our sample begins with the oldest available data, April 1996, and ends in December 1999. We end at that date because it is the first year-end that is (a) clearly after the end of the browser war and (b) well into the diffusion of commercially capable browsers.

Users browsing the UIUC Web site are not randomly drawn from the population of all Web users. For the most part, they are people interested in browsing engineering students' web pages. The advantage is that we study technology use over a period of time by a growing body of similar users. Users browsing at UIUC early and late in our sample are reasonably like one another: they are technically proficient. The disadvantage is that the users we sample are more likely to prefer new technology than the general population. They are likely to have lower costs of getting software onto their computers and downloading it. Thus, we are likely to overestimate the impact of technical improvements on choice of browser and to underestimate the impact of distribution relative to the general population.

Our sample dates vary across the four OSs. For each OS, we restrict the sample to the time period in which it is economically important, which we define as the time from the official introduction of the OS to the introduction of its replacement's

²³ We are grateful to Ed Kubaitis for archiving the logs and for technical advice.

replacement.²⁴ For Macintosh, user agent codes do not permit us to distinguish OS versions, so the entire period is included. The sample dates are shown in Table 4.

Those dates exclude users who run an OS before its official release date (beta testers and the like.) They also exclude users who got their computer long before browsing. Neither group seems likely to be representative of mass market users' tradeoff between distribution convenience and technical excellence.

Table 4 Observed operating systems and periods of use

OS	Macintosh	Windows 3.1	Windows 95	Windows 98
Sample Dates	4/96—12/99	4/96—7/98	4/96—12/99	8/98—12/99
No. of Months	45	28	45	17

Our data processing procedures for using the web server logs are documented in Appendix A. Each time a user accesses a web page, the web server's log records information about the browser and OS the user is running in the "user agent" field. That, plus other information recorded by the web server, forms the basis for our dataset.

An observation in our dataset is a browser / OS / month triple. We aggregate the UIUC data to the level of users running the same OS and then calculate the usage shares of each major browser version of each brand. The sample sizes to estimate these browser usage shares are substantial.²⁵ We report the descriptive statistics on browser shares used to generate our dependent variable in tables and graphs.

The first graphs we look at are the browser version diffusion curves within each brand. We examine Windows 95, since all the versions of all the browsers were available on that OS. In Figure 2 we show S_{bot} (the within-brand shares of a browser version b on an operating system o at time t) on Windows 95 for IE browser versions. Our sample period begins just as IE3 is beginning to diffuse on Windows 95, and just before the share of IE2 peaks. The within-brand diffusion paths that are visible (those of IE3, 4, and 5) are flattening over time.

²⁴ This is consistent with industry practice. Microsoft has a standing policy of only supporting the newest OS and its immediate predecessor, for example.

²⁵ The number of unique users rises from 35,757 in the first month of our data to 229,579 in the last month.

Figure 2 Market shares of major versions of Internet Explorer on Windows 95





Figure 3 Market shares of major versions of Netscape on Windows 95



We graph similar data for Netscape browser within-brand shares on Windows 95 in Figure 3. Netscape 2 was the newest version when Windows 95 was released, so it has the high (nearly 1) share at the beginning of our sample period. Once again, the diffusion paths of the three later browsers (versions 3, 4, and 5) tend to flatten. The similarity between Figure 3, which is based on our UIUC data, and Nielsen's (1998) Figure 1, based on commercial websites, suggests we can compare our UIUC data set to the broader commercial market for browsers.

A more familiar fact about browsers in this era is that brand leadership shifted from Netscape to Microsoft. Market shares of IE versus Netscape across all OSs in our data show this shift (Figure 4). Our figure, based on the UIUC data, is similar to those made from more commercial web browsing sources, such as Henderson (2000). This suggests that our study of technically-oriented web users may be representative of the broader population.

In his antitrust trial testimony, Richard Schmalensee (1999a) raised two objections to "hits" data like the data we use. He argued theoretically that positive feedback for a browser standard should be based on the number of users rather than the volume of usage. Since browser suppliers looked closely at hits data based on the observation that browser usage, not the number of users, is what matters to webmasters, the trial court rejected this theory, as do we. Schmalensee also raised the practical measurement problem that hits can be very hard to measure because of the practice of "caching" frequently viewed websites. It is unlikely that this is relevant to the UIUC websites, which do not host highly popular commercial websites but rather primarily the web pages of engineering students.

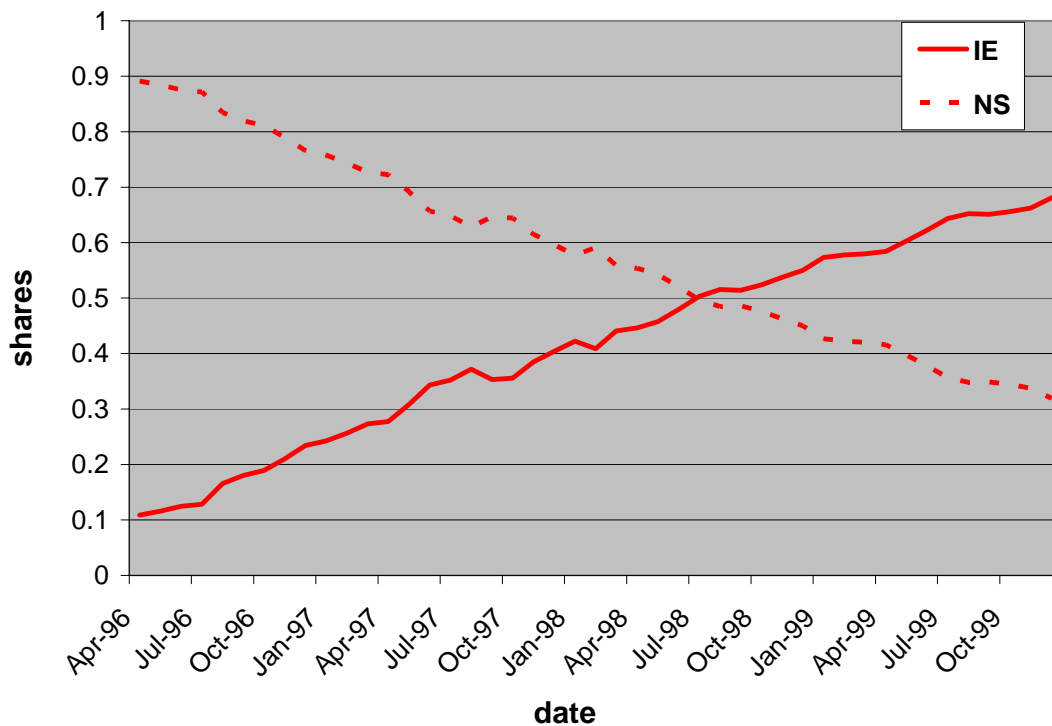
a) Regressors

Our distribution measures are based on the idea that convenience-oriented consumers will tend to use the software which came with their computer for a period of time. At a broad level, this amounts to looking at the shipments of new PCs as an explanation of browser utilization. By this means, we can distinguish the role of distribution convenience as distinct from adoption based only on technological progress in browsers. Consumers reacting to technical progress would respond to release of a

browser version (with some time lag), whereas consumers reacting to distribution convenience will tend to respond to variation in the shipments browsers with PCs.

Conceptually, our distribution measures should be based on the browser which came with users' computers (or ISP accounts). We do not have the ideal data: direct measures of which browsers were distributed with the PCs of our users. We can, however, construct measures of the probability that any particular browser came with users' PCs by conditioning on the date at which we observe the browser being used and on the PCs' OS.

Figure 4 Market shares of Internet Explorer and Netscape on all OSs



We proceed in two steps. The first step defines measures under the (false) assumption that the newest versions of both brands of browser were shipped with all new PCs. The second step takes account of the history of restrictions on distribution.

In the first step, we measure the probability that a user of operating system o at time t got browser b with their computer as the probability that a computer running operating system o at time t was sold when browser b was the newest of its brand. This

calculation does not use the UIUC data; it is based on the dates in Table 3 and on IDC and Microsoft data on the shipments of computers.²⁶

Let N_{om} be new PC shipments in month m running operating system o . If the depreciation rate for computers is δ , the stock of computers running operating system o at time t is $\sum_{m=0}^t N_{om} (1-\delta)^{t-m}$. Denote the dates (in Table 3) when browser b was newest of its brand on operating system o as the interval from m_1 to m_2 . The probability that a computer running operating system o observed at time t was shipped when browser b was the newest of its brand is²⁷

$$PCDISTR_{bot} \equiv \frac{\sum_{m=m_1}^{\max\{t, m_2\}} N_{om} (1-\delta)^{t-m}}{\sum_{m=0}^t N_{om} (1-\delta)^{t-m}}. \quad \text{Equation 2}$$

Here, the denominator is the stock of PCs running operating system o in use at time t . The numerator is the stock of PCs running operating system o in use at time t that were shipped when browser b was the newest of its brand.

For a given browser, $PCDISTR_{bot}$ varies both over time and across operating systems. A simple example will illustrate the definition and how it varies. The example, in Table 5, shows $PCDISTR_{bot}$ for IE version 1 and version 2 in late 1995 on Windows 95. As of November, IE1 was the only IE version which had been available for that OS, so it has $PCDISTR_{bot}$ of 1. In December, IE2 was introduced and, by the end of that month, 28% of Windows 95 computers ever distributed up to that time had been distributed in that month. Accordingly, $PCDISTR_{bot}$ for IE2 is 0.28 and $PCDISTR_{bot}$ falls for IE1.

²⁶ IDC, a leading IT market research firm, does not separately report monthly shipments by OS. It reports monthly shipments of all PCs and annual totals by OS (IDC 2000a, 2000c, 2000d, 2000e, 2000h (additional data from 1996-1998 were used)). Fortunately, Microsoft internal documents detail the rate at which new versions of its OS replace old ones in the market place. For example, the Microsoft “OEM Sales FY ’96 Midyear Review” gives the early history of Windows 95 vs. Windows 3.1 sales (Kempin 1998). This forms the basis for our allocation. We follow IDC by assuming 25% annual depreciation; lacking the retirements data they keep internally, we use a constant proportional depreciation assumption.

²⁷ for $t \geq m_1$; $PCDISTR_{bot} = \text{N.A.}$ otherwise.

Table 5 PCDISTR example

	$PCDISTR_{IE1 \text{ Win95 } t}$	$PCDISTR_{IE2 \text{ Win95 } t}$
November 1995	1	N.A.
December 1995	0.72	0.28

The time series variation in $PCDISTR_{bot}$ alone would not convincingly identify distribution convenience separately from “the shape of the S-curve.”²⁸ Across operating systems, $PCDISTR_{bot}$ varies in a more promising way. A user running an older operating system is likely to have bought their PC earlier. Therefore, a newer browser version will have lower $PCDISTR_{bot}$ on an older operating system than on a newer one. Consider $PCDISTR_{IE4 \text{ W95 } t}$ vs. $PCDISTR_{IE4 \text{ W98 } t}$ in the months shortly after Windows 98 replaced Windows 95. At that point, IE4 had been the newest version of IE for several months. Among Windows 95 users, only those who recently bought a computer would have had IE4 come with it, but all users who had a Windows 98 computer would have had IE4 distributed with it. Thus, for months in autumn 1998, $PCDISTR_{IE4 \text{ W98 } t} = 1$, but $PCDISTR_{IE4 \text{ W95 } t}$ is near 0.43. In diffusion specifications which include INTR, the variation in $PCDISTR_{bot}$ conditional on the other regressors is largely cross-section variation.

In the second step, we deal with restrictions on distribution. Not all users who bought a new PC received both brands of browser. Instead, contracts between Microsoft and PC manufacturers sometimes required distribution and display of the latest version of IE. Also, there was pressure from Microsoft on PC manufacturers not to distribute or display Netscape browsers.²⁹

A short history of the restrictions is in Appendix B – Restrictions on Distribution. Their dates are displayed in Table 11. That table shows the date at which “Must Carry” restrictions were in place requiring PC manufacturers to distribute and display IE. The

²⁸ Starting from the introduction date, $PCDISTR_{bot}$ first increases with t . After the browser version that succeeds “b” is introduced, $PCDISTR_{bot}$ begins to decrease. If the shipments of PCs were constant over time (which they are not) this time series variation, together with time since introduction plus the logistic transform of the dependent variable, would determine the “shape of the S-curve.”

²⁹ The contracts and the informal pressure were documented in the antitrust case, both in Microsoft internal documents and in its communications with outsiders such as OEMs. We base our regressors on these documents.

variable MC_{om} is the fraction of PC manufacturers shipping operating system o in month m who were contractually required to carry IE. Table 11 also shows the dates of “Must Not Carry” provisions under which Microsoft pressured PC manufacturers not to distribute Netscape Navigator. The variable EX_{om} is the fraction of PC manufacturers shipping operating system o at time t who were not distributing Netscape’s browser.

Combining the dates in Table 11 with information on the timing of PC shipments, we generate even sharper measures of the distribution of browsers with new computers. Descriptive statistics in Table 12 show the variation in these variables across OSs and over time.

The first variable measures the positive distribution advantage for Microsoft browsers from “Must Carry” provisions. We define $PCCARRY_{bot}$ as:

$$PCCARRY_{bot} \equiv \frac{\sum_{m=m_1}^{\max\{t, m_2\}} N_{om} MC_{om} (1-\delta)^{t-m}}{\sum_{m=0}^t N_{om} (1-\delta)^{t-m}} \quad \text{Equation 3}$$

$PCCARRY_{bot}$ is the probability that a computer running operating system o at time t (1) was shipped when browser b was the newest version of IE and (2) was shipped by a manufacturer subject to the “Must Carry” restrictions. It has the same definition as $PCDISTR_{bot}$ for IE except that we multiply the numerator by the “Must Carry” variable.

Since Microsoft always had “Must Carry” provisions in place for Windows 95 and Windows 98, $PCCARRY_{bot}$ is the same variable as $PCDISTR_{bot}$ for IE on those operating systems. These two variables differ on Windows 3.1, where $PCDISTR_{bot}$ is positive if small for some Internet Explorer browsers (which were new when late copies of Windows 3.1 were shipped) but $PCCARRY_{bot}$ is zero. They also differ on Macintosh, where restrictions were imposed only in the latter half of our sample.

In parallel, we define a variable, $PCEXCLU_{bot}$ for Netscape browsers based on EX_{om} . $PCEXCLU_{bot}$ is the probability that a computer running operating system o at time t was shipped (1) when browser b was the newest version of Netscape and (2) by a manufacturer that was pressured out of distributing Netscape browsers.

The cross-OS variation in $PCEXCLU_{bot}$ is broadly similar to that in $PCCARRY_{bot}$. Like $PCCARRY_{bot}$, $PCEXCLU_{bot}$ is distinct from $PCDISTR_{bot}$ for Netscape on the Macintosh OS, since the contract requiring Apple to make IE the default browser on the Macintosh took effect during our sample period. Like $PCCARRY_{bot}$, $PCEXCLU_{bot}$ is very different on Windows 3.1 than on other versions of Windows; it is always zero. The main difference between $PCCARRY_{bot}$ and $PCEXCLU_{bot}$ is that the “Must Not Carry” restrictions spread out over PC manufacturers over time, so that on Windows 95 and Windows 98, $PCEXCLU_{bot}$ is distinct from $PCDISTR_{bot}$. That said, most of the variation in $PCCARRY_{bot}$ and $PCEXCLU_{bot}$ in our data (conditional on other regressors) is the cross-OS variation.

There is controversy as to whether the distribution restrictions affected browser usage. The coefficients on $PCCARRY_{bot}$ and $PCEXCLU_{bot}$ permit us to test the hypothesis that the restrictions had no effect.

A second set of distribution variables are related to ISPs. ISPs, like PC manufacturers, saw a dramatic increase in their business in response to the widespread use of the Internet. ISPs distribute browsers and other network software to their customers. Just as some users may tend to use the software that came with their computers, they may also tend to use the network oriented software that came with their ISP subscription. The quantitative importance of this behavior will, of course, depend on the fraction of users who behave that way.

Accordingly, we define a distribution variable for ISPs, using ISP subscription data from IDC industry reports (IDC 2000b, 2000f, 2000g (additional data from 1996-1998 were used)). This variable, which we call $ISPDISTR_{bot}$ (distribution via ISPs) parallels $PCDISTR_{bot}$. It is defined by first calculating the stock of ISP subscribers and the flow of new subscribers. Then, for each browser on each OS at each time, we calculate the fraction of ISP users who were new subscribers when the browser was the newest of its brand available for that OS.

There is substantially less meaningful variation in $ISPDISTR_{bot}$ than in $PCDISTR_{bot}$, unfortunately. Both vary over time (when more computers are sold or more people get on the web). However, while $PCDISTR_{bot}$ also varies in cross section across operating systems, $ISPDISTR_{bot}$ varies only trivially across operating systems.

We also define a variable for Internet Explorer, $ISPTIED_{bot}$, based on the contracts between Microsoft and ISPs. These contracts required ISPs to distribute and display IE and not Netscape (for more details, see Appendix B – Restrictions on Distribution). The relationship between the $ISPTIED_{bot}$ variable and $ISPDISTR_{bot}$ is analogous to the relationship between $PCCARRY_{bot}$ and $PCDISTR_{bot}$: terms in the numerator are multiplied by the restriction variable.

A more detailed quantitative analysis of the ISP restrictions can be found in Fisher (1999). He compares ISPs bound by the restrictions to those not so bound and finds that the restrictions impact browser usage shares. Data limitations leave us unable to replicate his methods, because we cannot identify the (few) users who are using ISPs not bound by the restrictions. Another body of analyses of the ISP restrictions was carried out by Microsoft browser marketing personnel, who also concluded that the restrictions were effective.³⁰

Another adjustment to the timing of diffusion is idiosyncratic to software. Often, software is “prereleased” in test versions before it is officially released to the market. We see these prerelease versions in our data. While the rate of prerelease usage of IE browsers is low, prerelease usage of Netscape browsers can be as high as 12 percent of all browser usage on an OS. If the user is running a prerelease version of an OS as well as a browser, the observation is not in our sample (as explained above). Otherwise, we add a regressor, $PREREL_{bot}$, to capture the effect of prerelease browsers, and change the definition of $INTR_{bot}$ to start from the prerelease date. The prerelease date is the date of the most significant “beta” test version of the browser. These changes leave our results largely the same when we include the $PREREL_{bot}$ dummy (which gets a substantial negative coefficient, as expected). As a result, we do not show these analyses in tables.

5) Estimates

We discuss our estimates in three steps. We first examine predictors of the rate of diffusion of new browser versions of each brand. We then examine the direction of technical change by looking at the determinants of browser brand share. In the

³⁰ See U.S. Department of Justice et al. (1999) in the section entitled “Microsoft’s internal analyses evidence the impact of its restrictions” for quotations from these analyses.

subsequent section, we then use both those sets of estimates to analyze the quantitative importance of raw technological progress versus distribution and of the expanding field.

a) Within Brand Version Diffusion

We examine the diffusion of new versions of browsers within each brand. In this analysis, an observation is an OS and a month; e.g., browser users running Windows 95 in April, 1996.

Many of the ordinary specification issues of discrete choice models apply here. For example, to avoid the situation where the share of the newest version is 1 by default, we start with the second browser version of the brand available on the OS.³¹ Table 3 shows the dates for which each browser was newest on each OS.

All specifications include $PCDISTR_{b*ot}$, $INTR_{b*ot}$ and the interaction of $INTR_{b*ot}$ with the browser's version ($INTR_V_{b*ot} = INTR_{b*ot} \times V_{b*}$, where V_{b*} takes the values 2, 3, 4, and 5 for each major browser version).³²

Note that we include $INTR_{b*ot}$ and the interaction of $INTR_{b*ot}$ with V_{b*} but not V_{b*} itself. This is not an econometric oversight, but a specification decision. Including V_{b*} in Equation 1 would permit the long run penetration rate of newer versions of browsers to be lower or higher than older versions. While the long run penetration rate in a general diffusion study should vary (Griliches 1957) the diffusion of a new browser version should ultimately replace the previous version.

Observations for an OS are included only for the time periods shown in Table 4. We estimate by ordinary least squares after stacking the equations for each included browser on each OS. We have more observations ("N" in Table 7) for Netscape, since Netscape entered earlier and, especially early on, supplied more browser versions for more OSs. In addition to the coefficients of the estimating equation, Table 7 reports, in a line called Prob_Der_Mult, the multiplier that converts coefficients into probability

³¹ On the Macintosh OS in the Netscape columns the dependent variable is never the share of NS1; it is first the share of NS2 in the months when it is the newest browser, then of NS3, and so on. But for IE on the Macintosh, we drop IE2; IE1 was never offered for that OS. Similarly, we drop NS1 in our Windows 3.1 observations in the Netscape analysis, and we drop IE4 in our Windows 98 observations in the Internet Explorer analysis, since that was the first available IE browser. Table 3 shows the dates for which each browser was newest on each OS.

³² The use of cardinal values to represent different versions is a restrictive model of technical change. We also employ interaction terms between $INTR_{b*ot}$ and dummy variables for the browser versions. The estimated coefficients corroborate the magnitudes from employing $INTR_V_{b*ot}$, so we employ $INTR_V_{b*ot}$ in order to continue precise estimation of coefficients as we add more regressors.

derivatives at the mean values of the regressors (since the means of S_{b*ot} are similar for the two brands of browser, the two values of Prob_Der_Mult are close as well). Since the units of the distribution variables are defined as shares, they have a quantitative interpretation in our regression. Consider a probability derivative like $\partial \Pr(S_{b*ot|brand}) / \partial PCCARRY_{b*ot}$. That measures the tendency of users running a given brand of browser on operating system o to use the browser that came with it, thereby measuring the impact of PC manufacturer distribution.

Descriptive statistics of the dependent variables and regressors for the stacked samples may be found in Table 6. Estimates may be found in Table 7.

Table 6 Descriptive statistics for within-brand diffusion variables

	Mean	Std. Dev.	Min	Max
Netscape Browser Diffusion within Brand, n=131				
S_{b*ot}	0.349	0.204	0.014	0.946
$\ln[S_{b*ot}/(1 - S_{b*ot})]$	-0.825	1.201	-4.225	2.866
$PCDISTR_{b*ot}$	0.247	0.252	0.012	0.937
$PCEXCLU_{b*ot}$	0.181	0.238	0	0.875
$INTR_{b*ot}$	7.122	4.071	1	16
$INTR_V_{b*ot}$	29.069	19.711	3	75
$ISPDISTR_{b*ot}$	0.158	0.121	0	0.423
Internet Explorer Browser Diffusion within Brand, n=109				
S_{b*ot}	0.329	0.240	0.005	0.857
$\ln[S_{b*ot}/(1 - S_{b*ot})]$	-1.033	1.505	-5.323	1.794
$PCDISTR_{b*ot}$	0.260	0.223	0.002	0.752
$INTR_{b*ot}$	8.220	5.299	1	24
$INTR_V_{b*ot}$	30.780	21.342	3	96
$ISPDISTR_{b*ot}$	0.217	0.140	0	0.571
$ISPTIED_{b*ot}$	0.162	0.170	0	0.571

Table 7 Browser diffusion within brand

Dependent Variable: Logit of Share of Newest Version of Brand of Browser					
Regressor	Netscape			IE	
	(1)	(2)	(3)	(4)	(5)
C	-1.884 (0.174)	-1.932 (0.174)	-1.941 (0.197)	-3.125 (0.143)	-2.726 (0.128)
PCDISTR _{b*ot}	2.216 (0.360)	4.648 (1.239)	2.195 (0.363)	4.104 (0.350)	3.570 (0.304)
PCEXCLU _{b*ot}		-3.014 (1.472)			
INTR _{b*ot}	0.488 (0.063)	0.370 (0.084)	0.550 (0.119)	0.209 (0.051)	0.529 (0.059)
INTR_V _{b*ot}	-0.102 (0.013)	-0.073 (0.019)	-0.109 (0.017)	-0.023 (0.013)	-0.074 (0.014)
ISPDISTR _{b*ot}			-1.153 (1.885)		-9.850 (1.320)
ISPTIED _{b*ot}					5.192 (0.902)
Prob_Der_Mult	0.21	0.21	0.21	0.19	0.19
N	131	131	131	109	109
R ²	0.42	0.44	0.42	0.75	0.84
Standard errors in parentheses					

The base model is reported in columns (1) and (4) of Table 7 for Netscape and IE respectively. The pattern of the estimates is largely the same. In each, we have a large, positive, precisely estimated coefficient on PCDISTR_{b*ot}. This means that a higher rate of shipments of new computers increases the rate of diffusion of new browser versions (here, “new computers” refers to the ones that were shipped since the latest browser version was released).

The coefficients on PCDISTR_{b*ot} are economically large. At the mean of all the regressors, consider what the model predicts if we increase by 10% the fraction of users who got their PC since the newest version of their brand of browser was released. For IE,

the model predicts an 8% increase in use of the newest version.³³ For Netscape, the figure is 5%. These results show that there is a strong tendency of users to run the version of their brand of browser which came with their PC, particularly for IE users. That also means that the pace of sales of new PCs strongly influences the diffusion of the latest browser version.

Table 7 columns (1) and (4) also have implications for technical progress in browsers. Each specification includes $INTR_{b*ot}$ and $INTR_V_{b*ot}$. The interpretation in these two columns of $INTR_{b*ot}$ is that it measures all causes of the pace of diffusion other than distribution with new PCs. The positive coefficient of $INTR_{b*ot}$ and the negative coefficient of $INTR_V_{b*ot}$ show that the pace of diffusion of new browser versions, holding fixed the stock of PCs, is declining. For both brands, the decline is rapid.

We interpret the decline in the pace of diffusion as measuring a deceleration in technical progress (since distribution has been held constant). It is a broad measure of technical progress, including improvements in websites as well as improvements in browsers.

In percentage terms, the decline is somewhat faster for Netscape in column (1) than for IE in column (4), a pattern that persists across the other specifications in Table 7. This table, which is entirely about the pace of diffusion of browser versions within brand, thus confirms once again that during our sample period, imitator Microsoft was catching up to innovator Netscape. IE had more rapid technical progress than Netscape.

In addition to our base specification, we report two other specifications for Netscape and one more for IE in Table 7. Column (2) adds the PC manufacturer exclusion restrictions measured by $PCEXCLU_{b*ot}$. This permits us to sharpen our test of the distribution hypothesis by examining what happened when Microsoft imposed “must not carry” Netscape restrictions on PC manufacturers. These effects are estimated in a specification which permits a declining rate of technical progress in Netscape browsers, so they do not confuse distribution restrictions with technical progress.

³³ We use the probability derivative multipliers and make comparisons at the means of the data in column (1) and column (4). An increase of 10% in $PCDISTR_{b*ot}$ for the newest version of Internet Explorer leads to an increase in that newest version’s usage share in all Internet Explorer browsers of 7.8.% (the probability derivative is 0.19×4.104). For Netscape, an increase of 10% in $PCDISTR_{b*ot}$ leads to a smaller increase in the newest version’s share, 4.7% (0.21×2.216).

Column (2) differs from (1) in two regards. First, a large, negative, significant coefficient on $PCEXCLU_{b*ot}$ shows that the “Must Not Carry” restrictions substantially reduced distribution of Netscape browsers. Second, the coefficient on $PCDISTR_{b*ot}$ is much larger in the Netscape estimates in column (2) than in (1). In column (1), this coefficient is the average effect of sales of new computers, averaged over times with and without Microsoft-imposed restrictions on distribution. In column (2), this coefficient measures the effect of sales of new computers without the contractual restrictions. Economically, the difference means that PC distribution was an important channel for Netscape users and that the rate of diffusion of the newest Netscape browsers would have been much greater if not for the exclusion restrictions.

There is no column like (2) reported for Internet Explorer browsers in Table 7; if there were, it would have $PCDISTR_{b*ot}$ and $PCCARRY_{b*ot}$. As we noted earlier, there is not enough independent variation in $PCCARRY_{b*ot}$ conditional on $PCDISTR_{b*ot}$, $INTR_{b*ot}$, and $INTR_V_{b*ot}$ to estimate such a specification.

In columns (3) and (5) we add measures of ISP distribution. The ISP measures vary only across time, so they have much less meaningful variation in them. When we include a single $ISPDISTR_{b*ot}$ variable for Netscape in column (3), we get a negative coefficient, but not one that we can estimate precisely. The absence of $ISPTIED_{b*ot}$ as a regressor may lead the coefficient on $ISPDISTR_{b*ot}$ to reflect the exclusion of Netscape distribution through ISPs. However, adding $ISPTIED_{b*ot}$ puts too many regressors that vary together into that model to estimate any ISP effect precisely.

In column (5) we add both $ISPDISTR_{b*ot}$ and $ISPTIED_{b*ot}$ to an IE specification. The large positive coefficient on $ISPTIED_{b*ot}$ implies that Microsoft’s distribution restrictions on ISPs increased the diffusion rate for new versions of Internet Explorer. The IE marketing manager who obtained the distribution restrictions in order to avoid side by side product comparisons with Netscape browsers (Myhrvold 1999) appears to have been correct in his assessment. Absent the restrictions, an increase in ISP subscriptions leads to a decrease in the rate of IE version diffusion ($-9.85+5.192$) though the sum is not all that precisely estimated. This suggests either (a) that Internet-oriented PC users dislike Microsoft technologies, or (b) that the variation in the ISP variables may be picking up other effects which are moving around over time. There is considerable

reason to believe both of these stories and not enough information in the data to distinguish them.

We can precisely estimate the effects of distribution versus technical progress on diffusion of new versions of browsers when we have variation in both time series and across OSs, such as the PC manufacturer distribution variables. In other parts of our analysis, we have lower-quality variation in the regressors (as in the ISP variables which vary only over time), and our ability to measure their effects in these regressions is less. Nonetheless, the well-estimated coefficients of PC distribution and technical progress tell a clear story in which distribution is an important force.

b) Brand Shares: Direction of Diffusion

In Table 9, we examine the other demand dimension, browser brand shares. In all the columns of Table 9, the dependent variable is the share of IE browser usage in each month, and the denominator is Netscape plus IE browser usage. Descriptive statistics can be found in Table 8.

In columns (1) - (3), labeled “aggregate,” an observation is a month and the dependent variable is the aggregate share of IE browser usage by users running any of our four mass-market OSs, $S_{IE,t}$. In columns (4) - (7), labeled “each OS,” an observation is an OS / month, and the dependent variable is $S_{IE,ot}$, the share of IE on that OS. The number of observations increases by this disaggregation, but not by 4 times many observations because we do not observe all the OSs in all the months (see Table 4). Our sample also varies by whether we include or exclude the time period in which version 5 of IE is newest (long sample).³⁴

A literature (cited above) has already taken up the analysis of browser brand shares. We start with an analysis like the one presented by Schmalensee (1999a and 1999b) and Liebowitz and Margolis (1999). In column (1) we follow that earlier work as closely as possible, using the same time period and the same, single, predictor of Internet Explorer’s share. That predictor is $Rel_Qual_Jrnl_t$, the relative quality rating of browser

³⁴ In columns (4) - (7), estimation is by stacked OLS, whereas in columns (1)—(3), it is by OLS. In the short sample, we use a regressor based on Schmalensee (1999b). In the long sample, we extend his definition to version 5 browsers.

brands measure based on the average share of journalists' recommendations for and against IE at time t .³⁵

Table 8 Descriptive statistics for brand share variables

Variable					
	Sample				
		Mean	Std. Dev.	Min	Max
	Aggregate (n=45)				
$S_{IE,t}$		0.41	0.17	0.11	0.68
Rel_Qual_Jrnl _t		0.41	0.59	-1	1
D ₃ Dummy for version 3 newest ⁽¹⁾		0.31	0.47	0	1
D ₄ Dummy for version 4 newest		0.38	0.49	0	1
D ₅ Dummy for version 5 newest		0.22	0.42	0	1
	Each OS (n=135)				
$S_{IE,ot}$		0.306	0.264	0.002	0.941
PCCARRY _{b*ot}		0.251	0.228	0	1
(1) Our samples begin when version 2 browsers are newest, so that is the omitted category in the version dummies.					

Columns (2) and (3) continue to restrict the predictors of brand shares to relative quality measures, but use a larger sample size that includes version 5 browsers. Here we use two measures of relative quality. In column (2) we extend Rel_Qual_Jrnl_t to the era of IE5. In column (3), we include browser version dummies, a less restrictive treatment of relative quality.

In columns (4) and (6), we continue to follow the specification of the earlier literature, using only relative browser quality as a predictor of brand shares. Here, however, we use observations for each OS.

These specifications replicate the findings emphasized by Schmalensee (1999a and 1999b) and Liebowitz and Margolis (1999). First, increases in relative browser quality, measured by Rel_Qual_Jrnl_t, predict increases in browser brand market shares.

³⁵ Several features of Schmalensee (1999b) and Liebowitz and Margolis (1999) make it difficult to exactly replicate their work for purposes of statistical testing. Their work was graphical rather than statistical. Their work focuses on "leadership," i.e. a dummy that is one when the market share of a brand of browser (or other application) is largest. Leadership only changes once in browsers, so we replace it with share.

In these specifications, the coefficient is large and precisely estimated. Second, the predictive power of the model, measured by R^2 , is high, at least in the aggregate model.

Table 9 Browser brand shares

Dependent Variable: Share of Internet Explorer in IE + Netscape Browser Usage Models Differ by Definition of an Observation							
	IE Share Aggregate			IE Share Each OS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Constant	0.299 (0.013)	0.306 (.014)	0.119 (0.030)	0.226 (0.020)	0.088 (0.011)	0.238 (0.020)	0.111 (0.016)
Rel_Qual_Jrnl _t	0.224 (0.021)	0.266 (0.019)		0.188 (0.031)	0.108 (0.014)	0.239 (0.028)	0.171 (0.018)
Version 3 Newest Dummy			0.149 (0.034)				
Version 4 Newest Dummy			0.357 (0.033)				
Version 5 Newest Dummy			0.514 (0.035)				
PCCARRY _{b*ot}					0.623 (0.029)		0.584 (0.042)
N	35	45	45	105	105	135	135
R ²	0.78	0.82	0.89	0.26	0.87	0.36	0.74

Liebowitz and Margolis (1999) interpret their findings, which closely parallel those in columns (1)-(4) and (6) of Table 9, as evidence that “Good products win.”³⁶ They reject theories in which there are causes of brand leadership in software other than product quality, characterizing them as speculation. Their conclusion, they assert, comes from using superior methods to earlier scholars, particularly the “first systematic examination of real-world data from software applications markets. ... [their] most important finding is the close relationship between market share change and product quality.”³⁷ Schmalensee (1999a) agrees, writing “Netscape obtained fewer new users of Web-browsing software than Microsoft because its product did not keep pace with improvements in Microsoft’s IE, and because it made numerous business and technical

³⁶ Liebowitz & Margolis (1999), p. 135. They make similar analyses of several applications software markets.

³⁷ Liebowitz & Margolis (1999), p. 227.

mistakes.”³⁸ Schmalensee also contrasts his use of market outcome statistics to the government’s reliance on internal Microsoft correspondence and documents in the antitrust case. He rejects the government’s theory that distribution matters in favor of the quality theory on this basis.

The earlier literature just cited and the specifications just examined all have in common that relative quality is the only cause of brand market shares explicitly considered in the data analysis. However, since both quality and distribution are potential causes of brand shares, the appropriate empirical approach is to include regressors associated with both causes.

We can undertake this approach in our each-OS sample. In columns (5) and (7) of Table 9, we continue to use the Schmalensee measure of relative quality and also include $PCCARRY_{b*ot}$, the Internet Explorer distribution advantage variable. These columns are otherwise identical to columns (4) and (6), respectively.

For identification of both a quality and a distribution effect, it is essential that we have cross section variation across the OSs as well as time series variation.

$PCCARRY_{b*ot}$ varies across versions of Windows, and it varies over time on the Macintosh. As a result, in the each-OS sample, there is substantial independent variation in the distribution and quality measures. In column (7), for example, the correlation between $PCCARRY_{b*ot}$ and $Rel_Qual_Jrnl_t$ is 0.27. If we did not have cross-OS variation, however, we would not be able to pursue the analysis.

The result in both columns (5) and (7) is that the distribution advantage variable has a positive and precisely estimated coefficient. The estimates are economically significant. Increasing the percentage of OS users who obtained a computer on which Microsoft compelled the distribution and display of IE implies that the percentage use of IE rises by about 60% as much (0.62 or 0.58, depending on specification).

While quality continues to matter, it quantitatively matters less once we account for distribution. That is, specifications which look only at quality exaggerate its importance. Both quality and the distribution advantage for IE are trended upward, so the specification with only one regressor has omitted variable bias. Schmalensee (1999a and 1999b) and Liebowitz and Margolis (1999) examine only quality as a cause of market

³⁸ Schmalensee (1999b), slide C.

shares.³⁹ When we include a measure of distribution as well as their measure of quality, their restricted specification and its conclusion are rejected.

Our conclusion from the brand share analysis is the same as from the diffusion of new versions of the same brand. We find an important role for both technical progress and distribution. The two analyses are based on different phenomena. One comes from examining the diffusion of new versions of both brands of browsers, and the other comes from examining the brand shifts that occurred over time on different OSs. In each case, there is independent variation in the distribution variables. Our results show a common explanation. Users of computers respond to economic forces such as distribution convenience as well as to technical progress.

6) Quantitative Implications of the Estimates

Both the diffusion-within-brand results and the brand share results have estimates of the impact of distribution and of technical progress. We now examine the quantitative implications of the estimates for the relative importance of those two causes.

The first sense of the “relative importance” of technical progress is

$$\frac{\frac{\partial Y}{\partial T}}{\frac{\partial Y}{\partial \text{dist}}} \frac{\Delta T}{\Delta \text{dist}} \quad \text{Equation 4}$$

where Y is a demand behavior – either a brand share or a rate of adoption of new browsers, dist is a distribution measure, T is a technical progress measure, and Δdist and ΔT are comparable changes in distribution and in technical progress. For any given scenario, defined by Δdist and ΔT , this sense of “relative importance” asks whether a hypothetical change in distribution of size Δdist would have more or less impact than a hypothetical change in technical progress of size ΔT .

The second sense of “relative importance” is more historically grounded. Suppose that demand behavior is different at different times, as in the slowing rate of browser diffusion within brand over time or the rising brand share for IE over time. We can ask how distribution and technical progress changed between those times, and what contribution to the brand share and the pace of new version diffusion each makes.

³⁹ Liebowitz and Margolis (1999) also examine evidence they say relates to tipping, but not in the context of a model that also includes quality measures.

a) Brand Share

We begin with the brand market share estimates in Table 9. Here our definition of Y in Equation 4 is the share of Internet Explorer, S_{IE} .

For $\Delta dist$, we consider the counterfactual scenario in which IE had no distribution advantage. For ΔT , we consider the counterfactual scenario in which IE did not catch up to Netscape in quality at all. To estimate ΔT we note that historically, the value of $Rel_Qual_Jrnl_t$ grew from -1 (all journalists recommended Netscape) to 1 (all Microsoft). So we use $\Delta T = 2$, a change in $Rel_Qual_Jrnl_t$ of 2. To estimate $\Delta dist$, note that value of $PCCARRY_{b*ot}$ varied from 0 to 1, so we use $\Delta dist = 1$. These values for $\Delta dist$ and ΔT represent large changes in the regressors (though these changes are still within our sample), so one should interpret the prediction with care.

Using estimates of the slopes from column (7) of Table 9, we get an estimate of the relative importance of technical progress of just under 0.6 ($= 0.171/0.584 \times 2/1$). Technical progress is quantitatively less important as an explanation of brand shares than is distribution.⁴⁰

We now turn to the second sense of relative importance. How much did brand shares shift over time as a result of distribution and of technical progress? We have already picked appropriate values of $\Delta dist$ and ΔT to address this directly. However, quantitative interpretation of the numerator or the denominator of Equation 4 in this instance must be made carefully. To interpret them as the predicted market share change from a change in technological progress or in distribution is not obvious. Network effects mean that browser markets tip, so the predicted market share in a counterfactual experiment should be near either 0 or 1, given the size of $\Delta dist$ and ΔT .

The ratio of numerator to denominator has an interpretation whether or not there are network effects. It can be interpreted as showing the relative size of the impact of the historical change in IE product quality (technical progress) vs. the historical change in distribution advantages for IE. Our results imply a larger impact from the historical

⁴⁰ If we used column (5) of Table 9 instead, we would get a smaller estimate of the relative importance of technical progress. Both specifications lead to a much larger role for distribution than for technical progress.

distribution advantage than from the historical change in product quality – technical progress had about 0.6 as large an impact.

b) Within Brand Version Diffusion

We now turn to the relative importance of technology and distribution for the pace of diffusion within brand.

For this analysis, the relevant concepts of \dot{Y} , ΔT , and $\Delta dist$ are all time derivatives, so we label them with a dot above them. Our definition of \dot{Y} is the rate of growth of the share of the newest brand, \dot{S}_{b^*} . Our definition of ΔT is the difference in rates of progress between adjacent versions of the browser. That corresponds to a change in V of -1 (negative because we are comparing the higher rate of technical progress earlier to the lower rate of technical progress later.) For the numerator of Equation 4 we note that $(\beta_{INTR} + \beta_{INTR_V} \times V)$ is the model's prediction of \dot{S}_{b^*} . Accordingly, the numerator is $Prob_Der_Mult \times \beta_{INTR_V} \times (-1)$.

Our conceptual definition of $\Delta \dot{dist}$ is the monthly rate of growth of the stock of PCs in use, and we use the mean monthly rate of growth of $PCDISTR_{b^*ot}$ over all OSs and both brands in our sample; implicitly, we are comparing historical distribution to a counterfactual world of no distribution of browsers with new computers. In our sample, $PCDISTR_{b^*ot}$ is growing at 3.2% per month. The denominator of Equation 4 is $Prob_Der_Mult \times \beta_{PCDISTR} \times 0.032$.

We are now ready to calculate Equation 4 for Internet Explorer based on the estimates in column (4) of Table 7. Noting that the probability derivatives cancel, this is $\beta_{INTR_V} / \beta_{PCDISTR} \times -1/0.032 = -0.023/4.104 \times -1/0.032 = 0.18$. Quantitatively, the amount of technical progress represented by this measure of ΔT is less important than distribution measured by $\Delta \dot{dist}$.

Making a parallel calculation for Netscape browsers based on $\beta_{PCDISTR}$ in column (1) of Table 7 does not make economic sense. $PCDISTR_{b^*ot}$ in column (1) measures both the effect of more rapid distribution and the effect of limitations on distribution. A better set of estimates come from column (2) of Table 7, in which distribution is measured by $PCDISTR_{b^*ot}$ and restrictions on distribution are measured by $PXECLU_{b^*ot}$.

The first Netscape calculation is based on the coefficients of $PCDISTR_{b*ot}$ and $INTR_V_{b*ot}$ in column (2). This corresponds to the question, how important is distribution vs. technical progress for Netscape browsers when there are no restrictions on distribution of Netscape? This is a within-sample calculation, since there were times and OSs (such as on the early days of the Macintosh) in which there were no restrictions on Netscape browser distribution.⁴¹ Again using the same values for $\Delta \dot{dist}$ and $\Delta \mathcal{T}^{\&}$ and again letting the probability derivative terms cancel, this is $-0.073/4.648 \times -1/0.032 = 0.49$. Technical progress represented by this measure of $\Delta \mathcal{T}^{\&}$ is once again quantitatively the less important force.

There is a clear difference between the IE and Netscape results; the impact of technical progress is closer to that of distribution for Netscape (0.49 vs. 0.18). Given the underlying economic and technological situation, this result is not surprising. First, our measure of $\partial \mathcal{Y} / \partial \mathcal{T}^{\&}$ should be larger for Netscape. While both brands of browser show slowing technical progress across versions, in this era IE is catching up technically to Netscape, so the rate of decline of technical change for Netscape ($-\partial \mathcal{Y} / \partial \mathcal{T}^{\&}$) should be larger. Second, distribution should be more important for IE, since contractual restrictions requiring bundling of IE make distribution comparatively important for that brand. Both real-world differences go in favor of a larger impact of technical change relative to distribution for Netscape, which is what we find.

A second calculation for Netscape contrasts the effects of restrictions on distribution to technical progress, based on the coefficients of $PCEXCLU_{b*ot}$ and $INTR_V_{b*ot}$. This answers the economic question, which was a quantitatively more important predictor of the pace of adoption of Netscape browsers, technical progress or the restrictions on distribution imposed by Microsoft? Looking across the four eras in which the four major versions of Netscape browser were newest, we see monthly rates of growth of $PCEXCLU_{b*ot}$ (averaged across OSs) of 0.0%, 0.0%, 3.6%, and 3.3%. The obvious $\Delta \dot{dist}$ to use contrasts the early and late periods; we use $\Delta \dot{dist} = -0.035$ (negative to have it in units distribution rather than blockage of distribution).

⁴¹ Though there are requirements to distribute IE which may slow the distribution of Netscape. This will tend to bias downward this estimate relative to the concept we are attempting to quantify.

Accordingly, the value of Equation 4 is $-0.073/-3.014 \times -1/-0.035 = 0.69$. In this sense as well, the impact of technical progress is less than the impact of distribution (restrictions).

We now turn to the second sense of “relative importance,” and ask which forces are historically important in explaining the slowing pace of diffusion over time. For Internet Explorer, the slowing pace of diffusion is shown in Figure 2. We use estimates from Table 7, column (4). Here the explanation of the change over historical time is simple. The pace of technical progress is slowing for IE, $\beta_{INTR_V} < 0$. That slowing is not offset by the positive impact of more rapid distribution over time.⁴²

The slowing pace of diffusion of new Netscape browsers (shown in Figure 1) can also be explained using estimates from Table 7, particularly column (2). Here, too, one cause is slowing technical progress, $\beta_{INTR_V} < 0$. For Netscape browsers, two distribution coefficients matter. The coefficient of $PCDISTR_{b*ot}$ is 4.6; that of $PCEXCLU_{b*ot}$ is -3.0. The monthly rates of growth of $PCDISTR_{b*ot}$ during the era when our four major versions were the newest of their brand (averaged across OSs) were 4.4%, 3.7%, 3.0%, and 2.8%, respectively.⁴³ The monthly rates of growth of $PCEXCLU_{b*ot}$ during those three eras were 0.0%, 0.0%, 3.6%, and 3.3%. The trend in $PCDISTR_{b*ot}$ over time is slightly down, and the trend in $PCEXCLU_{b*ot}$ is upward. Accordingly, neither distribution coefficient works to offset the slowing pace of technical progress. Instead, the contribution of distribution forces is to further slow of the diffusion of Netscape browsers. Given the sizes of the coefficients and the trends in the Xs, more of this comes from the trend toward exclusion of Netscape browsers from distribution with new PCs than from slowing $PCDISTR_{b*ot}$ growth.

The last two subsections have examined the role of technical progress and of distribution from the perspective of our two models, and have found broadly similar conclusions: both forces matter, and the distributional forces quantitatively matter somewhat more.

⁴² Historically, during the eras when IE2, IE3, IE4, and IE5 were the newest of their brand, the rates of change of $PCDISTR$ (averaged across OSs) were 0.8%, 3.2%, 2.1% and 4.4%, respectively. That is an average rate of growth between versions of 0.8%, or 0.008. That implies a change in the rate of growth of the newest version of 0.00064 / month, so the upward trend in $PCDISTR$ is too small to offset slowing technical progress.

⁴³ The trends in $PCDISTR$ are different for the two brands because of weighting. IE in the later period has a bigger weight on Windows 98, which is rapidly growing.

c) Effect of PC Demand Growth on Browser Diffusion

In this subsection, we view these same quantitative findings from a different perspective. Since distribution with new PCs is quantitatively important, what role did the rapidly expanding field of diffusion for PCs play in the rate and direction of browser adoption? This analysis suggests aggregate implications of our estimates for economic growth and transformation.

The late 1990s saw a rapid expansion in the demand for information technology, brought about by the conversion of the Internet into a mass market technology for commercial applications. The social returns to that technology depended on the widespread diffusion of browsers suitable for mass market commercial applications. The private returns to that technology depended, in part, on which brand of browser would be in widespread use. Both the pace of diffusion of browsers and the direction in terms of brands were affected by the rapidly expanding field of diffusion, i.e., the growing installed base of PCs.

PC industry growth is typically rapid; it was extraordinary in the period we study. From January, 1995, to December, 1999, the installed base of personal computers doubled from 106 million to 213 million computers. How did that rapid growth affect the pace and direction of browser diffusion?

We compare that rapid growth to a replacement-demand scenario. We assume that one fifth of PCs go out of use each year, so that the gross flows of new PCs exceeds the net flow by enough to replace that.⁴⁴ Further, we assume that $PCDISTR_{b*ot}$, $PCCARRY_{b*ot}$, and $PCEXCLU_{b*ot}$ are all proportional to the rate of market growth. Finally, we assume that there is a new version of each brand of browser each year. We use derivatives from our browser diffusion model (Columns (2) and (4) of Table 7) at the mean of the data and fix technical progress at the level of a version 3 browser.

In the scenario of rapid growth found in our data during the sample period, the annual net rate of growth of PCs is 19% so the gross rate of new PCs is 39% (19% + 20% replacement demand). In the slower-growth scenario, the annual gross rate is 20%. These correspond to monthly growth rates of 2.8% and 1.5%, respectively.

⁴⁴ This is a steady-state assumption. In fact, during the historical era, retirements of PCs were less than a fifth of the stock because the stock consisted disproportionately of newer PCs.

First, as we can see by comparing the first two columns of Table 10, the impact of the more rapid growth of PC demand on Netscape version diffusion is small, raising the monthly diffusion rate from just under 4% (0.03 technology effect + 0.005 total distribution effect) to just over 4% (0.03 technology effect + 0.01 total distribution effect).

Table 10 The expanding field and browser diffusion

Effect on rate of diffusion of newest version	Netscape			IE	
	(1)	(2)	(3)	(4)	(5)
PC demand growth*	Fast	Slow	Fast, no Exclusion	Fast	Slow
<u>Technology</u> (INTR+INTR_V×3)	0.15	0.15	0.15	0.14	0.14
(xProb_Der_Mult)	0.032	0.032	0.032	0.027	0.027
<u>Expanding Field</u> (PCDISTR×PC growth)	0.13	0.07	0.13	0.11	0.06
<u>Exclusion</u> (PCEXCLU×PC growth)	-0.084	-0.045	0	0	0
<u>Total Distribution</u> ((Expanding Field + Exclusion) × Prob_Der_Mult)	0.010	0.005	0.027	0.021	0.011
* “Fast” = historical growth, “Slow”= replacement. Prob_der_mult = 0.21 for columns (1)-(3) and 0.19 for columns (4)-(5).					

In contrast, the same change (shown in the columns (4) and (5)) increases the Internet Explorer monthly diffusion rate by a percentage point, from 4% (0.03 technology effect + 0.01 total distribution effect) to 5% (0.03 technology effect + 0.02 total distribution effect). It is easy to understand the large difference between these results. Given the coefficient of $PCDISTR_{b*ot}$, an increase in the pace of new PC demand leads to an increase in the pace of distribution of the newest version of IE that is 83% larger. For Netscape, an increase in $PCDISTR_{b*ot}$ together with a proportionate increase in $PCEXCLU_{b*ot}$ leads to a much smaller effect.

A related result can be seen in the contrast between columns (1) and (3). Removing the impact of the exclusion restrictions substantially increases the pace of diffusion of Netscape; the distribution effect rises from .01 to .03. Again, the intuition is

simple. Exclusion substantially slowed the pace of diffusion of Netscape browsers in the historical world, despite the rapid pace of sales of new PCs.

The implicit calculation of the impact of the rapid growth of PC demand on brand shares compares the first (Netscape) to fourth (IE) columns. Given the high rates of new PC demand and the distribution advantages for IE, the stock of new browsers would soon become overwhelmingly IE. With slower growth (i.e., compare column 2 to the last column) the growth rates of the two brands would be much more similar.⁴⁵

Even for a rapidly progressing technology like the browser at the height of the competitive “browser war,” the role of the expanding diffusion field in overcoming the fixed costs of adoption was very important. Even with the restrictions on the distribution of one brand of browser – restrictions which had a quantitatively significant effect in the decline of Netscape– the expansion of the field was able to contribute to the widespread distribution of a new technology. Without the restrictions on OEM behavior imposed by Microsoft, the expanding field would have been even more powerful.

d) Implications

One quantitative conclusion is common to the brand share and newest-version diffusion estimates: both technical progress (in the broad sense) and distribution were important drivers of diffusion. The comparability of the economic impacts as measured in the brand share and diffusion models is not an artifact of our empirical specification. Those estimates use very different information in the data, suggesting that the comparatively large role for distribution is a robust finding.

The substantial role for distribution confirms what browser marketing executives in both Netscape and Microsoft learned from other kinds of quantitative evidence, such as surveys, and from their business experience. They viewed distribution as very important. (See Fisher (1999) and Bresnahan (2002) for quotations from internal documents.)

7) Our Economic Interpretation (and Its Limits)

The costs of adoption can significantly impact the diffusion of new technologies. In that case, users’ pace of adoption will be open to influence by distribution and

⁴⁵ This calculation is not the most complex one that could be made. Inside Microsoft, Haas (1998) made a somewhat more complex brand growth calculation leading the same conclusion, that distribution with new PCs would push IE into the lead.

marketing as well as by the attractiveness of new technologies. In the particular technology we study (browsers), the distribution forces had a large impact on the rate and direction of technical progress.

That conclusion is unlikely to be a statistical artifact. If anything, our estimates systematically overstate the importance of technical progress, counting the attractiveness of the entire Internet as part of the technical attractiveness of browsers.

Our conclusion is particularly relevant for mass market technologies and platform technologies where the benefits of newer versions of the technology do not always accrue directly to the individual user. In these industries, individual user adoption costs play a large role in the pace of adoption.

As time passes and the speed of connection to networks improve, one source of individual user's fixed costs of new software adoption falls: the time costs of downloading a new browser.⁴⁶ There are other sources of fixed costs of adoption, such as the costs of learning about new versions, their features, their problems (e.g., security) if any, and so on. Our estimates do not permit us to separate the sources of fixed costs. It is not clear from our estimates, therefore, whether we would get the same relative importance of distribution and technological progress later on in this same technology. Further work may shed light on this issue.

The pace of diffusion of narrowly "technical" technologies that do not have a substantial user cost of learning, adapting, adjusting to, or obtaining technologies from the frontier, will be less influenced by distribution. Such "technical" technologies, however, are not the ones for which information technology typically generates new applications with mass market commercial value.

8) Conclusion

Both technical and economic forces affect the diffusion of a new technology. We study these forces in the diffusion of new and improved versions of commercial web browsers in the late 1990s. By exploiting data on browser usage and documented changes in technical progress (browser versions) and distribution methods, we quantify

⁴⁶ Over our sample period, browsers grew larger while online connections grew faster apace, so that the time cost for an average user was not strongly trended.

the significance of these two types of forces. We find that browser distribution via new personal computers (PCs) had a larger effect on the rate and direction of technical change than technical browser improvements.

That is a controversial finding in one of our analyses, of brand choice, and we have been careful to explain why we get a different answer from some earlier data analyses. Our main change is that we include measures of both technical change and distribution in the model. Since we include both causes, our conclusion that one cause mattered more than the other is an advance over the earlier analyses.

We also find the larger role for distribution in the diffusion of new and improved (suitable for complex commercial applications) browser versions. Finding the same forces at work in diffusion and in brand choice reinforces our view of the underlying forces. Enough browser users value distribution convenience that distribution could impact the rate and direction of technical progress.

The role of distribution in overcoming the transactions costs of adoption was magnified, in the case of browsers, by the increasing demand for PCs. While that sounds very narrow and specific, it is in fact an important general economic conclusion.

In complex systems like the Internet, the invention of a new general purpose technology will typically spur growth in complements. Rapid growth in complements – in this case rapid growth in the demand for PCs – in turn feeds back to growth in the new invention – in this case the browser. One mechanism for this feedback is a rapidly expanding diffusion field. Inertial forces like the costs of switching to and adopting the latest technology are less important when the diffusion field is rapidly expanding. This mechanism is particularly important in systems industries with general purpose technologies, like the PC and the Internet.

The economic boom of the late 1990s involved both real exploitation of positive feedback effects and a speculative bubble. The debate over the boom has focused on efforts to measure the productivity impact of investments in new information technology capital and on the role of a stock-market bubble in encouraging over-investment. Our results, though limited in scope to particular markets, point to another set of economic forces at work. One component of the investment boom in IT capital in the late 1990s was that the IT-using side of the economy was moving closer to an expanding

technological frontier. This was due in part to positive feedback dynamics as analyzed in this paper, between new technologies like the browser and existing technologies like the PC. Demand for new technologies increased not only because they became more attractive but also because their adoption costs fell. Examination of the details of choices of particular technologies provides a microfoundation for understanding the overall economics of technical advance.

9) Appendix A – UIUC Data Details

When a computer accesses a page on a web server, it communicates to the server what browser and OS it is running. The server administrator can set the server to record that information in a log (see <http://www.ews.uiuc.edu/bstats/months/>). An example of the log results from September, 1996, is as follows:

Browser Versions - Top 40		
Specific Browser Version	Hosts	%
Mozilla/3.0 (Win95; I)	21598	8.4
Mozilla/2.0 (compatible; MSIE 3.0; Windows 95)	19515	7.6
Mozilla/2.02 (Win16; I)	10580	4.1
Mozilla/1.22 (compatible; MSIE 2.0; Windows 95)	9816	3.8

Web server logs record the Internet Protocol (IP) address of the computer accessing them. As a result, we use a unique IP address (or “host” in the odd lingo of server logs) as our definition of a single user. This will count multiple people who sit at the same computer as one user; the most important example of this is UNIX machines in university computer centers. This will count a user who gets a new IP address for each session as a single user and those who access the UIUC servers more than once in a month as multiple users; the most important example of this is dial-in users using ISPs. Users in offices, students in university dorms, and ISP subscribers who browse UIUC once a month will be counted correctly. In our first month of data, there were 35,757 unique hosts. This rose to 229,579 in the last month of our data.

When a computer accesses a web server, it passes a field called the “user agent” to the server. This field permits us to identify the browser and OS used by the accessing computer. We use the portion of the web log archive called "Browser Versions - All" which reports the number of “hosts” for each distinct user agent field. For each unique IP address (host) the archives record the last browser and OS used in that month.

There are thousands of distinct user agent fields, partly because there are many versions of browsers (especially early on, before conventions for the user agent field were set.) We aggregate all the distinct user agent codes linked to Microsoft or Netscape browsers to the major browser versions listed in Table 1 and Table 2.

Our algorithm for converting user agent codes into browsers and OSs begins with a Perl script used by Ed Kubaitis to make monthly statistical reports on the data we are using and with a number of browser detection software programs used on web servers. We make several changes to the script to catch oddities and exceptional cases.⁴⁷

Many user agent fields are fairly simple to parse. For example “Mozilla/2.0 (compatible; MSIE 3.02; Windows 95)” refers to Microsoft Internet explorer version 3.02 running on Windows 95 and is coded as IE3 on Windows 95 in our classification.

Others are more complicated. Some user-agent fields list multiple browsers and multiple OSs; often these appear as user-agent fields within user agent fields, demarcated by parentheses. For example, “Mozilla/3.0 (Windows 3.10;US) Opera 3.60b3 [en]” is a browser, Opera 3.6, pretending to be another browser, Netscape Navigator 3.0. This “spoofing” is common, since it lets a web server give pages to one browser as if it were another. All Microsoft Internet Explorer browsers, for example, begin their user agent field with the version of Netscape they imitate followed by “compatible.” In our algorithm below, we refer to this problem as “multiple browser names.”

The first step in our browser classification algorithm is to search for a list of names, like “Opera,” which are neither IE nor Netscape and to classify them as “other.” Regardless of where the name appears in the user-agent field, we classify these browsers as “other”.

Our second step is to classify all of the remaining user agent fields as IE or Netscape. If the browser name contains “MSIE,” we classify it as IE. This captures not only Microsoft-branded browsers, but co-branded browsers from AOL, etc., based on MSIE technology. Among the remaining browser names, they are classified as Netscape if they contain either “Netscape” or “Mozilla” but not “compatible.” If there are multiple browser names in the user agent field, we use the outermost (not in parentheses) and leftmost (first) one. Remaining browsers are then classified as “other.”

We also search the user-agent field for the OS that the computer is running. Just as multiple browser names are present, multiple OS names are sometimes present as well.

⁴⁷ Many of the improvements we make to the classification algorithm were not necessary to the purposes of contemporary web statistics. For example, there are versions of IE1 with “IE 4.40” in their user agent, years before the existence of IE4.

In addition to spoofing, sometimes a browser will identify itself as running on a list of OSs. We once again use the outermost and leftmost OS name.⁴⁸

All our econometric results are based on the Microsoft and Netscape browsers, but some users use other browsers. Also, there are some non-human “users”. Not every computer that accesses a web server is being used by a person; some are running automatic indexing programs or “spiders.” Also, a few users change the user agent code, typically to express their engineering individuality. We exclude the spiders and anomalous user-agent fields to the extent possible, keeping only the genuinely distinct browsers (Opera, Mosaic, etc.).

10) Appendix B – Restrictions on Distribution

We model some but not all of the distribution restrictions documented in the antitrust case.

MC_{om} measures the fraction of PC manufacturers selling operating system o who were contractually bound to distribute and display Internet Explorer in month m . PC manufacturers selling Windows 95 and Windows 98 were required to distribute the newest version of IE with new PCs. Thus, Table 11 shows MC_{om} (“must carry”) for the newest version of IE on those operating systems as 1 throughout the sample. PC manufacturers never had to distribute and display IE with Windows 3.1, but always had the MC_{om} restriction for later versions of Windows. For the Macintosh, the MC_{om} restriction activates with a contract that took effect in August, 1997. MC_{om} varies across the OSs in our data and, to a small degree, over time.⁴⁹

We omit any measure of the increasing stringency of the must carry restrictions. Throughout 1995, Microsoft compelled distribution of IE with Windows 95; beginning in early 1996, Microsoft enforced restrictions which compelled display of IE as well, including (for example) putting an IE icon on the Windows desktop, under the “Windows Experience” marketing label. This led to strife with PC manufacturers and monitoring of

⁴⁸ Our sample means depart from Ed Kubaitis’ statistics on this point. He assigns OSs based on a precedence system, classifying a user-agent field as coming from Windows 95 if that OS is named anywhere in the field, for example.

⁴⁹ According to Microsoft’s legal theory, there is no such thing as a browser so that the MC contract provisions simply required OEMs to take all of Windows. While this claim is incorrect, that does not affect our empirical analysis of regressors based on MC_{om} .

compliance by Microsoft. Microsoft bans on valuable technology, such as the special screens which appear when a PC is first used, also led to strife; they were imposed because the screens sometimes mentioned Netscape.

We also omit direct restrictions on end users. Starting with IE3, Microsoft went beyond limitations on PC manufacturers and made it harder for the end user to remove IE from their computer. With IE4, this was even more difficult for users. Similarly, there were increasingly tight bundling strategies for end users who bought Windows without a computer, for example, to upgrade.

There are a number of nuances in the MC_{om} provisions of PC manufacturer contracts not captured here. The most important of these are (a) that the MC_{om} restrictions on PC manufacturers grew more and more complex and onerous over time and (b) that technical restrictions also made it more difficult for the user to remove IE from her computer over time.

To some degree, the requirement to distribute and display IE led PC manufacturers to distribute only IE. Others distributed two browsers.

Furthermore, Microsoft used threats to block PC manufacturers from distribution and display of Netscape Navigator. This leads to a second measure reported in Table 11. The table shows EX_{om} , a variable for restrictions which blocked PC manufacturers from distributing Netscape browsers.⁵⁰ This variable measures the fraction of PC manufacturers who agreed not to distribute or not to display Netscape browsers with new PCs. The Windows 3.1 and Macintosh values are the same as MC_{om} – always zero for Windows 3.1, and changing from zero to one on the Macintosh. For Windows 95 and Windows 98, PC manufacturers agreed to these restrictions only over time. We use findings of fact from the antitrust case reporting on the status of PC manufacturers who had agreed to those restrictions, interpolating linearly between reported dates.⁵¹

The effect of all these restrictions was to lower the number of PC manufacturers carrying two browsers. Netscape had distribution agreements with OEMs in 1995. By

⁵⁰ For this variable, the view of Microsoft's legal defense would matter, since Microsoft denied that there were any such restrictions. Again, this claim is incorrect and does not affect our analysis. For our purposes, EX_{om} is simply the basis for a regressor.

⁵¹ See *Plaintiffs' Joint Proposed Findings of Fact*, US vs. Microsoft & New York vs. Microsoft (1999), Section VII, 364.4.1, i and 364.4.2, ii.

January, 1998, Microsoft had succeeded in blocking PC manufacturer distribution of Netscape browsers on almost all new computers.⁵²

Table 11 Distribution restrictions history

PC manufacturers OS	MC_{om}		EX_{om}	
Windows 3.1	0		0	
Windows 95	1		0	m < Jan. 1996
			11/17	May 1996
			rising linearly to	
			56/60	Jan. 1998
			constant thereafter	
Windows 98	1		56/60	
Macintosh	0	m < Aug. 1997	0	m < Aug. 1997
	1	m ≥ Aug. 1997	1	m ≥ Aug. 1997
ISPs	0	m < Jan. 1997	0	m < Jan. 1997
	1	m > Dec. 1996	1	m > Dec. 1996

Microsoft also imposed distribution restrictions on ISPs and online services (OLSs) like AOL. Starting in 1996, Microsoft sought contracts with these firms by which they agreed to distribute IE and to stringently restrict distribution of Netscape browsers.⁵³ AOL, for example, agreed in March 1996 to such a distribution deal that took effect with the launch of IE3 that summer. By early 1997, fourteen of the fifteen largest ISPs had signed similar contracts (the holdout was Erols).⁵⁴ Warren-Boulton (1998) dates the ISP restrictions program as being in place by late 1996.

⁵² See GX 421 (Kempin 1998).

⁵³ ISPs and online services could not offer or display any way for their customers to get Netscape browsers; if a customer demanded Netscape nonetheless, they could provide it, but not to more than 15% of total customers.

⁵⁴ See GX 93 (Hovstadius 1996) for remarks about IE distribution agreements circa Sept. 1996, GX 1833 (Sikka 1996) for a summary of all ISP browser distribution arrangements circa Dec. 1996, GX 440 (Chase 1997) for April 1997 remarks that ISPs object to the limitations on distribution, and GX 835 (Norberg 1997) for a late 1997 ISP distribution arrangements snapshot.

The restrictions on ISPs are simply represented by a date after which ISPs were required to distribute and display IE browsers and banned from displaying Netscape browsers or distributing them.

11) Complete Descriptive Statistics Table

Descriptive statistics are taken over months when an OS-browser combination exists according to Table 3 and the OS is within sample period as defined in Table 4.

Table 12 Descriptive Statistics for $PCDISTR_{bot}$, $PCCARRY_{bot}$, and $PCEXCLU_{bot}$

Browser	OS	N=	$PCDISTR_{bot}$				$PCCARRY_{bot}$	$PCEXCLU_{bot}$
			Mean	Std. Dev.	Min	Max	Mean	Mean
IE2	mac	45	0.108	0.030	0.017	0.150	0	0
	w31	27	0.075	0.016	0.024	0.088	0	0
	w95	45	0.257	0.199	0.089	0.742	0.257	0
NS2	mac	45	0.047	0.011	0.017	0.066	0	0
	w31	28	0.040	0.007	0.012	0.046	0	0
	w95	45	0.161	0.119	0.058	0.483	0	0.025
IE3	mac	36	0.108	0.035	0.012	0.151	0.041	0
	w31	19	0.064	0.025	0.017	0.087	0	0
	w95	41	0.416	0.143	0.104	0.694	0.416	0
NS3	mac	41	0.133	0.039	0.018	0.184	0	0
	w31	23	0.071	0.017	0.022	0.083	0	0
	w95	41	0.317	0.134	0.104	0.602	0	0.138
IE4	mac	24	0.198	0.114	0.015	0.374	0.198	0
	w31	5	0.005	0.002	0.002	0.008	0	0
	w95	27	0.412	0.126	0.072	0.536	0.412	0
	w98	17	0.694	0.303	0.248	1.000	0.694	0
NS4	mac	28	0.129	0.049	0.014	0.186	0	0.129
	w31	13	0.041	0.014	0.013	0.055	0	0
	w95	31	0.438	0.133	0.063	0.589	0	0.372
	w98	17	0.313	0.305	0.063	1.000	0	0.292
IE5	w95	10	0.096	0.046	0.022	0.158	0.096	0
	w98	10	0.520	0.199	0.150	0.752	0.520	0
NS5	mac	15	0.154	0.080	0.022	0.272	0	0.154
	w95	15	0.172	0.074	0.031	0.268	0	0.160
	w98	15	0.779	0.173	0.341	0.937	0	0.727

12) Bibliography

- Bank, David. (2001) *Breaking Windows: How Bill Gates Fumbled the Future of Microsoft*. New York: Free Press, 2001.
- Bresnahan, T. (2002) "The Economics of the Microsoft Case," working paper, Stanford University.
- Cusumano, Michael A. and David B. Yoffie. (1998) *Competing on Internet Time: Lessons from Netscape and Its Battle with Microsoft*. New York: Free Press, 1998.
- Chase, Brad. (1997) Email Re: Referral Server (GX 490) in *United States v. Microsoft Corporation, Civil Action No. 98-1232*.
- David, Paul A. and Shane Greenstein. (1990) "The Economics of Compatibility Standards: An Introduction to Recent Research." *Economics of Innovation & New Technology*. 1(1-2). pp. 3-41.
- Davis, Steven J. and Kevin M. Murphy. (2000) "A Competitive Perspective on Internet Explorer." *The American Economic Review* 90 (2): 184-187, May 2000.
- Downes, Tom and Shane Greenstein. (2002) "Universal Access and Local Internet Markets in the US." *Research Policy* 31: 1035-1052, 2002.
- Evans, Daniel S. and Richard L. Schmalensee. (2000a) "Be Nice to your Rivals: How the Government is Selling an Antitrust Case without Consumer Harm in *United States v. Microsoft*," in *Did Microsoft Harm Consumers?: Two Opposing Views*. Washington, D.C.: AEI Press, 2000.
- Evans, Daniel S. and Richard L. Schmalensee. (2000b) "Consumers Lose if Leading Firms Are Smashed for Competing," in *Did Microsoft Harm Consumers?: Two Opposing Views*. Washington, D.C.: AEI Press, 2000.
- Fisher, Franklin, (1999) Direct Testimony in *United States v. Microsoft Corporation, Civil Action No. 98-1232* (TPJ), January 1999.
<http://www.usdoj.gov/atr/cases/f2000/2057.htm>
- Fisher, Franklin M. and Daniel L. Rubinfeld. (2000a) "*United States v. Microsoft: An Economic Analysis*," in *Did Microsoft Harm Consumers?: Two Opposing Views*. Washington, D.C.: AEI Press, 2000.

- Fisher, Franklin M. and Daniel L. Rubinfeld. (2000b) "Misconceptions, Misdirection, and Mistakes," in *Did Microsoft Harm Consumers?: Two Opposing Views*. Washington, D.C.: AEI Press, 2000.
- Forman, Chris, Avi Goldfarb, and Shane Greenstein. (2005) Geographic Location and the Diffusion of Internet Technology. *Electronic Commerce Research and Applications* 4 (1): 1-10, Spring 2005.
- Goldfarb, Avi. (2004) "The (Teaching) Role of Universities in the Diffusion of the Internet." University of Toronto, Rotman School of Management, Working Paper, September, 2004.
- Goolsbee, Austan and Peter J. Klenow. (2002) "Evidence on Learning and Network Externalities in the Diffusion of Home Computers." *Journal of Law and Economics* 45: 317-343, October 2002.
- Gottardi, Giorgio. (2004) "Lo sviluppo di Internet. Un'analisi tramite modelli di diffusione (Growth of the Internet : A Diffusion Model Analysis. With English summary)." *L'Industria, Nuova Serie* 24 (April-June 2003): 223-248.
- Greenstein, Shane M. (2004) "The Economic Geography of Internet Infrastructure in the United States." Northwestern University, Center for the Study of Industrial Organization, Working Paper #0046, August 2004.
- Griliches, Z. (1957), "Hybrid Corn: An Exploration in the Economics of Technical Change" *Econometrics*, 25(4) pp.501-522.
- Griliches, Z. (1988), "Postscript on Diffusion," Chapter 3 in *Technology, Education, and Productivity*, Basil Blackwell, New York.
- Hall, Bronwyn H. (2003) "Innovation and Diffusion" Working Paper, University of California.
- Haas, Randy (1998) "Browser Model for 3YO," Email to Brad Chase and others, February 6th. Government Exhibit 310.
- Henderson, Rebecca E. (2000) Direct Testimony in *United States v. Microsoft Corporation*, Civil Action No. 98-1232, April 28, 2000.
- Hovstadius, Bjorn. (1996) Trial Exhibit: Email Re: ISP Marketing Update (GX93) in *United States v. Microsoft Corporation*, Civil Action No. 98-1232. September 9, 1996.

- IDC. (2000a) Client Operating Environments Market Forecast and Analysis, 2000-2004, IDC #22346, June 2000.
- IDC. (2000b) Consumer Internet Service Provider Market Share Update, 1999, IDC #22065, April 2000.
- IDC. (2000c) IDC's Quarterly PC Update: 1Q00 Review/2Q00 Outlook, IDC #22811, August 2000.
- IDC. (2000d) IDC's Quarterly PC Update: 2Q00 Review/3Q00 Outlook, IDC #23224, October 2000.
- IDC. (2000e) IDC's Quarterly PC Update: 4Q99 Review/1Q00 Outlook, IDC #22067, April 2000.
- IDC. (2000f) Internet Service Provider Market Review and Forecast, 1999-2004, IDC #21203, December 1999.
- IDC. (2000g) Network Insider: May 2000, IDC #22379, May 2000.
- IDC. (2000h) Operating Environments Market Forecast and Analysis, 2000-2004, IDC #22597, July 2000.
- Jenkins, Mark, Paul Liu, Rosa Matzkin, and Daniel L. McFadden. (2004) "The Browser War - Econometric Analysis of Markov Perfect Equilibrium in Markets with Network Effects." Working Paper. April 22, 2004, revised December 2004
- Kempin, Joachim. (1998) Trial Exhibit: Microsoft OEM Sales FY '98 Mid-Year Review (GX 421) in *United States v. Microsoft Corporation, Civil Action No. 98-1232.*
- Liebowitz, S. J. and Margolis, S.E. (1999) *Winners, Losers & Microsoft: Competition and Antitrust in High Technology*, The Independent Institute, Oakland, CA, 1998.
- Manes, Stephen and Paul Andrews. (1993) *Gates: How Microsoft's Mogul Reinvented an Industry—and Made Himself the Richest Man in America*. New York: Doubleday, 1993.
- McKean, Kevin. (2002) "Nine Timeless Tips for Tech Buyers." *PC World* (June 2002): p. 19.
- Mowery, David C. and Timothy Simcoe. (2002) "The Internet," in *Technological Innovation and Economic Performance*, eds. Benn Steil, David G. Victor, and Richard R. Nelson. Princeton, New Jersey: Princeton University Press, 2002.

- Myhrvold, Cameron. (1999) Direct Testimony in *United States v. Microsoft Corporation*, Civil Action No. 98-1232, February 10, 1999.
- Nielsen, Jakob (1998) “Alertbox: The Increasing Conservatism of Web Users: March 22, 1998 <http://www.useit.com/alertbox/980322.html>
- Norberg, Susan (preparer). (1997) Trial Exhibit: Netscape Competitive Analysis-ISP/OLS Channel Revenue Segment (GX 835) in *United States v. Microsoft Corporation*, Civil Action No. 98-1232.
- Nordhaus, W. (1997), “Do Real Output and Real Wage Measures Capture Reality? The History of Lighting Suggests No,” in Timothy Bresnahan and Robert Gordon, (eds.) *The Economics of New Goods*, University of Chicago Press, Chicago, pp. 29-70.
- Schmalensee, R. (1999a) Direct Testimony in *United States v. Microsoft Corporation*, Civil Action No. 98-1232 (TPJ), January 1999.
http://www.courtstv.com/archive/trials/microsoft/legaldocs/ms_wit.html
- Schmalensee, R. (1999b) Supplementary Exhibits to Direct Testimony (DX 2098) in *United States v. Microsoft Corporation*, Civil Action No. 98-1232 (TPJ), January 1999.
- Shapiro, Carl and Hal R. Varian. (1999) “The Art of Standards War.” *California Management Review* 41 (2): 8-32, Winter 1999.
- Sikka, Ajay. (1996) Trial Exhibit: Email Re: North American ISP XLS for Jekk Raikes (GX 1833) in *United States v. Microsoft Corporation*, Civil Action No. 98-1232.
December 19, 1996.
- Stango, Victor. (2004) “The Economics of Standards Wars.” *Review of Network Economics* 3 (1): 1-19, March 2004.
- Swisher, Kara. (1998) *Aol.com: How Steve Case Beat Bill Gates, Nailed the Netheads, and Made Millions in the War for the Web*. New York: Random House, 1998.
- U.S. Department of Commerce. (2002) Economics and Statistics Administration. National Telecommunications and Information Administration. *A Nation Online: How Americans are Expanding Their Use of the Internet*, February 2002.

- U.S. Department of Justice et al. (1999) “Plaintiffs Joint Proposed Findings of Fact – Revised” in Civil Action No. 98-1232, US v. Microsoft and State of New York (and others) v. Microsoft. September 10, 1999.
- Vogelsang, Ingo and Benjamin M. Compaine, eds. (2000) *The Internet Upheaval: Raising Questions, Seeking Answers in Communications Policy*. Cambridge, Massachusetts: MIT Press, 2000.
- Wallace, James. (1997) *Overdrive: Bill Gates and the Race to Control Cyberspace*. New York: John Wiley & Sons, 1997.
- Warren-Boulton, Frederick. (1998) Direct Testimony in *United States v. Microsoft Corporation*, Civil Action No. 98-1232 (TPJ), November 1998.
- Windrum, Paul (2000) “Back from the brink: Microsoft and the strategic use of standards in the Browser Wars” MERIT Working Paper, Universiteit Maastricht.
- Wiseman, Alan E. (2000) *The Internet Economy: Access, Taxes, and Market Structure*. Washington, D.C.: Brookings Institution Press, 2000.