

LEVERAGING MONOPOLY POWER BY DEGRADING INTEROPERABILITY:

Theory and evidence from computer markets*

CHRISTOS GENAKOS[†] KAI-UWE KÜHN[‡] AND JOHN VAN REENEN[§]

This Draft: August 2017

ABSTRACT. When will a monopolist have incentives to leverage his market power in a primary market to foreclose competition in a complementary market by degrading compatibility/interoperability of his products with those of her rivals? We develop a framework where leveraging extracts more rents from the monopoly market by “restoring” second degree price discrimination. In a random coefficient model with complements we derive a policy test for when incentives to reduce rival quality will hold. Our application is to Microsoft’s alleged strategic incentives to leverage market power from personal computer to server operating systems. We estimate a structural random coefficients demand system which allows for complements (PCs and servers). Our estimates suggest that there were incentives to reduce interoperability which were particularly strong at the turn of the 21st Century.

JEL CLASSIFICATION: O3, D43, L1, L4.

KEYWORDS: Foreclosure, Anti-trust, Demand Estimation, Interoperability

*We would like to thank Tim Besley, Cristina Caffarra, Francesco Caselli, Sofronis Clerides, Peter Davis, Ying Fan, Jeremy Fox, Georg von Graevenitz, Ryan Kellogg, Christopher Knittel, Mark McCabe, Aviv Nevo, Pierre Regibeau, Bob Stillman, Michael Whinston, two anonymous referees and participants at seminars at AEA, CEPR, Cyprus, FTC, LSE, Mannheim, Michigan, MIT, Pennsylvania and the NBER for helpful comments. Financial support has come from the ESRC Centre for Economic Performance. In the past Kühn and Van Reenen have acted in a consultancy role for Sun Microsystems. The usual disclaimer applies.

[†]Cambridge Judge Business School, AUEB, CEP and CEPR

[‡]University of East Anglia and CEPR

[§]Corresponding author: vanreene@mit.edu; CEP and MIT Department of Economics and Sloan.

1. INTRODUCTION

Many antitrust cases revolve around compatibility issues (called “interoperability” in software markets). For example, the European Microsoft case focused on the question of whether Microsoft reduced interoperability between its personal computer (PC) operating system - Windows, a near monopoly product - and rival server operating systems (a complementary market) to drive rivals out of the workgroup server market. Microsoft’s share of workgroup server operating systems rose substantially from 20% at the start of 1996 to near 60% in 2001 (see Figure 1) and the European Commission (2004) alleged that at least some of this increase was due to a strategy of making rival server operating systems work poorly with Windows. The possibility of such leveraging of market power from the PC to the server market was suggested by Bill Gates in a 1997 internal e-mail: “*What we’re trying to do is to use our server control to do new protocols and lock out Sun and Oracle specifically....the symmetry that we have between the client operating system and the server operating system is a huge advantage for us*”. Microsoft eventually lost the case leading to the largest fines in 50 years of EU anti-trust history.¹

Statements like those in Bill Gates’ e-mail could just be management rhetoric. Indeed, the rationality of such foreclosure strategies has been strongly challenged in the past by the “Chicago School” critique of leverage theory (e.g. Bork, 1978). For example, suppose one firm has a monopoly for one product but competes with other firms in a market for a second product, which is used by customers in fixed proportions with the first. The Chicago School observed that the monopolist in the first market did not have to monopolize the second market to extract monopoly rents. The monopolist will even benefit from the presence of other firms in this second market when there is product differentiation.² Following the Chicago tradition, there has been much work on trying to derive efficiency explanations for many practices that were previously seen as anti-competitive.³

More recently, studies of exclusive dealing (Bernheim and Whinston, 1998) and tying⁴ have

¹The initial interoperability complaint began in 1998 after beta versions of Windows 2000 were released. In 2004, the EU ordered Microsoft to pay €497 million for the abuse and supply interoperability information. In 2008, the EU fined Microsoft an additional €899 million for failure to comply with the earlier decision. See https://en.wikipedia.org/wiki/Microsoft_Corp_v_Commission).

²For a formal statement of this point, see Whinston (1990), Proposition 3.

³For example, Bowman (1957), Adams and Yellen (1976), Schmalensee (1984), McAfee, McMillan and Whinston (1989).

⁴See Whinston (1990), Farrell and Katz (2000), Carlton and Waldman (2002) among others.

shown that rational foreclosure in markets for complements is possible in several models.⁵ Most of these models have the feature that exclusionary strategies are not necessarily profitable in the short run. However, such strategies through their impact on investment, learning by doing, etc., can make current competitors less effective in the future, making the exclusionary strategy profitable in the long run.

This paper makes several contributions. We propose a new theory of foreclosure through interoperability degradation and apply it to the market for PCs and servers. The theory suggests a relatively straightforward policy-relevant test for foreclosure incentives that can be used in many contexts. To implement the test we develop a structural econometric approach using detailed market level data (quarterly data from the US PC and server markets between 1996 and 2001), which requires extending a random coefficient model to allow for complementary products. We find strong and robust incentives for Microsoft to degrade interoperability as the competition authorities alleged.⁶

In our theory, the reduction of competition in the complementary (server) market allows the PC monopolist to more effectively price discriminate between customers with heterogeneous demand. If customers with high elasticity of demand for PCs also have low willingness to pay for servers, server purchases can be used for second degree price discrimination. A monopolist both of PC and server operating systems would lower the price for the PC operating system and extract surplus from customers with inelastic PC demand by charging higher server operating system prices. Competition on the server market will limit the ability to *price discriminate* in this way. By reducing interoperability, the monopolist can reduce competition on the server market, re-establishing the ability to price discriminate.

Although the incentive can exist in theory, whether it is relevant in practice depends on the interplay between two effects. The PC operating system monopolist benefits from reducing interoperability because he gains share in the server market. But because interoperability lowers the quality of rival servers, some customers will purchase fewer PCs, and this reduces his profits from the PC operating system monopoly. Our test quantifies the magnitude of this difference.

⁵See Rey and Tirole (2007) for a comprehensive review of this literature and Whinston (2001) for an informal survey in relation to some aspects of the U.S. vs. Microsoft case.

⁶Hence, our static motivation complements dynamic theories, for example those based on applications network effects, that have been shown to generate anti-competitive incentives to extend monopoly (e.g. Carlton and Waldman, 2002). These dynamic effects would only make our static foreclosure incentives stronger.

For the argument we are making, modelling the heterogeneity between buyers is essential for generating foreclosure incentives. Modelling customer heterogeneity in a flexible way is also a central feature of recent approaches for estimating demand systems in differentiated product markets. We therefore first develop the theory on the basis of a discrete choice model with random coefficients as used in demand estimations by Berry, Levinsohn, and Pakes (1995, henceforth BLP). We extend this approach to allowing complementarity between two markets and compare our results to those from existing approaches such as Gentzkow (2007) and Song and Chintagunta (2006). We show theoretically and empirically how different assumptions on complementarity will affect foreclosure incentives. For example, we show how overly strong restrictions on the assumed form of complementarity (e.g. not allowing a PC only purchase) can cause the econometrician to underestimate the scope for foreclosure.

One caveat to our approach is that although we can test whether there is a foreclosure incentive, our methodology cannot unequivocally resolve the question of whether foreclosure occurs and how much of the change in market shares was due to interoperability constraints.⁷

The paper is structured in the following way. Section 2 gives the basic idea and presents the core theoretical results relating to foreclosure. Section 3 details the econometrics, section 4 the data, section 5 the results and section 6 concludes. In the online Appendices we give more details of proofs (A), data (B), estimation techniques (C and E), derivations (D) and robustness (F).

2. MODELLING FORECLOSURE INCENTIVES

Our basic approach is to measure, at any point in time, the incentives for a monopolist in a primary market to reduce the quality of a rival in a complementary market through changes in the interoperability features of its monopoly product. In our application we examine whether Microsoft had an incentive to degrade interoperability between its PC operating system and rival server operating systems in order to foreclose competition in the server operating system market. In this section we give an overview of how we identify this incentive.

⁷Demonstrating that foreclosure took place through interoperability degradation requires a more in-depth market investigation. In particular, we cannot separately identify whether the introduction of a new Microsoft operating system only enhanced the quality of Microsoft servers relative to others, or whether the decreases in interoperability also decreased the effective quality of rival server operating systems. In the anti-trust case, the European Commission (2004) claimed that changes in Windows technology did seriously reduce interoperability. The evidence in this paper is consistent with the claim that Microsoft had an incentive to do this.

2.1. The Test for Foreclosure Incentives. To outline our approach, we first introduce some notation that will be maintained for the rest of the paper. There are J different types of PCs offered in the market. A buyer of PC j has to pay the price \hat{p}_j for the hardware and ω for the operating system of the monopolist.⁸ In the data, we observe the vector of PC prices $\mathbf{p}_j = \hat{\mathbf{p}}_j + \omega \cdot \mathbf{1}$ with element $p_j = \hat{p}_j + \omega$. For servers we observe the corresponding vector of hardware/software total system prices $\mathbf{p}_k = \hat{\mathbf{p}}_k + \boldsymbol{\omega}_k$ with element $p_k = \hat{p}_k + \omega_k$, where \hat{p}_k is the hardware price of server k and ω_k is the price for the operating system running on that server.⁹ We use the notation $\omega_k = \omega_M$ when the server product k uses the PC monopolist's server operating system. We parameterize the degree of interoperability of the operating system of server k with the monopolist's PC operating system as $a_k \in [0, 1]$. We set $a_k = 1$ for all servers that run the server operating system of the monopolist and $a_k = a \leq 1$ for servers with competing operating systems.

Given the price vectors we can define demand functions for total demand for PCs, $q(\mathbf{p}_j, \mathbf{p}_k, a)$, and the total demand for the monopolist's server operating system as $q_M(\mathbf{p}_j, \mathbf{p}_k, a)$. Total profits of the monopolist are given by:

$$\Pi(\mathbf{p}_j, \mathbf{p}_k, a) = (\omega - c)q(\mathbf{p}_j, \mathbf{p}_k, a) + (\omega_M - c_M)q_M(\mathbf{p}_j, \mathbf{p}_k, a), \quad (1)$$

where c and c_M are the corresponding marginal costs for the monopolist's PC and server operating system respectively.¹⁰

We are interested in the incentive of the monopolist to decrease the interoperability parameter a . By the envelope theorem, there will be such an incentive if:

$$(\omega - c) \left. \frac{dq(\mathbf{p}_j, \mathbf{p}_k, a)}{da} \right|_{\omega, \omega_M} + (\omega_M - c_M) \left. \frac{dq_M(\mathbf{p}_j, \mathbf{p}_k, a)}{da} \right|_{\omega, \omega_M} < 0 \quad (2)$$

The demand derivatives with respect to the interoperability parameter are total derivatives of the respective output measures holding the monopolist's operating system prices constant. This

⁸Over our sample period Apple, Linux and others had less than 5% of the PC market, so that Microsoft could be considered the monopoly supplier. We therefore do not use subscripts for PC operating system prices. We use "M" as the subscript when we refer to prices and quantities of sales of the monopolist's server operating system.

⁹Note that we treat two servers with different operating system as different server products even if the hardware is identical.

¹⁰The marginal cost can be thought of as being very close to zero in software markets.

derivative contains the direct effect of interoperability on demand as well as the impact of the price responses to a change in interoperability by all rival software and hardware producers. Total demand for PCs will increase with greater interoperability because of complementarity between PCs and servers. Greater interoperability means that some customers start purchasing more PCs as the monopolist’s rival servers have become more attractive. At the same time we expect the demand for the monopolist’s server operating system, q_M , to decrease when interoperability increases because some customers will switch to a rival server operating system. The relative impact on server and PC operating system demand from interoperability degradation will therefore be critical to the incentives to foreclose. Rearranging terms we obtain that there is an incentive to decrease interoperability at the margin if:

$$\frac{\omega_M - c_M}{\omega - c} > - \frac{\left. \frac{dq(\mathbf{p}_j, \mathbf{p}_k, a)}{da} \right|_{\omega, \omega_M}}{\left. \frac{dq_M(\mathbf{p}_j, \mathbf{p}_k, a)}{da} \right|_{\omega, \omega_M}} \quad (3)$$

On the left hand side of equation (3) we have the “**relative margin effect**”. Interoperability degradation will only be profitable if the margin on the server operating system of the monopolist ($\omega_M - c_M$) sufficiently exceeds the margin on the PC operating system ($\omega - c$). We call the expression on the right hand side of (3) the “**relative output effect**” as it measures the relative impact of a change in interoperability on demand for the PC operating system (the numerator increases with interoperability) and the monopolist’s server operating system (the denominator decreases with interoperability) respectively.

Our estimation approach is designed to verify whether the strict inequality (3) holds in the data. Why is this a good test for foreclosure incentives when one might expect an optimal choice of interoperability by the monopolist to lead to a strict equality? First, it is costly to change operating systems to reduce the degree of interoperability and there are time lags between the design of the less interoperable software and its diffusion on the market. Second, non-Microsoft server operating system vendors such as Novell and Linux sought to overcome the reduction in interoperability through a variety of measures such as developing “bridge” products, redesigning their own software, reverse engineering, etc. Third, there are many reasons why it will be impossible for a monopolist to reduce all interoperability to zero (i.e. making rival

server operating systems non-functional). One reason is that there are different server market segments. For example, in European Commission (2004) it was claimed that Microsoft had an incentive to exclude rivals in workgroup server markets (the market which we focus on), but not in the markets for web servers or enterprise servers.¹¹ The protocols that achieve interoperability for web servers may, however, provide some interoperability with workgroup servers thus preventing full interoperability degradation. This means the monopolist would want to reduce quality of the server rivals further if he could. Finally, since the late 1990s, anti-trust action in the US and EU may have slowed down Microsoft’s desire to reduce interoperability. All these reasons suggest that in the presence of foreclosure incentives we should find a strict incentive to foreclose at the margin, which is why we focus our analysis on estimating the relative margin and output effects.

2.2. Measuring the Relative Margin Effect. The margins on PC and server operating systems are very hard to observe directly. For our econometric estimations we only have prices of PCs and servers bought inclusive of an operating system. While there do exist some list prices of operating systems that allow us to infer an order of magnitude, we have to estimate the operating system margins from the data. For this estimation we therefore have to impose a specific model of price setting. Given the complementarity between software and hardware as well as between PCs and servers, the move order in price setting is important for determining the pricing incentives for the monopolist. We assume that the hardware and software companies set their prices simultaneously so that the price the software company charges is directly added to whatever price the hardware company charges for the computer. This assumption seems consistent with what we observe in the market as Microsoft effectively controlled the price of the software paid by end users through licensing arrangements.¹² Maximizing equation (1) with respect to the PC operating system price ω and the monopolist’s server operating system price ω_M yields the first order conditions:

¹¹Enterprise servers are high-end corporate servers that manage vast amounts of “mission critical” data in large corporations. They need very high levels of security and typically use custom written software. Web servers host the web-sites of companies and are also used for e-commerce.

¹²Our assumption greatly simplifies the analysis of the monopolist’s problem. While the optimal software price does depend on the expected prices for the hardware, we do not have to solve for the pricing policies of the hardware producers to analyze relative margin effect. If the software company would move first setting prices and the hardware company second, the software company would have to take into account the price reactions of the hardware company.

$$q + (\omega - c) \frac{\partial q}{\partial \omega} + (\omega_M - c_M) \frac{\partial q_M}{\partial \omega} = 0 \quad (4)$$

$$q_M + (\omega - c) \frac{\partial q}{\partial \omega_M} + (\omega_M - c_M) \frac{\partial q_M}{\partial \omega_M} = 0 \quad (5)$$

Denoting $\frac{\partial q}{\partial \omega} \frac{1}{q} = \varepsilon_\omega$ as the semi-elasticity of the impact of a change in operating system price (ω) on quantity demanded (q), we can solve equations (4) and (5) for the PC monopolist's profit margins:

PC monopolist's operating system margin:

$$(\omega - c) = -\frac{1}{\varepsilon_\omega} \left(\frac{1 - \frac{q_M}{q} \frac{\varepsilon_\omega^M}{\varepsilon_{\omega_M}^M}}{1 - \frac{\varepsilon_{\omega_M}}{\varepsilon_\omega} \frac{\varepsilon_\omega^M}{\varepsilon_{\omega_M}^M}} \right) \quad (6)$$

PC monopolist's server operating system margin:

$$(\omega_M - c_M) = -\frac{1}{\varepsilon_{\omega_M}^M} \left(\frac{1 - \frac{q}{q_M} \frac{\varepsilon_{\omega_M}}{\varepsilon_\omega}}{1 - \frac{\varepsilon_{\omega_M}}{\varepsilon_\omega} \frac{\varepsilon_\omega^M}{\varepsilon_{\omega_M}^M}} \right) \quad (7)$$

There are four relevant semi-elasticities: the own-price elasticity of the operating systems of PCs (ε_ω), the own-price elasticity of the monopolist's server operating system ($\varepsilon_{\omega_M}^M$), the cross price elasticity of the monopolist's servers with respect to PC operating system prices (ε_ω^M) and the cross price elasticity of PCs with respect to the monopolist's server operating system prices (ε_{ω_M}). The semi-elasticities that determine the right hand side of these two equations can be estimated from PC and server sales and price data. The operating systems margins and the relative margin effect can therefore be inferred from estimating the parameters of an appropriate demand system.

A first remark on equations (6) and (7) is that the price cost margins differ from the standard single product monopoly margins due to the ratios of cross- to own-price elasticities of PC and server operating system demands, $\frac{\varepsilon_{\omega_M}}{\varepsilon_\omega}$ and $\frac{\varepsilon_\omega^M}{\varepsilon_{\omega_M}^M}$. In general, mark-ups will be affected both by the degree of competition and by the degree of complementarity. As a benchmark case, suppose that PCs and servers are perfect complements which means that customers buy servers and PCs in fixed proportions (i.e. exactly w PCs for every server purchased). With competition between different server operating systems we should generally expect $|\varepsilon_\omega^M| < |w\varepsilon_{\omega_M}^M|$: the

demand response of the monopolist's server operating system should be greater for an increase in the server operating system price (ω_M) than the PC operating system price (ω), because the latter leads to a price increase for all servers and therefore does not lead to substitution between servers due to relative price changes. In the limit, as the server operating system market becomes perfectly competitive, i.e. $\varepsilon_{\omega_M}^M \rightarrow -\infty$ and $\frac{\varepsilon_{\omega}^M}{\varepsilon_{\omega_M}^M} \rightarrow 0$, the PC operating system margin of the monopolist goes to the single product monopoly margin, i.e. $(\omega - c) \rightarrow -\frac{1}{\varepsilon_{\omega}}$. At the same time the server operating system margin goes to zero, i.e. $(\omega_M - c_M) \rightarrow 0$. Hence generally, we would expect $\frac{\varepsilon_{\omega}^M}{\varepsilon_{\omega_M}^M}$ to decrease as competition in the market for server operating systems increases. A further implication of this discussion is that a naive estimation of PC operating system margins that ignored the complementarity between PCs and servers (as the literature typically does) will systematically generate incorrect results for estimated margins. Generally, we would expect operating system margins to be over-stated by the failure to recognize this complementarity. This could be why estimates of PC operating system margins on the basis of elasticities appear to be much higher than their actual empirical values (e.g. Werden, 2001, Schmalensee, 2000 and Reddy et al., 2001).

2.3. Measuring the Relative Output Effect. While the direct impact of a uniform quality reduction of all rivals on demand can be deduced directly from the demand estimates, the total output effect needs to take into account the pricing reactions of rival server operating system and hardware producers. To measure this indirect effect of a quality change on relative output we impose the assumption of profit maximizing behavior also for all rival software and hardware companies. A server with lower quality will command lower prices in equilibrium. Furthermore, if PC demand is reduced as a result of lower server qualities, PC hardware sellers will also partly accommodate by reducing their prices in order to increase demand. These equilibrium price adjustments are crucial to measure the size of the relative output effect. We therefore compute the equilibrium pricing response of each hardware and software producer to a common change of quality in non-Microsoft servers given the estimated demand function assuming a Nash equilibrium in prices. These price responses can then be used to compute the relevant demand derivatives to determine the relative output effect (see Appendix D for details).

To check the robustness of our results we also estimate reduced form equations for PC server operating systems that depend only on quality indices and the estimated price cost margins of the monopolist. The derivatives of this reduced form demand with respect to the quality indices can then be used directly to calculate the relative output effect. This approach avoids the strong structural assumptions we have to make in the first approach, but has more ambiguities of interpretation. We show that the qualitative conclusions of the two approaches are essentially the same (see Appendix F for details).

2.4. Implications of customer heterogeneity for incentives to degrade interoperability. In this sub-section we show how different types of heterogeneity map into foreclosure incentives. The mechanisms that generate foreclosure incentives are all based on theories in which competition in the server market interferes with a (privately) optimal price discrimination strategy by the PC monopolist. By foreclosing the server market, the monopolist can increase rent extraction by using the price of the PC and server operating systems to target different types of customers.

The sign of the relative margin effect is determined by the sign of the server margin in equation (7), which in turn depends only on:

$$1 - \frac{q}{q_M} \frac{\varepsilon_{\omega_M}}{\varepsilon_{\omega}} = -\frac{1}{\varepsilon_{\omega}} \int \left[\frac{q(\theta, w)}{q} - \frac{q_M(\theta, w)}{q_M} \right] [\bar{\varepsilon}_{\omega} - \varepsilon_{\omega}(\theta, w)] dP(\theta) d\Upsilon(w) \quad (8)$$

where “the aggregate elasticity of demand” (i.e. across all consumers) is $\bar{\varepsilon}_{\omega} = \int \varepsilon_{\omega}(\theta, w) dP(\theta) d\Upsilon(w)$ and $P(\theta)$ and $\Upsilon(w)$ are the population distribution functions of θ and w .¹³ It follows that the server operating system margin will be positive if the own price semi-elasticity of the PC operating system, $-\varepsilon_{\omega}$, is positively correlated with $\left(\frac{q(\theta, w)}{q} - \frac{q_M(\theta, w)}{q_M} \right)$. This means that on average buyers with relatively more elastic PC demand (a more negative ε_{ω} than $\bar{\varepsilon}_{\omega}$) have higher market shares in PC purchases than in server purchases from the monopolist. It follows that the server margin will be zero if there is no heterogeneity - the monopolist does best by setting the price of the server at marginal cost and extracting all surplus through the PC operating system price. In this case there is no incentive to foreclose rival servers because all rent can be extracted through the PC price:

¹³The derivation of (8) can be found in Appendix A.2, equation (31).

Proposition 1. *If there is no demand heterogeneity in the parameter vector (θ, w) , then the “one monopoly profit” theory holds. The PC operating system monopolist sells the server operating system at marginal cost and extracts all rents with the PC price. The monopolist has no incentive to degrade interoperability.*

Proof. See Appendix A.3 ■

In order to generate foreclosure incentives the optimal extraction of surplus for the monopolist involves making a margin on the server product. In that case, competition among server operating systems reduces the margin that can be earned on servers and thus restricts the ability of the PC monopolist to extract monopoly rents. By limiting interoperability, rival server quality is reduced and the monopolist “restores” a server margin.

Foreclosure incentives require sorting customers with inelastic PC demand into buying the monopolist’s server operating system, and customers with more elastic PC demand into not buying the server. Our central model generates this feature by assuming customers do not necessarily need to buy a server in order to gain value from PCs. By contrast, servers are complements to PCs in the sense that they only have value when they are consumed with PCs. We call this the *imperfect complementarity* case.¹⁴

To see this, consider a limited type of heterogeneity where buyers have different marginal valuations of server quality γ_i , which can be either $\bar{\gamma}$ or 0. Server quality is denoted y_k . We also assume that there is no other heterogeneity across consumers with respect to the server product and without loss of generality assume only one brand of PC and one brand of server with the monopolist’s server operating system. Given this we can prove:

Proposition 2. *Suppose that $\gamma_i \in \{0, \bar{\gamma}\}$. Then the monopolist sets the server operating system price strictly above marginal cost. Then there exists $\bar{y}_k > y_M$, such that for all $y_k \in (y_M, \bar{y}_k)$ it is optimal for the monopolist to foreclose a competitor by fully degrading the quality of a competing server operating system.*

Proof. See Appendix A.3 ■

¹⁴We also consider a more general case of “free complementary” in section 5.3, where servers can be purchased stand alone.

Proposition 2 states that when the monopolist has a server margin, for some values of rival server quality above the monopolist's server quality ($\bar{y}_k > y_M$), the monopolist will degrade rival quality. The positive value of the server margin follows from equation (8). Although all groups have positive PC sales, the server market share of customers with high value for the server ($\gamma_i = \bar{\gamma}$) is 1 whereas that of customers with low value ($\gamma_i = 0$) is zero. This means that there is a positive correlation between the elasticity of demand of the type γ and the relative importance of that type in PC sales.

Now consider competition on the server operating system market. By standard Bertrand arguments, competition between the two server products will compete down the price of the lower quality product to no more than marginal cost, and the higher quality firm can extract (at most) the additional value provided by its higher quality. If the rival server does not have too much higher quality it will extract all of the quality improvement over the monopolist's product in the server price. This means that a PC monopolist with the lower quality server will generate the same profit as setting the server price at marginal cost without competition and setting the conditionally optimal PC price. Fully foreclosing the competitor is therefore optimal even if the competitor has arbitrarily better quality than the monopolist. Note also that for a rival firm with *lower* server quality the monopolist cannot extract the full value of its own server quality but only the improvement over the quality of the rival. Hence, reducing the quality of the rival slightly will increase the ability to extract surplus.

Proposition 2 holds because competition limits the ability of the PC operating system monopolist to optimally extract surplus. If customers with high elasticity of PC demand sort away from servers, then server sales can be used as a second degree price discrimination device, allowing the monopolist to extract more surplus from high server value/low elasticity customers.¹⁵

¹⁵Note that even where this foreclosure effect exists, it is not always the case that there are *marginal incentives* to foreclose. For example in the above model there are no marginal incentives to foreclose when the rival has higher quality. A small reduction in the quality has no effect on the profits of the monopolist in that case. Only a reduction below the quality of the incumbent will increase profits. In more general models there can even be a *negative marginal incentive to foreclose* when there are *global incentives to foreclose*. This arises from a vertical product differentiation effect. Locally a small increase in the quality of a higher quality rival can lead to higher profits for the monopolist by relaxing price competition as in a Shaked and Sutton (1992) style product differentiation model. Nevertheless, there may be incentives to dramatically reduce quality of the rival in order to increase profits even further. Our focus in the empirical analysis of the marginal incentives to foreclose may therefore lead to an underestimation of the true foreclosure incentives.

3. ECONOMETRIC STRATEGY

3.1. The Model of Demand. The previous section showed the key theoretical mechanisms at play. We now turn to how we econometrically model the problem by considering individual demand as a discrete choice of “workgroup” purchases. A buyer i of type w has demand for a PC workgroup which consists of w PCs and one server. We assume that each buyer can connect his workgroup to at most one server.¹⁶ As before, there are J producers of PCs and K producers of servers indexed by j and k respectively.¹⁷ The index $j = 0$ refers to a purchase that does not include PCs while $k = 0$ refers to an option that does not include a server. A buyer i with workgroup size w who buys the PCs from producer j and the server from producer k has conditional indirect utility:

$$u_{ijk}(w_i) = w_i \left[x_j \beta_i + a_k y_k \gamma_i - \lambda_i \left[p_j + \frac{1}{w_i} p_k \right] + \xi_j + \xi_k + \xi_{jk} + \epsilon_{ijk} \right] \quad (9)$$

The total price for the workgroup is given by $w_i p_j + p_k$ ¹⁸ and the income sensitivity of utility of buyer i is measured by λ_i . The characteristics of PC j are captured by the vector x_j and the characteristics of server hardware software k are represented by the vector y_k . The vectors β_i and γ_i represent the marginal value of these characteristics to buyer i . We normalize quality by assuming that the interoperability parameter $a = 1$ whenever server producer k has the Windows operating system installed. We assume that $a_k = a \leq 1$ is the same for all non-Microsoft servers. In the case of $j = 0$, (x_j, p_j) is the null vector, while in the case of $k = 0$, (y_k, p_k) is the null vector. These represent “workgroup” purchase without a server or without PCs respectively.

The models we estimate differ in whether these choices are allowed, which captures different

¹⁶Assuming that the purchase decisions are only about the setup of a whole “workgroup” implies some important abstractions from reality. If server systems are used for serving one workgroup we effectively assume that the whole system is scalable by the factor $1/w$. Effectively, we are capturing all potential effects of pre-existing stocks of servers and PCs (together with their operating systems) by the distribution of ϵ_{ijk} in equation (8). Since we are assuming that this distribution is invariant over time, we are implicitly assuming that (*modulo* some time trend) the distribution of stocks of computers is essentially invariant. Also note that scalability of workgroups implies that we are not allowing for any difference in firm size directly. All such differences will be incorporated into the distribution of the ϵ_{ijk} and the parameters $(\beta_i, \gamma_i, \lambda_i)$ including a (heterogenous) constant. The idea is to make the relationship between size and purchases as little dependent on functional form as possible.

¹⁷For notational simplicity we are associating one producer with one PC or server hardware type. In the empirical work we, of course, allow for multi-product firms.

¹⁸We can allow for two part tariffs by having p_k take the form $p_k(w) = p_{k1} + w p_{k2}$. This can allow for typical pricing structures in which there is a fixed price for the server operating system and a pricing component based on the number of users (i.e. w “Client Access Licences” have to be purchased). We can accommodate such pricing without any problems in our approach. All that is really important for the pricing structure is that there is some fixed component to the pricing of the monopolist’s server operating system. For simplicity we will exposit all of the analysis below ignoring licenses based on client user numbers.

assumptions about the degree of complementarity. The terms ξ_j and ξ_k represent unobserved quality characteristics of the PC and server respectively, while ξ_{jk} represents an interaction effect between a specific PC and server type.

The term ϵ_{ijk} represents a buyer specific shock to utility for the particular workgroup selected. Assumptions on the distribution of this term among customers will model the degree of horizontal product differentiation between different workgroup offerings. Given that we make ϵ_{ijk} workgroup specific, the variables ξ_{jk} and ϵ_{ijk} capture all of the potential complementarities between the PCs and the servers in a workgroup. In the empirical section we generally assume that $\xi_{jk} = 0$ except for one model version in which ξ_{jk} is a common shift variable for utility whenever a buyer consumes PCs and servers together.

Following BLP, we allow random coefficients on the parameter vector $\theta_i = (\beta_i, \gamma_i, \lambda_i)$ as well as heterogeneity in the size of work groups w_i (captured by a random coefficient on the server price, $\lambda_i^S \equiv \lambda_i/w_i$).¹⁹ We derive demand from the above utility function in the standard way (see Appendix A), the key assumptions being that ϵ_{ijk} comes from a double exponential distribution and that (θ_i, w_i) are drawn from a multivariate normal distribution.²⁰ We can then calculate market shares, s_{ij} for buyer i of PC j as:

$$s_{ij} = e^{\delta_j + \mu_{ij}} \sum_{k=0}^K \frac{e^{\delta_k + \mu_{ik}}}{1 + \sum_{j=1}^J \sum_{k=0}^K e^{\delta_j + \mu_{ij} + \delta_k + \mu_{ik}}} \quad (10)$$

and for servers as:

$$s_{ik} = e^{\delta_k + \mu_{ik}} \sum_{j=1}^J \frac{e^{\delta_j + \mu_{ij}}}{1 + \sum_{j=1}^J \sum_{k=0}^K e^{\delta_j + \mu_{ij} + \delta_k + \mu_{ik}}} \quad (11)$$

where the mean utilities are:

$$\delta_j = x_j \beta - \lambda^{PC} p_j + \xi_j^{PC}; \quad \delta_k = a_k y_k \gamma - \lambda^S p_k + \xi_k^S \quad (12)$$

¹⁹Hence, the ratio of the estimated price coefficients of PCs' to servers' can provide us with an estimate of the implied workgroup size.

²⁰Although the multivariate normal is the most popular choice (e.g. BLP; Nevo, 2001), other possibilities have also been explored (e.g., Petrin, 2002). One could object to the assumption that w_i is normally distributed because strictly speaking it is a count data process (the number of PCs is a positive integer). However, a typical workgroup contains a large number of PC clients (see data section), so viewing λ_i/w_i as distributed normally seems like a reasonable approximation.

and the “individual effects” are:

$$\mu_{ij} = \sigma^{PC} x_j \nu_i^{PC} + \sigma_p^{PC} p_j \nu_{ip}^{PC}; \quad \mu_{ik} = \sigma^S y_k \nu_i^S + \sigma_p^S p_k \nu_{ip}^S$$

The $(\nu_i^{PC}, \nu_{ip}^{PC}, \nu_i^S, \nu_{ip}^S)$ is a vector of the normalized individual effects on the parameters and $(\sigma^{PC}, \sigma_p^{PC}, \sigma^S, \sigma_p^S)$ is the vector of standard deviations of these effects in the population. Notice that the individual effects, μ_{ij} and μ_{ik} , depend on the interaction of customer specific preferences and product characteristics.²¹

3.2. Baseline Imperfect Complementarity Model. Our baseline imperfect complementarity demand model can be empirically implemented in the standard fashion of BLP demand models. We allow customers to select either w PCs or a “workgroup” of w PCs and one server. The indirect utility of the outside option is $u_{i00} = \xi_0^{PC} + \xi_0^S + \sigma_0^{PC} \nu_{i0}^{PC} + \sigma_0^S \nu_{i0}^S + \epsilon_{i00}$, where the price of the outside good is normalized to zero. Since relative levels of utility cannot be identified, the mean utility of one good has to be normalized to zero so we set $\xi_0 = 0$. The terms in ν_{i0} accounts for the outside alternatives’ unobserved variance.

To connect the empirical framework with the theoretical model, we model the interoperability parameter (a) as a multiplicative effect that customers derive from having a Microsoft (M) server:

$$\delta_k = y_k \gamma_1 + \gamma_2 M + \gamma_3 (M y_k) - \lambda^S p_k + \xi_k \tag{13}$$

where M is a dummy variable equal to one if the server runs a Microsoft operating system and zero otherwise. In this way, the interoperability parameter is captured by a combination of the estimated coefficients and therefore we can calculate the “relative output effect” in one step (see Appendix D for details). Given this parameterization, the relationship between the mean utility for servers in equation (12) and the estimates is that $\gamma_3 = \gamma(1 - a)$ and $\gamma_1 = a\gamma$, where $0 \leq a \leq 1$ is the interoperability parameter.²² If there were no interoperability limitations

²¹Hence, our modelling approach is to fix the random draws on the heterogeneity of consumers’ preferences and to compare the results of different assumptions on the strength of the complementarity between servers and PCs (imperfect complementarity vs. free complementarity). An alternative approach would have been to allow consumers to have some correlation in their preferences across PCs and servers as well as an idiosyncratic draw that could directly be estimated from the data. The aggregate nature of our data did not allow us to estimate such more flexible models.

²²We allow γ_2 to be freely estimated as it could reflect the higher (or lower) quality of Windows compared to other operating systems. Alternatively, γ_2 could also reflect interoperability limitations. We examine this

between Microsoft and non-Microsoft operating systems ($a = 1$), then γ_3 , the coefficient on the interaction variable in equation (13), would be estimated as zero.

3.3. Estimation. Our estimation strategy closely follows the spirit of the BLP estimation algorithm, but modifies it so that “multiple product categories” (i.e. PCs and servers) can be accommodated. In essence, the algorithm minimizes a nonlinear GMM function that is the product of instrumental variables and a structural error term. This error term, defined as the unobserved product characteristics, $\xi = (\xi_j^{PC}, \xi_k^S)$, is obtained through the inversion of the market share equations after aggregating appropriately the individual customer’s preferences. However, the presence of multiple product categories means that we need to compute the unobserved term, ξ , via a category-by-category contraction mapping procedure (for a detailed description of the algorithm followed see Appendix C).

Implementing the contraction mapping for PCs and servers is consistent with BLP, but a concern is that feedback loops between the two categories could alter the parameter estimates. There is no theoretical proof we know of for a contraction mapping for random coefficient models with complements, but we made two empirical checks on the results involving further iterations of the algorithm across the two product categories. These both lead to very similar results to the ones presented here (see discussion in Appendix C).

The weighting matrix in the GMM function was computed using a two-step procedure. To minimize the GMM function we used both the Nelder-Mead nonderivative search method and the faster Quasi-Newton gradient method based on an analytic gradient.²³ We combine all these methods to verify that we reached a global instead of a local minimum.

Standard errors corrected for heteroskedasticity are calculated taking into consideration the additional variance introduced by the simulation.²⁴ In our benchmark specification we draw a sample of 150 customers, but we also experiment with more draws in our robustness section. Confidence intervals for nonlinear functions of the parameters (e.g., relative output and relative

possibility in a robustness exercise.

²³In all contraction mappings, we defined a strict tolerance level: for the first hundred iterations the tolerance level is set to 10E-8, while after every 50 iterations the tolerance level increases by an order of ten.

²⁴We do not correct for correlation in the disturbance term of a given model across time because it turns out to be very small. Two features of our approach appear to account for this finding: First, firm fixed effects are included in the estimation. Second, there is a high turnover of products, with each brand model observation having a very short lifecycle compared to other durables like autos.

margin effects) were computed by using a parametric bootstrap. We drew repeatedly (2,000 draws) from the estimated joint distribution of parameters. For each draw we computed the desired quantity, thus generating a bootstrap distribution.

3.4. Identification and instrumental variables. Identification of the population moment condition is based on an assumption and a vector of instrumental variables. Following BLP we assume that the unobserved product level errors are uncorrelated with the observed product characteristics. In other words, that the location of products in the characteristics space is exogenous. In our view, this is realistic for both the PC and server manufacturers since most R&D and components built are developed and produced by other firms, not the PC or server manufacturers themselves. Given this exogeneity assumption, characteristics of other products will be correlated with price, since the markup for each model will depend on the distance from its nearest competitors. To be precise, for both PCs and servers we use the number of products produced by the firm and the number produced by its rivals, as well as the sum of various characteristics (PCs: speed, RAM, hard drive; servers: RAM, whether the server is rack optimized, number of racks, number of models running Unix) of own and rival models. One of the many contributions of Gandhi and Houde (2015) is the insight that the power of the instruments can be significantly improved by choosing a sub-set of “close” characteristics of rival products. Following this idea, we also constructed all instruments in the PC equation separately for desktops and laptops.²⁵

As emphasized by the important contribution of Berry and Haile (2014), we also examine the robustness of our results by varying the type of instruments used. First, we experimented using alternative combinations of computer characteristics. Second, we use hedonic price series of computer inputs, such as semi-conductor chips, which are classic cost shifters. The results are robust to these two alternative sets of instruments, but they were less powerful in the first stage. Finally, we followed Hausman (1996) and Hausman et al (1994) and used model-level prices in other countries (such as Canada, Europe or Japan) as alternative instruments. These instruments were powerful in the first stage, but there was evidence from the diagnostic tests that they were invalid (see Genakos, 2004 and Van Reenen, 2004, for more discussion).

²⁵See also Bresnahan, Stern and Trajtenberg’s (1997) study of the PC market.

Finally, one important limitation of using aggregate data is that we cannot separate true complementarity (or substitutability) of goods from correlation in customers' preferences (see Gentzkow, 2007). Observing that firms that buy PCs also buy servers might be evidence that the two product categories in question are complementary. It might also reflect the fact that unobservable tastes for the goods are correlated - that some firms just have a greater need for "computing power". However, notice that for our purposes such a distinction does not make a major difference to the theoretical results - so long as there is a correlation between customers' heterogeneous preferences for PCs and their probability of buying servers, the incentive to foreclose can exist.

4. DATA

Quarterly data on quantities and prices between 1996Q1 and 2001Q1 was taken from the PC Quarterly Tracker and the Server Quarterly Tracker, two industry censuses conducted by International Data Corporation (IDC). The Trackers gather information from the major hardware vendors, component manufacturers and various channel distributors and contains information on model-level revenues and transaction prices.²⁶ The information on computer characteristics is somewhat limited in IDC, so we matched more detailed PC and server characteristics from several industry datasources and trade magazines. We concentrate on the top fourteen computer hardware producers with sales to large businesses (over 500 employees) in the US market to match each observation with more detailed product characteristics.²⁷ We focus on large businesses as these are the main customers who clearly face a choice to use servers (see Genakos, 2004, for an analysis of other consumer segments).

For PCs the unit of observation is distinguished into form factor (desktop vs. laptop), vendor (e.g. Dell), model (e.g. Optiplex), processor type (e.g. Pentium II) and processor speed (e.g. 266 MHZ) specific. In terms of characteristics we also know RAM (memory), monitor size and whether there was a CD-ROM or Ethernet card included. A key PC characteristic is the performance "benchmark" which is a score assigned to each processor-speed combination based

²⁶Various datasets from IDC have been used in the literature (Davis and Huse, 2009; Foncel and Ivaldi, 2005; Van Reenen, 2006; Pakes, 2003; Genakos, 2004).

²⁷These manufacturers (in alphabetical order) are: Acer, Compaq, Dell, Digital, Fujitsu, Gateway, Hewlett-Packard, IBM, NEC, Packard Bell, Sony, Sun, Tandem and Toshiba. Apple was excluded due to the fact that we were unable to match more detail characteristics in the way its processors were recorded by IDC.

on technical and performance characteristics.

Similarly, for servers a unit of observation is defined as a manufacturer and family/model-type. We also distinguish by operating system, since (unlike PCs) many servers run non-Windows operating systems (we distinguish six other categories: Netware, Unix, Linux, VMS, OS390/400 and a residual category). For servers key characteristics are also RAM, the number of rack slots,²⁸ whether the server was rack optimized (racks were an innovation that enhanced server flexibility), motherboard type (e.g. Symmetric Parallel Processing - SMP), and chip type (CISC, RISC or IA32). Appendix B contains more details on the construction of our datasets.

Potential market size is tied down by assuming that firms will not buy more than one new PC for every worker per year. The total number of employees in large businesses is taken from the US Bureau of Labour Statistics. Results based on different assumptions about the potential market size are also reported.

Table A1 provides sales weighted means of the basic variables for PCs that are used in the specifications below. These variables include quantity (in actual units), price (in \$1,000), benchmark (in units of 1,000), memory (in units of 100MB) as well as identifiers for desktop, CD-ROM and Ethernet card. Similarly, Table A2 provides sales weighted means of the basic variables that are used for servers. These variables include quantity (in actual units), price (in \$1,000), memory (in units of 100MB), as well as identifiers for rack optimized, motherboard type, each operating system used and number of racks. The choice of variables was guided by technological innovation taking place during the late 1990s, but also developments and trends in related markets (e.g. Ethernet for internet use or CD-ROM for multimedia).

Table A1 shows that there was a remarkable pace of quality improvement over this time period. Core computer characteristics have improved dramatically exhibiting average quarterly growth of 12% for “benchmark” and RAM. New components such as the Ethernet cards that were installed in only 19% of new PCs at the start of the period were standard in 52% of PCs by 2001. CD-ROM were installed in 80% of new PCs in 1996 but were ubiquitous in 2001. Furthermore, technological progress is accompanied by rapidly falling prices. The sales-weighted average price of PCs fell by 40% over our sample period (from \$2,550 to under \$1,500).²⁹

²⁸Rack mounted servers were designed to fit into 19 inch racks. They allow multiple machines to be clustered or managed in a single location and enhance scalability.

²⁹There is an extensive empirical literature using hedonic regressions that documents the dramatic declines in

Similar trends hold for the server market. Core characteristics, such as RAM, exhibit an average quarterly growth of 12% over the sample period, the proportion of servers using rack-optimization rose from practically zero at the start of the period to 40% by the end. The average price of servers fell by half during the same period (from \$13,523 to \$6,471). More importantly, for our purposes, is the dramatic rise of Windows on the server from 20% at the start of the sample to 57% by the end. As also seen in Figure 1, this increase in Windows’ market share comes mainly from the decline of Novell’s Netware (down from 38% at the start of the sample to 14% by the end) and, to a lesser extent of the various flavors of Unix (down from 24% to 18%). The only other operating system to have grown is open source Linux, although at the end of the period it had under 10% of the market.³⁰

5. RESULTS

5.1. Main Results. We first turn to the demand estimates from a simple logit model (Table 1 for PCs and Table 2 for servers) and the full random coefficients model (Table 3), before discussing their implications in terms of the theoretical model. The simple logit model (i.e. $\mu_{ij} = \mu_{ik} = 0$) is used to examine the importance of instrumenting the price and to test the different sets of instrumental variables discussed in the previous section for each product category separately. Table 1 reports the results for PCs obtained from regressing $\ln(s_{jt}) - \ln(s_{0t})$ on prices, brand characteristics and manufacturer identifiers. The first two columns include a full set of time fixed effects, whereas the last four columns include only a time trend (a restriction that is not statistically rejected). Column (1) reports OLS results: the coefficient on price is negative and significant as expected, but rather small in magnitude. Many coefficients have their expected signs - more recent generations of chips are highly valued as is an Ethernet card or CD-ROM drive. But a key performance metric, RAM, has a negative and significant coefficient, although the other quality measure, performance “benchmark”, has the expected positive and significant coefficient. Furthermore, the final row of Table 1 shows that the vast majority of products (85.5%) are predicted to have inelastic demands, which is clearly unsatisfactory.

the quality adjusted price of personal computers. See, for example, Berndt and Rappaport (2001) and Pakes (2003).

³⁰Even Linux’s limited success, despite being offered at a zero price, is mainly confined to server functions at the “edge” of the workgroup such as web-serving rather than the core workgroup task of file and print and directory services (see European Commission, 2004, for more discussion). Web servers have been considered outside the relevant market in the European Commission decision.

Column (2) of Table 1 uses sums of the number of products and their observed characteristics offered by each firm and its rivals as instrumental variables. Treating price as endogenous greatly improves the model - the coefficient on price becomes much more negative and most other coefficients have now the expected signs.³¹ Most importantly, under 1% of models now have inelastic demands. Columns (3) and (4) report the same comparison between the OLS and IV results when we include a time trend instead of a full set of time dummies. Again, as we move from OLS to IV results, the coefficient on price becomes much more negative leaving no products with inelastic demands and all the other coefficients on PC characteristics have the expected sign. For example, both benchmark and RAM have now positive and significant coefficients and virtually all products have now elastic demands. In terms of diagnostics, the first stage results (reported in full in Table A3) indicate that the instruments have power: the F-statistic of the joint significance of the excluded instruments is 9 in column (2) and 27 in column (4). In the last two columns we restrict the number of instruments dropping hard disks in column (3) and also speed in column (4). Focusing on a sub-set of the more powerful instruments further improves our results. In the last column, for example, the first stage F-test is 41, moving the price coefficient further away from zero, leaving no PC with inelastic demand.

Table 2 reports similar results from the simple logit model for the server data. In columns (1) and (2) the OLS and IV results are again reported based on regressions that include a full set of time fixed effects, whereas the latter four columns include instead a time trend (a statistically acceptable restriction). The price terms are significant, but with a much lower point estimate than PCs. Consistent with the PC results, the coefficient on server price falls substantially moving from OLS to IV (e.g. from -0.040 in column (3) to -0.179 in column (4)).

Columns (4)-(6) of Table 2 experiment with different instrument sets (first stages are reported in full in Table A4). Empirically, the most powerful set of instruments were the number of models by the firm, the number of models produced by rivals firms and the sum of RAM by rivals (used in columns (2) and (6)). We use these instruments in all columns and also include the official series for quality-adjusted prices for semi-conductors and for hard-disks (two key inputs for servers) in columns (4) and (5). In addition, column (5) includes sums of rivals' char-

³¹The only exception is monitor size which we would expect to have a positive coefficient whereas it has a small negative coefficient. This is likely to arise from the introduction of more advanced and thinner monitors of the same size introduced in 1999-2001. These are not recorded separately in the data.

acteristics (rack-optimized servers, numbers of racks and use of Unix). Although the parameter estimates are reasonably stable across the experiments, the F-test of excluded instruments indicates that the parsimonious IV set of column (6) is preferred, with a F-statistic of 12.9. In these preferred estimates we find that RAM, the number of racks (an indicator of scalability) and type of chip appear to be significantly highly valued by customers. Most importantly, the estimated proportion of inelastic model demands in the final row falls from over 80% in column (3) to 22% in column (6). Notice also that the coefficient on the interaction of Windows and RAM is always positive and significant in the IV results which is consistent with the idea of some interoperability constraints.³²

A sense-check on the results is to use the fact that the implied workgroup size (the number of PCs per server) can be estimated from the ratio of the price coefficients in the PC equation to those in the server equations. In the IV specifications, the implied workgroup size ranges from 10.6 in column (9) to 18.9 in column (2) which are plausible sizes of workgroups (e.g. European Commission, 2004; International Data Corporation, 1998). Our estimates of PC hardware brand-level elasticities are within the typical range of those estimated in the literature, but are relatively inelastic probably because we focus on large firms rather than on households.³³ One diagnostic problem is that the Hansen-Sargan test of over-identification restrictions reject throughout Tables 1 and 2, a common problem in this literature. There is some improvement as we move to the preferred more parsimonious specifications, but it is a concern for the instruments.

Results from the baseline random coefficients model are reported in column (1) of Table 3. The first two panels (A and B) report the mean coefficients for PCs and servers respectively. Almost all mean coefficients are significant and have the expected sign. The lower rows (C and D) report the results for the random coefficients. We allow random coefficients only on price and one other basic characteristic in our baseline specification - performance benchmark for PCs and RAM for servers.³⁴ Our results indicate that there is significant heterogeneity in the price

³²We also estimated models allowing other server characteristics to interact with the Microsoft dummy. These produced similar evidence that these characteristics were less highly valued when used with a non-Microsoft server. The other interactions were not significant, however, so we use the RAM interaction as our preferred specification.

³³For various estimates of computer demand elasticities see Foncel and Ivaldi (2005), Goeree (2008), Hendel (1999), Ivaldi and Lorincz (2008) or Stavins (1997).

³⁴We also estimated models allowing a random coefficient on the interaction of RAM with Microsoft. This was

coefficient for PCs and servers. For PCs, the random coefficient for the performance benchmark has a large value and although insignificant in column (1) is significantly different from zero in the other columns (see below). The robust finding of heterogeneity in the PC price coefficient is important for our theory as this drives the desire to price discriminate which, according to the model, underlies the incentive to foreclose the server market.

As a cross-check on the plausibility of the estimates it is important that the implied hardware margins from the baseline model seem realistic for both PCs and servers. Assuming multi-product firms and Nash-Bertrand competition in prices for PC and server hardware firms, our derived median margin is 16% for PCs and 34% for servers. This is in line with industry reports at the time that put the gross profit margins of the main PC manufacturers in the range of 10%-20% and for server vendors in the range of 25%-54%.³⁵ Furthermore, the implied mean workgroup size of 10.2 (the ratio of the mean coefficients on PC vs. server prices) is also reasonable.

Figure 2 plots the calculated relative output and operating system margin effects based on these coefficients (together with the 90% confidence interval).³⁶ Server operating system margins are higher than PC operating system margins (as indicated by relative margins well in excess of unity). Note that the operating system margin differences are similar to some crude industry estimates.³⁷ The higher margin on servers than PC operating systems reflects both greater customer sensitivity to PC prices, but more interestingly the finding that there is significant heterogeneity in the effects of price on demand across customers. According to proposition 2 this heterogeneity creates incentives for the PC monopolist to use the server market as a price discrimination device by charging a positive server margin. The positive value of the relative output effect indicates that reducing interoperability has a cost to Microsoft which is the loss of PC demand (due to complementarity). The shaded area in Figure 2 indicates where we estimate

insignificant and the implied overall effects were similar so we keep to the simpler formulation here.

³⁵See International Data Corporation (1999a,b). The numbers are also consistent with other results in the literature. For example, Goeree (2008), using a different quarterly US data set for 1996-1998, reports a median margin of 19% for PCs from her preferred model, whereas ours is 16%.

³⁶Figure A1 in the Appendix plots the calculated relative output and margin effects together with the 95% confidence interval.

³⁷Large businesses will enjoy more discounts than individuals, so we cannot simply look at list prices. IDC (1999, Table 1) estimate server operating environment revenues for Windows as \$1,390m million and license shipments for Windows NT were as 1,814 (Table 4). This implies a “transaction price” for a Windows server operating system (including CALs) as \$766. Similar calculations for PC operating systems are around \$40, suggesting a relative margin 19 to 1 similar to Figure 1.

that Microsoft has significant incentives to degrade interoperability.

Four key findings stand out in Figure 2. First, looking at the period as a whole, the relative margin effect exceeds the output effect from the end of 1996 onwards indicating incentives to degrade interoperability. Second, the two effects trend in opposite directions with the relative output effect decreasing and the relative margin steadily increasing. By the end of our sample period in 2001, the difference between the two effects takes its largest value with the relative margin effect clearly dominating the relative output effect. Third, the two lines diverge around the end of 1999 and beginning of 2000. These dates coincide with the release of the new Microsoft operating system (Windows 2000). The European anti-trust case hinged precisely on industry reports that Windows 2000 contained interoperability limitations that were much more severe than any previous version on Windows (European Commission, 2004). Fourth, the relative output effect is very small and very close to zero from 1999 onwards. This means that foreclosure inducing activities would have had cost Microsoft very little in terms of lost PC operating system sales. As we will show later these four findings are robust to alternative empirical models of complementarity and a battery of robustness tests.

The increase in the relative server-PC margin is mainly driven by the increase in the absolute value of the PC own price elasticity. This is likely to be caused by the increasing “commodification” of PCs over this time period linked to the increasing entry of large numbers of PC brands by low cost manufacturers (e.g. Dell and Acer) as the industry matured and cheaper production sites in Asia became available. The relative output effect is declining primarily because the aggregate number of servers sold was rising faster than the number of PCs, which is related to the move away from mainframes to client-server computing (see Bresnahan and Greenstein, 1999). Thus, a marginal change in interoperability had a smaller effect on loss of PC quantity (relative to the gain in servers) in 2001 than in 1996.

5.2. Robustness. Columns (2)-(8) of Table 3 report various robustness tests of the baseline model (reproduced in Figure 3A to ease comparisons) to gauge the sensitivity of the results to changes in our assumptions. We show that the basic qualitative result that there were incentives to degrade interoperability is robust. First, we vary the number of random draws following the Monte Carlo evidence from Berry, Linton and Pakes (2004) for the BLP model. In column (2)

we increase the number of draws to 500 (from 150 in the baseline model). The estimated results are very similar to our baseline specification, the only exception being that the PC benchmark now has a significant random coefficient. Unsurprisingly, the calculated relative output and margin effects in Figure 3B exhibit the same pattern as in Figure 3A.

In column (3) and (4) of Table 3 we make different assumptions about the potential market size. In column (3) we assume that firms will only make a purchase decision to give all employees a computer every two years, essentially reducing the potential market size by half. In column (4) we assume that the potential market size is asymmetric, whereby firms purchase a PC every year whereas they purchase a server bundle every two years. In both experiments the estimated coefficients are hardly changed and Figures 3C and 3D are similar.

In columns (5) and (6) of Table 3 we reduce the number of instruments used for both the PCs and servers. On the one hand, using the most powerful instruments increases the absolute value of the coefficients. For example, the mean coefficient on PC price increases from -3.301 in the baseline model to -3.622 and -5.598 in columns (5) and (6) respectively. On the other hand, using fewer instruments means that we are reducing the number of identifying restrictions and this is reflected in higher standard errors. As a result very few coefficients are significant in column (6). Despite these differences, Figures 3E and 3F reveal a qualitative similar picture as before.

In columns (7) and (8) of Table 3 we experiment using different random coefficients. In column (7), we add a random coefficient on the constant in both equations. The estimated coefficients indicate no significant heterogeneity for the outside good at the 5% level for either PCs or servers. In column (8) we reduce the number of estimated random coefficients by allowing only a random coefficient on server price. As before, both the estimated coefficients and calculated effects in Figures 3G and 3H look similar to our baseline specification: at the beginning of our sample the relative output effect dominates the relative margin effect, but by the end of 2000 the ordering is clearly reversed indicating strong incentives from Microsoft's perspective to reduce interoperability.

5.3. Free complementarity model. Our baseline model is more restrictive in that complementarity between PCs and servers is built in rather than estimated, as customers are assumed

to buy either a PC, a bundle of a server and PC or the outside good. This choice was driven both by our understanding of how the market for “workgroup” purchases operates (firms buy servers not to use them on a stand alone basis but to coordinate and organize PCs). However, we also analyze a more general model that allows the data to determine the degree of complementarity or substitutability between the two products that we call the “free complementarity” model.

Under the free complementarity model a bundle, indicated by (j, k) , can include either a server, or a workgroup of PCs, or both. Denote d^{PC} as an indicator variable that takes the value of one if any PC is purchased and zero otherwise; similarly we define d^S to be the indicator for servers. Each customer i , maximizes utility by choosing at each point in time, t , the bundle of products, (j, k) , with the highest utility, where utility is given by:

$$u_{ijk} = \delta_j + \mu_{ij} + \delta_k + \mu_{ik} + \Gamma(d^{PC}, d^S) + \epsilon_{ijk} \quad (14)$$

This is identical to the baseline except we have included an additional term, $\Gamma(d^{PC}, d^S)$, that it is not affected by a choice of particular brand once (d_{PC}, d_S) is given and does not vary across customers. This utility structure allows us to model complementarity or substitution at the level of the good, i.e. PC or server, via a "free" parameter, $\Gamma^{PC,S}$, that captures the extra utility that a customer obtains from consuming these two products together over and above the utility derived from each product independently. When $\Gamma^{PC,S}$ is positive we call PCs and servers complements (and if negative substitutes). This model borrows directly from the work of Gentzkow (2007), who was the first to introduce a similar parameter in a discrete setting. Our utility model is more general in that we allow for random coefficients on the model characteristics and prices (Gentzkow does not have price variation in his data). More importantly, our model is designed to be estimated with aggregate market level data. We identify the $\Gamma^{PC,S}$ parameter in the standard way, by using aggregate time series variation in server prices (in the PC demand equation) and time series variation in PCs prices (in the server demand equation).³⁸ Further

³⁸Song and Chintagunta (2006) also build on Gentzkow to allow for a common complementarity/substitution parameter and apply it on store level data for detergents and softeners. We differ from Song and Chintagunta in three ways: (i) we specify a different brand and consumer part of the utility that is closer to the original BLP specification, (ii) we use a different set of instruments to address the issue of price endogeneity and (iii) we implement a more robust estimation method.

model and estimation details are given in Appendix E.

The “free complementarity” model is presented in the last column of Table 3 where we allow customers to purchase standalone servers (as well as standalone PCs, bundles of PCs and servers or the outside good), and let complementarity to be freely estimated through the parameter $\Gamma^{PC,S}$. The estimated $\Gamma^{PC,S}$ parameter is positive and significant, confirming our previous assumption and intuition that PCs and servers are complements. The mean and random coefficients all exhibit similar patterns to the baseline results with evidence of significant heterogeneity in price (for servers and PCs) and significant heterogeneity in customers’ valuation of PC quality (benchmark) but not server quality (RAM). Figure 4 plots the relative output and margin effects and their confidence interval. The relative margin is somewhat lower than under the baseline model of column (1). This is because second degree price discrimination is less powerful because customers who buy servers now come from a group with both low PC demand elasticity and no demand for PCs at all. Nevertheless, we still find incentives to degrade interoperability towards the end of the sample period. Relative output effects remain small. Given that this is a much more demanding specification, the consistency of results with our baseline case is reassuring.³⁹

So far, to measure the relative output effect we relied heavily on our structural demand model to compute the equilibrium pricing response of each hardware and software producer to a common change of quality in non-Microsoft servers. As a final robustness test we also consider an alternative approach to estimating the relative output effect, which considers only the “reduced form” residual demand equations for Microsoft servers and PCs. Since these will be a function of non-Microsoft quality (and other variables), we can use the coefficients on these to calculate the output effects of degradation directly. Appendix F gives the details and shows that relative margins continue to lie far above the relative output effect. The relative output effect is 4 or less, far below the mean relative margin estimated at around 20 in Figure 2. So even this simpler, less structural approach, suggests strong incentives for Microsoft to reduce interoperability.

³⁹The reason why we do not use this model as our baseline is because estimation of the free complementarity was significantly slower to converge and more sensitive to starting values (causing convergence problems). Since identification of both the random coefficients and the $\Gamma^{PC,S}$ parameter come solely from time variation, these problems are hardly surprising given the limited time span of our data.

6. CONCLUSIONS

In this paper we examine the incentives for a monopolist to degrade interoperability in order to monopolize a complementary market. These type of concerns are very common in foreclosure cases such as the European Commission’s landmark 2004 Decision against Microsoft. Structural econometric approaches to examining the plausibility of such foreclosure claims have generally been unavailable. This paper seeks to provide such a framework, developing both a new theory and a structural econometric method based upon this theory.

In our model, the incentive to reduce rival quality in a secondary market comes from the desire to more effectively extract rents from the primary market that are limited *inter alia* by the inability to price discriminate. We detail a general model of heterogeneous demand and derived empirically tractable conditions under which a monopolist would have incentives to degrade interoperability. We then implement our method in the PC and server market, estimating demand parameters with random coefficients and allowing for complementarity. According to our results it seemed that Microsoft had incentives to decrease interoperability at the turn of the 21st century as alleged by competition authorities. In our view, the combination of theory with strong micro-foundations and detailed demand estimation advances our ability to confront complex issues of market abuse.

There are many limitations to what we have done and many areas for improvement. First, our model is static, whereas it is likely that dynamic incentives are also important in foreclosure (e.g. Carlton and Waldman, 2002). An important challenge is how to effectively confront such dynamic theoretical models with econometric evidence (see, for example, Lee, 2013). In the context of the Microsoft case, it is likely that the dynamic effects would strengthen the incentive to foreclose as the European Commission (2004) argued. Second, we have used only market-level data but detailed micro-information on the demand for different types of PCs and servers could lead to improvements in efficiency (see Bloom, Draca and Van Reenen, 2016, for examples of such detailed IT data). Although we have gone some of the way in the direction of endogenising one characteristic choice (interoperability decisions) there is still a long way to go.

REFERENCES

- [1] Abel, J., Berndt, E. and White, C. (2007) “Price Indices for Microsoft’s Personal Computer Software Products” in Berndt, E. and Hulten, C. (eds) *Hard to Measure Goods and Services*, Chicago: Chicago University Press.
- [2] Adams, W. and Yellen, J. (1976) “Commodity Bundling and the Burden of Monopoly”, *Quarterly Journal of Economics*, 90, 475 – 498.
- [3] Angrist, J. and S. Pischke (2009) *Mostly Harmless Econometrics*, Princeton: Princeton University Press
- [4] Bajari, P. and L. Benkard (2005) “Demand Estimation with Heterogeneous Customers and Unobserved Product Characteristics: A Hedonic Approach”, *Journal of Political Economy*, 113(6), 1239-1276.
- [5] Bernheim, D. and Whinston, M. (1998) “Exclusive Dealing”, *Journal of Political Economy*, 106(1), 64-103.
- [6] Berndt, E., and Rappaport (2001) “Price and Quality of Desktop and Mobile Personal Computers: a Quarter Century Overview”, *American Economic Review*, 91(2), 268-71.
- [7] Berry, S., Levinson, J., and Pakes, A. (1995) “Automobile Prices in Market Equilibrium”, *Econometrica*, 63(4), 841-90.
- [8] Berry, S., Linton, O., and Pakes, A. (2004) “Limit Theorems for Estimating the Parameters of Differentiated Product Demand Systems”, *Review of Economic Studies*, 71(3), 613-654.
- [9] Berry, S. and P. Haile (2014) “Identification in Differentiated Product Markets using market level data” *Econometrica*, 82(5) 1749-1797
- [10] Bloom, N., Draca, M. and Van Reenen, J. (2016) “Trade Induced Technical Change”, *Review of Economic Studies*, 83(1): 87-117
- [11] Bork, R. (1978) *The Antitrust Paradox*, New York: Free Press
- [12] Bowman, W. (1957) “Tying Arrangements and the Leverage Problem”, *Yale Law Journal*, 67, 19-36
- [13] Bresnahan, T., S. Stern and M. Trajtenberg (1997) “Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the late 1980s”, *RAND Journal of Economics*, 28, S17-S44.
- [14] Bresnahan, T. and Greenstein, S. (1999) “Technological Competition and the Structure of the Computer Industry”, *Journal of Industrial Economics*, 47(1), 1-40.
- [15] Carlton, D. and Waldman (2002), “The Strategic Use of Tying to Preserve and Create Market Power in Evolving Industries”, *RAND Journal of Economics*, 33(2), 194-220.
- [16] Davis, P. and Huse. C. (2009) “Estimating the Co-ordinated Effects of Mergers”, *mimeo*, Stockholm School of Economics.
- [17] European Commission (2004) “Commission Decision of 23.04.04 relating to the Case COMP/C-3/37.792 Microsoft”
- [18] Farrell, J. and Katz (2000) “Innovation, Rent Extraction and Integration in Systems”, *Journal of Industrial Economics*, 48(4), 413-432
- [19] Foncel, J. and M. Ivaldi (2005) “Operating System Prices in the Home PC Market”, *Journal of Industrial Economics*, 53(2), 265-297.
- [20] Ganghi, A. and J-F. Houde (2015) “Measuring Substitution Patterns in Differentiated Products Industries”, Cornell mimeo

- [21] Genakos, C. (2004) “Differential Merger Effects: The Case of the Personal Computer Industry”, *working paper*, STICERD EI 39.
- [22] Gentzkow, M. (2007) “Valuing New Goods in a Model with Complementarity: Online Newspapers”, *American Economic Review*, 97(3), 713-744.
- [23] Goeree, A. S. (2008) “Limited Information and Advertising in the US Personal Computer Industry”, *Econometrica*, 76(5), 1017-1074.
- [24] Hausman J. A., G. Leonard and J. D. Zona (1994) “Competitive Analysis with Differentiated Products”, *Annales D’ Economie et de statistique*, 34, 159-180.
- [25] Hausman J. A. (1996). “Valuation of New Goods Under Perfect and Imperfect Competition” in T. F. Bresnahan and R. Gordon, eds., *The Economics of New Goods*. Chicago: National Bureau of Economic Research.
- [26] Hendel, I. (1999) “Estimating Multiple Discrete Choice Models: an Application to Computerization Returns”, *Review of Economic Studies*, 66(2), 423-446.
- [27] International Data Corporation (1998) *IDC Server Market Methodology, Database Structure and Definitions* IDC Report No. 15909R.
- [28] International Data Corporation (1999a) *Server Operating Environment and Platform Report*.
- [29] International Data Corporation (1999b) *Client Operating Environments: 1999 Worldwide markets and trends* IDC Report No. 19477.
- [30] Ivaldi, M. and S. Lörincz (2011) “Implementation Relevant Market Tests in Antitrust Policy: Applications to Computer Servers”, *Review of Law and Economics*, 7(1), 31-73.
- [31] Lee, R. (2013) “Vertical Integration and Exclusivity in Platform and Two-Sided Markets”, *American Economic Review*, 103(7), 2960-3000.
- [32] McAfee, J., McMillan, J. and Whinston, M. (1989) “Multiproduct Monopoly, Commodity Bundling, and Correlation of Values”, *Quarterly Journal of Economics*, 104(2), 371-383.
- [33] Nelder, J. A. and R. Mead (1965) “A Simplex Method for Function Minimization”, *Computer Journal*, 7, 308-313.
- [34] Nevo, A. (2001) “Measuring Market Power in the Ready-to-Eat Cereal Industry”, *Econometrica*, 69(2), 307-342.
- [35] Nevo, A. (2003) “New Products, Quality Changes and Welfare Measures Computed from Estimated Demand Systems” *Review of Economics and Statistics*, 85(2), 266-275.
- [36] Oyer, P (1998) “Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality”, *Quarterly Journal of Economics*, 113(1), 149-185.
- [37] Pakes, A. (2003) “A Reconsideration of Hedonic Price Indices with an Application to PC’s”, *American Economic Review*, 93(5), 1578-1596.
- [38] Petrin A. (2002) “Quantifying the Benefits of New Products: The Case of the Minivan”, *Journal of Political Economy*, 110(4), 705-729.
- [39] Reddy, B., Evans, D., Nichols, A. and Schmalensee, R. (2001) “A Monopolist would still charge more for Windows: A Comment on Werden”, *Review of Industrial Organization*, 18, 263-268.
- [40] Rey, P. and Tirole, J. (2007) “A Primer on Foreclosure”, *Handbook of Industrial Organization Volume III*, Armstrong, M. and Porter, R. (eds) 2145-2220.
- [41] Schmalensee, R. (1984) “Gaussian demand and commodity bundling”, *Journal of Business*, 57(1) 211-30.

- [42] Schmalensee, R. (2000) “Anti-trust issues in Schumpeterian Industries”, *American Economic Review*, 90(2), 192-196.
- [43] Song, I. and Chintagunta, P. (2006) “Measuring cross-category price effects with aggregate store data”, *Management Science*, 52(10), 1594-1609.
- [44] Shaked, A. and Sutton, J. (1992) “Relaxing price competition through product differentiation”, *Review of Economic Studies*, 49 (1), 3-13.
- [45] Stavins, J. (1997) “Estimating Demand elasticities in a differentiated product industry: The Personal Computer market”, *Journal of Economics and Business*, 49, 347-36
- [46] Van Reenen, J. (2004) “Is there a market for workgroup servers?”, Centre for Economic Performance Discussion Paper No.650.
- [47] Van Reenen, J. (2006) “The Growth of Network Computing: Quality adjusted price changes for network servers”, *Economic Journal*, 116, F29-F44.
- [48] Werden, G. (2001) “Microsoft’s pricing of Windows and the economics of derived demand monopoly”, *Review of Industrial Organization*, 18, 357-262.
- [49] Whinston, M. (1990) “Tying, Foreclosure, and Exclusion”, *American Economic Review*, 80(4), 837-59.
- [50] Whinston, M. (2001) “Exclusivity and Tying in U.S. v. Microsoft: What We Know and Don’t Know”, *Journal of Economic Perspectives*, Spring, 63-80.

TABLE 1 - RESULTS FOR SIMPLE LOGIT FOR PC MARKET SHARE

	(1)	(2)	(3)	(4)	(5)	(6)
Estimation method	OLS	IV	OLS	IV	IV	IV
Dependent variable	$\ln(S_{jt})-\ln(S_{0t})$	$\ln(S_{jt})-\ln(S_{0t})$	$\ln(S_{jt})-\ln(S_{0t})$	$\ln(S_{jt})-\ln(S_{0t})$	$\ln(S_{jt})-\ln(S_{0t})$	$\ln(S_{jt})-\ln(S_{0t})$
Price	-0.336*** (0.037)	-1.400*** (0.281)	-0.404*** (0.037)	-2.085*** (0.204)	-2.275*** (0.239)	-2.488*** (0.258)
Benchmark	0.305*** (0.108)	0.953*** (0.211)	0.388*** (0.095)	1.153*** (0.160)	1.239*** (0.180)	1.336*** (0.190)
RAM	-0.458*** (0.101)	0.339 (0.246)	-0.333*** (0.105)	0.920*** (0.220)	1.062*** (0.239)	1.221*** (0.262)
CD-ROM	0.226** (0.095)	0.257** (0.112)	0.188* (0.096)	0.278** (0.130)	0.288** (0.136)	0.299** (0.143)
Ethernet	0.140* (0.077)	0.354*** (0.103)	0.105 (0.077)	0.463*** (0.109)	0.504*** (0.116)	0.549*** (0.123)
Desktop	0.375*** (0.070)	-0.406* (0.213)	0.273*** (0.071)	-0.908*** (0.169)	-1.042*** (0.193)	-1.192*** (0.208)
5th Generation	1.068*** (0.244)	1.814*** (0.364)	0.894*** (0.229)	2.520*** (0.379)	2.704*** (0.410)	2.911*** (0.432)
6th Generation	0.889*** (0.268)	2.314*** (0.496)	0.954*** (0.252)	3.652*** (0.472)	3.957*** (0.523)	4.299*** (0.556)
7th Generation	1.112*** (0.395)	2.037*** (0.526)	1.084*** (0.385)	3.087*** (0.561)	3.313*** (0.595)	3.568*** (0.626)
Monitor Size	-0.066*** (0.008)	-0.086*** (0.009)	-0.066*** (0.008)	-0.097*** (0.010)	-0.101*** (0.010)	-0.105*** (0.011)
Trend			-0.051*** (0.013)	-0.368*** (0.041)	-0.404*** (0.047)	-0.444*** (0.051)
Time Dummies (21)	yes	yes	no	no	no	no
Overidentification test		60.383 [0.000]		65.425 [0.000]	50.836 [0.000]	27.114 [0.000]
1 st Stage F-test		8.8 [0.000]		27.21 [0.000]	30.40 [0.000]	40.620 [0.000]
Own Price Elasticities						
Mean	-0.73	-3.04	-0.88	-4.52	-4.94	-5.40
Standard deviation	0.31	1.28	0.37	1.90	2.07	2.27
Median	-0.68	-2.83	-0.82	-4.21	-4.60	-5.03
% inelastic demands	85.51%	0.70%	71.44%	0.03%	0.00%	0.00%

Notes: Demand estimates from a simple logit model based on 3,305 observations. "Benchmark" are numbers assigned to each processor-speed combination based on technical and performance characteristics. "Generation" dummies indicate common technological characteristics shared among central processing units. All regressions include a full set of nine hardware vendor dummies. Columns (2) and (4) use BLP-type instruments: the number of the same form factor own-firm products, the number of the same form factor products produced by rival firms, the sum of the values of the same characteristics (speed, RAM and hard disk) of other products of the same form factor offered by the same firm and the sum of values of the same characteristics of all same factor products offered by rival firms. In the last two columns, we restrict the number of instruments dropping hard disks in column (3) and also speed in column (4). Full first stage results can be found in Table A3 of the Appendix. The Hansen-Sargan test of overidentification and the first-stage Angrist-Pischke (2009) F test of excluded instruments with the p-values in square parentheses below are reported. Robust standard errors are reported in parenthesis below coefficients: *significant at 10%; **significant at 5%; ***significant at 1%.

Source: Authors' calculations based on the IDC Quarterly PC Tracker data corresponding to sales and prices of PC models for large business customers matched to more detailed PC characteristics from several industry datasources and trade magazines, US market (1996Q1-2001Q1).

TABLE 2 - RESULTS FOR SIMPLE LOGIT FOR SERVER MARKET SHARE

	(1)	(2)	(3)	(4)	(5)	(6)
Estimation method	OLS	IV	OLS	IV	IV	IV
Dependent variable	$\ln(S_{kt})-\ln(S_{0t})$	$\ln(S_{kt})-\ln(S_{0t})$	$\ln(S_{kt})-\ln(S_{0t})$	$\ln(S_{kt})-\ln(S_{0t})$	$\ln(S_{kt})-\ln(S_{0t})$	$\ln(S_{kt})-\ln(S_{0t})$
Price	-0.040*** (0.003)	-0.075*** (0.020)	-0.040*** (0.003)	-0.179*** (0.031)	-0.201*** (0.035)	-0.234*** (0.041)
RAM	0.002 (0.007)	0.031* (0.019)	0.002 (0.008)	0.116*** (0.032)	0.133*** (0.036)	0.161*** (0.042)
Windows	-0.861*** (0.113)	-0.567*** (0.196)	-0.867*** (0.114)	0.305 (0.282)	0.484 (0.310)	0.766** (0.357)
Windows × RAM	0.013 (0.011)	0.025* (0.014)	0.012 (0.011)	0.060*** (0.023)	0.068*** (0.025)	0.079*** (0.029)
Symmetric Parallel Processor	0.474*** (0.081)	0.705*** (0.156)	0.474*** (0.081)	1.388*** (0.224)	1.528*** (0.246)	1.748*** (0.284)
Rack Optimized	0.455*** (0.110)	0.337** (0.134)	0.458*** (0.110)	-0.005 (0.182)	-0.076 (0.197)	-0.187 (0.225)
Number of Racks	-0.009 (0.010)	0.006 (0.013)	-0.008 (0.010)	0.051** (0.022)	0.060** (0.024)	0.074*** (0.028)
Linux	0.037 (0.413)	0.542 (0.484)	-0.033 (0.392)	1.995*** (0.605)	2.307*** (0.647)	2.795*** (0.715)
Unix	-0.675*** (0.166)	0.351 (0.589)	-0.681*** (0.168)	3.393*** (0.907)	4.019*** (1.000)	5.000*** (1.176)
OS390/400	-1.750*** (0.204)	-0.711 (0.611)	-1.717*** (0.204)	2.390** (0.936)	3.020*** (1.037)	4.008*** (1.218)
VMS	-1.961*** (0.255)	-1.620*** (0.330)	-2.009*** (0.257)	-0.610 (0.574)	-0.396 (0.639)	-0.059 (0.734)
Other OS	-2.088*** (0.222)	-1.094* (0.596)	-2.070*** (0.222)	1.874** (0.900)	2.480** (0.992)	3.429*** (1.163)
Trend			-0.030*** (0.007)	-0.144*** (0.025)	-0.161*** (0.028)	-0.189*** (0.032)
Time Dummies (21)	yes	yes	no	no	no	no
Overidentification test		64.409 [0.000]		35.389 [0.000]	20.061 [0.000]	12.03 [0.002]
1 st Stage F-test		18.53 [0.000]		5.82 [0.000]	8.70 [0.000]	12.87 [0.000]
Own Price Elasticities						
Mean	-0.63	-1.18	-0.63	-2.84	-3.18	-3.71
Standard deviation	0.62	1.15	0.62	2.79	3.12	3.64
Median	-0.44	-0.83	-0.44	-2.01	-2.25	-2.62
% inelastic demands	81.13%	57.40%	80.89%	32.02%	28.08%	22.11%

Notes: Demand estimates from a simple logit model based on 2,967 observations. All regressions include a full set of nine hardware vendor dummies. Columns (2) and (6) use BLP-type instruments: the number of own-firm products, the number of products produced by rival firms, the sum of RAM of products offered by rival firms. In columns (4) and (5) we also experiment with additional instruments based on server characteristics (sum of Rack and Rack Optimized of products offered by rival firms and sum of Unix own-firm models) and input prices (quality adjusted indices for semi-conductors and hard disks). Full first stage results can be found in Table A4 of the Appendix. The Hansen-Sargan test of overidentification and the first-stage Angrist-Pischke (2009) F test of excluded instruments with the p-values in square parentheses below are reported. Robust standard errors are reported in parenthesis below coefficients: *significant at 10%; **significant at 5%; ***significant at 1%.

Source: Authors' calculations based on the IDC Quarterly Server Tracker data corresponding to sales and prices of server models matched to more detailed server characteristics from several industry datasources and trade magazines, US market (1996Q1-2001Q1).

TABLE 3 - RESULTS FROM ALTERNATIVE COMPLEMENTARITY MODELS

Estimation method	(1) GMM Imperfect complementarity - baseline model	(2) GMM sample of 500 consumers	(3) GMM potential mkt size reduced in half	(4) GMM asymmetric potential mkt size	(5) GMM reduced number of instruments	(6) GMM reduced number of instruments	(7) GMM include random coef. on constant	(8) GMM reduce random coef. on servers	(9) GMM "free" complementarity
PANEL A: PC - Means									
Price	-3.301*** (0.629)	-2.795*** (0.501)	-3.002*** (0.411)	-3.353*** (0.604)	-3.622*** (0.676)	-5.598 (3.882)	-2.768*** (0.555)	-3.350*** (0.635)	-3.314*** (0.592)
Benchmark	0.021 (1.243)	-2.503 (1.572)	0.070 (0.493)	0.088 (1.114)	-0.786 (2.355)	-1.388 (5.770)	-1.971* (1.152)	0.020 (1.229)	-0.153 (1.176)
RAM	0.760** (0.316)	1.088*** (0.284)	0.837*** (0.278)	0.747** (0.303)	0.639* (0.348)	0.244 (0.568)	0.753** (0.320)	0.765** (0.312)	0.801*** (0.303)
CD-ROM	0.275** (0.130)	0.321*** (0.140)	0.261** (0.132)	0.267** (0.129)	0.304** (0.135)	0.315 (0.193)	0.316** (0.133)	0.275** (0.131)	0.278** (0.131)
Ethernet	0.423*** (0.134)	0.490*** (0.134)	0.486*** (0.127)	0.410*** (0.130)	0.403*** (0.149)	0.305 (0.228)	0.443*** (0.136)	0.424*** (0.132)	0.444*** (0.131)
5th Generation	2.783*** (0.395)	2.955*** (0.461)	2.811*** (0.401)	2.766*** (0.391)	2.869*** (0.400)	3.128*** (0.547)	3.153*** (0.491)	2.795*** (0.394)	2.821*** (0.399)
6th Generation	4.053*** (0.574)	4.517*** (0.607)	4.103*** (0.517)	4.007*** (0.558)	4.154*** (0.616)	4.296*** (0.547)	4.619*** (0.718)	4.066*** (0.569)	4.132*** (0.567)
7th Generation	2.709*** (0.606)	3.034*** (0.738)	2.757*** (0.652)	2.663*** (0.597)	2.529*** (0.700)	1.858 (1.335)	3.005*** (0.759)	2.702*** (0.605)	2.733*** (0.629)
Constant	-3.426*** (0.708)	-3.528*** (0.936)	-2.708*** (0.709)	-3.451*** (0.704)	-3.269*** (0.653)	-2.402 (2.671)	-6.319** (2.818)	-3.379*** (0.708)	-3.368*** (0.704)
PANEL B: Server - Means									
Price	-0.282*** (0.089)	-0.352*** (0.133)	-0.288*** (0.094)	-0.258*** (0.085)	-0.249*** (0.081)	-0.298** (0.131)	-0.352*** (0.113)	-0.281*** (0.086)	-0.674*** (0.155)
RAM	0.173*** (0.057)	0.203*** (0.058)	0.177*** (0.052)	0.162*** (0.051)	0.161*** (0.045)	0.180*** (0.060)	0.220** (0.096)	0.174*** (0.049)	0.208*** (0.066)
Windows	0.794* (0.451)	1.069* (0.556)	0.828* (0.460)	0.688 (0.431)	0.683** (0.342)	0.888 (0.737)	1.342** (0.590)	0.781* (0.436)	1.543*** (0.483)
Windows × RAM	0.077** (0.034)	0.092*** (0.041)	0.078** (0.034)	0.072** (0.032)	0.074** (0.033)	0.085* (0.047)	0.102** (0.041)	0.076** (0.032)	0.133*** (0.039)

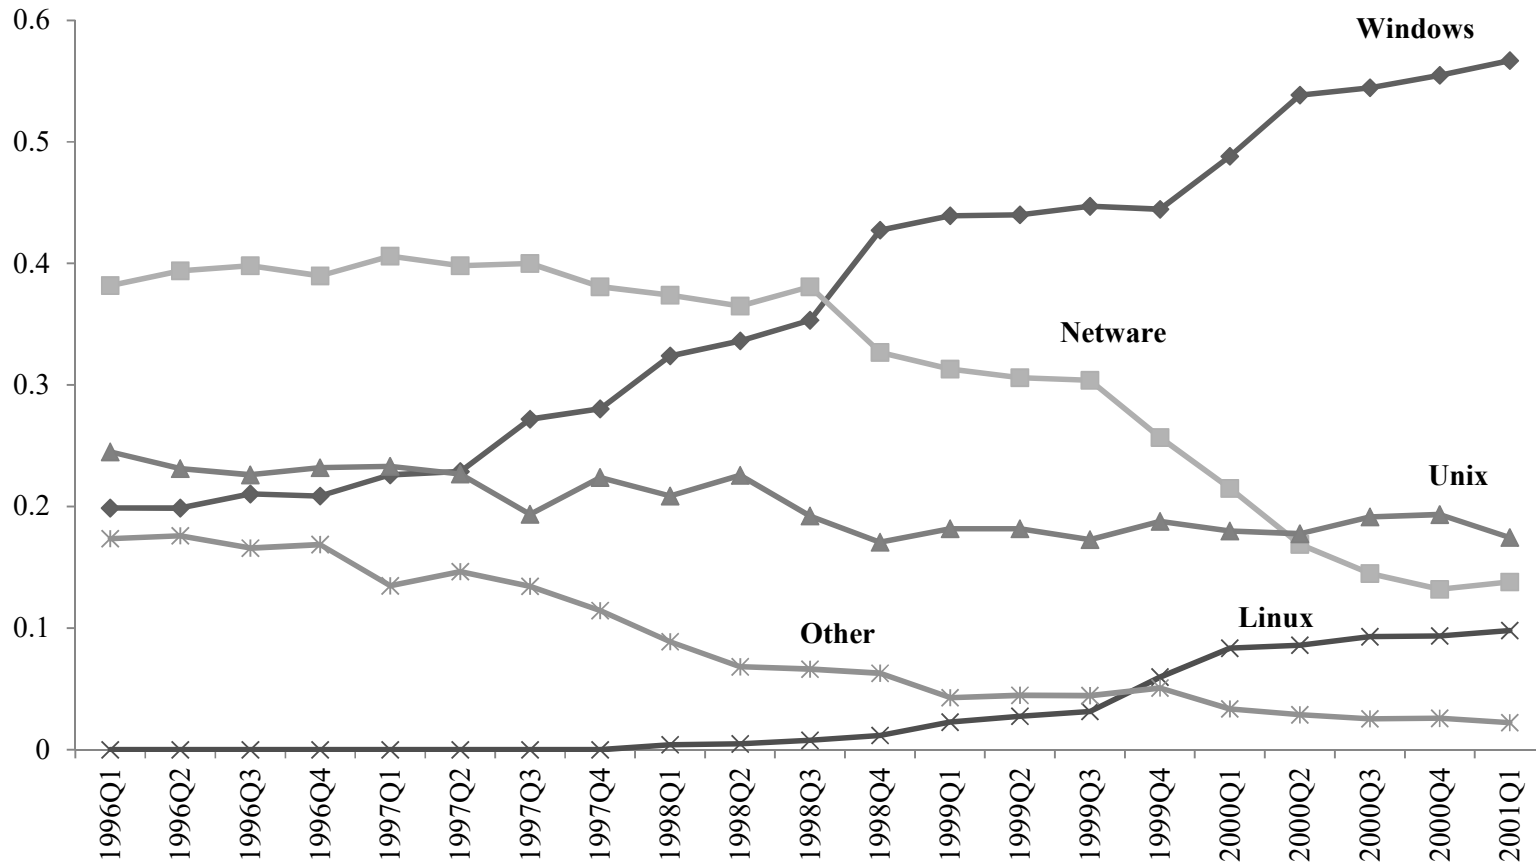
TABLE 3 - RESULTS FROM ALTERNATIVE COMPLEMENTARITY MODELS - cont.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Estimation method	GMM	GMM	GMM	GMM	GMM	GMM	GMM	GMM	GMM
Symmetric Parallel Processor	1.787*** (0.390)	2.015*** (0.498)	1.810*** (0.399)	1.698*** (0.369)	1.690*** (0.307)	1.858*** (0.665)	2.234*** (0.486)	1.773*** (0.358)	2.620*** (0.408)
Rack Optimized	-0.185 (0.234)	-0.266 (0.267)	-0.199 (0.237)	-0.145 (0.223)	-0.154 (0.203)	-0.208 (0.329)	-0.441 (0.296)	-0.176 (0.227)	-0.373 (0.273)
Number of Racks	0.060 (0.039)	0.084 (0.050)	0.063 (0.040)	0.056 (0.036)	0.055 (0.035)	0.074 (0.058)	0.094* (0.052)	0.060 (0.037)	0.140*** (0.034)
Constant	-5.814*** (0.228)	-5.807*** (0.260)	-5.809*** (0.228)	-5.128*** (0.223)	-5.896*** (0.294)	-5.748*** (0.269)	-6.197*** (0.566)	-5.816*** (0.228)	-8.096*** (0.669)
PANEL C: PC - Standard Deviations									
Price	0.916** (0.363)	0.520* (0.283)	0.758*** (0.273)	0.955*** (0.346)	1.140*** (0.413)	2.292 (1.777)	0.795*** (0.270)	0.938*** (0.362)	0.902*** (0.338)
Benchmark	1.321 (0.822)	2.794** (1.102)	1.610*** (0.444)	1.282* (0.771)	1.938 (1.532)	2.690 (3.199)	2.658*** (0.910)	1.332 (0.812)	1.450* (0.752)
Constant				2.569			(1.839)		
PANEL D: Server - Standard Deviations									
Price	0.048** (0.024)	0.062** (0.031)	0.048* (0.025)	0.042* (0.025)	0.035 (0.037)	0.054* (0.030)	0.049* (0.028)	0.049** (0.024)	0.162*** (0.042)
RAM	0.014 (0.104)	0.011 (0.090)	0.011 (0.106)	0.010 (0.103)	0.000 (0.159)	0.007 (0.312)	0.017 (0.145)		0.027 (0.091)
Constant				0.930*			(0.528)		
$\Gamma^{PC,S}$ parameter									2.647** (1.271)
GMM Objective (df)	75.613 (10)	68.583 (10)	71.783 (10)	88.356 (10)	54.146 (5)	46.723 (3)	56.899 (8)	79.292 (12)	57.493 (9)

Notes: Demand estimates from the baseline model as described in section 4.1 in the text based on 6,272 observations. Parameters were estimated via a two-step GMM algorithm described in the estimation subsection 4.2. In all columns (except columns (5) and (6)) the instruments used are the same as in the baseline model in column (1). In column (2) we increase the number of draws relative to the baseline model to 500. In column (3) we assume that firms make a purchase decision every two years. In column (4), we assume that firms purchase a PC every year, whereas a server bundle every two years. In column (5) we reduce the number of instruments used to the ones corresponding to column (5) in Tables 1 and 2 for PCs and servers respectively. In column (6) we further reduce the instruments used for PCs to the ones corresponding to column (6) in Table 1. In column (7) we include a random coefficient on the constant for both PCs and servers. In column (8), we only allow for a random coefficient on price on servers. The last column reports results from the "free complementarity" model estimated via a two-step GMM algorithm as described in Appendix E. All specifications include all the characteristics in Tables 1 and 2, i.e. for PCs: desktop, monitor size, CD-ROM, firm dummies and time trend; for servers: full set of operating system and firm dummies and time trend. The standard errors take into account the variance introduced through the simulation by bootstrapping fifty times the relevant component of the variance in the moment conditions. Robust standard errors are reported in parenthesis below coefficients: *significant at 10%; **significant at 5%; ***significant at 1%.

Source: Authors' calculations based on the IDC Quarterly PC and Server Tracker data matched to more detailed PC and Server characteristics from several industry datasources and trade magazines, US market (1996Q1-2001Q1).

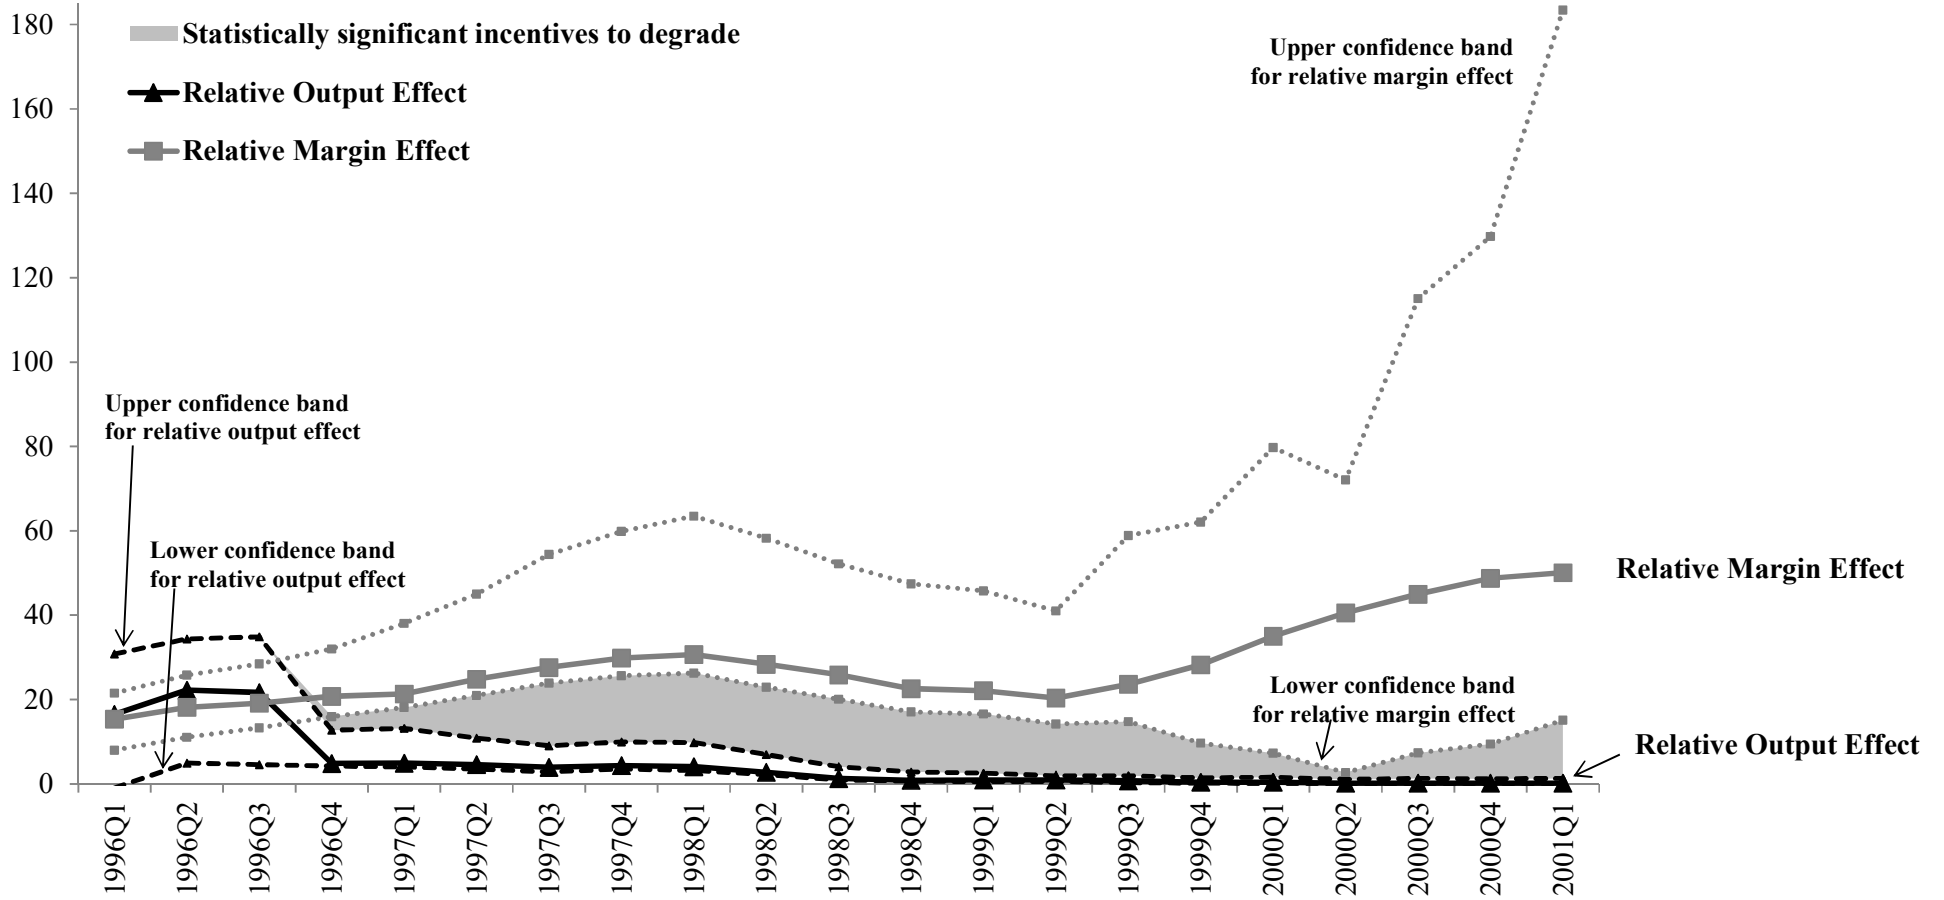
FIGURE 1: EVOLUTION OF MARKET SHARES FOR SOFTWARE VENDORS IN US (units)



Notes: This plots the evolution of shares for different operating systems. Shares are measures in unit volumes. Other includes operating systems include IBM's OS390/400, Compaq's VMS and some other smaller non-Unix operating systems.

Source: International Data Corporation (IDC) Quarterly Server Tracker survey, US market (1996Q1-2001Q1).

FIGURE 2: RELATIVE MARGIN AND OUTPUT EFFECTS (IMPERFECT COMPLEMENTARITY)



Notes: This plots the evolution of calculated relative output and margin effects based on the formulas provided in Appendix D and their 90% confidence interval. The shaded area highlights the period where the relative margin effect is statistically higher than the relative output effect, hence Microsoft had significant incentives to degrade according to our model.

Source: Authors' calculations based on the IDC Quarterly PC and Server Tracker data matched to more detailed PC and Server characteristics and the estimated coefficients in column (1) of Table 3.

FIGURE 3: RELATIVE MARGIN AND OUTPUT EFFECTS (ROBUSTNESS)




FIGURE 3A: COLUMN (1), TABLE 3




FIGURE 3B: COLUMN (2), TABLE 3




FIGURE 3C: COLUMN (3), TABLE 3

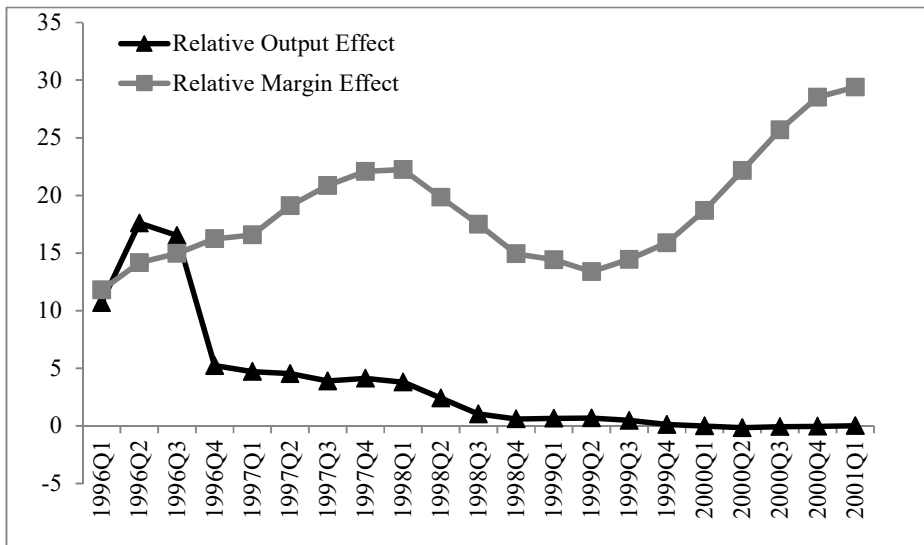


FIGURE 3D: COLUMN (4), TABLE 3

Source: Authors' calculations based on the IDC Quarterly PC and Server Tracker data matched to more detailed PC and Server characteristics and the estimated coefficients from Table 3.

Notes: These plot the evolution of calculated relative output and margin effects based on the formulas provided in Appendix D. When the relative margin effect lies above the relative output effect we predict that there are incentives to degrade interoperability.

FIGURE 3: RELATIVE MARGIN AND OUTPUT EFFECTS (ROBUSTNESS) - cont.




FIGURE 3E: COLUMN (5), TABLE 3




FIGURE 3F: COLUMN (6), TABLE 3




FIGURE 3G: COLUMN (7), TABLE 3





FIGURE 3H: COLUMN (8), TABLE 3

Source: Authors' calculations based on the IDC Quarterly PC and Server Tracker data matched to more detailed PC and Server characteristics and the estimated coefficients from Table 3.

Notes: These plot the evolution of calculated relative output and margin effects based on the formulas provided in Appendix D. When the relative margin effect lies above the relative output effect we predict that there are incentives to degrade interoperability.

FIGURE 4: RELATIVE MARGIN AND OUTPUT EFFECTS (FREE COMPLEMENTARITY)



Notes: This plots the evolution of calculated relative output and margin effects based on the formulas provided in Appendix D and their 90% confidence interval. The shaded area highlights the period where the relative margin effect is statistically higher than the relative output effect, hence Microsoft had significant incentives to degrade according to our model.

Source: Authors' calculations based on the IDC Quarterly PC and Server Tracker data matched to more detailed PC and Server characteristics and the estimated coefficients in column (9) of Table 3.