

המכון למחקר כלכלי בישראל
על-שם מוריס פאלק בע"מ (חל"צ)
The Maurice Falk Institute for
Economic Research in Israel Ltd.



The Dynamics of Technology Adoption and Vertical Restrictions: An Empirical Analysis

by

Alon Eizenberg, Andras Pechy and Michelle Sovinsky

Discussion Paper No. 17.01

May 2017

בנין פרץ נפתלי, קמפוס האוניברסיטה העברית, הר הצופים, ירושלים 9190501
The Hebrew University Campus, MT. Scopus, 9190501 Jerusalem, Israel
www.falk.huji.ac.il

The Dynamics of Technology Adoption and Vertical Restraints: An Empirical Analysis

By

Alon Eizenberg, Andras Pechy and Michelle Sovinsky*

Abstract

This paper studies the impact of vertical restraints in the x86 processor industry, where a dominant upstream supplier (Intel) competes with a smaller contender, Advanced Micro Devices (AMD). During the studied period, Intel's strategy included a controversial program, "Intel Inside", through which it offered downstream clients rebates and subsidies that were conditioned on the volume purchased from it and, sometimes, on the volume purchased from AMD. We document the manner by which such restraints interact with the dynamic process of downstream technology adoption. Our results indicate, first, that Intel's restraints were binding: restrictions imposed on a downstream client reduced the rate of its AMD adoption. Nonetheless, we also find that the adoption of the AMD technology by a given downstream client was negatively affected by restrictions imposed on other clients. Furthermore, adoption was an increasing function of both the intensity of antitrust litigation against Intel, and AMD's production capacity. These results highlight that a downstream client considers whether to adopt AMD's technology, in part, based on its perception regarding this supplier's ability to expand its capacity and meet high levels of demand in the future. The client may therefore be discouraged from adopting the AMD technology as a consequence of vertical restraints imposed on other clients, if those imply that AMD's cash flow and investment opportunities would be limited. The client may, analogously, be more likely to adopt if mounting litigation implies that Intel's restraints may soon be removed. Taken together, our results suggest that competition authorities need to pay particular attention to market dynamics and to their implications for the channels via which vertical restraints affect competition.

The Maurice Falk Institute for Economic Research in Israel Ltd.

Jerusalem, May 2017 • Discussion Paper No. 17.01

* Eizenberg: The Hebrew University of Jerusalem and CEPR; Pechy: Accenture; Sovinsky: University of Mannheim and CEPR. Toby Webber provided excellent research assistance. We are grateful for helpful comments from Jan Boone, Greg Crawford, Brett Gordon, Saul Lach, Konrad Stahl, Yona Rubinstein, and to seminar participants at DICE, European Economics Association meetings (Geneva), EARIE (Lisbon), ETH, German IO meetings (2016), Pompeu Fabra, Barcelona Applied IO Workshop, IDC Herzlia, Asian IO Meetings (Melbourne), Tel Aviv, Mannheim, the Hebrew University of Jerusalem, Edinburgh, Hamburg, and Tilburg. We also thank Brett Gordon for sharing his data with us for earlier years. Correspondence: alon.eizenberg@mail.huji.ac.il (Eizenberg), mssovinsky@econ.uni-mannheim.de (Sovinsky), apechy@gmail.com (Pechy). This project was supported by the Maurice Falk Institute for Economic Research in Israel.

1 Introduction

“There is perhaps no aspect of competition policy that is as controversial or has been as inconsistent over time and across jurisdictions as policy towards restraints between upstream firms and their downstream retailers” Lafontaine and Slade (2008)

Upstream manufacturers often impose exclusive dealing contracts on their retailers, which may result in foreclosure of a competing brand. On the other hand, there may be procompetitive effects of exclusive dealing. For example, exclusivity could enhance market performance by inducing a retailer to focus its promotional activities on the manufacturer’s products, and improve customer service. It could also secure investments made by the manufacturer (such as quality assurance and advertising) by preventing the retailer from “free-riding” on these investments.² Due to its potential procompetitive effects, the *per se* illegality characterization of exclusive dealing was rejected in *Standard Oil Co. v. United States* (Standard Stations), 337 U.S. 293, 305 -06 (1949) and the US courts treat exclusive dealing under the “rule of reason” legality rule.³ Motivated by these conflicting aspects of exclusive dealing contracts, a vast theoretical and empirical literature has sought to identify their impact given various market conditions.

One of the challenging tasks within this research agenda has been to empirically identify the foreclosure effect, i.e., the impact of vertical restraints imposed by one upstream firm on the sales of a competing upstream firm. An example of such work is Ater’s (2015) documentation of the negative effect of exclusive dealings on upstream rivals’ market shares in the fast-food industry. In general, empirical evidence on this issue is scarce. Our paper contributes to this literature while emphasizing an aspect that has not been addressed to date, to the best of our knowledge: the effect of vertical restraints on the dynamics of downstream technology adoption.

² Marvel (1982) and Besanko and Perry (1993) provide more discussion of this issue.

³ Potential pro- and anti-competitive effects of exclusive dealing and the history of its legal statues in the US are discussed in Areeda and Kaplow (1997) and Sullivan and Hovenkamp (2003). Exclusive dealing may violate the Clayton Act (Section 3) and the Sherman Act (Section 2). The rule of reason approach was reaffirmed in *Tampa Elec. Co. v. Nashville Coal Co.*, 365 U.S. 320 (1961)

Consider an upstream industry that sells critical components to downstream customers. The upstream industry features a dominant incumbent firm, and a smaller entrant, whose stated goal is to challenge the dominant position of the incumbent. Suppose further that this upstream industry is characterized by rapid innovation, large sunk investments, and capacity constraints. Downstream clients, for their part, crucially depend on timely shipments from the upstream industry.

In this setup, exclusive dealing contracts between the incumbent and major downstream clients may curb the downstream adoption of the rival's technology, effectively inhibiting its growth. Being excluded from selling to the major downstream customers, the smaller upstream supplier may struggle to expand its customer base and production volumes. This, in turn, may limit its ability to finance investments in Research and Development (R&D) and capacity expansion, perpetuating its inferior position. This mechanism is likely to affect clients' expectations regarding the smaller supplier's ability to deliver timely and reliable shipments, lowering the perceived long-run benefits from adopting its technology.

This paper estimates dynamic panel models that document such patterns in the semiconductor industry. Specifically, we study competition in the x86 microprocessor market between Intel, traditionally controlling about 80% of the market, and its rival AMD that controls much of the remaining 20% market share. Our analysis covers the years 2002-2009 and focuses on the role played by the "Intel Inside" program. Mostly throughout 2002-2005, this program funneled substantial payments from Intel to its downstream clients, i.e., PC manufacturers such as Dell, HP and Toshiba, in the form of rebates and advertising subsidies. As revealed in court cases and regulatory investigations, such payments were sometimes explicitly conditioned on the extent of the client's purchases from Intel's rival, AMD (see Lee, Pechy and Sovinsky 2013, hereafter LPS). Intel's arrangement with HP, for example, specified that the share of HP's business line PCs using AMD's chips was not to exceed 5%, while manufacturers such as Dell refrained from using AMD's chips altogether.

The "Intel Inside" program illustrates the basic dilemma faced by regulators with respect to exclusive dealings. On the one hand, excluding AMD from downstream manufacturers' product lines clearly raises anticompetitive concerns. On the other hand, the program may have had procompetitive aspects: it provided downstream manufacturers with incentives to

focus their entire production process on Intel’s technology, possibly creating economies of scale with some of the efficiency gains being passed to consumers. Fierce downstream competition may have, indeed, allowed consumers to enjoy a substantial share of such efficiency gains, and rebates on CPU purchases may have similarly ended up benefitting consumers. This program may have also facilitated Intel’s ability to capitalize on its investments, thus promoting innovation. Nonetheless, LPS find evidence that the advertising subsidies received by some of Intel’s large customers (notably Dell) were *predatory* in the sense that their magnitude appeared inconsistent with profit maximization. The controversial nature of the program was manifested in a series of complaints and lawsuits filed by AMD with antitrust authorities and courts worldwide, leading to active investigations and lawsuits filed by regulators. This process gained strong traction in the years 2004-2006. Ultimately, the bulk of these legal proceedings were settled, with Intel agreeing to roll back the controversial aspects of this program.

Ideally, empirical work should measure the social benefits and the social costs of the Intel Inside program. Such measurement is, however, outside the scope of the current paper, and our goal is different: we study the impact of vertical restraints on the downstream adoption of a rival’s technology within a dynamic environment. Our goal is not merely to quantify this impact, but rather to shed light on the channels via which this impact is realized.

Our empirical analysis combines several datasets. We document the downstream adoption of AMD’s technology using market level data on PC brands’ prices, characteristics and sales, indicating the brand-level share of PCs that had an AMD chip installed over time. We also collected data on the evolution of the upstream firms’ technology and capacity. Further, we draw on court cases and additional sources to construct variables that describe various aspects of the Intel Inside program, such as the volume of payments made to individual downstream clients, and specific restrictions dictated by Intel on their use of the AMD technology. Finally, we construct indices capturing the extent of litigation mounted by AMD and competition authorities worldwide in connection with the Intel Inside program. The joint variation of such variables over time allows us to identify the dynamic impact of technology, vertical restraints, and litigation, on the downstream adoption of AMD’s technology. Sharp variation in some of these measures, such as Intel’s decision to roll back much of the

controversial aspects of its Intel Inside program midway through the sample period, is very helpful in separately identifying these effects.

We estimate linear and nonlinear dynamic panel models in which the unit of observation is an individual downstream product line by market segment (e.g., HP’s “Pavilion” desktop for the home market), and where the dependent variable is the share of this product line’s sales that have an AMD chip installed. Our results indicate that the adoption of AMD’s technology by a given downstream client responds negatively to the extent of payments via “Intel Inside” to the downstream client itself, and to specific restrictions on the extent of usage of AMD’s technology imposed on the client by the Intel Inside arrangement. Such restrictions were, therefore, binding. Importantly, we also find that the rate of AMD adoption by a given client responds negatively to restrictions imposed on *other clients*. Two primary factors that have a positive effect on the rate of AMD adoption are AMD’s production capacity, and the extent of anti-“Intel Inside” litigation.

These findings reflect the importance of dynamics in the technology adoption process. Institutional details suggest that this decision is inherently dynamic. First, as is typical with technology adoption, the extent of current adoption affects the future costs of using the technology. Second, the extent of current adoption of the AMD technology may affect both current and future benefits granted to the client via the “Intel Inside” program. For example, increased adoption of AMD today may cause Intel to retaliate by withholding rebates and other benefits, or by being less responsive to the customer’s needs, and this effect may persist into future periods. Third, clients rely on timely shipments from their upstream suppliers, and may, therefore, be weary of increased adoption of a supplier’s technology, if they expect that supplier to have low capacity levels in the future. Ultimately, downstream customers need to weigh the potential benefits from adopting AMD’s technology against its costs, taking into account their expectations regarding two strategic variables: AMD’s production capacity, and the future viability of Intel’s vertical restraints.

Our results, as described above, are consistent with the importance of these dynamic factors. First, the documented negative response of a client’s AMD adoption to restrictions imposed by Intel on other clients indicates the role played by clients’ expectations. Although the exact contracts signed between Intel and a specific downstream client were not common

knowledge at the time, other clients may have had some information about them. For example, firms could easily observe the quantity of PCs with AMD processors sold by other firms, where this quantity was zero for major market players over many periods. If a client concludes or observes (albeit imperfectly) that other firms accept substantial restrictions on their adoption of the AMD technology, it may revise downward its expectations regarding AMD's future market position, investment and performance, resulting in the negative adoption effect that we document. Competition authorities should be aware of this subtle effect of exclusive contracts in dynamic environments, i.e., the fact that restrictions on one client may affect decisions by other clients.

Second, the finding of a positive effect of AMD's production capacity on the rate of its adoption is also consistent with the dynamic mechanisms illustrated above. Finally, the positive effect of increased litigation on the adoption of the AMD technology is consistent with downstream clients revising upward their expectations regarding the benefits of adoption given an increased likelihood that the restrictions will soon be lifted.

The rest of the paper is organized as follows. After a brief literature review, section 2 describes the data. Section 3 explains our empirical strategy, i.e., the applicability of dynamic panel models to the question at hand. Section 4 provides our results, and section 5 concludes with some discussion of limitations and avenues for future research.

Related literature. Our paper belongs in a small but growing empirical literature on exclusive dealing and vertical contracts (see Lafontaine and Slade 2008 for an overview). Asker (forthcoming) examines the effect of exclusive dealing on entry in the Chicago beer market. He finds that rivals do not have higher costs when facing competitors who sell under exclusive dealing agreements. Sass (2005) also studies the beer market and finds that exclusive dealing is more prevalent in smaller markets, in contrast to the predictions of foreclosure theory models. Nurski and Verboven (2016) estimate a structural model of demand with product and spatial differentiation and dealer exclusivity applied to the European automobile market. They find that exclusive dealing has served as a mild barrier to entry against Asian competitors, but with considerable consequences on consumers' domestic welfare because of reduced spatial coverage.

Our paper is closest to Ater (2015) who empirically quantifies the effect of exclusive deal-

ing contracts on sales in the fast food industry. He finds that exclusive dealing reduces sales, and concludes that this is inconsistent with efficiencies, so that exclusive dealing must be used for anti-competitive reasons. Our paper contributes to this line of research by considering a technology market where a dynamic relationship arises between exclusive dealings, downstream technology adoption, and upstream capacity investments.⁴

Our paper is also related to a large literature on the PC and CPU industries. Several papers study the nature of innovation in the x86 microprocessor industry. Some examples that rely on static structural models include Song (2007), who quantifies the benefits from such innovation, and Eizenberg (2014), who studies the impact of CPU innovation on the variety of downstream PC configurations. Gordon (2009) uses a dynamic demand model to study consumer replacement cycles, and Goettler and Gordon (2011) estimate a dynamic model in which innovation by Intel and AMD is endogenously determined, and use it to predict the impact of innovation from a hypothetical exclusion of AMD from the market. Our work differs from these papers by studying the dynamic adoption process of the AMD technology by downstream PC makers. Our empirical approach relies on dynamic panel methods rather than on structural modeling. While our approach limits our ability to analyze out-of-sample scenarios, it allows us to avoid some of the strong assumptions required in the estimation of dynamic games. In particular, our approach allows us to account for rich product-level characteristics without running into large state space concerns.

Another related paper is LPS who use a structural approach to estimate the marginal benefit to Intel from a dollar invested in the “Intel Inside” program. As this marginal benefit appears substantially lower than the marginal cost, LPS conclude that the subsidizing of advertising via this program has been predatory in nature. LPS do not consider PC firms’ product-line choices, focusing instead on Intel’s advertising decisions. The current paper, by contrast, explicitly addresses product-line choices.

Finally, a vast theoretical literature examines exclusive contracts and other vertical restraints. Exclusive dealing was initially considered an anticompetitive barrier to entry. However, the Chicago critique maintained that an exclusive deal is not in the joint interest of

⁴ Additional contributions include Slade (2000) and Suzuki (2009).

the upstream and downstream firm, thereby dismissing the view that exclusive contracts could be used to exclude a rival.⁵ The post-Chicago theoretical literature, in contrast, has identified conditions under which socially harmful exclusive contracts may arise. These are situations in which an incumbent and a retailer may have a joint incentive to contract on exclusive dealing as a means of foreclosing entry. The main insight is that such contracts imply externalities on other players not accounted for in the Chicago critique (see, e.g., Spector (2011) and Yehezkel (2008)).⁶

2 Data

We use data from several sources, containing information on PC and CPU sales and attributes, PC firms' advertising expenditures, measures of processor quality, and processor makers' production capacity. We also use information on the vertical restraints imposed by Intel, and on legal action taken against these restraints.

2.1 Sales, Attributes and Advertising

We use quarterly data on PC sales in the US home and business sectors available from the Gartner Group, covering the years 2002-2009.⁷ A unit of observation is defined as a combination of PC vendor (e.g., Dell), PC vendor brand (e.g., Inspiron), market segment

⁵ First, if offering a second brand increases the retailer's profit, then the manufacturer can charge the retailer higher franchise fees. Therefore, if a manufacturer finds it profitable to foreclose a competing brand then it has to be that this brand is not profitable to begin with. Second, if for whatever reason a manufacturer wishes to foreclose a competing brand, then the manufacturer can choose between imposing exclusive dealing on the retailer, or offering him quantity discounts to induce the retailer to choose not to carry the competing brand. Either way, the manufacturer needs to compensate the retailer for the foregone profits from offering the competing brand. Thus it is not clear why exclusive dealing is any better from the manufacturer's viewpoint than quantity discounts that are less open to antitrust scrutiny. Third, the fact that the retailer has the option to carry the competing brand will force the manufacturer to offer discounts that the retailer is likely to pass on, at least partially, to consumers. In that sense, the competitive pressure from the competing brand holds even in the presence of exclusive dealing. As Gilbert (2000) points out, the arguments made by the "Chicago School" parallels a more tolerant approach by US courts towards exclusive dealing.

⁶ For related theoretical contributions, see for example Inderst and Shaffer (2007) and Lommerud, Straume and Sogard (2006).

⁷ We do not include servers as server sales were not recorded in the Gartner dataset prior to 2005.

(e.g., Home), CPU vendor (e.g., Intel), CPU family (e.g., Pentium 4) and quarter. We focus our analysis on the home and business segments.⁸ We exclude Apple products as those exclusively used IBM's chips during much of the sample period (using Intel's chips afterwards).

	Obs	Mean	Std. Dev.	Min	Max
Percentage AMD Sold	3508	0.13	0.25	0	1
PC Characteristics					
Price PC (1000\$)	3508	1.02	0.45	0.24	3.69
Brand Advertising (M\$)	3508	1.01	3.10	0.00	29.14
Firm Advertising (M\$)	3508	8.16	16.46	0.00	87.80
CPU Characteristics					
AMD CPU benchmark/dollar, if non-zero	1234	7.31	4.54	1.75	22.65
Intel CPU benchmark/dollar, if non-zero	3349	4.69	3.85	0.84	32.66
Quarters Brand/AMD family available, if non-zero	1234	5.42	3.20	1	19
Quarters Brand/Intel family available, if non-zero	3349	6.42	3.76	1	30
Capacity Related Variables					
Free Cash (100M\$)	3508	8.31	3.12	3.97	19.05
AMD Capacity Index	3508	8.08	3.59	3	13
Intel Capacity Index	3508	31.93	6.71	23	44
Exclusionary Restriction/Antitrust Related Variables					
Exclusionary Restriction Index	3508	1.03	2.03	0	6
Intel Payments to Dell (M\$)	3508	147.38	154.15	13.37	603.05
Intel Payments to each PC Firm (M\$)	3508	36.83	95.24	0.00	603.05
Cumulative Antitrust Cases Against Intel	3508	4.19	2.15	1	7
Pending Antitrust Cases Against Intel	3508	3.45	1.62	1	6

Table 1: Descriptive Statistics (PC Brand Segment Quarter Level)

Table 1 shows descriptive statistics at the PC brand-segment-quarter level. The total number of observations is 3,508. As the table indicates, the rate of utilization of AMD's chips, averaged across observations, is 13 percent. PC prices display significant variation, ranging from \$240 to \$3,690. We also create a variable to account for how long a PC brand has been available equipped with AMD chips (respectively Intel), counting the number of quarters.

⁸ All variables expressed in monetary terms were deflated using the quarterly consumer price index of the Bureau of Labor Statistics, basis set at the year 2000 USD.

The evolution of the principal variable, AMD’s market share, is shown in Figure 1 (overall market share across PC brands and segments, the remaining share garnered by Intel). The rise in market share during the year 2006 from 10 percent to 20 percent is of interest. It seems indeed to be closely timed with a loosening of the restrictions of the Intel Inside program, discussed below and illustrated in Figure 4. The gain in market share is however not maintained as it declines after 2008 to reach 12.5 percent by 2009.

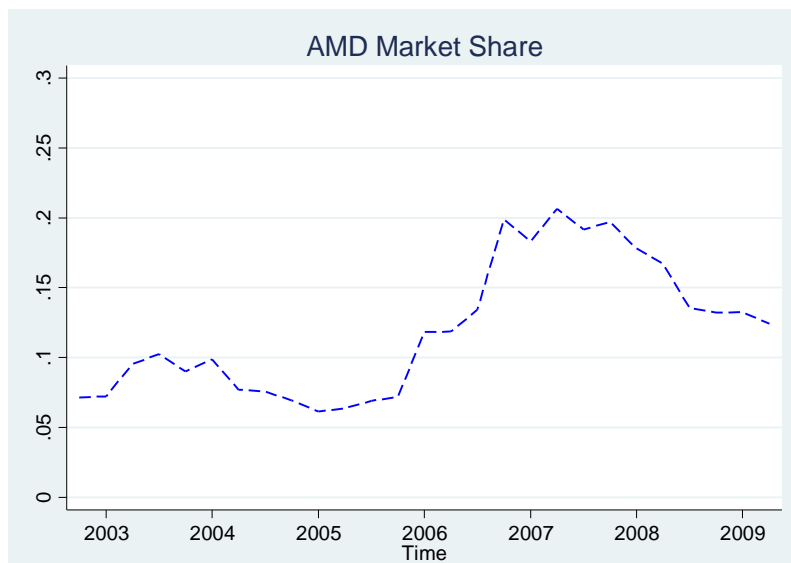


Figure 1: AMD Market Shares

Figure 2 provides client-specific information regarding the adoption of the AMD technology over time. It shows AMD’s market share within selected PC firms (the remainder being Intel’s market share) and the share of these PC firms in the market. We display mean (over quarters) market shares, before and after the first quarter of 2006, when most major restrictions were lifted by Intel. The figure reveals that the increase in AMD’s market share did not only come from those PC firms that were directly subject to Intel’s exclusionary restrictions. For example, AMD’s adoption by Acer, that was not subject to these restraints, went up considerably. This descriptive evidence is consistent with one of our formal findings: that adoption of the AMD technology by a given downstream client was affected not only by restrictions imposed directly on it, but also by restrictions imposed on other clients.

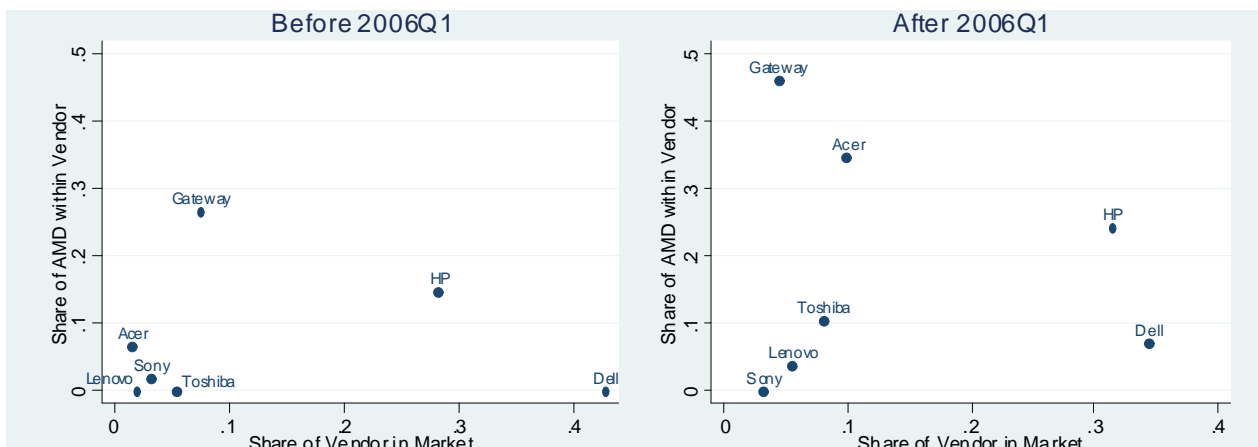


Figure 2: PC Vendor Specific Market Shares

PC advertisement data come from the Kantar Media Group. These data are important because many rebates given by Intel were a function of advertising done by PC firms. These consist of PC brand-specific ad expenditures (e.g. Acer Aspire) and PC firm level ad expenditures (e.g. Acer) where we match sales and advertising data across brands.⁹ Table 1 shows that expenditures for brand-specific ads averaged \$1 million while firm-specific expenditures averaged \$8.2 million (the observed maximum of \$87 million was due to exceptionally large expenditures on TV advertisements).

2.2 CPU Quality

We obtain data on the quality of the CPU from Passmark’s CPU Mark publications.¹⁰ This company collects measurements on CPU tests from users around the world, and creates a database of CPU performance at the CPU model level. It provides a benchmark score: a continuous measure of performance for each CPU model. These contain a continuous quality measure (“benchmark”) for each CPU model. We use this together with CPU prices, which we gathered from published list prices or obtained from Instat, to construct a benchmark

⁹ For a few PC firms, the Kantar brands were available at a more aggregate level than the Gartner sales data. Thus, the match occurred at the Kantar brand level.

¹⁰ www.cpubenchmark.net

per dollar spent index.¹¹ We view this as a natural measure of the value delivered to customers by both processor makers. To be clear, these measures capture the sales weighted average of the benchmark per dollar for AMD and Intel. The weighting across different Intel (respectively, AMD) chips is done based on the percentage of sales of the specific PC product line that have those chips installed. A challenge then arises with respect to the calculation of the AMD benchmark-per-dollar measure in product lines that were exclusive with Intel. We solve this issue by substituting the average AMD benchmark-per-dollar across all product lines. Complete details regarding the construction of those series are provided in the Appendix.

The descriptive statistics provided in Table 1 were computed on observations with positive adoption of AMD only. The Table shows that AMD CPUs have, on average, a 35% higher benchmark per dollar measure than Intel CPUs. Figure 3 shows the evolution over time of the sales weighted average of the benchmark per dollar measure for both firms. AMD's quality measure was consistently better than Intel's throughout much of the sample period, both due to lower AMD prices and higher AMD benchmark values.

Nonetheless, as of 2006, Intel's benchmark-per-dollar measure experienced much faster growth than AMD's, so that by the end of the sample period, both companies were neck-to-neck in terms of this indicator. This development has had a lot to do with the introduction of new generations of Intel chips (specifically, the Centrino platform and its successors) that offered substantial improvements over incumbent generations. This pattern is of interest, as it suggests that Intel has been enjoying different competitive advantages during different sample periods. While in the first part of the sample it enjoyed the ability to impose restraints via the Intel Inside program, in the second part of the sample it has largely shifted away from such restraints, but at the same time started to enjoy a substantial technological advantage (noting that Intel's pricing was consistently higher than that of AMD's, a tie in terms of benchmark per dollar implies a benchmark advantage for Intel).

¹¹ Instat "Intel Rosetta Stone: Intel Processor Shipments, Forecasts, Technology and Roadmaps" November 2005.

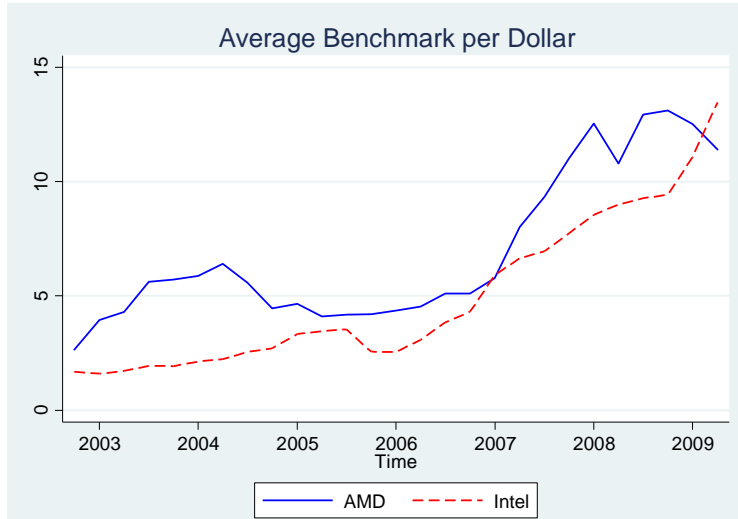


Figure 3: CPU Quality Per Dollar

2.3 CPU Production Technology and Capacity

To measure the processor makers’ production technology and capacity, we use information from their annual reports. These report the number of fabrication units available for making microprocessors (FABs), the silicon wafer size of each FAB (the larger the wafer, the more CPUs can be printed simultaneously), and the precision in nanometers available at each FAB for “printing” the integrated circuits (the smaller the precision, the more CPUs can be printed, and additionally, CPU power efficiency is improved). We construct a measure that captures both the capacity and the quality associated with Intel’s (respectively, AMD’s) production process, the *capacity index*, as the sum over all FABs of (ordered) wafer size and (ordered) IC process. Table 1 shows that the mean capacity index of Intel is, on average, four times larger than AMD’s. This is indicative of Intel’s large capacity advantage, often described as one of its main strategic advantages over AMD. We also obtain a measure of AMD’s cash flow that is available for investment at the beginning of each quarter from SEC quarterly reports.¹² Table 1 reveals that the free cash available for AMD in each quarter is on average \$831 million. We focus on AMD’s cash flows as it is well accepted that AMD — and not Intel — encountered challenges in financing investments in new production

¹² The quarterly filings were accessed on September 18, 2014 from <http://ir.amd.com/>

facilities. We provide more institutional details on this issue, and its importance, when discussing below the results of our formal regression analysis.

Table 2 shows the evolution of the processor makers' production technology and capacity over time. AMD usually lags behind Intel regarding the IC process and the wafer size. The variation in Intel's number of FABs is due to technological upgrades and relocations.

Year	Number of Fabs		Mean ICP in nm		Mean wafer in mm		Capacity index	
	AMD	Intel	AMD	Intel	AMD	Intel	AMD	Intel
2002	1	10	130	150	200	220	3	28
2003	1	7	130	130	200	229	3	23
2004	1	7	130	113	200	243	3	27
2005	2	6	90	78	250	300	9	33
2006	2	5	90	75	250	300	9	28
2007	2	5	78	57	250	300	10	32
2008	2	7	65	60	300	300	12	44
2009	2	6	55	50	300	300	13	41
2010	2	8	45	48	300	300	14	56

Notes 1. ICP stands for Integrated Circuit Process
2. The Fab capacity index is computed by ranking IC process (largest to smallest) and wafer size (smallest to largest) , then summing these points over all fabs

Table 2: Evolution of Capacity Variables

2.4 Exclusive Restrictions

We examined case files from the Department of Justice and the European Commission that listed the types of exclusive restrictions Intel imposed on downstream PC firms. The restrictions span a wide variety of instruments via which Intel could discourage clients from using AMD's chips. These include: caps on the amount sold of AMD-based products, or restrictions on sales in specific segments/product lines; restrictions on the distribution channels that could be used to sell AMD-based products; provision of rebates in exchange for selling certain amounts of Intel-based machines; limitations on the marketing PC firms could undertake for AMD-based products; guarantees of preferred supply of Intel CPUs; and restrictions imposed on bidding on contracts using AMD-based products. The documents also detail threats made by Intel to certain PC firms to remove funding, divert funding to rivals, or other retaliation, as a punishment for selling more than the specified amount of AMD-based PCs.

We form two indices to capture the presence of these restrictions. The first measure is the number of restrictions imposed on a downstream client. The second measure only counts restrictions we define as extreme: excessive rebates, demands to exclude AMD from certain product lines completely, threats, or a promise to increase Intel market shares provided to the PC firm.

Table 1 reports statistics on the first measure, and reveals that, on average across brand-segment-quarters, 1.03 restrictions were in place while the maximum was 6. Figure 4 shows the number of brand-segment data cells that were affected by various values of the exclusionary restrictions index in each quarter. As can be seen, most of the restrictions took place before 2007, and were rapidly eliminated thereafter. This sharp variation will be useful in identifying the effect of these restrictions in our econometric model.

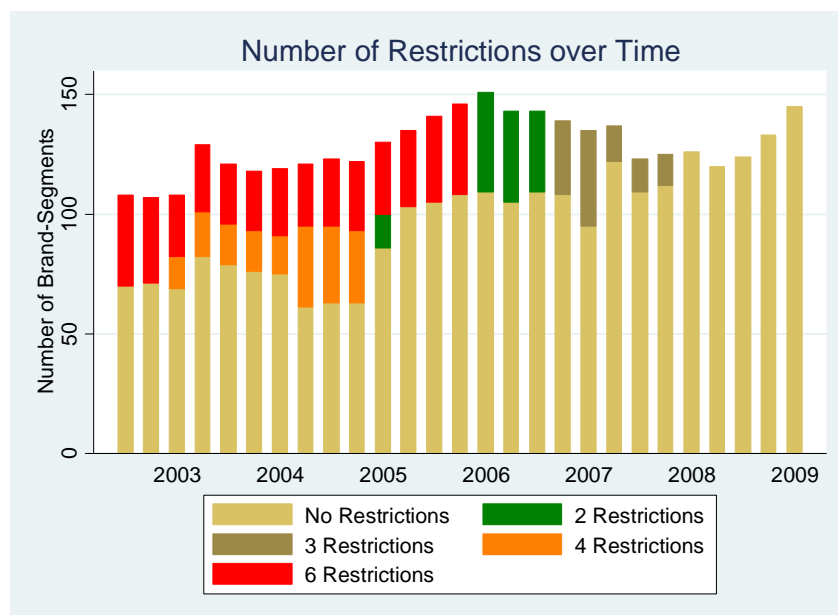


Figure 4: Evolution of Restrictions Imposed by Intel

We also observe rebates offered by Intel to PC firms via the “Intel Inside” program. The exact amounts paid to Dell were available from the court decision SEC vs. Dell Inc. July 22 2010, covering the years 2003-2007. These documents indicate that the payments to Dell were substantial and deviated considerably from the official description of the Intel Inside

payments provided on Intel’s website. To calculate payments to Dell after 2007, we refer to this official description, which stated that 3 percent of the CPU costs will be rebated to the PC maker to finance ads for PC models equipped with Intel CPUs.¹³ Hence, for the period after 2007, the “Intel Inside” payments are computed as 3 percent of Dell’s CPU costs, computed using Gartner sales data and the price dataset we described previously. An identical approach, using a 3 percent rebate, was used to compute the “Intel Inside” payments received by other PC firms throughout the sample period. The variable is defined at the firm level and summed over all brands and segments.

Figure 5 provides an overview of these “Intel Inside” payments. Payments to Dell are displayed on the left axis and the average payments to the other PC firms are on the right axis. Payments to Dell during 2003-2007 were close to 100 times the payments the company was supposed to receive based on the advertised 3 percent rebate. Examining payments to other PC firms, it appears that the “Intel Inside” program conveyed non-negligible benefits to the thin-margin PC makers: these payments averaged between \$2.3 and \$5.5 million per quarter per firm.

2.5 Legal Action

We compiled information on investigations or legal action taken against Intel relating to its practices against AMD. We construct two variables that capture this legal action, as reflected in lawsuits filed by the Federal Trade Commission, European Commission, Korean Fair Trade Commission, Japanese Fair Trade Commission, and the State of New York. The source for this information is Intel and AMD’s shareholder reports. We include measures of the cumulative number of antitrust cases/investigations brought against Intel as of 2001 and the number of pending antitrust cases/investigations in process against Intel. These variables may capture the extent to which market participants revise their beliefs regarding the future viability of Intel’s restraints. Table 1 shows that, on average across all observations, there were 3.45 pending cases, and 4.19 cumulative cases.

¹³ For further details about the program, see for example Lee, Péchy and Sovinsky (2013).

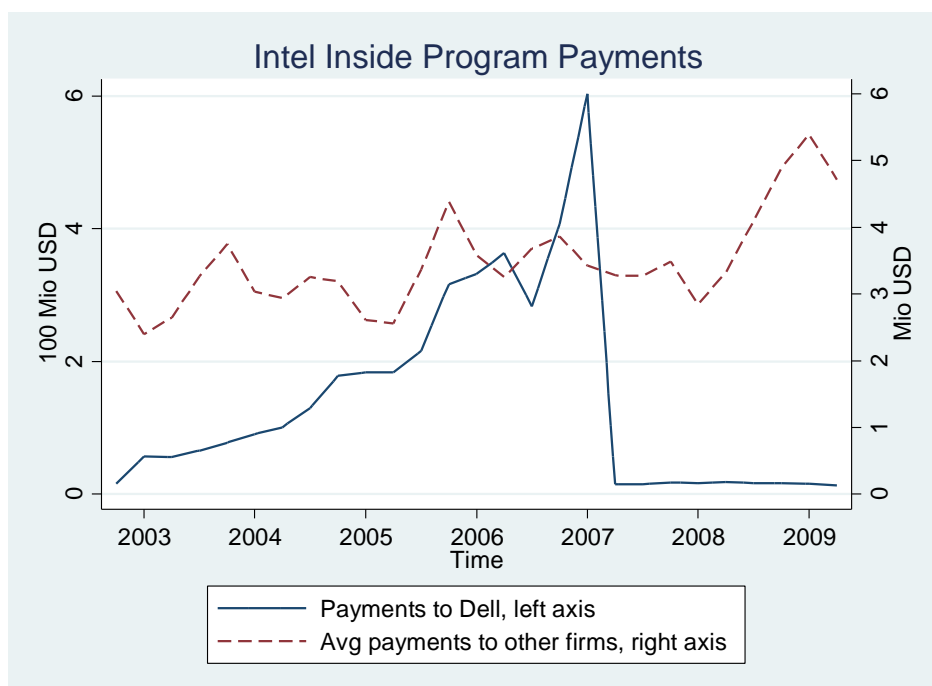


Figure 5: Evolution of Intel's Payments to PC Firms

3 Econometric Model

3.1 Preliminaries

Our goal is to identify the effect of vertical restraints imposed by an upstream supplier (Intel) on downstream input adoption decisions. To this end, we set up an econometric model that treats a PC brand-segment-quarter combination as the unit of observation, and defines the extent of the usage of AMD chips as the dependent variable. Specifically, the dependent variable is the share of units sold in the PC brand-segment data cell that have an AMD chip installed. Defining the dependent variable in this fashion allows us to link our analysis to the nature of Intel's vertical restraints, which often specified a cap on the extent of AMD chips used in each segment as a condition for eligibility to the Intel Inside program benefits. Key explanatory variables of interest are those that capture the extent of Intel's restrictions, the extent of litigation against those restrictions, AMD's production capacity, and the benchmark-per-dollar performance of the two processor vendors.

A major challenge to the analysis is the dynamic nature of technology adoption decisions, stemming from several reasons. First, beginning to use AMD chips (or expanding the extent of their adoption) requires a certain degree of investment by the PC maker. In particular, the PC maker must learn how to configure the hardware to AMD's specifications, and set up (or expand) a production line that installs AMD chips. Since AMD and Intel chips are not "pin-compatible," it is not possible to simply plug one of them in instead of the other, and certain adjustments need to be made. Such an investment in the learning and internal organization is not likely to be a static decision, but rather a cumulative process. Importantly, the infrastructure created in a given quarter (where by "infrastructure" we mean the accumulated know-how, experience and physical aspects of an AMD-based production line) is likely to reduce the cost of employing the AMD technology in future periods.

The adoption decisions are best viewed, then, as a dynamic process in which current investment decisions are taken given expectations regarding future market conditions, formed based on observing current values of state variables, some of which are endogenous. The state-dependent nature of these decisions manifests itself in several fashions. In particular, the endogenous adoption decision at time $t - 1$ affects the adoption decision at time t , motivating the consideration of dynamic panel methods.

Another source of dynamic links therefore involves the expectations formed by the PC maker at time t regarding market conditions in future periods. One such market condition is the evolution of supplier capacity. For a PC maker to be willing to expand its reliance on AMD chips, it must believe that AMD would be able to meet its level of demand. PC makers rely on thin inventories, making this issue crucial for the creation of strategic dependence on AMD's chips. Capacity constraints are, in general, an important aspect of the integrated circuit industry. Intel enjoyed a substantial production capacity advantage.

Constructing chip production facilities is extremely expensive, and with low sales, AMD lacked the financial capabilities to finance such investments. Institutional details provide ample support for this possibility. AMD's struggle to finance capacity investments are vividly described, for instance, in a quote from an AMD executive who left the company over disagreements related to the construction of AMD's FAB in Dresden, completed in 2000: "(t)he trouble in the entire economic model was that AMD did not have enough capital to

be able to fund fabs...(t)he point at which I had my final conflict was that (they) started the process of building a new fab with borrowed money prematurely. We didn't need a fab for at least another year. If we had done it a year later, we would have accumulated enough profits to afford the fab in Germany. He (referring to AMD's CEO at the time) laid the foundation for a fundamentally inefficient capital structure that AMD never recovered from."¹⁴

Similar views are present in blogs that cover the industry in the relevant years. For example, the following opinion was posted by a blogger following the industry in 2002: "...I think that AMD knows that if they do only what they have announced in terms of their capacity expansion road map, they will allow Intel to retreat into the part of the market AMD can't supply, lick their wounds, and buy/or finish developing technology that can compete with AMD in a year or two."¹⁵ In other words, it seems that AMD's insufficient capacity may have prevented it from fully enjoying the benefits from its ability to introduce innovative chips that, at least at some point, may have provided superior value to that of Intel. Recalling the evolution of our benchmark-per-dollar measures from Figure 3, the blogger's prediction may have, in fact, materialized: while AMD did offer better value than Intel in 2002, Intel was able to eventually recuperate and regain its technological edge in later years. AMD's lack of production capacity may have contributed to this development.

The market's expectations regarding AMD's ability to expand its capacity, therefore, may have played a crucial role in clients' adoption decisions. We address this issue by including measures of AMD's (Intel's) capacity as explanatory variables. This allows important events, such as the opening of a new AMD production facility, to affect current decisions by downstream clients. To capture the effect of the ability of Intel and AMD to provide value to customers, we also include our benchmark-per-dollar measures for these companies as explanatory variables.

Intel's vertical restraints play a major role on the right-hand-side of our econometric model. First, they directly affect the extent of adoption of AMD chips. Second, they also

¹⁴ Source: "The rise and fall of AMD: How an underdog stuck it to Intel," arstechnica.com, April 2013 (accessed on March 9th, 2017). Text in paranthesis added by the authors.

¹⁵ Source: "AMD's Future Fab Capacity," a January 2002 post by ValueNut on the online community "The Motley Fool (<http://www.lnksrv.com/community/pod/2002/020122.htm>, accessed on March 9th 2017).

may affect expectations. If Intel aggressively increases payments via the program, it may signal to the market that it is determined to block AMD from expanding. If this signal is credible, it may discourage current adoption of AMD chips. We therefore include as explanatory variables not only the vertical restraints imposed on the PC firm in question, but also measures that capture the overall scope and magnitude of Intel’s restraints and subsidies. The rationale is that the rate of AMD adoption by one PC firm can be affected by restrictions imposed on other PC firms, as those affect the firm’s expectations regarding the future standing of AMD in the market.

Finally, mounting litigation against Intel over its Intel Inside practices also affected expectations. It may have signaled to the market that Intel will soon be forced to pull back the exclusionary restrictions, increasing the expected value of adopting the AMD technology. This expected value increases both because fear of Intel’s retaliation against the customer is alleviated, and because the easing of the pressure on AMD increases its ability to invest in capacity expansion and perform as a reliable and effective supplier in the future. We therefore include our measures of legal action against Intel as explanatory variables. We lag these variables by one quarter, as it may take a certain amount of time for the market to respond to news regarding such litigation.

3.2 Formal econometric setup

Our unit of observation, indexed by i , is the PC brand-quarter-segment data cell (e.g, Acer’s Aspire Home Market), where the fraction of usage of AMD chips at time t is the dependent variable, W_{it} . Our model for the PC firm’s choice is given by

$$W_{it} = \alpha W_{it-1} + \beta x_{it} + \lambda c_t + \eta m_{it} + \delta r_{it} + \gamma l_{it} + \mu_i + \varepsilon_{it}. \quad (1)$$

Following the above discussion, W_{it-1} , the lagged percentage of segment-brand i sold with an AMD processor, serves as an explanatory variable to help us capture the dynamic link between current and future decisions. The μ_i term represents brand fixed effects and ε_{it} is an idiosyncratic iid error term.

Time varying observed characteristics of PC brand i are captured by x_{it} .¹⁶ These include the price of brand i , brand-level advertising, and firm-level advertising. Given that often more than one product is available within the brand, when the component of x_{it} varies across such products we compute the weighted average within the brand.

Variables relating to CPU manufacturing capacity constraints are included in c_t . To capture upstream production technology investments (along with their effect on expectations), we include a capacity index for AMD and for Intel, which is a function of the number of FABs, wafer size and IC process (as described in section 2), as well as the (lagged) amount of free cash available to AMD for investment. Our assumption, supported by institutional details, is that Intel was much less subject to cash constraints than AMD. Measures of capacity refer to worldwide capacity for production, while our variable of interest relates to sales in the US, which mitigates endogeneity concerns.

Several possible strategies present themselves with respect to including lags and leads of these variables. One could argue that leads of capacity should be included, since the knowledge that AMD is about to expand its capacity may also affect decisions. On the other hand, one may argue that announcing that the new IC process would become operational on a given date, and announcing that it is in fact operational on a given date, are very different pieces of information, which may justify using lagged values but not leads. For this reason, we estimate a variety of specifications with leads and lags of capacity variables (noting however that, while the cashflow variable varies quarterly, the capacity index is computed based on Annual reports so quarterly leads and lags deliver limited variation).

CPU related variables are given in m_{it} . These include the extent of technological progress as measured by the benchmark per dollar indices for Intel and AMD, respectively, and the number of quarters the segment-brand-CPU family combination for those families sold by AMD (resp., Intel) has been available.¹⁷

The vector r_{it} captures upstream restrictions and can be conceptually decomposed into

¹⁶ We do not include product characteristics that are not time varying within brands (such as the platform/form factor) because these are absorbed by the brand fixed effect term.

¹⁷ We do not have information prior to 2002 so the number of quarters available is counted starting from the first quarter in 2002.

two parts: the part related to exclusive restrictions on PC sales imposed by Intel, and the part related to payments via the Intel Inside program. Identifying the impact of the upstream restrictions on product line choice is the main goal of our analysis. Specifically, r_{it} contains total payments Intel made to the relevant PC maker, payments made to other downstream firms, indices that capture the intensity of exclusionary restrictions imposed by Intel (as discussed in section 2), and the number of such restrictions imposed on rival firms.

The vector l_{it} captures legal issues that may impact whether AMD is sold. These include the cumulative number of antitrust cases/investigations brought against Intel as of 2001, and the pending number of such cases.

Identification. Our goal is to identify the causal effect of upstream vertical restraints on downstream product line choices. Our identification strategy takes advantage of the dynamic panel aspect of our data. This approach is valuable for a number of reasons. First, a dynamic panel can help us address unobserved firm heterogeneity in the decision to adopt AMD’s chips. Unobserved heterogeneity may arise if some firms are fundamentally better-suited to gain from using AMD’s chips. For example, some firms are positioned as “value PC” makers, and may thus be more likely to offer PCs based on the cheaper AMD chips. Firms may also differ in the flexibility of their production processes: for example, some firms may enjoy large economies of scale from using a single type of chip, making them less likely to adopt the AMD technology in addition to that of Intel’s.

A particularly important issue is the endogeneity of Intel’s vertical restraints. Again, the panel structure of the data is crucial for identification in this context. It enables the employment of a fixed-effect estimator, removing time-invariant unobserved firm heterogeneity. Thus, the effect of the vertical restraint on the outcome variable, (as captured in δ) is causal and is identified by time-series variation. Note that we need time-series variation in the use of exclusive constraints, which is present in the data.

We estimate specifications that address different aspects of the data. We estimate both linear and nonlinear dynamic panel versions of the relationship in equation (1). These methods allow us to explicitly account for the dynamic role played by the past-adoption decision W_{it-1} , and to rely on fixed effects to alleviate endogeneity concerns. As explained below, the linear approach allows us to additionally include instrumental variables. The

nonlinear approach, in contrast, allows us to address the mass point at zero characterizing the dependent variable, arising from situations where firms buy exclusively from Intel.

Linear specifications. Our linear model treats μ_i as a fixed effect to be differenced out following Arellano and Bond (1991). We obtain a first-differences regression of the form:

$$\Delta W_{it} = \alpha_L \Delta W_{it-1} + \beta_L \Delta x_{it} + \delta_L \Delta r_{it} + \lambda_L \Delta c_t + \eta_L \Delta m_{it} + \Delta \varepsilon_{it}. \quad (2)$$

One could assume that the variables represented by $y_{it} \equiv \{x_{it}, r_{it}, c_t, m_{it}\}$ are strictly exogenous. That is, that ε_{is} is independent of y_{it} for any t, s conditional on μ_i . Strict exogeneity implies Δy_{it} is exogenous and hence Δy_{it} can serve as its own instrument. However, Δy_{it} may be a weak instrument.¹⁸ Another solution to potentially weak instruments is to assume that ε_{it} is independent over time, in which case the (two period) lag of the endogenous variable(s) is a valid instrument.¹⁹

However some of the regressors, even if independent of current disturbances, may be influenced by past ones. These regressors are then not strictly exogenous but rather exhibit sequential exogeneity where $E(\varepsilon_{it} | y_{is}, \mu_i) = 0$ for $s \leq t$. Relaxing the assumption of strict exogeneity implies Δy_{it} is endogenous. In this case one can use W_{it-2} and y_{it-1} as IV in the first-differenced equation.²⁰ We estimate specifications allowing for sequential exogeneity.

Note that estimators using too many lags of explanatory variables (relative to the number of observations) as instruments are known to have poor finite sample properties (Arellano and Bover (1995), and Blundell and Bond (1998)). In practice it is best not to use lags back to $t = 1$. We follow this approach and consider specifications using only lags of three periods.

Nonlinear specification. The linear specifications above allowed us to control for both state dependence, and for individual heterogeneity, in explaining the rate of adoption of AMD’s chips. These specifications, however, did not address the “corner solution” issue:

¹⁸ One can test the null hypothesis that the overidentifying restrictions hold using a Wald statistic which is valid under heteroskedasticity and clustering. The critical value is $\chi^2(l)$, where l is the degree of overidentification.

¹⁹ In this case it is important to test the null hypothesis that ε_{it} is independent over time. We implement an autocovariance test of the null hypothesis of no autocorrelation in the idiosyncratic error term.

²⁰ We can test for weak instruments using the standard first stage regression results: if y_{it-1} are not weak instruments then they should affect W_{it-1} conditional on y_{it} . Again, we test for serial correlation in the errors.

many product lines, at different times, chose not to use AMD chips at all.

To address this issue, we follow Wooldridge (2002, 2005) in specifying a dynamic non-linear model that builds on Chamberlain (1984). This allows us to include random effects that capture time-constant heterogeneity, as well as a lagged dependent variable capturing state dependence. However, this framework does not allow us to relax a strict exogeneity assumption. Both the linear and the nonlinear models therefore have their specific strengths and weaknesses, and considering both approaches allow us to provide a more complete picture of the data patterns of interest. As we shall see, our main conclusions hold across both analyses.

This specification treats the dependent variable W_{it} as a continuous measure with a mass point at zero given by

$$W_{it} = \max(0, \alpha_{NL}W_{it-1} + \beta_{NL}x_{it} + \delta_{NL}r_{it} + \lambda_{NL}c_t + \eta_{NL}m_{it} + \mu_i + u_{it}) \quad (3)$$

$$u_{it} | (\bar{y}_i, W_{i,t-1}, \dots, W_{i0}, \mu_i) \sim N(0, \sigma_u^2), \quad (4)$$

where, as in the linear specification, we denote by y_{it} the collection of all the explanatory variables in all time periods. The mean (over time) of these variables is \bar{y}_i .

One issue concerns the initial value of W_{i0} . One possibility is to treat it as nonrandom, which would imply that μ_i and W_{i0} are independent. However, this may not necessarily be the case, and so we follow the suggestion in Wooldridge (ibid.) and specify the density of the fixed effect conditional on the initial value. That is, we specify the fixed effects as

$$\mu_i = \psi + \xi_0 W_{i0} + \bar{y}_i \xi + a_i, \quad a_i | (W_{i0}, \bar{y}_i) \sim N(0, \sigma_a^2). \quad (5)$$

The fixed effects can then be integrated to yield the likelihood function of the random effects Tobit model with time- t , observation- i explanatory variables: $(y_{it}, W_{i,t-1}, W_{i0}, \bar{y}_i)$. That is, \bar{y}_i and w_{i0} are controlled for in each time period. This likelihood function is used to obtain estimates of the parameters $(\alpha_{NL}, \beta_{NL}, \delta_{NL}, \lambda_{NL}, \eta_{NL}, \psi, \xi_0, \xi, \sigma_a^2)$.

4 Results

Table 3 presents results from the linear specification. Recall that an observation corresponds to a PC brand-quarter-segment combination, and that the dependent variable is the fraction of sales of that brand-quarter-segment that have an AMD processor installed.

The different specifications correspond to different included combinations of variables capturing Intel's exclusionary restrictions, and the litigation prompted by those litigations. Across these specifications, the estimates of α , the state-dependence factor, are on the order of 0.7-0.75, noting that values between zero and one are considered valid.

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Lagged Percentage AMD Sold	0.7303*** (0.0701)	0.7248*** (0.0708)	0.6994*** (0.0737)	0.7409*** (0.0678)	0.6987*** (0.0741)	0.7128*** (0.0753)	0.7204*** (0.0756)	0.7440*** (0.0657)	0.7438*** (0.0659)	0.7590*** (0.0721)	0.7010*** (0.0717)	0.6999*** (0.0722)
PC Characteristics												
Price PC (100\$)	-0.0041 (0.0041)	-0.0046 (0.0042)	-0.0068 (0.0047)	-0.0062 (0.0040)	-0.0097** (0.0047)	-0.0052 (0.0048)	-0.0059 (0.0047)	-0.0029 (0.0039)	-0.0031 (0.0040)	-0.0024 (0.0042)	-0.0064 (0.0045)	-0.0066 (0.0046)
Brand Advertising (10M\$)	-0.0059 (0.0036)	-0.0063 (0.0038)	-0.0069 (0.0044)	-0.0035 (0.0039)	-0.0060 (0.0051)	-0.0075* (0.0041)	-0.0067 (0.0041)	-0.0047 (0.0034)	-0.0044 (0.0034)	-0.0065* (0.0034)	-0.0071 (0.0044)	-0.0069 (0.0044)
Firm Advertising (10M\$)	-0.0023* (0.0014)	-0.0036** (0.0016)	-0.0024 (0.0015)	-0.0033** (0.0014)	-0.0029* (0.0016)	-0.0024* (0.0014)	-0.0011 (0.0014)	-0.0027* (0.0014)	-0.0028** (0.0014)	-0.0030** (0.0015)	-0.0024 (0.0015)	-0.0024 (0.0015)
CPU Characteristics												
AMD CPU benchmark/dollar (10,000\$)	-0.0037** (0.0018)	-0.0038** (0.0018)	-0.0038** (0.0019)	-0.0050*** (0.0021)	-0.0053** (0.0021)	-0.0043** (0.0019)	-0.0048** (0.0019)	-0.0035* (0.0018)	-0.0036** (0.0018)	-0.0048*** (0.0018)	-0.0038* (0.0020)	-0.0039** (0.0020)
Intel CPU benchmark/dollar (10,000\$)	-0.0025 (0.0020)	-0.0026 (0.0020)	-0.0034 (0.0022)	-0.0037* (0.0022)	-0.0049** (0.0022)	-0.0027 (0.0022)	-0.0031 (0.0022)	-0.0021 (0.0019)	-0.0022 (0.0020)	-0.0018 (0.0021)	-0.0033 (0.0021)	-0.0034 (0.0021)
Number Quarters Brand/AMD family available	0.0122*** (0.0033)	0.0123*** (0.0033)	0.0131*** (0.0034)	0.0116*** (0.0033)	0.0127*** (0.0033)	0.0127*** (0.0034)	0.0123*** (0.0034)	0.0119*** (0.0032)	0.0119*** (0.0032)	0.0119*** (0.0033)	0.0130*** (0.0034)	0.0131*** (0.0034)
Number Quarters Brand/Intel family available	-0.0055*** (0.0016)	-0.0057*** (0.0017)	-0.0065*** (0.0018)	-0.0055*** (0.0016)	-0.0068*** (0.0019)	-0.0060*** (0.0019)	-0.0059*** (0.0019)	-0.0052*** (0.0015)	-0.0052*** (0.0015)	-0.0046*** (0.0017)	-0.0065*** (0.0018)	-0.0065*** (0.0018)
Capacity Related Variables												
Lagged Free Cash (1000M\$)	0.0171 (0.0108)	0.0169 (0.0108)	0.0158 (0.0105)	0.0203* (0.0105)	0.0197* (0.0103)	0.0203* (0.0104)	0.0199* (0.0106)	0.0168 (0.0112)	0.0153 (0.0117)	0.0202* (0.0114)	0.0163 (0.0110)	0.0149 (0.0115)
AMD Capacity Index	0.0025** (0.0012)	0.0026** (0.0012)	0.0028** (0.0012)	0.0036*** (0.0012)	0.0033*** (0.0012)	0.0002 (0.0012)	0.0016 (0.0012)	0.0032*** (0.0012)	0.0032*** (0.0012)	0.0015 (0.0012)	0.0029** (0.0012)	0.0029** (0.0012)
Intel Capacity Index	-0.0009 (0.0006)	-0.0009 (0.0006)	-0.0009 (0.0006)	-0.0013** (0.0006)	-0.0011* (0.0006)	-0.0007 (0.0006)	-0.0007 (0.0006)	-0.0012* (0.0006)	-0.0011* (0.0006)	-0.0007 (0.0006)	-0.0010 (0.0007)	-0.0010 (0.0007)
Exclusionary Restriction/Antitrust Related Variables												
Exclusionary Restriction Index	-0.0028*** (0.0010)											
Extreme Exclusionary Restrictions Index		-0.0071*** (0.0024)								-0.0122** (0.0025)		
Intel Payments to PC Firms (M\$)			-0.0014*** (0.0004)		-0.0015*** (0.0005)	-0.0018*** (0.0004)	-0.0021*** (0.0004)				-0.0014*** (0.0004)	
Intel Payments to Dell (100M\$) (For Non-Dell Firms)				-0.0053*** (0.0019)	-0.0052*** (0.0018)							
Number Exclusionary Restrictions on Other Firms						-0.0027*** (0.0010)						
Number Extreme Restrictions on Other Firms							-0.0084*** (0.0023)			-0.0076*** (0.0024)		
Lagged Number Pending Antitrust Cases Against Intel								0.0011 (0.0031)			0.0009 (0.0031)	
Lagged Number of Antitrust Cases Against Intel									-0.0009 (0.0022)	-0.0008 (0.0021)		-0.0008 (0.0022)

Notes: Robust Clustered Standard Errors in Parentheses. Each regression contains a constant and a quarterly trend. *** denotes significance at the 1% level, ** at the 5% level, and * at the 10% level. Each regression contains 3236 observations.

Table 3: Arellano-Bond Regressions

Considering first the effect of PC characteristics, the PC price appears to have a weak (often insignificant) negative relationship with the intensity of AMD adoption. There is also a negative relationship of the AMD adoption with the intensity of brand and firm-level advertising. These patterns are consistent with AMD processors being attractive for use in PC product lines that target value-seeking consumers.

The effects of CPU characteristics are, at a first glance, surprising. As expected, the rate of adoption of AMD chips decreases with Intel's benchmark-per-dollar index. At the same time, controlling for this index of Intel's delivered value, the rate of AMD adoption appears to *decrease* with the AMD benchmark-per-dollar index. One would expect, of course, to see AMD's adoption increasing with the value it offers its clients. This negative relationship is, however, consistent with institutional details. Recalling the displayed evolution in Figure 3, AMD's largest benchmark-per-dollar advantage over Intel obtained in the first part of the sample period, in the years 2003-2005. Recalling Figure 4 and additional discussion in Section 2, these years were also characterized by highly intensive application of the Intel Inside program, leading to extreme restrictions on the adoption of the AMD technology. AMD's benchmark-per-dollar advantage seems therefore to have been negated by Intel's ability to engage in vertical restraints. AMD enjoyed some growth in its market share in later periods, but these are periods when its benchmark-per-dollar advantage over Intel actually eroded. This negative relationship seems to be driving the negative coefficient on the AMD CPU benchmark-per-dollar variable.

This result merits some attention. It suggests that AMD's ability to pull ahead by innovating and offering competitive chips did little to expand its market share. This situation is consistent with an industry where technological leadership is *not necessarily* the primary driver of market share growth. Instead, per our additional results described below, market share growth is strongly affected by capacity investments and the ability to engage in vertical restraints and exclusive deals — two areas where Intel enjoyed a fundamental incumbency advantage. The implication that delivering higher value than the rival is not sufficient for market share growth is consistent with substantial anticompetitive forces stemming from the dynamics of this industry, and, specifically, from the effect of vertical constraints and their interaction with the capacity investment issue.

We next observe that the capacity-related variables have the signs predicted by our discussion in Section 3. Namely, AMD’s adoption increases in the extent of (lagged) free cash it has for investment, and in its capacity index (the former effect is only marginally statistically-significant, or insignificant across specifications). Both effects are consistent with the notion that higher values of these variables send a positive message to clients regarding AMD’s ability to serve as a viable substitute to Intel as a supplier in present and future periods. Higher AMD capacity (and cash flow that enables investment in capacity) suggests that a client should be less concerned about shifting towards increased reliance on AMD. One of the reasons is that if Intel retaliates by limiting supply of its own chips to the customer, AMD would be more able to pick up the slack. Intel’s capacity index, for its part, reduces the adoption of AMD chips, once again consistent with Intel’s capacity advantage being an important factor in its ability to retain its high market shares.

Continuing to move down Table 3 we reach our main results, concerning the impact of Intel’s vertical restraints. Both the restriction index and the extreme restriction index are found to have a negative and statistically significant effect on the rate of AMD’s adoption. Simply put: restrictions placed by Intel on the intensity of AMD utilization of a given downstream client indeed bind, and reduce the fraction of AMD-based machines produced by the client. Payments to the client from the “Intel Inside” program also have a negative relationship with AMD adoption, which is not surprising given that these payments depend positively on the utilization of a substitute, the Intel technology. These results confirm that Intel’s exclusionary practices, as measured by our indices, were binding, and indeed led to reduced adoption of AMD’s technology.

The more interesting results, however, are those associated with the variables “Intel payments to Dell (for non-Dell firms)”, “number of exclusionary restrictions on other firms,” and “number of extreme restrictions on other firms.” Controlling for the payments made by Intel to the downstream client, all these variables have a negative effect on the client’s adoption of the AMD technology. These results shed light on an important and little-explored aspect of vertical restraints: in a dynamic environment, they may affect downstream clients’ expectations regarding the future viability of the competing supplier. As discussed in Section 3 in detail, a client observing that Intel exerts substantial vertical restraints on another client

may revise its own assessment of AMD’s future ability to invest in production capacity and innovation.

This mechanism implies that vertical restraints imposed on other clients can cause the client to perceive AMD as a weaker competitor, thus increasing the client’s reliance on Intel. Importantly, this mechanism may operate *even if the client in question was not subject to Intel’s vertical restraints at all*. These findings are consistent with descriptive evidence presented in Figure 2 above: companies like Gateway and Acer, that were less affected by Intel’s vertical restraints, displayed some of the biggest jumps in AMD adoption following the removal of these restraints. This mechanism presents a policy-relevant issue to be considered by regulators: in evaluating the consequences of exclusive deals, their impact on clients that were not directly subject to them should potentially be considered — in particular, in markets characterized by a dynamic process of technology adoption.

Finally, we note that the lagged antitrust cases measures have an insignificant effect on the rate of AMD adoption in the linear specifications. Our analysis suggested that these should have a positive effect on adoption, as they shift expectations: more litigation implies higher chances that the restraints will be removed, making retaliation from Intel for increased AMD adoption less likely (while also improving the expected value of buying from AMD, as this firm will be expected to have more cash flow to invest in capacity expansion and product innovation). As we shall see below, these variables would indeed have the expected positive sign in our nonlinear specifications, and we shall discuss possible reasons for that.

To conclude our presentation of results from the linear specifications, we note that the results described above are not driven by one particular market segment. Tables 4 and 5 present the same specifications as in Table 3, applied to the Business and Home segments of the PC market, separately. These tables effectively obtain the same conclusions as those derived in Table 3 that considered both segments jointly.

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Lagged Percentage AMD Sold	0.7900*** (0.0719)	0.7836*** (0.0740)	0.7545*** (0.0764)	0.8162*** (0.0711)	0.7657*** (0.0800)	0.7760*** (0.0787)	0.7869*** (0.0798)	0.8064*** (0.0661)	0.8051*** (0.0658)	0.8263*** (0.0782)	0.7550*** (0.0754)	0.7539*** (0.0752)
PC Characteristics												
Price PC (100\$)	-0.0021 (0.0037)	-0.0026 (0.0036)	-0.0046 (0.0042)	-0.0035 (0.0036)	-0.0073* (0.0042)	-0.0026 (0.0044)	-0.0032 (0.0042)	-0.0008 (0.0035)	-0.0011 (0.0036)	-0.0003 (0.0038)	-0.0045 (0.0041)	-0.0047 (0.0041)
Brand Advertising (10M\$)	-0.0054 (0.0036)	-0.0058 (0.0039)	-0.0069 (0.0046)	-0.0035 (0.0036)	-0.0063 (0.0054)	-0.0070* (0.0041)	-0.0064 (0.0041)	-0.0041 (0.0033)	-0.0040 (0.0034)	-0.0054 (0.0035)	-0.0069 (0.0046)	-0.0068 (0.0046)
Firm Advertising (10M\$)	-0.0015 (0.0015)	-0.0027 (0.0018)	-0.0016 (0.0017)	-0.0023 (0.0015)	-0.0019 (0.0018)	-0.0016 (0.0016)	-0.0002 (0.0016)	-0.0019 (0.0016)	-0.0019 (0.0016)	-0.0020 (0.0016)	-0.0016 (0.0017)	-0.0016 (0.0017)
CPU Characteristics												
AMD CPU benchmark/dollar (10,000\$)	-0.0034 (0.0022)	-0.0034 (0.0022)	-0.0035 (0.0023)	-0.0047** (0.0022)	-0.0049* (0.0025)	-0.0040* (0.0022)	-0.0044** (0.0022)	-0.0032 (0.0022)	-0.0033 (0.0022)	-0.0045** (0.0022)	-0.0035 (0.0024)	-0.0036 (0.0024)
Intel CPU benchmark/dollar (10,000\$)	-0.0018 (0.0021)	-0.0019 (0.0021)	-0.0028 (0.0021)	-0.0028 (0.0021)	-0.0042** (0.0021)	-0.0019 (0.0023)	-0.0022 (0.0022)	-0.0014 (0.0020)	-0.0015 (0.0021)	-0.0010 (0.0022)	-0.0027 (0.0021)	-0.0028 (0.0021)
Number Quarters Brand/AMD family available	0.0093*** (0.0035)	0.0094*** (0.0036)	0.0103*** (0.0036)	0.0079** (0.0037)	0.0093** (0.0039)	0.0097*** (0.0037)	0.0091** (0.0037)	0.0088*** (0.0034)	0.0089*** (0.0034)	0.0078** (0.0037)	0.0103*** (0.0036)	0.0103*** (0.0036)
Number Quarters Brand/intel family available	-0.0039** (0.0018)	-0.0041** (0.0019)	-0.0049** (0.0020)	-0.0036** (0.0017)	-0.0050** (0.0021)	-0.0041* (0.0022)	-0.0040* (0.0021)	-0.0035** (0.0021)	-0.0036** (0.0021)	-0.0029 (0.0020)	-0.0049** (0.0020)	-0.0049** (0.0020)
Capacity Related Variables												
Lagged Free Cash (1000M\$)	0.0133 (0.0145)	0.0131 (0.0144)	0.0123 (0.0142)	0.0161 (0.0139)	0.0158 (0.0138)	0.0170 (0.0140)	0.0165 (0.0143)	0.0129 (0.0150)	0.0116 (0.0157)	0.0167 (0.0155)	0.0125 (0.0147)	0.0112 (0.0154)
AMD Capacity Index	0.0026* (0.0014)	0.0027* (0.0014)	0.0027* (0.0015)	0.0036*** (0.0014)	0.0032** (0.0014)	0.0001 (0.0019)	0.0016 (0.0014)	0.0032** (0.0014)	0.0032** (0.0014)	0.0015 (0.0014)	0.0028* (0.0015)	0.0028* (0.0015)
Intel Capacity Index	-0.0010 (0.0008)	-0.0009 (0.0008)	-0.0009 (0.0008)	-0.0013* (0.0008)	-0.0010 (0.0008)	-0.0007 (0.0007)	-0.0007 (0.0007)	-0.0012 (0.0008)	-0.0011 (0.0008)	-0.0007 (0.0008)	-0.0009 (0.0008)	-0.0009 (0.0008)
Exclusionary Restriction/Antitrust Related Variables												
Exclusionary Restriction Index	-0.0024** (0.0011)									-0.0113*** (0.0030)		
Extreme Exclusionary Restrictions Index		-0.0061** (0.0026)										
Intel Payments to PC Firms (M\$)			-0.0012*** (0.0004)		-0.0013*** (0.0005)	-0.0016*** (0.0004)	-0.0018*** (0.0005)				-0.0012*** (0.0004)	
Intel Payments to Dell (100M\$) (For Non-Dell Firms)					-0.0053** (0.0026)							
Number Exclusionary Restrictions on Other Firms						-0.0027** (0.0013)						
Number Extreme Restrictions on Other Firms							-0.0084*** (0.0030)			-0.0082** (0.0032)		
Lagged Number Pending Antitrust Cases Against Intel								0.0006 (0.0037)			0.0003 (0.0036)	
Lagged Number of Antitrust Cases Against Intel									-0.0009 (0.0028)	-0.0009 (0.0027)		-0.0010 (0.0027)

Notes: Robust Clustered Standard Errors in Parentheses. Each regression contains a constant and a quarterly trend. *** denotes significance at the 1% level, ** at the 5% level; and * at the 10% level. Each regression contains 1726 observations.

Table 4: Arellano-Bond Results for Business Market

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Lagged Percentage AMD Sold	0.6818** (0.1067)	0.6767*** (0.1075)	0.6546*** (0.1096)	0.6869*** (0.1033)	0.6483*** (0.1095)	0.6645*** (0.1100)	0.6703*** (0.1109)	0.6975*** (0.1009)	0.7025*** (0.1013)	0.7031*** (0.1084)	0.6571*** (0.1049)	0.6576*** (0.1070)
PC Characteristics												
Price PC (100\$)	-0.0066 (0.0062)	-0.0070 (0.0062)	-0.0090 (0.0069)	-0.0087 (0.0069)	-0.0115 (0.0070)	-0.0078 (0.0070)	-0.0084 (0.0069)	-0.0053 (0.0064)	-0.0047 (0.0062)	-0.0047 (0.0064)	-0.0087 (0.0070)	-0.0082 (0.0069)
Brand Advertising (10M\$)	-0.0044 (0.0067)	-0.0048 (0.0069)	-0.0050 (0.0077)	-0.0014 (0.0073)	-0.0038 (0.0086)	-0.0059 (0.0073)	-0.0049 (0.0074)	-0.0033 (0.0065)	-0.0034 (0.0062)	-0.0058 (0.0063)	-0.0052 (0.0078)	-0.0054 (0.0075)
Firm Advertising (10M\$)	-0.0035 (0.0021)	-0.0048* (0.0025)	-0.0036 (0.0023)	-0.0045** (0.0023)	-0.0041* (0.0024)	-0.0036 (0.0023)	-0.0023 (0.0023)	-0.0038* (0.0023)	-0.0037* (0.0022)	-0.0042* (0.0024)	-0.0035 (0.0024)	-0.0035 (0.0023)
CPU Characteristics												
AMD CPU benchmark/dollar (10,000\$)	-0.0040 (0.0030)	-0.0040 (0.0030)	-0.0041 (0.0031)	-0.0054* (0.0032)	-0.0057 (0.0035)	-0.0046 (0.0031)	-0.0051 (0.0031)	-0.0037 (0.0030)	-0.0039 (0.0029)	-0.0050* (0.0030)	-0.0040 (0.0032)	-0.0042 (0.0032)
Intel CPU benchmark/dollar (10,000\$)	-0.0033 (0.0030)	-0.0034 (0.0030)	-0.0040 (0.0032)	-0.0045 (0.0031)	-0.0052 (0.0033)	-0.0034 (0.0033)	-0.0037 (0.0033)	-0.0029 (0.0031)	-0.0028 (0.0031)	-0.0025 (0.0032)	-0.0038 (0.0031)	-0.0037 (0.0032)
Number Quarters Brand/AMD family available	0.0148*** (0.0052)	0.0149*** (0.0052)	0.0156*** (0.0053)	0.0145*** (0.0052)	0.0156*** (0.0054)	0.0152*** (0.0053)	0.0150*** (0.0054)	0.0144*** (0.0050)	0.0142*** (0.0050)	0.0140*** (0.0053)	0.0155*** (0.0052)	0.0155*** (0.0053)
Number Quarters Brand/Intel family available	-0.0071*** (0.0024)	-0.0072*** (0.0024)	-0.0081*** (0.0027)	-0.0072*** (0.0023)	-0.0084*** (0.0027)	-0.0076*** (0.0027)	-0.0075*** (0.0027)	-0.0066*** (0.0023)	-0.0065*** (0.0023)	-0.0063*** (0.0025)	-0.0080*** (0.0026)	-0.0079*** (0.0026)
Capacity Related Variables												
Lagged Free Cash (1000M\$)	0.0216 (0.0163)	0.0215 (0.0162)	0.0201 (0.0158)	0.0254 (0.0159)	0.0245 (0.0156)	0.0247 (0.0157)	0.0244 (0.0160)	0.0213 (0.0169)	0.0195 (0.0175)	0.0245 (0.0171)	0.0208 (0.0167)	0.0192 (0.0173)
AMD Capacity Index	0.0025 (0.0020)	0.0026 (0.0020)	0.0029 (0.0021)	0.0037* (0.0019)	0.0035* (0.0020)	0.0031 (0.0025)	0.0017 (0.0020)	0.0033* (0.0019)	0.0033* (0.0019)	0.0016 (0.0020)	0.0030 (0.0021)	0.0031 (0.0021)
Intel Capacity Index	-0.0009 (0.0010)	-0.0009 (0.0010)	-0.0010 (0.0010)	-0.0013 (0.0010)	-0.0013 (0.0010)	-0.0007 (0.0010)	-0.0007 (0.0010)	-0.0012 (0.0010)	-0.0012 (0.0010)	-0.0007 (0.0010)	-0.0011 (0.0010)	-0.0011 (0.0010)
Exclusionary Restriction/Antitrust Related Variables												
Exclusionary Restriction Index	-0.0030* (0.0016)											
Extreme Exclusionary Restrictions Index		-0.0076** (0.0038)								-0.0129*** (0.0041)		
Intel Payments to PC Firms (M\$)			-0.0015** (0.0007)		-0.0016** (0.0007)	-0.0020*** (0.0007)	-0.0023*** (0.0007)				-0.0015** (0.0007)	-0.0015** (0.0006)
Intel Payments to Dell (100M\$) (For Non-Dell Firms)												
Number Exclusionary Restrictions on Other Firms						-0.0028* (0.0014)						
Number Extreme Restrictions on Other Firms							-0.0089** (0.0035)			-0.0074** (0.0034)		
Lagged Number Pending Antitrust Cases Against Intel								0.0014 (0.0051)			0.0015 (0.0052)	
Lagged Number of Antitrust Cases Against Intel									-0.0008 (0.0035)	-0.0009 (0.0034)		-0.0007 (0.0036)

Notes: Robust Clustered Standard Errors in Parentheses. Each regression contains a constant and a quarterly trend. *** denotes significance at the 1% level, ** at the 5% level, and * at the 10% level. Each regression contains 1510 observations.

Table 5: Arellano-Bond Results for the Home Market

Table 6 reports results from the nonlinear (Tobit) specification. The results are largely consistent with those obtained from the linear specifications. The effects of PC and CPU characteristics are qualitatively the same as those described above. The only notable exception is that PC firm advertising now has a positive relationship with AMD adoption, as opposed to the negative relationship obtained in the linear models of Table 3. PC brand advertising, nonetheless, still has a negative effect as before.

Importantly, the results concerning the role played by capacity and by Intel’s exclusive restrictions, including “Intel Inside” payments, are all preserved in the nonlinear specifications. Of note, the effect of our antitrust litigation measure “Lagged Number Pending Antitrust Cases Against Intel” becomes positive and significant at the 10 percent level (whereas it was positive but statistically insignificant in the linear specifications of Table 3). It could be that the main effect of litigation was to convince companies that were exclusive with Intel (notably, Dell and Toshiba) to switch into positive adoption of the AMD technology, suggesting that the Tobit specification, that allows for a mass point at zero for the dependent variable, may be better suited to capture this effect.

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lagged Percentage AMD Sold								
Price PC (100\$)	0.722*** (0.033)	0.702*** (0.033)	0.701*** (0.033)	0.702*** (0.033)	0.706*** (0.033)	0.705*** (0.033)	0.698*** (0.033)	0.701*** (0.033)
PC Characteristics								
Brand Advertising (10M\$)	-0.025*** (0.005)	-0.026*** (0.005)	-0.026*** (0.005)	-0.026*** (0.005)	-0.026*** (0.005)	-0.026*** (0.005)	-0.025*** (0.005)	-0.025*** (0.005)
Firm Advertising (10M\$)	-0.031 (0.029)	-0.037 (0.029)	-0.037 (0.029)	-0.037 (0.029)	-0.023 (0.029)	-0.026 (0.030)	-0.029 (0.029)	-0.032 (0.030)
AMD CPU benchmark/dollar (10,000\$)	0.009** (0.004)	0.004 (0.004)	0.004 (0.004)	0.004 (0.004)	0.012** (0.005)	0.003 (0.004)	0.011** (0.005)	0.002 (0.004)
Intel CPU benchmark/dollar (10,000\$)	-0.012*** (0.002)	-0.012*** (0.002)	-0.011*** (0.002)	-0.011*** (0.002)	-0.011*** (0.002)	-0.011*** (0.002)	-0.013*** (0.002)	-0.013*** (0.002)
Number Quarters Brand/AMD family available	-0.000 (0.002)	-0.002 (0.002)	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)
Number Quarters Brand/Intel family available	0.024*** (0.002)	0.024*** (0.002)	0.024*** (0.002)	0.024*** (0.002)	0.024*** (0.002)	0.024*** (0.002)	0.024*** (0.002)	0.024*** (0.002)
Capacity Related Variables								
Lagged Free Cash (1000M\$)	-0.018*** (0.002)	-0.018*** (0.002)	-0.018*** (0.002)	-0.018*** (0.002)	-0.018*** (0.002)	-0.018*** (0.002)	-0.018*** (0.002)	-0.018*** (0.002)
AMD Capacity Index	0.031 (0.025)	0.011 (0.025)	0.019 (0.026)	0.018 (0.026)	0.029 (0.026)	0.023 (0.026)	0.049* (0.027)	0.034 (0.026)
Intel Capacity Index	0.011** (0.005)	0.010** (0.005)	0.010** (0.005)	0.010** (0.005)	0.007 (0.005)	0.010** (0.005)	-0.004 (0.006)	0.008 (0.005)
Exclusionary Restrictions/Antitrust Related Variables								
Intel Payments to PC Firms (M\$)	-0.004*** (0.001)	-0.004*** (0.001)	-0.004*** (0.001)	-0.004*** (0.001)	-0.004*** (0.001)	-0.004*** (0.001)	-0.002 (0.001)	-0.003** (0.001)
Exclusionary Restriction Index	-0.011*** (0.003)	-0.011*** (0.003)	-0.012*** (0.003)	-0.011*** (0.003)	-0.012*** (0.003)	-0.012*** (0.003)	-0.011*** (0.003)	-0.011*** (0.003)
Extreme Exclusionary Restrictions Index	-0.019*** (0.005)	-0.019*** (0.005)	-0.019*** (0.005)	-0.019*** (0.005)	-0.019*** (0.005)	-0.019*** (0.005)	-0.027*** (0.006)	-0.027*** (0.006)
Number Exclusionary Restrictions on Other Firms	-0.032** (0.014)	-0.032** (0.014)	-0.032** (0.014)	-0.032** (0.014)	-0.032** (0.014)	-0.032** (0.014)	-0.032** (0.014)	-0.032** (0.014)
Number Extreme Restrictions on Other Firms	0.005 (0.007)	0.005 (0.007)	0.005 (0.007)	0.005 (0.007)	0.005 (0.007)	0.005 (0.007)	0.005 (0.007)	0.005 (0.007)
Lagged Number Pending Antitrust Cases Against Intel	0.016* (0.009)	0.016* (0.009)	0.016* (0.009)	0.016* (0.009)	0.015* (0.009)	0.015* (0.009)	0.015* (0.009)	0.015* (0.009)
Lagged Number of Antitrust Cases Against Intel								

Notes: Robust Clustered Standard Errors in Parentheses. Regressions include quarterly trends, a constant, and mean regressors over time. *** denotes significance at the 1% level, ** at the 5% level; and * at the 10% level. Each regression contains 3280 observations.

Table 6: Tobit Estimates of Percent of AMD sold

5 Conclusions

In this paper, we examine the impact of exclusionary restrictions put in place by Intel on PC firms in the semiconductor industry. We investigate the manner by which such restraints interact with the dynamic process of downstream technology adoption. To do so we use rich data on PC and CPU sales to estimate dynamic panel models that document such patterns in this market. Our analysis covers the years 2002-2009 and focuses on the role played by the “Intel Inside” program.

Our results shed light on important mechanisms that have not, to the best of our knowledge, received attention in the empirical and theoretical literature on vertical restraints. We show that, not only do restraints imposed on a given downstream client reduce its adoption of the rival’s technology, but this client is also less likely to adopt this technology when restraints are imposed on *other clients*. This is consistent with an important role for dynamics and client expectations regarding the future value of adopting the rival’s technology.

These mechanisms are reinforced by capacity constraints. When the vertical restraints are imposed by an incumbent upstream firm that enjoys a substantial capacity advantage, as is the case in our empirical application, these restrictions may interact with the capacity issue in an interesting fashion. If they involve exclusive, or near-exclusive deals with major clients (such as Dell and Toshiba in our case), the rival may lack the volume of sales and cashflow that would allow it to invest in capacity expansion. Limited capacity by the rival can cause downstream clients to be weary of adopting its technology, as they cannot rely on it to deliver large, timely shipments on which they depend. This causes clients to stick with the large, incumbent supplier. Our results indeed show that capacity is an important factor driving the adoption decision.

Our analysis has some limitations. In particular, we do not formally model the adoption decision. Indeed, in future work we plan on extending the analysis to consider a full-blown structural dynamic model of technology adoption in this industry. Such an analysis would enable us to quantify the costs associated with adopting the AMD technology, including the component of these costs that is directly driven by Intel’s restraints. This would allow us to better identify the separate costs and benefits associated with adopting AMD’s technology.

It would also allow us to consider policy counterfactuals - i.e., to consider the implications of different policies regarding these restraints. Counterfactual analysis would also enable us to determine the extent to which a faster buildup of production capacity by AMD would have allowed it to negate the impact of Intel's restraints on its growth.

6 Appendix: Data Details

Advertisement Variables. We describe the creation of the advertisement variables based on a dataset from the Kantar Media Group. We identify three types of ad expenditures in the Kantar data: PC brand level advertising (Ad1), PC advertising categorized as business-to-business (Ad2) and PC firm level promotions (Ad3). We define two advertisement variables: brand specific advertising and firm level advertising. For brand specific advertising, we create the variable differently depending on whether the observation is in the home or non-home segment. Indeed, while Ad1 expenditures are likely to influence choices on both segments (households or firms), the Ad2 expenditures should only affect the non-home segment.

The firm level advertising are identically defined on both segments and consist of Ad3. The definitions of the two variables are summarized in Table A1. For those observations of the advertisement data whose brand could not be matched with the Gartner data, the expenditures were accounted for as firm level expenditures. Finally, the above described ad variables were matched to the Gartner data at the brand level. As the Kantar data contains less details about PC brands than the Gartner dataset, the match occurred based on the Kantar brands.

Variable / Segment	Home Segment	Non-home segment
Brand Advertising	Ad1	Ad1+Ad2
Firm Advertising	Ad3	Ad3

Table A1: Segment-specific Definition of Advertisement Variables

CPU Quality. We measure CPU quality in terms of CPU benchmark per dollar. In what follows we describe first the creation of the CPU family level price measures, and then the CPU family level benchmark measures.

To our knowledge, a comprehensive CPU price database for the US in the time period of interest is not available. We thus create our own CPU price dataset. We use four different sources: Instat estimated Intel CPU core prices (D1), Instat forecasted Intel CPU core prices (D2), Intel list prices (D3) and AMD list prices (D4). Table A2 offers an overview of their

respective coverage. The level of aggregation of these datasets differs from one to the other and from that of the Gartner data. In what follows, we describe how each of these datasets was merged to the Gartner data to obtain a consistent dataset at the CPU family-quarter level. In the case of Intel, we also discuss how the different sources (Instat and List prices) are merged to generate a unified dataset.²¹

	2003	2004	2005	2006	2007	2008	2009
AMD	D3 List Prices						
Intel	D1 Instat Estimate			D2 Instat Forecast		D4 List Prices	

Table A2: Time Coverage of the Price Data Sources

We first describe the treatment of the Intel prices for 2002Q3-2005Q4 (D1). These are computed based on information from Instat’s "Rosetta Stone" report on CPU core prices. We follow the methodology described in Lee, Pechy and Sovinsky (2013). A given CPU core is often marketed under different family names depending on which features are available. For example, the CPU core “Northwood” is used in both “Pentium 4” and “Mobile Celeron” CPU families. Moreover, the CPU core used in a CPU family can change over time. Taking these into consideration, the CPU cores are matched to the CPU families of the PC data at the platform group (whether desktop or mobile)/type (mainstream/value/ultraportable)/family/speed/quarter level.²²

Table A3 provides the product cross-referencing. Table A4 provides an overview of the variation of the prices of these CPU model price estimates at the family level. The most famous Intel families, Celeron and Pentium 4, have more than a hundred price observations. Prices vary significantly within a family. The Pentium D model has only two observations

²¹ The Gartner sales data also records a few CPU families which are neither Intel nor AMD produced (Cru, Eff, ViaC7). These observations are dropped due to lack of price information.

²² For the CPUs not matched at first attempt, the type is dropped of the matching criteria. When unmatched, the data are matched based on family/marketing name of a CPU, CPU speed, year, and quarter, ignoring platform group. When the data are not matched, we try matching based on platform group, family/marketing name of a CPU, CPU speed, ignoring time. For observations still not matched, we take the averages of price estimates of CPUs of the same marketing name, year and quarter.

as it is introduced at the end of the sample.

Platform	CPU Core	Family Name	Speed (Frequency: MHz)	
Desktop	Mainstream	Willamette	1300 - 2000	
		Northwood	1600 - 3400	
		Prescott	2260 - 3800	
		Smithfield*	Pentium D	2667 - 3200
	Value	Tualatin	Pentium III	1000 - 1400
			Celeron	900 - 1400
		Willamette Northwood	Celeron	1500 - 2000
				1600 - 2800
	Prescott	Celeron D	2133 - 3460	
Mobile	Mainstream	Northwood	Mobile Pentium 4-M	1200 - 2600
		Prescott	Mobile Pentium 4	2300 - 3460
		Banias Dothan	Pentium M	1200 - 1800 1300 - 2267
	Value	Tualatin	Mobile Celeron	1000 - 1330
			Mobile Pentium III-M	866 - 1333
		Northwood	Mobile Celeron	1400 - 2500
		Banias Dothan	Celeron M	1200 - 1500
				1200 - 1700
	Low-Power	Tualatin LV Tualatin ULV	Mobile Pentium III-M	733 - 1000
				700 - 933
		Tualatin LV Tualatin ULV	Mobile Celeron	650 - 1000
				650 - 800
		Banias LV Banias ULV Dothan LV Dothan ULV	Pentium M	1100 - 1300
				900 - 1100
				1400 - 1600
				1000 - 1300
		Banias ULV Dothan ULV	Celeron M	600 - 900
	900 - 1000			

Notes: * Dual-core processor
 Low-power mobile PCs are mini-notebook, tablet, and ultraportables.
 (LV: low-voltage; ULV: ultra-low-voltage)

Table A3. Cross-Reference from CPU Core to Family Name in 2002Q3-2005Q4

The price datasets from list prices of Intel and AMD (D3 and D4) were created as follows. Intel prices were collected in the form of Intel's price catalogues (unit price in case of 1000

CPU Firm	CPU Family	Median	Std Dev	Min	Max	Obs
	Cel	66	7	49	77	140
	Cel M	94	32	87	203	12
Intel	P3	128	46	49	170	36
	P4	176	17	130	202	171
	PD	245	4	242	247	2
	PM	219	29	190	317	51

Table A4: Descriptive Statistics of CPU Instat Estimated Prices by CPU Family in \$

units) from a large variety of websites.²³ AMD prices (unit price in case of 1000 units) were collected from the corporate website list price publications using waybackmachine.com, a website storing (many) historical saves of given webpages. These list prices are published and observed at the CPU model level (e.g. AMD Athlon 64 2800+) with variable frequency and must be merged to the Gartner market share dataset at the CPU family level (e.g. AMD Athlon 64) observed quarterly.

Figure A1 provides an overview of the availability of list prices resulting from our data sources. The left panel depicts the number of publications observed in a quarter.²⁴ It reveals that for the majority of quarters, more than one price publication is observed. For both AMD and Intel, there is one period where no publication could be collected: 2008Q3 and 2007Q1 respectively. The treatment of these two periods is discussed below. In 2008Q3, nine price publications are observed for Intel. The right panel shows the quarterly average number of CPU models per publication per quarter: the price of a family is based on average on more than ten CPU models. For AMD in particular, this number declines over time. This is explained by changes in the product portfolio. Before 2003Q4 only one or two families were marketed, but these proposed many different models. Afterwards, more families were gradually introduced (seven families observed in 2008Q2) with lower number of models per family.

We propose an overview of the differences across families in these model-quarter level

²³ Complete list available from the authors upon request.

²⁴ We note that Intel list prices could not be collected prior to 2005Q4. Most likely, the company was not publishing list prices in PDF format on the web prior to this quarter.

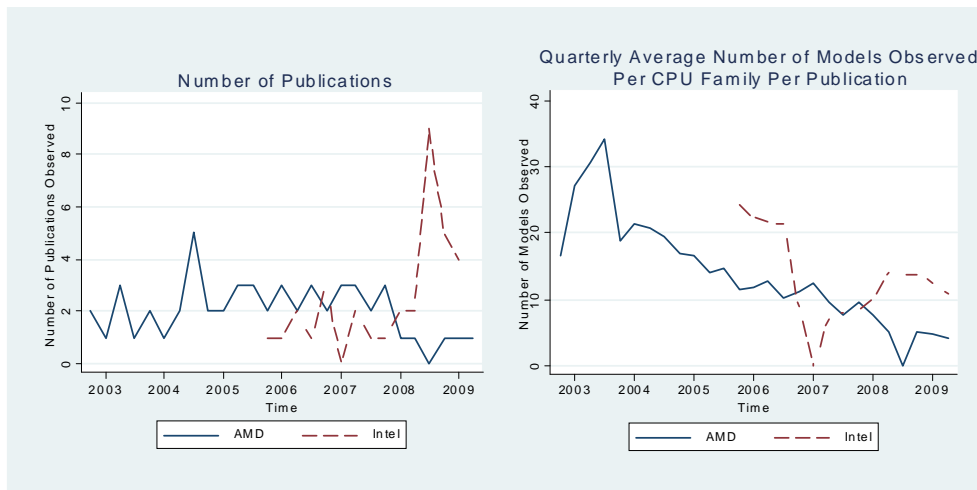


Figure A1: Availability of List Price Data

prices in Table A5. Most families have at least 50 observations and some of them more than 500. Model prices are widely spread, which is explained by high prices at the introduction and a strong decline over time.

We now describe the aggregation of the list prices. The procedure is identical for AMD and Intel CPUs. The CPU prices at the publication date-model level are aggregated to the quarter-family level by taking the median over models. Second, the obtained price dataset is merged with the Gartner market share data at the family quarter level to verify price data availability for each quarter of a family’s market share sequence. Out of 164 Intel CPU family-quarter observed in the Gartner data, 124 have a match in our Intel list price dataset. These numbers are respectively 164 and 148 for AMD CPU family-quarters.

For the periods of a sequence where price data is not available, we proceed as follows. When the price is missing in the middle of the sequence, it is approximated with kernel density interpolation at the family level. For prices missing in the first quarters of the sequence, the first observed price is used. These new introductions have usually very small market shares and high prices, which are preserved by this approximation. For prices missing in the last periods of the sequence, the last observed price is used. In a few cases where for a CPU family no price at all is observed the observations are dropped from the dataset (the related market shares are negligible, as these mostly concern server CPUs). These necessary

CPU Firm	CPU Family	Median	Std Dev	Min	Max	Obs
AMD	Ath	97	66	51	588	583
	Ath 64	223	210	36	1'031	579
	Ath 64 X2	178	193	62	1'001	269
	Dur	62	14	42	89	12
	Phe II X4	195	28	175	245	5
	Phe X3	122	23	101	195	24
	Phe X4	173	29	142	251	38
	Sem	86	26	30	145	617
	Tur 64	184	63	145	525	239
	Tur 64 X2	220	60	154	354	93
Intel	Atom	40	37	20	135	179
	Cel	65	26	30	134	479
	Cel M	107	34	45	161	178
	Cel X2	83	3	80	86	4
	Core 2 Duo	262	196	113	999	1'226
	Core 2 Quad	316	434	163	1'499	336
	Core 2 Solo	262	9	241	262	16
	Core Duo	294	129	209	706	184
	Core Solo	241	25	209	278	51
	Core i7	562	298	284	999	26
	P4	218	186	55	999	179
	PD	178	176	74	999	80
	PDC	64	9	64	86	137
	PM	304	114	130	702	409

Table A5: Descriptive Statistics of CPU Model List Prices by CPU Family in \$

inter/extrapolations are listed in Table A6 for AMD and Table A7 for Intel. In the end, we obtain from the list prices, a dataset at the family quarter level with the following coverage: 2005Q4 until 2009Q1 for Intel, and 2002Q3 until 2009Q1 for AMD.

At this point, the Intel price data stem from two different sources: Instat for 2002Q3-2005Q4 (D1) and list prices for 2005Q4-2009Q1 (D3). To obtain a consistent measure of CPU prices, we define a correction coefficient. We take the mean of the “Instat price / list price” ratios at the CPU family level for periods where both types of prices are available. As this is only fulfilled in period 2005Q4, we propose a second correction coefficient on periods 2005Q4-2006Q4 using the Instat CPU core price predictions for year 2006 (D2). To obtain prices at the CPU family level, a cross-referencing between Gartner and Instat is executed as previously described except for speed information, which is not available in the Gartner

data for 2006.

As for some CPU families more than one core is matched, the mean price over cores is retained. The cross-referencing is provided in Table A8. Using these prices, the second correction coefficient can be computed. The two price correction coefficients are summarized in Table A9. The Instat prices are on average 22% below the value of the list prices (14% when the Instat predicted prices (D2) are also included). As expected, the standard deviation of the coefficient which is computed over both estimated and predicted Instat prices is larger. The observed Min and Max values are to our understanding due to CPU model introductions being predicted too early/late, thus leading to a large value of the price difference for this CPU family. Based on these two price correction coefficients two different variables for Intel CPU family prices are defined. For robustness, we run our regressions using each variable.

	CPU Family Name	Quarter
First Value	Ath 64 X2	2005Q1
Interpolation	Ath 64, Ath 64 X2, Phe X3 Phe X4, Sem	2008Q3 2008Q3
Last Value	Ath Tur 64 Tur 64 X2	2005Q3-2006Q1 2007Q3-2008Q2 2008Q3-2009Q1
Dropped Obs	-	-

Table A6: AMD List Price Corrections

Benchmark. CPU benchmark information is gathered from Passmark publications.²⁵ This company collects measurements on CPU tests from users around the world, and creates a database of CPU performance at the CPU model level. We now discuss the treatment of the CPU benchmark information. The benchmark level of a given CPU family in a given period is built with two different approaches, exploiting the best information available in each quarter.

²⁵ Source: <https://www.cpubenchmark.net/>

	CPU Family Name	Quarter
First Value	Cel X2	2008Q2-2008Q4
	Core 2 Quad	2007Q1
	PDC	2007Q1-2007Q2
Interpolation	Cel, Cel M, Core 2 Duo	2007Q1
	Core Duo, Core Solo, P4, PD	2007Q1
Last Value	Core Duo	2008Q3-2009Q1
	Core Solo	2008Q3-2009Q1
	P4	2008Q3-2008Q4
	PD	2007Q3-2008Q1, 2008Q3-2009Q1
	PM	2006Q4-2008Q1, 2008Q3-2009Q1
Dropped Obs	A110	2007Q3-2008Q2

Table A7: Intel List Price Corrections

Platform		CPU Core	Family Name	
Desktop	Mainstream	Conroe*	Celeron	
		Conroe*	Core 2 Duo	
		Prescott	Pentium 4	
		Presler*	Pentium D	
		Gallatin	Xeon	
	Value	Cedar Mill	Celeron D	
		Cedar Mill	Pentium 4	
		Prescott	Celeron D	
	Mobile	Mainstream	Yonah*	Core Duo
			Dothan	Pentium M
Value		Dothan	Celeron M	
		Yonah	Celeron M	
		Yonah	Core Solo	
Low-Power		Dothan LV	Pentium M	
		Dothan ULV	Celeron M	
		Dothan ULV	Pentium M	
		Yonah LV	Xeon	
		Yonah ULV	Celeron M	
		Yonah ULV	Core Solo	

Notes: * Dual-core processor
Low-power mobile PCs are mini-notebook, tablet, and ultraportables.
(LV: low-voltage; ULV: ultra-low-voltage)

Table A8: Cross-Reference from CPU Core to Family Name in 2006Q1-Q4

Used Instat Prices	Instat Price/List Price Ratio					
	Overlapping Quarters	Mean	Std. Dev.	Min	Max	Obs
Estimated Instat Prices	2005Q4	0.78	0.24	0.58	1.15	5
Estimated and Predicted Instat Prices	2005Q4-2006Q4	0.86	0.45	0.40	2.30	33

Table A9: CPU Price Correction Coefficients

In the first approach, we rely on Gartner data information and match CPU benchmark to the Gartner data at the CPU family-CPU speed-platformgroup level (let us call this approach Gartner based) following Lee, Pechy and Sovinsky (2013). In the second approach, the availability of CPU models over time is inferred from our list price dataset described above (let us call this approach List Price based): those CPU models which are available in the period according to the list price information are those which define the value of the benchmark of that family in that period. The matches between the CPUs of the benchmark and the list price data are achieved by taking the best of 3 different matching criteria (in order of preference): family/model code/speed, family/model code, family/speed.²⁶ Then, to obtain the level of observation of the Gartner dataset after 2005 (without speed information), a CPU family quarter, we take the median of the benchmark level over CPU models in each quarter.²⁷

Table A10 offers an overview of the approach used in each time period for each CPU firm. For AMD, the Gartner based approach is used from the beginning of the sample until 2005Q1, as in this quarter, speed information is not available anymore and thus the List Price based approach is preferred. For Intel, the Gartner based approach is used until 2005Q4, as this is the first period where Intel list prices are observed and thus the List Price

²⁶ Note that this last criteria is required in a minority of cases only. It can potentially aggregate very different benchmark levels (aggregating benchmarks of CPUs available in 2005 with some of 2008). To exclude these cases, we only use observations where the min and the max benchmarks are distant by less than 10%.

²⁷ For observations where benchmark information is missing, we use the same procedure as described above for prices (interpolation, first observed benchmark, last observed benchmark) since the benchmark data availability is corresponding to price data availability.

based approach is applied that period onwards.

	2003	2004	2005	2006	2007	2008	2009
AMD	<i>Gartner based</i>			<i>List Price based</i>			
Intel	<i>Gartner based</i>			<i>List Price based</i>			

Table A10: Methodologies used To Proxy Benchmark Level

Table A11 offers a summary of the benchmark scores of each family in the Gartner based approach. Table A12 offers a summary of the benchmark scores in the List Price based approach. There are clear differences across families. For example, Athlon models have low scores, while Phenom models are top performers. Differences within a family are less large but show that, as expected, various benchmark levels are proposed within a family.

CPU Firm	CPU Family	Median	Std Dev	Min	Max	Obs
AMD	Ath	410	119	200	610	14
	Ath 64	527	100	418	764	10
	Dur	272	17	243	272	3
	Sem	426	11	412	434	4
Intel	Cel	258	54	186	409	27
	Cel M	342	86	231	437	4
	P3	243	42	162	296	12
	P4	311	149	133	641	26
	PD	905	-	905	905	1
	PM	356	130	226	596	9

Table A11: Descriptive Statistics of CPU Benchmark Scores by CPU Family in the *Gartner based Approach*

Benchmark per Dollar Measure. Before getting to the merger of the price and the benchmark information, we provide a summary of the CPU landscape in Figure A2. The median price and the median benchmark (horizontal and vertical axis respectively) are shown for the CPU families of both Intel and AMD for two periods 2004Q1 and 2008Q2 (left and right panel respectively). As can be seen, the number of families on the market

CPU Firm	CPU Family	Median	Std Dev	Min	Max	Obs
AMD	Ath	428	35	341	454	45
	Ath 64	559	249	445	1'597	388
	Ath 64 X2	1'264	197	805	1'781	266
	Phe II X4	3'602	329	3'100	3'941	5
	Phe X3	1'938	135	1'655	2'095	24
	Phe X4	2'585	259	2'168	3'047	38
	Sem	441	38	362	604	541
	Tur 64	467	64	387	616	239
	Tur 64 X2	894	138	768	1'273	93
Intel	Atom	304	120	163	634	155
	Cel	556	227	321	1'227	460
	Cel M	425	69	221	482	135
	Cel X2	1'220	54	1'173	1'267	4
	Core 2 Duo	1'547	488	587	2'652	1'049
	Core 2 Quad	3'575	478	2'976	4'606	253
	Core 2 Solo	316	84	311	502	16
	Core Duo	843	159	544	1'144	170
	Core Solo	402	86	280	514	46
	Core i7	6'123	547	5'555	7'022	26
	P4	548	86	180	688	134
	PD	809	83	672	1'000	80
	PDC	1'249	289	907	1'944	137
	PM	448	90	248	596	409

Table A12: Descriptive Statistics of CPU Benchmark Scores by CPU Family in the *List Price based Approach*

is much larger in 2008Q2. There is significant variation across families as some are low-end (low price and low benchmark level) while others are high-end (ex: Phe X4, Core 2 Duo).²⁸

To account for this information in our model, we define a benchmark per dollar variable. For each CPU family quarter, the ratio of benchmark per price is computed. Then the data is merged with the Gartner sales data, and we compute the average of this ratio weighted by market share of all PC models in a PC brand-segment. We obtain the variables of interest, the benchmark per dollar for AMD (resp. Intel) at the PC brand-segment level.

²⁸ We note that on a given benchmark level, a given CPU manufacturer provides various families at different price levels. The existence of the more expensive families is explained by the fact that beside the benchmark, other CPU characteristics influence the price (ex: power consumption). These could not be accounted for here due to lack of data.

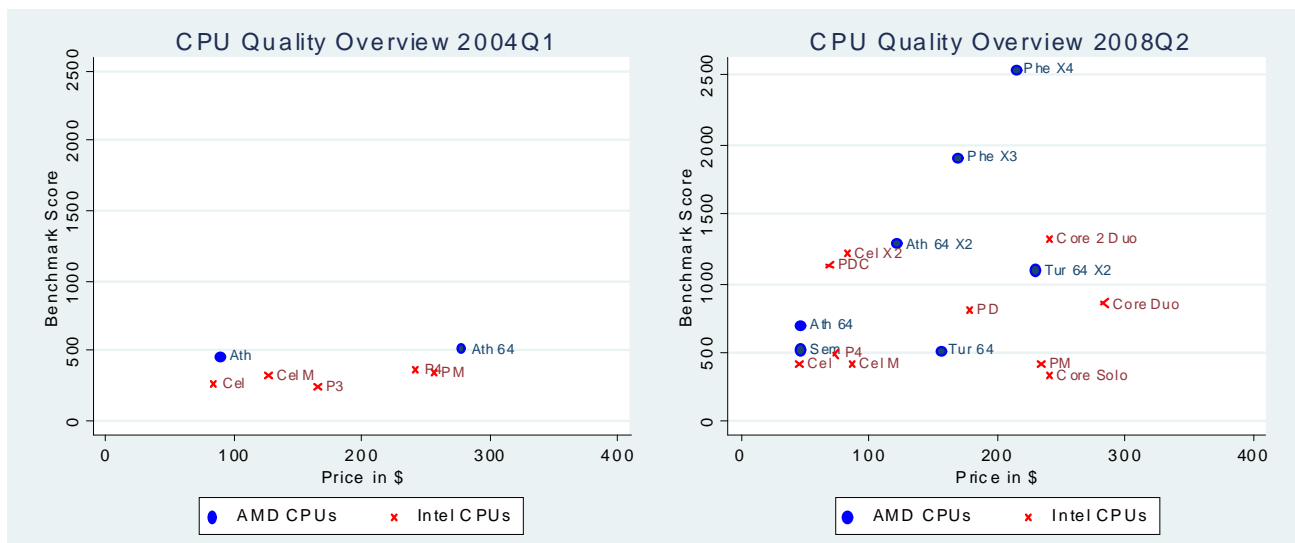


Figure A2: CPU Quality At The Family Level

References

- Arellano, M., and O. Bover (1995) “Another look at the Instrumental Variables Estimation of Error Components Models,” *Journal of Econometrics* 68: 29-51.
- Asker, J. (forthcoming) “Diagnosing Foreclosure due to Exclusive Dealing,” *Journal of Industrial Economics*.
- Ater, I. (2015) “Vertical Foreclosure in Shopping Malls: Evidence from Shopping Malls,” *Journal of Economics and Management Strategy*, 24: 620 - 642.
- Areeda, P., and L. Kaplow (1997). “*Antitrust Analysis*,” Aspen Publishes, Inc.
- Arellano, M. and S. Bond (1991) “Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations,” *Review of Economic Studies* 58: 277-97.
- Blundell, R. and S. Bond (1998) “Initial Conditions and Moment Restrictions in Dynamic Panel Data Models,” *Journal of Econometrics* 87: 11-143.
- Besanko, D., and M.K. Perry (1993). “Equilibrium Incentives for Exclusive Dealing in a Differentiated Products Oligopoly,” *RAND Journal of Economics*, 24(4): 646-67
- Chamberlain, G. (1984) “Panel Data,” in Z. Griliches and M.D. Intriligator (eds.), *Handbook of Econometrics*, Volume 2. Amsterdam: North Holland, 1247-1318.

- Eizenberg A. (2014) “Upstream Innovation and Product Variety in the United States Home PC Market,” *The Review of Economic Studies*, 81(3): 1003-1045
- Inderst, R., and G. Shaffer (2007). “etail Mergers, Buyer Power and Product Variety,” *The Economic Journal*, 117(516): 45-67
- Lafontaine, F. and M. Slade (2008) “Exclusive Contracts and Vertical Restraints: Empirical Evidence and Public Policy,” in Paolo Buccirossi, ed., Handbook of Antitrust Economics, Cambridge: MIT Press.
- Lee, H., Pechy, A., and M. Sovinsky (2013) “Is Intel’s Marketing Campaign Predatory?” Working Paper, University of Zurich.
- Lommerud, K.E., Straume, O.R., and L. Sorgard, (2006). "National versus international mergers in unionized oligopoly," *RAND Journal of Economics*, 37(1): 212-233
- Marvel, H.P. (1982), “Exclusive Dealing,” *Journal of Law and Economics*, 25: 1-25
- Goettler, R. and B. R. Gordon (2011) “Competition and innovation in the microprocessor industry: Does AMD spur Intel to innovate more?” *Journal of Political Economy*, 119(6), 1141-1200
- Gordon, B. (2009) “A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry,” *Marketing Science*, 28(5): 846–8.
- Lafontaine, F. and M. Slade (2008) Exclusive Contracts and Vertical Restraints: Empirical Evidence and Public Policy," in Paolo Buccirossi, ed., Handbook of Antitrust Economics, Cambridge: MIT Press, 2008, pp. 391-414.
- Moon, Y. E. and Darwall, C. (2005) “Inside Intel Inside,” *Harvard Business School Case* 502-083
- Nurski, L. and F. Verboven (2016) “Exclusive Dealing as an Entry Barrier - Evidence from the Car Market,” *The Review of Economic Studies*, forthcoming
- Prince, J. T. (2008), “Repeat Purchase Amid Rapid Quality Improvement: Structural Estimation of Demand for Personal Computers,” *Journal of Economics and Management Strategy* 17(1), 1–33.
- Salgado, H. (2008a) “Dynamic Firms Conduct and Market Power: The Computer Processors Industry under Learning-by-Doing,” Working Paper, University of California at Berkeley
- Salgado, H. (2008b) “Brand Loyalty, Advertising and Demand for Personal Computer Processors: The Intel Inside Effect,” Working Paper, University of California at Berkeley

- Sass, T.R. (2005). "The competitive effects of exclusive dealing: Evidence from the U.S. beer industry," *International Journal of Industrial Organization*, 23: 203-25
- Slade, M. E. (2005) "Product Rivalry with Multiple Strategic Weapons: An Analysis of Price and Advertising Competition," *Journal of Economics & Management Strategy*, 4(3): 445-476
- Sullivan, E. T., and H. Hovenkamp (2003). "Antitrust Law, Policy and Procedure: Cases, Materials, Problems," 5th edn, Charlottesville, Va.: LexisNexis.
- Song, M. (2007) "Measuring Consumer Welfare in the CPU Market: An Application of the Pure Characteristics Demand Model," *RAND Journal of Economics*, 38, 429-446.
- Sovinsky Goeree, Michelle (2008) "Limited Information and Advertising in the US Personal Computer Industry," *Econometrica*, 1017-1074.
- Spector, D. (2011). "Exclusive contracts and demand foreclosure," *RAND Journal of Economics*, 42(4): 619-38
- Wooldridge, J.M. (2002) "Econometric Analysis of Cross Section and Panel Data," *MIT Press: Cambridge, US*
- Wooldridge, J.M. (2005) "Simple Solutions to the Initial Conditions Problem in Dynamic, Nonlinear Panel Data Models with Unobserved Heterogeneity," *Journal of Applied Econometrics*, 20: 39-54
- Yehezkel, Y. (2008). "Retailer's choice of product variety and exclusive dealing under asymmetric information," *RAND Journal of Economics*, 39(10): 115-43