The Dynamics of Technology Adoption and Vertical Restraints: an Empirical Analysis*

June 2018

Work in progress, please do not quote without permission

Abstract

This paper examines an aspect of exclusive deals that has not, to the best of our knowledge, been addressed by the literature: their interrelated, dynamic effects among downstream customers. We study these effects in the x86 processor industry, where Intel, a dominant upstream supplier, competes with a smaller contender, Advanced Micro Devices (AMD). Several regulatory agencies worldwide have asserted that Intel offered downstream clients rebates and subsidies that were sometimes conditioned on their purchases from AMD. We use publicly available documents from the relevant antitrust cases to create indices that capture the scope and extent of such exclusionary restrictions. Combining these with market data, we analyse the impact of the restrictions on the downstream adoption of the AMD technology. Estimated dynamic panel models imply that adoption 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's exclusionary practices and AMD's production capacity. Taken together, these results illustrate the role played by exclusionary restrictions in shaping downstream expectations regarding the payoff from adopting a smaller competitor's technology.

Alon Eizenberg
The Hebrew University of Jerusalem and CEPR

Michelle Sovinsky
University of Mannheim and CEPR

Andras Pechy
Accenture

*Toby Webber provided excellent research assistance. We are grateful for helpful comments from Steve Bond, Jan Boone, Greg Crawford, Brett Gordon, Phil Haile, Saul Lach, Vincent Mak, Konrad Stahl, Steve Tadelis, Yona Rubinstein, and to seminar participants at APIOS (Asia-Pacific Industrial Organization Conference), Barcelona Applied IO Workshop (2016), Cambridge, CREST Paris, Dusseldorf (DICE), EARIE (Lisbon), Edinburgh, EEA (Geneva), European University (EUI), ETH Zurich, German IO meetings (2016), Gothenborg, Hamburg, Hebrew University of Jerusalem, IDC Herzlia, IIOC (2016), Mannheim, Oxford, Pompeu Fabra, Tel Aviv, and Tilburg. We also thank Brett Gordon for sharing his data with us for earlier years. This project was supported by the Falk Institute for Economic Research (Eizenberg) and the European Research Council Grant #725081 FORENSICS (Sovinsky).

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 firms often engage in exclusive deal contracts with their downstream customers. While such deals can result in foreclosure of a competitor, they may also have procompetitive effects. For example, they may induce a downstream customer to improve customer service, or help secure upstream investments by preventing downstream free-riding. As a consequence of these conflicting motives, the US courts treat exclusive dealing under a "rule of reason" approach.[1] The issue has motivated a vast literature that studies the impact of exclusive deals under various market conditions.[2]

This paper explores the possibility that exclusive deals can interact with industry dynamics in a manner that has not, to the best of our knowledge, been addressed in the literature. Specifically, we consider an upstream industry featuring a large incumbent supplier, facing a smaller entrant. In this environment, decisions by downstream customers to adopt the technology of the small upstream entrant are of an interrelated, dynamic nature. As a consequence, exclusive deals between the dominant upstream supplier and certain downstream customers can affect the incentives of *other downstream customers* to adopt the smaller supplier's technology. This effect may operate via both demand and supply side channels.

On the demand side, limiting adoption by some customers can slow down the process by which the smaller supplier's technology is "legitimized" with final consumers, thus reducing the perceived value for other downstream firms from adopting it. On the supply side, several mechanisms may be at play. First, exclusive deals between the incumbent and any single customer limit the smaller supplier's sales and cashflow. This may signal to other downstream customers that this supplier would fail to finance Research and Development and capacity expansion, once again lowering the perceived payoff from adopting its technology.

Second, downstream competition in the final product market may drive some downstream customers to either "imitate" by lowering, or "differentiate" by expanding their adoption of

[1] 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). Their *per se* illegality has been rejected in Standard Oil Co. v. United States (Standard Stations), 337 U.S. 293, 305 -06 (1949). The rule of reason approach was reaffirmed in Tampa Elec. Co. v. Nashville Coal Co., 365 U.S. 320 (1961)

[2] See, for example, Marvel (1982), Mathewson and Winter (1987), Bernheim and Whinston (1998), Besanko and Perry (1993), Yehezkel (2008) and Spector (2011).

the smaller supplier's technology when adoption by any downstream customer is restricted by exclusive dealing with the incumbent. Third, the upstream incumbent may retaliate against downstream customers who adopt the small supplier's technology. A downstream customer may expect the scope of retaliation to be linked to the degree to which other customers are exclusive with the incumbent, again reinforcing a negative effect of exclusive deals with some customers on the willingness of other customers to adopt the entrant's technology.

Taken together, these potential channels suggest that the effect of exclusive deals on downstream choices may be quite subtle. In this paper, we empirically explore such mechanisms in the x86 microprocessor market, where Intel, traditionally controlling about 80 percent of the market, competes with a smaller rival, Advanced Micro Devices (AMD), that serves much of the remaining 20 percent. We focus on a time period when Intel's relationship with prominent downstream customers was characterized by a strong degree of exclusive deals and vertical restraints. We estimate dynamic panel models that capture the effect of such restraints on the downstream adoption of AMD's technology, exploiting sharp variation in the deployment of such restraints.

The microprocessor market offers an appropriate context for the study of such mechanisms. It features an upstream industry that sells critical components to downstream customers — personal computer (PC) makers — who depend on the timely supply of these components. The upstream industry is characterized by rapid innovation, large sunk investments, and capacity constraints. Intel's capacity advantage over AMD has been of first-order importance: shifting a substantial portion of their demand from Intel to AMD would have required downstream customers to rely on AMD to produce timely, large shipments. Excluding AMD from selling to major downstream customers reduced its sales, restricting its ability to finance the necessary capacity expansion.[3] Comments by contemporary industry sources also suggest that blocking AMD's sales to major downstream customers slowed down the process of "legitimizing" its technology in the marketplace, and that downstream customers chose to delay the introduction of AMD-based PCs until other customers do so. For all these reasons, Intel's exclusive deals with some downstream customers may have lowered the incentives of other customers to buy from AMD.

Whether or not such forces were at play in this market, as well as their quantitative importance, are empirical questions. We study these issues using data covering the years from 2002 to 2010. During the studied period, Intel has provided substantial benefits to its

[3] As discussed below, the difficulty to finance new production facilities was perceived as a major obstacle for AMD, as reflected in discussions among its own top executives.

downstream clients, i.e., PC manufacturers such as Dell, HP and Toshiba, in the form of rebates and advertising subsidies via the "Intel Inside" program. This practice has attracted considerable legal attention in the form of lawsuits and regulatory investigations, focusing on allegations that the benefits were sometimes explicitly conditioned on the extent of the client's purchases from Intel's rival (see Lee, Pechy and Sovinsky 2013, hereafter LPS). For example, regulators have asserted that Intel's arrangement with HP specified that the share of HP's business line PCs using AMD's chips was not to exceed 5%, while Intel's arrangements with other customers such as Dell were conditioned on full exclusivity.

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 by regulators. This process gained strong traction during our sample period. Ultimately, many of the proceedings were settled. In 2009, the EU fined Intel 1.06 billion Euro on account of anticompetitive behavior. This decision was upheld by the European General Court on 2014, and is being appealed by Intel. Throughout these proceedings, and until the present time, Intel consistently rejects the allegations that it engaged in anticompetitive practices.

Our analysis does not address the legal question of whether Intel's actions were exclusionary or anticompetitive. Nor do we attempt to measure the overall welfare effect associated with exclusive deals. Indeed, the practice of buying chips exclusively or near-exclusively from Intel may have been associated with procompetitive effects.[4] Our focus, instead, is on exploring the interrelated, dynamic channels via which exclusive deals with some customers affect other customers' choices.

Our empirical analysis begins with the construction of an extensive dataset, combining several data sources. We observe market level data on PC 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 production capacity. We further draw on publicly available legal documents associated with the lawsuits and investigations mentioned above, as well as financial reports, to construct variables that capture the exclusionary nature of Intel's contracting with downstream firms, as asserted

[4] Exclusive deals could have 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. Exclusivity could have also facilitated Intel's ability to capitalize on its own investments, thus promoting innovation.

by both AMD and regulators. Finally, we construct quantitative indices of the extent of antitrust activity targeted at Intel's alleged practices. The joint variation of these variables allows us to identify the dynamic impact of technology, vertical restraints, and the legal environment, on the rate of adoption of AMD's technology. Identification is facilitated by sharp data variation with both time-series and cross-sectional dimensions.

Dynamic panel analysis. We estimate 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. The explanatory variables of primary interest are measures that capture the extent of restrictions placed by Intel on AMD purchases by the client in question, as well as restrictions placed on other clients. Incorporating dynamics into the econometric model is in line with our research questions, and with institutional details. First, the extent of current adoption of AMD's technology affects the future costs of using it. Second, the extent of current adoption may affect future benefits granted to the client from Intel. Third, current adoption depends on the customer's current beliefs regarding AMD's future viability in the market.

Results. Our results indicate, first, that the adoption of AMD's technology by a given downstream client responds negatively to the extent of rebates and subsidies paid to this client by Intel, and to the extent of exclusionary restrictions characterizing the client's relationship with Intel. Second, we find that, controlling for the extent of restrictions and payments directed at the client by Intel, restrictions and payments directed *at other clients* further decrease the client's adoption of AMD's technology. Finally, we find that both AMD's production capacity, and the extent of antitrust legal activity in connection with Intel's practices, had a positive effect on the rate of AMD adoption.

The documented negative response of a client's AMD adoption to restrictions imposed on other clients is consistent with our discussion above of the (demand and supply) channels via which reduced adoption of AMD's technology by one customer may diminish incentives by other customers to do so. Conversely, higher production capacity by AMD causes an upward revision in clients' expectations regarding its future viability as a supplier, consistent with its documented positive effect on current adoption levels. Mounting legal pressure on Intel to curb its exclusionary practices may similarly improve clients' view of AMD, as they signal that the restrictions may soon be lifted. Our results, therefore, support an important role for dynamics in shaping the industry's response to exclusive deals.

The rest of the paper is organized as follows. Following a literature review, section 2 describes the data. Section 3 explains our empirical strategy. 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). Sass (2005) studies the beer market and finds that exclusive dealing is more prevalent in smaller markets, in contrast to the predictions of foreclosure theory models. Asker (2016) 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. 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, and also resulted in reduced spatial coverage.

Ater (2015) empirically quantifies the effect of exclusive dealing 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.[5] A closely related issue concerns the foreclosure versus efficiency motives in the context of vertical integration. In this context, Hortaçsu and Syverson (2007) find little evidence for anti-competitive effects in the cement industry.

Our paper contributes to this line of research by studying a *dynamic relationship* between exclusive dealings, downstream technology adoption, and upstream capacity investments. It therefore connects with a small empirical literature that investigates the role played by client expectations in shaping firm growth trajectories. Atalay et al. (2011) demonstrate that signals of supplier financial distress discourage purchases of durable goods as consumers revise downward their expectations of receiving future services in the event of supplier bankruptcy. Our framework features downstream customers who revise their assessment of the future viability of a supplier, thus capturing a similar mechanism.

More broadly, our paper contributes to studies of the dynamic process that entrants undergo en route to closing the size gap versus established incumbents. Providing a complete survey of this literature is beyond the scope of this paper, but we provide a few examples. Jovanovich and Lach (1989) consider the dynamic process of entry, exit and S-shape diffusion in the presence of learning by doing. Financing constraints, which play an important role in our context, are modeled by Cabral and Mata (2003). Foster, Haltiwanger and Syverson (2016) emphasize the role of informational or reputation frictions that must dissipate for

[5] Additional contributions include Slade (2005) and Suzuki (2009).

customers to embrace a newcomer.[6] Our paper speaks to these issues and argues that exclusive deals between incumbents and key downstream clients can frustrate an entrant's effort to expand its customer base and close the market share gap.

Our paper also belongs in 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 its impact on the variety of downstream product configurations offered to consumers. 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 on innovation of a hypothetical exclusion of AMD from the market. Our work differs from these previous contributions in terms of both questions and methods. By relying on dynamic panel methods rather than on structural modeling we restrict our ability to analyze out of sample scenarios. Yet, this choice allows us to account for rich product-level factors without running into large state space concerns.

2 Data

We use data from several sources, containing information on PC and CPU sales and attributes, PC firms' advertising expenditures, measures of CPU quality, CPU makers' production capacity, the restrictions characterizing Intel's vertical contracts, and the scope of the legal action taken in connection with these restrictions.

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 2002-2010.[7] Following Sovinsky Goeree (2008), we define our unit of observation as a combination of PC vendor (e.g., Dell), PC vendor brand (e.g., Inspiron), market segment (e.g., Home), CPU vendor (e.g., Intel), CPU family (e.g., Pentium 4) and quarter, yielding 3,280 observations. We exclude Apple products as those exclusively used IBM's chips during much of the sample period (using Intel's chips afterwards).

[6] See Bar-Isaac and Tadelis (2008) for a review of the literature on such mechanisms.

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

The Gartner data reveal the number of units sold, and revenue.[8] We can therefore compute the average price by dividing revenue by the number of units. Furthermore, we use the number of units sold to compute, for each combination of PC brand-segment-quarter, the percentage of units sold with an AMD processor installed. We refer to such brand-segment combinations as PC product lines.

Table 1 shows descriptive statistics, with Panel A pertaining to PC characteristics. PC prices display significant variation, ranging from $241 to $3,521. Statistics regarding downstream advertising expenditures are also provided. The advertising data come from the Kantar Media Group, and 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] As the table shows, brand-specific ad expenditures averaged $1 million while firm-specific expenditures averaged $8 million. Importantly, the "Intel Inside" program provided rebates that were a function of advertising by PC firms.

As reported in Panel A of Table 1, the rate of utilization of AMD's chips, averaged across observations, is 13 percent. Figure 1 goes beyond this statistic and displays the evolution of AMD's market share over time. The rise in market share, around the years 2005-2006, from 10 percent to 20 percent is of interest. We discuss below several developments that took place during the relevant time period, involving both technological trends and strategic changes in the competitive arena between Intel and AMD. In particular, the later part of the sample may have been associated with a loosening of Intel's restrictions on its downstream clients. We shall explain how our data ultimately allow us to disentangle the effect of those restrictions from these other factors. As evident in the figure, the market share gain was temporary, with AMD's share declining to 12.5 percent by 2009.

Figure 2 displays AMD's market share within selected customers, as well as these customers' shares of the PC market. Mean (over quarters) market shares are displayed, before and after the first quarter of 2006. The figure shows that most downstream makers have increased their adoption of the AMD technology. For some, like Dell and Toshiba, the initial level was zero, as they have been exclusive with Intel in the earlier part of the sample. Overall, the rate of increase in AMD adoption varies substantially across firms. The joint variation of those changes and of the changes in the nature of the contracts between Intel and these firms are helpful in identifying the effects of Intel's restraints.

[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.

[9] For a few PC firms, the Kantar brands were available at a more aggregate level than the Gartner sales data. In these cases, we performed the match at the Kantar-reported brand level.

2.2 CPU quality

An important determinant of the adoption rate of the AMD technology is the quality of its products vis-a-vis those of Intel. We therefore obtain data on CPU quality from Passmark's CPU Mark publications.[10] This company collects data from CPU tests conducted by users around the world, and creates a "benchmark" score: a continuous measure of performance for each CPU model.

We gain some perspective on the value delivered by various CPUs by computing a benchmark-per-dollar measure, i.e., by dividing the benchmark measure by the CPU price. This required us to obtain data on CPU prices. We obtained such price data by combining information from published list prices for Intel and AMD with price data obtained from Instat.[11] For PC product lines that use both Intel and AMD chips, we compute both an AMD- and an Intel-based benchmark-per-dollar index for each PC product line. To compute the AMD-based index for a given product line, we aggregate over the various AMD chips used by the product line by taking the sales-weighted average index over such chips, and analogously for Intel. If the PC product line exclusively uses Intel chips, we substitute the average AMD benchmark-per-dollar across all product lines for that observation. Complete details regarding the construction of those variables are provided in the Appendix.

Panel B of Table 1 reports descriptive statistics. On average, the AMD benchmark scores were as much as 45 percent higher than the Intel ones. We emphasize that these calculations do not compare AMD chips against Intel chips directly, but rather in a way that takes into account the rate of their utilization within downstream product lines.

These indices were subject to important temporal variation, displayed in Figure 3. AMD's quality measure was consistently higher than Intel's, especially in the earlier sample years. As of 2006, however, Intel's quality 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 has had much to do with the introduction of new generations of Intel chips — specifically, Intel Core product family — that offered substantial improvements over incumbent generations (noting that AMD's lower prices imply that a tie in terms of benchmark per dollar is equivalent to a quality advantage for Intel). These observed patterns are consistent with statements from relevant antitrust cases.[12]

[10] Accessed from www.cpubenchmark.net.

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

[12] For example, the European Commission 2009 decision against Intel cites a 2002 internal HP presentation stating that AMD's Athlon desktop processor "had a unique architecture" and was "more efficient on many

Considering both Figure 1 and Figure 3, we obtain a preliminary view of a surprising, weak correlation between AMD's product quality and its market share. AMD's market share was relatively low in the earlier part of the sample, when it seemed to have had a better value proposition than that of Intel's, and its market share actually went up around the time that Intel started to regain a technological edge. Indeed, product quality may not have been the only, or the most important determinant of competitive advantage in this industry. Intel's ability to engage in exclusive deals, as well as its inherent advantage in production capacity, may have allowed it to prevail through periods in which AMD offered superior product quality. Nonetheless, controlling for these shifting patterns of technological leadership is imperative for identifying the impact of these other factors.

Finally, we also create a variable to account for the time (in quarters) in which a PC brand has been available equipped with AMD chips (respectively Intel). As reported in Panel B of Table 1, on average across observations, an AMD-based (Intel-based) PC brand configuration lasted 5.6 (6.8) quarters on the shelf, respectively.

2.3 CPU production technology and capacity

Upstream production capacity plays an important role in the microprocessor industry. Intel and AMD's annual financial reports indicate their number of fabrication units (FABs), the silicon wafer size at each FAB (the larger the wafer, the more CPUs can be printed simultaneously), and the Integrated Circuit process at each FAB (capturing precision in nanometers — the smaller the precision, the more CPUs can be printed, and additionally, CPU power efficiency is improved).

The multiple dimensions of the production process — FABs, wafer sizes and IC processes — motivate us to develop a production capacity index. We define a processor maker's *capacity index* by a weighted sum of its FABs, where the weight of each FAB is the sum of its ranked wafer size and its ranked IC process. That is: we rank FABs by their IC process (largest to smallest) and wafer size (smallest to largest), and then sum these ranks over all FABs of the manufacturer in question. As the information is gleaned from annual reports, the capacity indices are measured at an annual frequency.

Panel C of Table 1 provides summary statistics on these indices, showing that Intel's mean (over time) capacity index is four times larger than AMD's. Table 2 shows the evolution

tasks," adding that AMD offers "no-compromise performance at superior value." The same EC decision states that "... Intel has made references to having recently 'caught up' with AMD following the launch of its new generation of CPUs based on the 'Core' micro-architecture" (COMP/C-3 /37.990 - Intel).

of the processor makers' production technology and capacity over time. AMD usually lags behind Intel regarding the IC process and the wafer size.

We also use AMD's quarterly financial reports to obtain the cash it had available for investment at the beginning of each quarter.[13] Panel C of Table 1 reveals that the free cash available for AMD in each quarter was on average $833 million. To provide perspective, the cost to build an advanced FAB in 2007 was about $5 billion (Brown and Linden, 2009). 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 facilities.

AMD encountered challenges in financing the construction of production facilities. An AMD executive who left the company stated that "(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]

Other sources emphasize the significance of Intel's production capacity advantage. A blogger following the industry commented in 2002 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] 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 2009 State of New York case against Intel stated that "(a)ll major computer manufacturers depend on Intel in a variety of ways and are reliant on it for microprocessors, since AMD is, and in the foreseeable future will remain, unable to fulfill more than a small share of their requirements."[16]

[13] The quarterly filings were accessed on September 18, 2014 from http://ir.amd.com/

[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 parenthesis 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).

[16] Source: https://www.intel.com/pressroom/legal/docs/NY_AG_v._Intel_COMPLAINT.pdf

2.4 Exclusive deals and other vertical restraints

We next describe our construction of data regarding Intel's exclusionary restrictions. Several challenges are associated with this task. First, no single, authoritative source exists with respect to these restrictions. Rather, what is observed are legal documents from a number of legal proceedings conducted in connection with Intel's practices. These include lawsuits by AMD, and cases brought by the EU, the Japanese Fair Trade Commission, the Korean Fair Trade Commission, the State of New York, and the US Federal Trade Commission. Table 3 provides some basic information regarding the timing of these proceedings.

Second, while these documents provide very rich information regarding Intel's restraints, they are by no means comprehensive. The text of the cases brought against Intel often states that they include examples, rather than a complete list, of Intel's practices. Finally, Intel consistently rejects the assertions brought about in the aforementioned cases.[17]

Nonetheless, these cases contain very detailed information, much of it based on internal documents retrieved by regulators from the companies involved, that we wish to exploit in our econometric study of the dynamic aspects of vertical restraints. The restraints span a wide variety of instruments via which the adoption of AMD's technology could have been affected. We use this information to define, at the customer-quarter level, binary indicators that take the value 1 if the specific instrument were used, and zero otherwise.

Specifically, we construct indicators for the following restrictions on the customer's use of AMD's technology: caps on the amount sold of AMD-based PCs; exclusion of AMD from certain product lines or delayed launch of specific AMD-based machines; 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; restrictions imposed on bidding on contracts using AMD-based products; threats to remove funding, divert funding to rivals, or other retaliation, as a consequence of selling AMD-based PCs; and guarantees of preferred supply of Intel CPUs.[18]

[17] Intel chose not to contest the charges brought by the Japanese Fair Trade Commission in 2005. Several cases ended in settlements in which Intel did not admit the charges. In particular, Intel paid AMD 1.25 billion US dollars in 2009 to settle AMD's lawsuits. The Korean Fair Trade Commission ruled against Intel in 2008 and upheld that decision against Intel's appeal in 2013. The European Commission fined Intel by 1.06 billion Euro in 2009, a decision which was upheld by the court in 2014, but is still being appealed by Intel.

[18] The latter benefit was particularly valuable to downstream customers, for, as noted by one industry insider, "(s)urvival practically depends on being able to get allocations of the newest chips which are always in short supply coming out of the gate." See http://www.zdnet.com/article/so-dells-not-a-wintel-lapdog-it-

Given the multi-faceted nature of the restraints, we again resort to an index approach to quantify their presence and temporal evolution. Our first measure counts the number of restrictions imposed on a downstream client within a given quarter. A second measure only counts restrictions we define as extreme: a zero cap on AMD-based machines; exclusion from certain product lines or delayed launch of specific AMD-based machines; threats or retaliation; or a commitment to increase Intel's share of the PC producer's CPU purchases.

Naturally, we can only observe the restrictions through the lens of what was asserted by the plaintiffs in legal proceedings. Our indices are therefore subject to measurement error. At the same time, the plaintiffs had an incentive to provide detailed information on what they believed to be the important aspects of Intel's restrictions, and often based their case on detailed internal documents. Moreover, if the measurement error causes an attenuation bias, then one may interpret our findings as providing lower bounds on the true effect of the restraints. Noting that we find considerable effects, this strengthens our results.

Panel D of Table 1 reports statistics on the first measure, and reveals that, on average across brand-segment-quarters, 1.6 restrictions were in place while the maximum was 6. Figure 4 shows the number of brand-segment data cells that were affected by different numbers of restraints over the sample period. The recorded restrictions took place until 2007, with no observed activity afterwards. This sharp variation, controlling for other factors such as the processor quality offered by the upstream firms over time, is helpful in identifying the effect of the restrictions in our econometric model.

One natural concern is that our indices are biased by truncation error: while our product market data runs through 2010, many of the legal documents we rely on are dated prior to that year, suggesting that there may have been restrictions after 2007 that we simply do not observe. This possibility does not appear to be supported by institutional details. In a May 2010 document, the FTC specifies information on Intel's actions, some of which taking place in 2008 and 2009.[19] However, those pertain to products such as GPUs, chipsets and compilers, rather than to CPUs. The document does not specify CPU related restrictions after 2007. So while there are some claims (e.g., in the case brought about by the State of New York) that the restrictions we cover continued beyond 2007, there is no specific information about this possibility. It therefore does not appear that the variation in our measures is driven by a truncation problem in the process that generates the data.[20]

should-probably-thank-amd/accessed November 1, 2017.

[19] See https://www.ftc.gov/sites/default/files/documents/cases/2010/05/100517intelmemoopphpmoquash.pdf.

[20] The European Commission 2009 decision (ibid.) also lends support to the notion that the restrictions by Intel were more intensive in the earlier part of our sample period: "Most of the individual abuses concerned

The restrictions indices also display useful cross-sectional variation. Differential restrictions applied to different customers, as well as to different market sectors. For example, HP was subject to restrictions on the amount of AMD-based machines it was permitted to sell in the business sector. As a consequence, we observe both product line level, and firm level variation in the number of restrictions in place.

One of our research hypotheses would be that adoption of the AMD technology by downstream clients was affected by restrictions on other clients. This begs the question of whether one client was actually aware of restrictions on other clients, given that contracts between Intel and its downstream customers were not public. We argue that clients were able, and in fact did monitor the process of adoption of the AMD technlogy by rivals. We provide such evidence in section 3.1 below.

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. The payments to Dell were substantial and deviated considerably from the official description of the program provided on Intel's website. This official description 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.[21] For the period after 2007, we compute Intel's payments to Dell as 3 percent of Dell's CPU costs, computed using Gartner sales data and the price dataset described previously. We apply the 3 percent approach to compute the payments to 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 payments. Payments to Dell are displayed on the left axis (in 100 millions) and the average payments to the other PC firms are on the right axis (in millions). 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, the average per-quarter per-firm payment varied over time between \$2.3 and \$5.5 million.

An index of antitrust activity. Finally, we construct variables that capture the scope and magnitude of the legal activity described above. Our measures count both the cumulative number of antitrust cases brought against Intel by AMD and regulators worldwide as of 2001, and the number of pending cases (see the appendix for additional details — TBC). We

are concentrated in the period ranging from 2002 to 2005, whilst, after the end of 2005, at most two individual abuses occur simultaneously at any given point in time."

[21] For further details about the program, see for example Lee, Péchy and Sovinsky (2013, LPS). As reported there, at some point, payments from Intel represented a substantial portion of Dell's income.

shall examine below how this legal activity affects the dynamic adoption process of AMD's technology. The use of the cumulative number of cases allows for the possibility that market participants take note not only of pending cases, but also the critical mass generated by the overall extent of antitrust activity. Panel D of Table 1 shows that, on average across all observations, there were 3.47 pending cases, and 4.22 cumulative cases. The temporal evolution of these measures is described in Table 4.

2.5 Data patterns: a summary

Upon conclusion of our data section, we wish to highlight the main data patterns presented above, and the manner with which they motivate our research hypotheses.

Several important phenomena seem to have taken place as of 2006-2007. First, the rate of adoption of AMD's technology picked up significantly (but then decreased again towards the end of the sample period). Second, AMD's benchmark-per-dollar advantage began to erode as Intel regained a technological edge. Finally, still around that point in time, there appears to be a decline in the deployment of exclusion restrictions by Intel, as captured in the antitrust cases.

A hypothesis that emerges from this joint data variation is that Intel's restrictions have slowed down the rate of adoption of AMD's technology. This hypothesis is supported by the notion that, despite Intel's technological advantage picking up midway through the sample period, AMD's adoption started to rise around that point in time. Moreover, the rise in AMD's share took place at different rates among downstream customers whose transactions with Intel were characterized by differential degrees of vertical restraints. This variation will help us test another hypothesis: that restrictions placed on other clients were also important to the decision making of a given client, consistent with dynamic effects.

Other data patterns indicate that Intel enjoyed a capacity advantage throughout the sample period, and that antitrust activity in connection with its practices has been expanding. We return to the significance of these issues in the next section, where we spell out our hypotheses, present our formal econometric model, and discuss identification.

3 Econometric Model

3.1 A dynamic game framework with implied research hypotheses

We derive our research hypotheses by casting the choices made by downstream PC manufacturers in a dynamic game environment, which we describe in an informal fashion. In this game, downstream clients face two upstream input suppliers: Intel, a large incumbent, and AMD, a smaller contender. The downstream client maximizes its discounted stream of payoffs by adjusting a single, continuous control variable: the fraction of inputs purchased from the smaller supplier. This fraction is chosen in each period given the values of state variables, some of which are endogenous.

The intensity of purchases from the smaller supplier affects the customer's flow payoff function via multiple channels. First, the smaller supplier charges lower input prices, allowing the customer to reduce its marginal cost. Second, this supplier offers a technology which is differentiated from that of the incumbent supplier, affecting the utility and willingness to pay of final consumers. Third, spillovers could arise from the product market interaction among downstream customers. For example, the introduction of AMD-based PCs by one downstream firm may increase the visibility of the AMD technology in the marketplace, affecting the payoff to other firms from offering such machines. Those other firms' incentives may be further complicated by equilibrium effects that we discuss below.

Dynamics and state variables. While increasing the fraction of inputs purchased from the small supplier can shore up margins and boost the downstream customer's flow payoff by reducing marginal costs and offering more variety to final consumers, it also involves several types of costs. First, the customer may be giving up economies of scale associated with buying all or most of the input from a single supplier. In our empirical context, we note that Intel and AMD chips are not "pin compatible." As a consequence, beginning to use AMD chips, or expanding their use, requires a certain degree of investment: the client must learn how to configure the hardware to AMD's specifications, set up a production line that installs AMD chips, and develop working relations with the supplier.

Such investment is likely to involve a cumulative, dynamic process. 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 endogenous fraction of inputs that the client purchased from the small supplier in period $t-1$ is therefore one of the state variables affecting its adoption decision in period t.

Some other costs associated with expanding the use of the smaller supplier's inputs are related to the effect on the client's relationship with the incumbent supplier. The client's flow payoff may be adversely affected if, by reducing the extent of purchases from the incumbent, rebates and other benefits from this supplier are forgone. More broadly, both current and future payoffs may be affected if the expanded use of the small supplier's input violates restrictions placed by the incumbent supplier. In our context, the client may lose favorable terms of trade with Intel, such as guaranteed preferred supply of top of the line chips, lose payments provided by Intel (which, as in the case of Dell, may be substantial), or worse, see those benefits being diverted to rivals. As a consequence, the restraints placed by Intel on the client are another state variable affecting its choices.[22]

The extent of antitrust activity taken against the deployment of such restraints is another state variable: this activity affects the perceived future viability of Intel's restraints, and therefore, the perceived future costs of purchasing from AMD today.

The benefits from the expansion of purchases from AMD are functions of the value embedded in this supplier's input, versus the value offered by Intel. Empirically, we capture these values via the benchmark per dollar measures described in section 2.2. The higher is AMD's benchmark per dollar versus that of Intel, the more its technology can help the customer provide better quality to final consumers per dollar spent on CPU purchases. These measures therefore affect the customer's margins and are relevant state variables.

Yet another state variable is the smaller supplier's production capacity. As reviewed in detail in section 2.3, Intel enjoyed a substantial production capacity advantage which may have allowed it to prevail through periods in which AMD enjoyed a technological edge. For a downstream client to be willing to expand its reliance on AMD, it must believe that this supplier would be able to meet its current and future levels of input demand. Downstream PC manufacturers rely on thin inventories and their survival depends on the ability to receive large shipments of chips within short time frames. Developing a strategic reliance on AMD's chips is therefore more likely given expectations that it would be able to deliver large production volumes in the future. This is especially true if the client expects that supply from Intel could be jeopardized on account of retaliation by Intel. AMD's capacity, as well as the cash available for its investment activities, are therefore both state variables that affect the willingness to adopt its technology in the current period.

[22] Defining the client's continuous control as the fraction of chips purchased from AMD accords with the nature of Intel's vertical restraints, as asserted in the case files, which sometimes specified a cap on the extent of AMD chips used in each segment as a condition for eligibility to benefits from Intel.

Restrictions imposed by the incumbent on *other downstream clients* are also a relevant state variable in the client's decision. This holds true for a number of reasons. First, lower adoption of the AMD technology by other downstream customers will exacerbate the challenges of the smaller supplier in financing future investments in research and development and capacity expansions, making it less attractive to adopt its technology today.[23]

Second, the adoption of the AMD technology by some downstream PC makers, especially if they sold the AMD-based machines to large corporate customers, could "legitimize" this technology as a valid substitute to that of Intel.[24] Restrictions that limit the adoption of the AMD technology by specific downstream PC makers can therefore reduce the perceived benefits to other PC makers from using AMD chips.

Third, equilibrium effects suggest that restricting the AMD adoption by one customer may increase or decrease the incentives of other customers to adopt the technology. Other customers may choose to either imitate that customer by reducing their own adoption, or to differentiate themselves in the final consumer market by expanding it. Nonetheless, the case files discuss another equilibrium effect which operates unequivocally to reduce adoption given a restriction on other customers.

The State of New York case discusses downstream incentives as a "...'prisoner's dilemma': If all of the OEMs had been willing to deal with AMD without Intel-imposed restrictions, the resulting strengthened competition would have benefited them all, as well as consumers, by lowering their microprocessor costs. Nevertheless, there were strong — often overwhelming — incentives for any individual OEM to accept the pay-offs – and avoid the punishments – which Intel dealt out." It goes on to describe IBM executives as having "grave concerns" that "...if IBM were first to market with Opetron-based server products, IBM would be

[23] The European Commission 2009 decision (Ibid.) states: "The emergence of AMD as a competitive threat to Intel was dependent on the availability of investors willing to finance risky investments in research and development as well as AMD production facilities. Such investments are only undertaken when there is a prospect of an adequate return if the research and development is successful and well implemented. Given Intel's conduct, AMD's products did not reach final customers in the volumes that their quality and price would have justified had competition been exclusively on the merits."

[24] The European Commission 2009 decision states that "Intel itself expressed concern that success for AMD with HP corporate desktops would lead to a "spill –over possibility of ...products into corporate space 'legitimizing' AMD platforms." Another quote from the decision asserts that "...the largest OEMs have a greater ability to legitimise (that is to create consumer trust in the capabilities of a new product) a new x86 CPU in the market, and hence provide an important springboard for a x86 CPU supplier that wants to significantly increase its penetration in the market." The State of New York case quotes an HP executive who noted in 2002 that "Intel's worst fear will be a sufficient ramp of commercial Athlon [an AMD microprocessor product] such that it becomes legitimized for commercial markets. Once AMD is out of the box, Intel cannot put it back in." Quotes in parentheses appear in the original text.

particularly exposed to Intel." The European Commission 2009 decision quotes a 2004 e-mail in which a senior Acer executive assures Intel of Acer's commitment that "...Acer will stop both flyers and advertisements for any Acer sub-brand K8 notebook worldwide from now on, until any other major brand, such as HP, Toshiba, Sony, Fujitsu and Fujitsu-Siemens, or similar class, announces their K8 notebook."

In other words, it appears that a restriction that ensured that some downstream customers limit their AMD purchases would have operated to diminish AMD purchases by other downstream customers. Note also that the quotes above show that downstream customers monitored the adoption of AMD by their rivals, so that even if they were not aware of the precise restrictions on those rivals, their actions could still be affected by such restrictions.

Hypotheses. The above description characterizes a downstream client's *policy function*, i.e., the function that determines how the client adjusts the fraction of chips purchased from AMD given the values of state variables. It also implies predictions for the effects of these state variables on that policy function. Those are summarized in Table 5 below.

We do not formally solve the model or take it to the data in a structural estimation exercise. Rather, we test our predictions for the policy function by estimating the relationship between the continuous control (fraction of chips purchased from AMD) and the state variables. We do this via dynamic panel techniques that allow us to acknowledge the cumulative nature of the technology adoption process, while employing instruments to deal with endogenous state variables.

Before proceeding to the econometric implementation, we complete this section with a couple of comments on the dynamic framework described above. First, our description has focused entirely on choices made by downstream clients in response to actions (restraints, innovation, pricing, capacity expansion) on part of upstream suppliers. A more complete characterization would also model these upstream choices as endogenous. Importantly, our empirical implementation does not treat these choices as exogenous.

Second, our discussion has emphasized channels via which downstream input choices can be interrelated across clients. These interdependencies can lead to multiple equilibria: in one equilibrium, adoption of AMD may be low, and no individual customer would have an incentive to expand it. In another equilibrium, downstream adoption may be widespread. What we observe in the data appears to be the first, rather than the second equilibrium.

It follows that in a dynamic environment characterized by large sunk investment and capacity constraints, exclusive deals can divert the industry into a less competitive equilibrium. This need not result in complete foreclosure: AMD did not exit the market, but it failed to surpass a 20 percent market share level throughout the studied period.

3.2 Formal econometric setup

We observe a cross-section of product lines (defined as PC brand-segment combination, e.g., Acer's Aspire for the Home Market), indexed by i, over time periods (quarters) $t = 1, ..., T$. Our dependent variable is the fraction of product line i's sales that have an AMD chip installed at time t, denoted w_{it}. Our model for the PC firm's choice is given by

$$w_{it} = \alpha w_{it-1} + \beta x_{it} + \lambda z_t + \eta m_{it} + \delta r_{it} + \gamma l_t + c_i + \varepsilon_{it.} \tag{1}$$

The lagged percentage of segment-brand i sold with an AMD processor, w_{it-1}, captures the state dependence between current and past decisions. The c_i term captures unobserved heterogeneity at the brand-segment level, while ε_{it} is an idiosyncratic iid error term.

Time varying observed characteristics are captured by x_{it}, taking a sales-weighted average of those characteristics across the individual products offered within the product line.[25] Included in x_{it} are the PC price, brand level advertising, and firm level advertising. Those characteristics are included since the AMD technology may be differentially attractive across product lines that differ in such characteristics. In particular, a downstream producer may find it attractive to use the cheaper AMD chips in thin-margin "value" products.

Variables relating to CPU manufacturing capacity are included in z_t. Both Intel and AMD capacity indices are included, where those are functions of the number of FABs, wafer size and IC process (see section 2).[26] We also include the (lagged) amount of free cash available to AMD. Our assumption, supported by institutional details, is that Intel was much less subject to cash constraints than AMD.

CPU related variables are contained in m_{it}. These include the extent of technological progress as measured by the benchmark per dollar indices for Intel and AMD, respectively. We additionally capture the age of the CPUs used by the product line via the number of quarters in which the segment-brand-CPU family combination has been available.[27]

[25] Time-constant product line characteristics are not included as those would be absorbed by c_i.

[26] The relationship between AMD's production capacity and its sales is not a mechanical one, even if this supplier were to operate at capacity. Our measure of AMD's capacity refers to its overall production operations, used to generate worldwide sales, while our dependent variable relates to the fraction of downstream demand directed to AMD's product in the US alone. This alleviates endogeneity concerns and suggests that the capacity variable may indeed capture the subtle effect on client expectations discussed above.

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

The vector r_{it} captures upstream vertical restrictions, and is therefore the main explanatory variable of interest. We include indices that capture the intensity of exclusionary restrictions imposed on the relevant firm, and the number of such restrictions imposed on rival firms. We also include such indices that capture only the number of extreme restrictions (see section 2). We also include total payments Intel made to the relevant PC maker, and payments made to Dell. Finally, the vector l_t captures our legal environment variables: the cumulative and pending numbers of antitrust cases brought against Intel.

Identification. Our primary goal is to identify the causal effect of upstream vertical restraints on downstream input choices at the product line level. A particularly important issue is the endogeneity of these restraints: one would expect them to be correlated with the terms c_i and ε_{it}. Specifically, Intel may have been setting these restraints in response to unobserved factors affecting the downstream client's demand for AMD's chips.

Our identification strategy takes advantage of the panel aspect of our data in addressing unobserved client heterogeneity in the incentives to adopt AMD's technology. Unobserved heterogeneity may arise if some firms are fundamentally better-suited to gain from using AMD's chips, and may be reflected in aspects of the demand faced by different downstream PC makers, or in the flexibility of their production processes. For example, some firms enjoy large economies of scale from using a single type of chip, and their incentive to adopt the AMD technology in addition to that of Intel may be relatively low. Nevertheless, simply accounting for the time-invariant unobserved heterogeneity term c_i within the model would fall short of addressing endogeneity concerns.

Examining equation (1) we note that a standard fixed effect estimator, while allowing for arbitrary correlation of the unobserved heterogeneity with the explanatory variables, would require a strict exogeneity assumption with respect to the variables $y_{it} \equiv \{w_{it-1}, x_{it}, z_t, m_{it}, r_{it}, l_t\}$. That is, consistent estimation would require us to assume that ε_{it} is mean independent of y_{is} for every $t, s = 1, ..., T$. Such as assumption is, however, immediately violated by the presence of the lagged dependent variable, since y_{it+1} contains w_{it} which is necessarily correlated with ε_{it}. Furthermore, if a current positive shock to the rate of AMD adoption in brand-segment i causes Intel to impose stricter restrictions in the next period, then ε_{it} should be correlated with $r_{it+1,}$ again violating strict exogeneity.[28]

Arellano and Bond (1991) propose instrumenting the differenced variables that are not strictly exogenous with all their available lags in levels. Arellano and Bover (1995) note that lagged levels can be poor instruments for first differences if the variables are close to a random

[28] See, for example, Wooldridge (2002, ch. 10).

walk. They suggest adding the original equation in levels to the system, making it possible to incorporate additional instruments and increase efficiency. The modification they propose is to include lagged levels as well as lagged differences. The Arellano-Bover/Blundell-Bond estimator, known as the system GMM estimator, is our leading specification.

Concretely, first differencing results in:

$$\Delta w_{it} = \alpha \Delta w_{it-1} + \beta \Delta x_{it} + \lambda \Delta z_t + \eta \Delta m_{it} + \delta \Delta r_{it} + \gamma \Delta l_t + \Delta \varepsilon_{it}. \tag{2}$$

Strict exogeneity would imply that Δy_{it} is exogenous and hence can serve as its own instrument. 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 $E(\varepsilon_{it} \mid y_{is}, \mu_i) = 0$ for $s \leq t$. It is possible to use w_{it-2} and y_{it-1} as instruments in the first-differenced equation, which yields an overidentified system.[29] The system GMM estimator adds additional moments generated by the levels equation (1). Following Arellano and Bover (1995) we can use lagged differences Δw_{it-1} as instruments for the lagged regressor w_{it-1} in the levels equation. Blundell and Bond (1995) show that these additional moment restrictions hold under the assumption that the mean of the initial observation of w equals its steady state.[30]

In our application of this system GMM estimator, we incorporate additional instruments for the endogenous PC price. These are motivated by the observation from Berry, Levinsohn and Pakes (1995) that markups, and hence prices, should be correlated with characteristics of rivals products. Specifically, we use the number of brands offered by rival firms in the quarter as an instrument. In Appendix (6.2) we present the first-stage estimation results showing that these instruments are not weak.

Additional specifications. In addition to the system GMM estimates, we also present below results from a standard fixed effects estimator, as well as from a pooled OLS estimator. While the discussion above suggests that those specifications do not generate consistent estimators in our application, we shall demonstrate that our main economic conclusions hold under these two alternative models as well.

The system GMM specification allows us to control for both state dependence, and for

[29] See Arellano and Bond (1991). We 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}. We also conduct AR(3) tests to make sure there is no remaining serial correlation in the errors.

[30] While formally testing this assumption appears far from trivial, recall from Section 2 that the mean rate of AMD's adoption was maintained in the 10 to 20 percent range, suggesting that the steady-state assumption may be reasonable.

individual heterogeneity. it does not allow us, however, to address a "corner solution" issue: some product lines, at different times, used no AMD chips at all. To address this issue, we estimate a nonlinear, dynamic Tobit-like model following Wooldridge (2002, 2005) whose model builds on insights from Chamberlain (1984). This approach treats the time-constant heterogeneity c_i as random effects. Its main drawback is that, unlike the linear approach outlined above, it does not allow us to relax a strict exogeneity assumption.

The dependent variable w_{it} is modeled as having a mass point at zero:

$$w_{it} = max\big(0, \alpha w_{it-1} + \beta x_{it} + \lambda z_t + \eta m_{it} + \delta r_{it} + \gamma l_t + c_i + u_{it}\big) \tag{3}$$

$$u_{it}|(\overline{y}_i, w_{i,t-1}, ...w_{i0}, c_i) \sim N(0, \sigma_u^2), \tag{4}$$

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 for a cross-sectional unit (i.e., brand-segment combination) i is $\overline{y}_i$.

With respect to the initial value of w_{i0}, we follow Wooldridge (ibid.) and specify the density of the unobserved heterogeneity terms conditional on w_{i0}:

$$c_i = \psi + \xi_0 w_{i0} + \overline{y}_i \xi + a_i,\ \ a_i|\,(w_{i0}, \overline{y}_i) \sim N(0, \sigma_a^2). \tag{5}$$

The heterogeneity terms c_i can therefore be integrated out to yield the likelihood function of the random effects Tobit model with time-t, observation-i explanatory variables: $(y_{it}, w_{i,t-1}, w_{i0}, \overline{y}_i)$. That is, $\overline{y}_i$ and w_{i0} are controlled for in each time period. This likelihood function is used to obtain estimates of the parameters $(\alpha, \beta, \delta, \lambda, \eta, \psi, \xi_0, \xi, \sigma_a^2, \sigma_u^2)$.

4 Results

This section presents our empirical findings. Section 4.1 provides system GMM estimates that allow us to test the predictions from our theory, as summarized in Table 5. Section 4.2 then follows with a discussion of economic magnitudes and additional specifications.

4.1 Leading specification: System GMM estimates

We present results for our leading specification, i.e., system GMM estimates of the model in equation (2), in three steps. First, in Table 6, we present specifications that include PC and CPU characteristics (x_{it}, m_{it}), capacity-related variables (z_t), and variables that capture Intel's vertical restraints (r_{it}). Second, in Table 7, we remove the vertical restraints (r_{it}) and include instead the legal environment variables (l_t). Finally, in Table 8, we include all the above variables, i.e., $(x_{it}, m_{it}, z_t, r_{it}, l_t)$. This step-by-step analysis reveals how our variables of main interest — the vertical restraints variables (r_{it}) and the legal environment variables (l_t) — perform both individually, and together with all other variables. Our theoretical predictions from section 3.1 will be tested throughout these three steps.

We begin with Table 6, where the columns present results corresponding to different combinations of included restrictions-related variables (where column (1) includes no such variables). The coefficients in columns (2) through (4) imply that the number of restrictions imposed on the client, the number of such extreme restrictions, and the payments made by Intel to the client, all have a negative effect on the rate at which the client utilizes AMD's technology. These findings confirm our theoretical predictions (1) and (2) (recall Table 5) and show that the restraints were effective in restricting purchases from AMD.

The findings in columns (5) through (7) address the impact on the client of restrictions on *other clients*. Column (5) shows that, conditional on the restrictions imposed on the client, the number of restrictions imposed on other clients has a negative, yet statistically insignificant effect on the client's willingness to purchase from AMD. Column (6) shows that, controlling for the number of extreme restrictions imposed on the client, the number of extreme restrictions imposed on other clients has a negative effect on the client's purchases from AMD, with statistical significance at the 10 percent level. Finally, column (7) indicates that, controlling for the payments the client receives from Intel, it is discouraged from buying AMD's chips by payments made to Dell, with statistical significance at the 1 percent level.

The results therefore confirm our theoretical predictions (3) and (4) from Table 5: vertical restraints imposed on other clients can cause the client to reduce its adoption of AMD's technology. Above we have discussed both demand and supply side mechanisms that may give rise to this pattern, which presents a policy-relevant conclusion: in evaluating the consequences of exclusive deals, their impact on clients that were not directly subject to them should potentially be considered.

Throughout all specifications in Table 6, the signs on the capacity-related variables confirm prediction (5) from Table 5. Namely, AMD's adoption increases with the (lagged) free cash it has for investment, and with its capacity index. 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 in present and future periods. This makes clients less concerned about shifting their input demand towards AMD. In particular, a client may consider that if Intel retaliates against it by limiting the supply of its own chips, a higher-capacity AMD would be more able to pick up the slack. Intel's capacity index, for its part, reduces the adoption of AMD chips, consistent with Intel's capacity advantage being a strategic factor affecting its market position.

Continuing to examine the results in Table 6, we next consider our theoretical predictions (6) and (7): that Intel's (AMD's) benchmark-per-dollar measure should negatively (positively) affect the rate of AMD's adoption by downstream clients. Neither one of these variables has a statistically significant effect across the various specifications. We argue that this surprising result is, in fact, consistent with institutional details and with the dynamic mechanism stated by our theory.

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, these years were also characterized by highly intensive application of Intel's vertical restraints. AMD's benchmark-per-dollar advantage seems therefore to have been negated by Intel's ability to engage in exclusionary restrictions. 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.

These descriptive patterns suggest that AMD's ability to innovate and offer competitive chips was very weakly correlated with its ability to expand its market share, consistent with our statistically insignificant estimates. This suggests that technological leadership is *not necessarily* the primary driver of market share growth in the industry in question. Instead, per our additional results described above, 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.

To complete the discussion of results from Table 6, the estimates of α, the state dependence factor, are on the order of 0.7-0.8, noting that values between zero and one are considered valid. The PC price appears to have a weak (and statistically insignificant) negative relationship with the intensity of AMD adoption. There is also a negative relationship of the AMD adoption rate with the intensity of brand and firm advertising, consistent with AMD being particularly attractive for use in PCs that targeted value-seeking consumers.

To summarize, the estimation results confirm predictions (1)-(5) from Table 5 regarding the role of Intel's vertical restraints and of the capacity-related measures. Predictions (6)-(7) regarding the role of the benchmark-per-dollar variables were not supported in the sense that those variables were found to have a statistically insignificant effect — but we have argued that this, too, is in line with the basic premise of our theory.

It remains to test prediction (8) that pertains to the role of legal measures taken against Intel's restraints. To reduce clutter, we do this in two steps: first, Table 7 presents specifications that, relative to Table 6, remove the restraints variables r_{it} and include the legal environment variables l_t instead. Then, Table 8 presents specifications including both r_{it} and l_t.[31] Our theory suggested that antitrust cases brought against Intel should have a positive effect on adoption. They imply that retaliation from Intel for increased AMD adoption becomes less likely, while also improving AMD's expected ability to make sales and have more cash to invest in capacity expansion and product innovation. We include the lagged numbers of such cases, as it will likely take downstream customers at least one quarter to respond to news regarding such litigation in terms of their input choices.

The results in Table 7 demonstrate that neither the lagged number of pending antitrust cases (LPC) nor the lagged number of cumulative cases (LCC) have a significant impact on AMD offerings (columns 1 and 2). In columns 2 through 8, however, we examine how these cases affect firms that were less subject to strong exclusivity (or near-exclusivity) with Intel. The motivation is that firms that were largely restricted from using AMD chips were probably not able to respond to these antitrust cases by increasing the rate of AMD adoption, and certainly not in the short run.

We therefore interact both LPC and LCC with indicators that take the value 1 for firms that are not Dell, HP, or Toshiba, as these firms were either exclusive with Intel during much of the sample period (as in the case of Dell and Toshiba) or strongly restricted from using AMD chips above certain levels (as in the case of HP). Following a similar logic, we also interact LPC and LCC with an indicator that takes the value 1 for all clients other than Dell, and with an indicator that takes the value 1 for firms that used a positive amount of AMD chips in the previous period. As columns 2 through 8 demonstrate, all these interactions are positive and significant. Litigation, therefore, leads to an increase in the rate of adoption of AMD's technology, confirming prediction (8) of our theory.

Finally, Table 8 adds back the variables pertaining to Intel's exclusionary restrictions, r_{it}, in addition to the legal environment variables and all other variables suggested by our theory. These specifications lead to exactly the same conclusions as above, confirming our

[31] In these analyses, we continue to control for PC characteristics, CPU characteristics, capacity related variables, a constant and a time trend, exactly as in Table 6 above — and the findings for these variables are qualitatively the same as reported in that Table and discussed above.

theoretical predictions. Including the legal environment variables actually strengthens the results regarding the role of the extreme exclusionary restrictions: Columns (2) through (5) show that the negative effect of extreme restrictions on other clients (controlling for the extreme restrictions on the client itself) is now always statistically significant at either the 5 percent or 1 percent levels, whereas in Table 6, where the legal environment variables were not included, significance at the 10 percent level was demonstrated.

Taken together, the results of our baseline analysis therefore suggest that our predictions for the role played by exclusionary restrictions, and for that played by the legal environment, are supported whether we include just one of those factors at a time, or both of them together.

4.2 Economic magnitudes and additional specifications

Economic magnitudes. The above results confirmed our predictions in terms of signs and statistical significance of the explanatory variables included in the model. We next turn to examine economic significance: that is, the magnitudes of the effects of interest, expressed as marginal (or partial) effects. For concreteness, we focus on the results presented in Table 8, where the full set of explanatory variables are included.

While the model was estimated using its differenced version, equation (2), for the purpose of interpreting coefficient magnitudes it is best to look at the model before the differencing takes place — i.e., equation (1). We are interested in the quantity:

$$\frac{\partial E w_{it} | w_{i,t-1}, x_{it}, z_t, m_{it}, r_{it}, l_t, c_i}{\partial r_{itk}}, \tag{6}$$

where r_{itk} is the k^{th} element of r_{it}, (e.g., the number of extreme restrictions imposed at time t on the downstream producer). Given the linearity, this effect is consistently estimated by $\hat{\delta}_k$, i.e., the coefficient on r_{itk}. Considering column (1) of Table 8, the estimated effect of an additional restriction is -0.0060. It is interpreted as the short-run effect of an additional restriction, implying that one additional restriction reduces the firm's AMD adoption rate by 0.0060 percentage points. Given that the mean AMD adoption rate (i.e., the mean fraction of AMD-based machines across product lines) is 0.13 (see Table 1), this implies that one such additional restriction reduces adoption by 0.0060/0.13, i.e., by 4.6 percent. Columns (2) through (5) similarly imply that one additional extreme restriction reduces the rate of adoption of AMD chips by 13.3 to 18.6 percent. Intel's restraints on a specific client's use of AMD's chips therefore have a quantitatively important effect on that client's choices.

Examining columns (2)-(5) again, we note that one additional extreme restriction on *another client* reduces the adoption rate of a given client by 4.8-6.0 percent. These sizable effects demonstrate that, to fully appreciate the impact of upstream restraints, one needs to go beyond examining their effects on the client that is directly subject to them. According to our conceptual framework, such restraints affect choices by other clients as well, and our estimates suggest that those indirect effects, while smaller than the direct ones (which is to be expected), are still of notable quantitative importance.

Next, we note that the inertia embedded in the dynamic panel model implies that long-run effects of such restrictions may be larger. Concretely, an increase of size τ in current adoption w_{it} implies a cumulative effect of $\tau + \alpha\tau + \alpha^2\tau + ... = \tau/(1-\alpha)$, where α is the coefficient on the lagged dependent variable. Revisiting column (1), one additional restriction imposed on a rival today reduces a client's AMD adoption, in the long run, by 0.0060/(1-0.7949)=0.029. Dividing by 0.13, the mean adoption rate, implies that the long-run effect is the reduction of adoption by 22.5 percent. Repeating this for an additional extreme restriction imposed on a rival, per columns (2) through (5), implies a long run reduction effect of 21.2 to 27.9 percent. Clearly, the long-run effects of the restrictions are of considerable magnitudes.

Turning to the quantitative effect of an additional antitrust case brought against Intel's practices, and following similar calculations, we learn (columns 1 and 2) that the short (long) run effect of one additional lagged pending case are an increase of 12.4-12.8 (59.3-60.4) percent in the rate of adoption of AMD's chips by a client that was not subject to an exclusive contract in the previous period. The increase in adoption by clients other than Dell (column 3) are 3.9 percent in the short run, and 17.4 percent in the long run. These results lend strong support to the notion that antitrust litigation against exclusionary restrictions can have a substantial effect on expectations and choices by downstream customers.

Additional specifications. Following discussion in Section 3.2, we examine estimates from two alternatives to the system GMM approach: a fixed effects estimator, and a Pooled OLS estimator. Table 9 shows these estimates, along with estimates from our system GMM baseline models, reproduced from Tables 6 and 7.

The estimates of the coefficient α on the lagged dependent variable from the system GMM lie between the estimates from the fixed effects and pooled OLS estimators. This is consistent with theoretical observations: Nickell (1981) shows that the coefficient on the lagged dependent variable is biased downward in the fixed effects specification. The pooled OLS estimator of the same coefficient would display an upward bias if there are systematic differences across product lines in their propensity to offer an AMD chip. Our main hypotheses regarding the impact of Intel's restrictions and the impact of antitrust cases are supported by the two alternative models.

We next examine whether the results described in our baseline analysis are driven by a particular market segment. To that end, we repeat the System GMM estimation, described in Section 4.1 above, but apply it separately to the Home and Business segments. The results are reported in Tables 10 and 11. These tables effectively deliver very similar conclusions as those derived when considering both segments jointly. Interestingly, the effect of extreme restrictions on other clients is barely statistically significant (or insignificant) in the analysis of the Home segment, but is strongly significant in the Business segment. In some of the antitrust cases, claims were made that it was in the lucrative Business segment that Intel was most concerned about AMD's expansion. One may therefore indeed expect to see a stronger impact of the restrictions in that market segment.

Finally, we turn to the potential non-linearity brought about by the fact that the rate of adoption of the AMD technology, our dependent variable, is bounded between 0 and 1. We therefore estimate a (dynamic) Tobit specification along the lines described in Section 3.2, though, as discussed there, such a model does not allow us to instrument for endogenous variables, such as Intel's restrictions.

Table 12 reports results. The Tobit results replicate the qualitative message of some, but not all of our findings from the baseline linear model. We still find a statistically significant, negative effect of restrictions, extreme restrictions, and payments from Intel on the rate of adoption of the AMD technology. Extreme restrictions on other customers no longer have a significant effect on the customer's adoption rate (columns 2, 4 and 5) but payments to Dell do continue to reduce other customers' adoption rates in a statistically significant fashion. The impact of AMD's free cash is no longer significant, yet the impact of its capacity index is still positive and, in two of the specifications, significant at the 10 percent level. Finally, the positive, significant effect of antitrust cases on the AMD adoption rate is maintained. To conclude, the Tobit specification maintains several of our previous findings, despite its inability to address the endogeneity of key explanatory variables.

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, interrelated process of downstream technology adoption. To do so we use rich data from multiple sources to estimate dynamic panel models.

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*. The number of antitrust cases brought against Intel, as well as AMD's production capacity, have a positive effect on the downstream adoption of the AMD technology. These results are consistent with an important role for dynamics and client expectations regarding the future value of adopting the rival's technology.

Our results indicate that technological leadership is *not necessarily* the primary driver of market share growth. Instead, we find that 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 incumbent is not sufficient for market share growth is indicative of the complex nature of competition in a market characterized by a dynamic process of technology adoption and large sunk investments.

Our approach has limitations. In particular, we do not formally model the adoption decision. A full-blown structural dynamic model of technology adoption would have allowed 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 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 could 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. We leave such goals for future research.

References

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.

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. (2016) "Diagnosing Foreclosure due to Exclusive Dealing," *Journal of Industrial Economics*, 64(3): 375-410.

Atalay, E., A. Hortaçsu, J. Roberts, C. Syverson (2011) "Network structure of production" *Proceedings of the National Academy of Sciences,* 108 (13): 5199-5202.

Ater, I. (2015) "Vertical Foreclosure in Shopping Malls: Evidence from Shopping Malls," *Journal of Economics and Management Strategy*, 24: 620 - 642.

Bar-Isaac, H., S. (2008) "Seller Reputation", Foundations and Trends in Microeconomics: Vol. 4: No. 4, pp 273-351.

Berry, S., J. Levinsohn and A.Pakes (1995) "Automobile Prices in Market Equilibrium," *Econometrica* 63(4): 841-890.

Bernheim, B.D., and M.D. Whinston (1998). "Exclusive Dealing," *The Journal of Political Economy,* 106(1): 64-203

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

Blundell, R. and S. Bond (1998) "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models," *Journal of Econometrics* 87: 11-143.

Brown, C. and G. Linden (2009) "Chips and Change: How Crisis Reshapes the Semiconductor Industry" *MIT Press: Cambridge, US.*

Cabral, L., and J. Mata (2003) "On the Evolution of the Firm Size Distribution: Facts and Theory" *American Economic Review*, 93(4): 1075-1090.

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

Foster, L., Haltiwanger, J. and Syverson, C. (2016) "The Slow Growth of New Plants: Learning about Demand?" *Economica*, 83: 91–129.

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

Hortaçsu, A., and C. Syverson (2007) "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices," *Journal of Political Economy*, 115(2): 250–301.

Inderst, R., and G. Shaffer (2007) "Retail Mergers, Buyer Power and Product Variety," *The Economic Journal,* 117(516): 45-67

Jovanovic, B., and S. Lach (1989) "Entry, Exit, and Diffusion with Learning by Doing," *American Economic Review* 79: 690-99.

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.

Lee, H., Pechy, A., and M. Sovinsky (2013) "Is Intel's Marketing Campaign Predatory?" Working Paper, University of Zurich.

Mathewson, G.F. and R.A. Winter (1987). "The Competitive Effects of Vertical Agreements: Comment," *The American Economic Review*, 77(5):1057-62

Marvel, H.P. (1982) "Exclusive Dealing," *Journal of Law and Economics,* 25: 1-25

Nickell, S. (1981) "Biases in Dynamic Models with Fixed Effects," *Econometrica*, 49, 1417-1426.

Nurski, L. and F. Verboven (2016) "Exclusive Dealing as an Entry Barrier - Evidence from the Car Market," *The Review of Economic Studies*, forthcoming

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

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, M. (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

Sullivan, E. T., and H. Hovenkamp (2003). “Antitrust Law, Policy and Procedure: Cases, Materials, Problems,” 5th edn, Charlottesville, Va.: LexisNexis.

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)

6 Appendix

6.1 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 business segment. Indeed, while Ad1 expenditures are likely to influence choices on both segments (households or firms), the Ad2 expenditures should only affect the business 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 fewer details about PC brands than the Gartner data, we matched 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 Advertisment 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 with 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.[32]

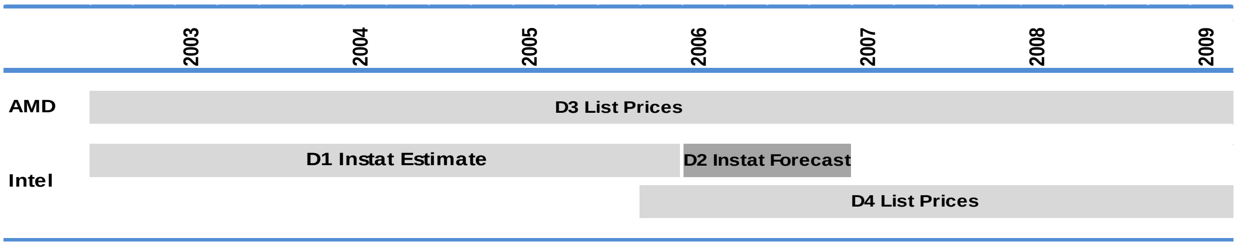

Table A2: Time Coverage of the Price Data Sources

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 this into consideration, the CPU cores are matched to the CPU families of the PC data at

[32] 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.

the platform group (whether desktop or mobile)/type (mainstream/value/ultraportable)/ family/speed/quarter level.[33]

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 as it is introduced at the end of the sample.

[33] For the CPUs not matched at first attempt, the type is dropped from 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.

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

CPU Firm	CPU Family	Median	Std Dev	Min	Max	Obs
Intel	Cel	66	7	49	77	140
	Cel M	94	32	87	203	12
	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 $

The price datasets from list prices of Intel and AMD (D3 and D4) were created as follows. Intel prices were collected from Intel's price catalogues (unit price for 1000 units) from a large variety of websites.[34] AMD prices (unit price for 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 are merged with the Gartner market share data 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.[35] It reveals that for the majority of quarters, more than one price publication is observed.[36] 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 contained many different models. Over time, more families were gradually introduced (e.g., seven families were present in 2008Q2) with fewer models per family.

[34] Complete list available from the authors upon request.

[35] 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.

[36] For 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.

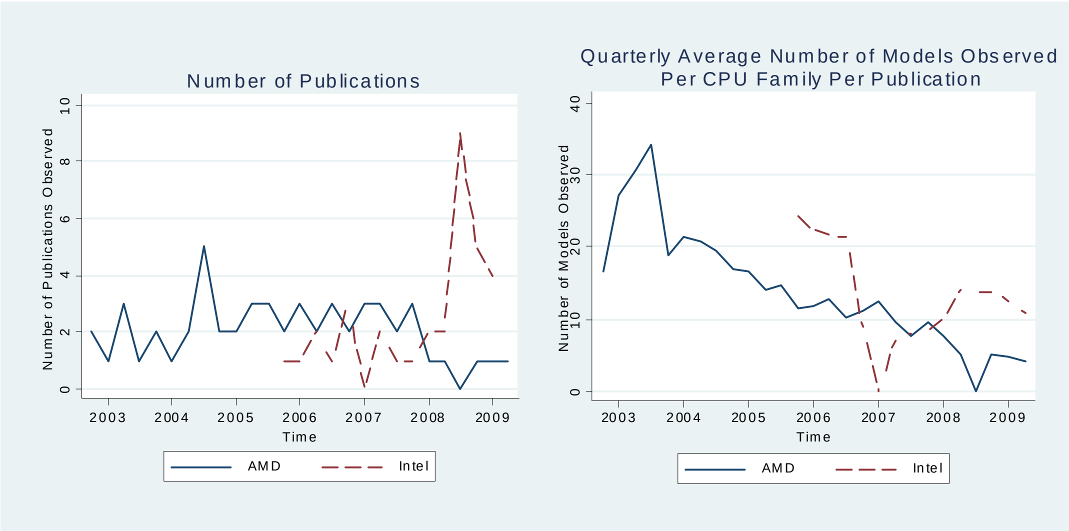

Figure A1: Availablity of List Price Data

Table A5 shows differences across families in model-quarter level prices. Note that the majority of families have at least 50 observations and some have more than 500. Model prices have high variation due to high introductory prices which decline over time.

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 $

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 are 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 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.

For some CPU families more than one core is matched, in this case we retain the mean price over cores. 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 minimum 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	2008Q3
	Phe X4, Sem	2008Q3
Last Value		
	Ath	2005Q3-2006Q1
	Tur 64	2007Q3-2008Q2
	Tur 64 X2	2008Q3-2009Q1
Dropped Obs		
	-	-

Table A6: AMD List Price Corrections

	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

	Instat Price/List Price Ratio					
Used Instat Prices	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

Benchmark. CPU benchmark information is gathered from Passmark publications.[37] 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

[37] Source: https://www.cpubenchmark.net/

period is built with two different approaches, exploiting the best information available in each quarter.

In the first approach, we use Gartner data and match CPU benchmark to the Gartner data at the CPU family-CPU speed-platform-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.[38] 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.[39]

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 based approach is applied that period onwards.

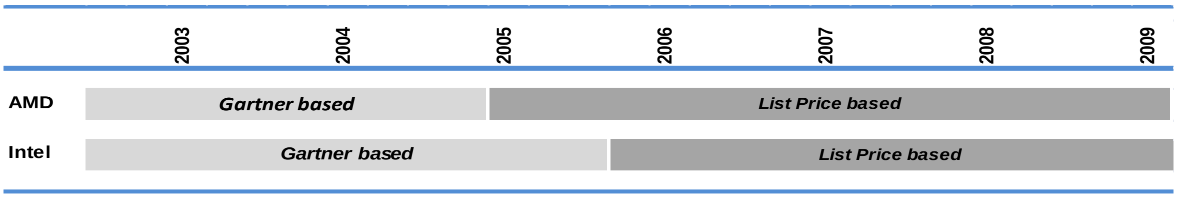

Table A10: Methodologies used To Proxy Benchmark Level

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

[38] 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%.

[39] 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.

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

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

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 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).[40]

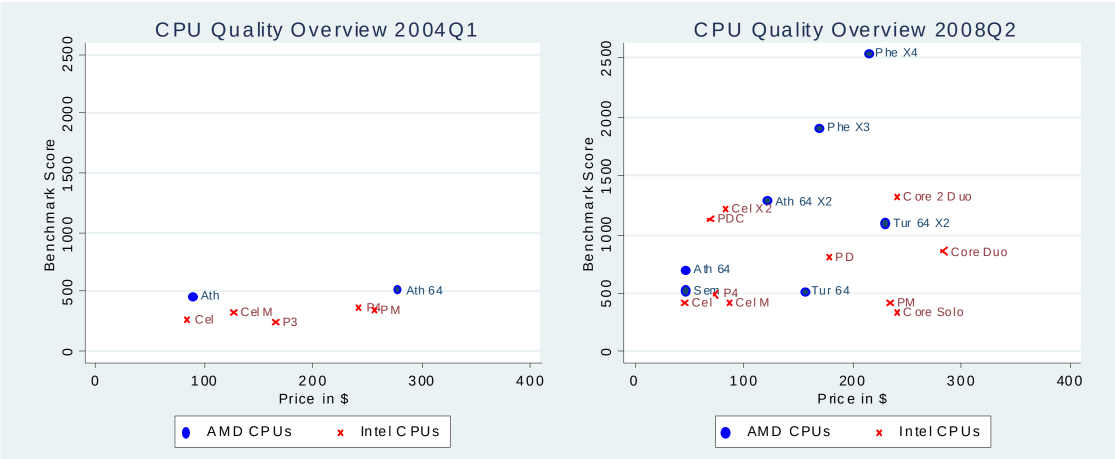

Figure A2: CPU Quality At The Family Level

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 (respectively Intel) at the PC brand-segment level.

6.2 First-Stage Regression Results for Endogeneity of PC Price

In our application of this system GMM estimator, we incorporate additional instruments for the endogenous PC price. These are motivated by the observation from Berry, Levinsohn

[40] 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 (e.g., power consumption). These could not be accounted for here due to lack of data.

and Pakes (1995) that markups, and hence prices, should be correlated with characteristics of rivals products. Specifically, we use the number of brands offered by rival firms in the quarter as an instrument. Table B1 below presents the first-stage estimation results which show that these instruments are not weak in our setting.

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Instrument							
Number of Brand/Rivals	-0.0292***	-0.0296***	-0.0285**	-0.0322***	-0.0365***	-0.0382***	-0.0250**
	(0.0111)	(0.0110)	(0.0111)	(0.0109)	(0.0104)	(0.0103)	(0.0106)
Other Regressors							
Brand Advertising (10M$)	0.0101	0.0179	0.0063	0.0088	0.0656	0.0608	0.0022
	(0.1902)	(0.1862)	(0.1909)	(0.1896)	(0.1873)	(0.1929)	(0.1898)
Firm Advertising (10M$)	0.0497	0.0560	0.0601	0.0081	0.0662	0.0549	0.0079
	(0.0663)	(0.0662)	(0.0673)	(0.0655)	(0.0659)	(0.0678)	(0.0657)
AMD CPU benchmark/dollar (10,000$)	-0.0409**	-0.0420**	-0.0388*	-0.0374*	-0.0154	0.0209	-0.0583**
	(0.0207)	(0.0209)	(0.0209)	(0.0210)	(0.0217)	(0.0235)	(0.0249)
Intel CPU benchmark/dollar (10,000$)	-0.0679***	-0.0675***	-0.0686***	-0.0815***	-0.0711***	-0.0613***	-0.0864***
	(0.0171)	(0.0171)	(0.0171)	(0.0177)	(0.0173)	(0.0171)	(0.0178)
Lagged Free Cash (1000M$)	-1.4131***	-1.4100***	-1.4058***	-1.5325***	-1.6901***	-1.6838***	-1.3957***
	(0.1604)	(0.1591)	(0.1623)	(0.1567)	(0.1630)	(0.1639)	(0.1556)
AMD Capacity Index	-0.5229***	-0.5277***	-0.5161***	-0.4998***	-0.3678***	-0.3957***	-0.4779***
	(0.0387)	(0.0415)	(0.0405)	(0.0383)	(0.0388)	(0.0363)	(0.0384)
Intel Capacity Index	0.0326***	0.0341***	0.0299**	0.0339***	0.0119	0.0092	0.0247**
	(0.0118)	(0.0122)	(0.0121)	(0.0119)	(0.0128)	(0.0125)	(0.0106)
Number of Exclusionary Restrictions		-0.0264			0.0386		
		(0.0474)			(0.0446)		
Number of Extreme Exclusionary Restrictions			0.1254			0.3213***	
			(0.0992)			(0.0873)	
Payments Received from Intel (M$)				-0.0896***			-0.0849***
				(0.0154)			(0.0154)
Intel Payments made to Dell (100M$)							-0.0919***
							(0.0264)
Number of Exclusionary Restrictions on Other Firms					0.1079***		
					(0.0202)		
Number of Extreme Restrictions on Other Firms						0.3706***	
						(0.0547)	
Observations	3,508	3,508	3,508	3,508	3,508	3,508	3,508
R-squared	0.6613	0.6614	0.6617	0.6664	0.6665	0.6723	0.6681

Notes: Robust Clustered Standard Errors in Parentheses. *** denotes significance at the 1% level, ** at the 5% level; and * at the 10% level. Each regression contains 3508 observations.

Table B1: First-Stage Results

6.3 Tables and Figures

Table 1: Descriptive statistics

	Obs	Mean	Std. Dev.	Min	Max
A. PC Characteristics					
Contains AMD CPU	3280	0.13	0.24	0	1
Price PC (100$)	3280	10.23	4.43	2.41	35.21
Brand Advertising (M$)	3280	1.03	3.15	0	29.14
Firm Advertising (M$)	3280	8.11	16.43	0	87.80
B. CPU Characteristics					
AMD CPU benchmark/dollar*	1181	7.38	4.59	1.75	22.65
Intel CPU benchmark/dollar*	2099	4.09	3.39	0.84	32.66
Num. Quarters Brand/AMD family available*	1181	5.62	3.14	1	19
Num. Quarters Brand/Intel family available*	3144	6.78	3.61	1	30
C. Capacity and cashflow					
Free Cash (100M$)	3280	8.33	3.10	3.97	19.05
AMD Capacity Index	3280	8.10	3.56	3	13
Intel Capacity Index	3280	31.93	6.64	23	44
D. Exclusionary Restrictions & antitrust					
Exclusionary Restrictions Index	3280	1.60	1.82	0	6
Extreme Restrictions Index	3280	0.61	0.75	0	3
Intel Payments to Dell (M$)	3280	150.0	155.23	13.37	603.1
Intel Payments to (non-Dell) PC Firm (M$)	2658	3.50	4.08	0	17.73
Num. Cumulative Antitrust Cases Against Intel	3280	4.22	2.12	1	7
Num. Pending Antitrust Cases Against Intel	3280	3.47	1.60	1	6

Notes: Descriptive statistics over 3,280 observations as defined in the text. *AMD (Intel) statistics are reported over AMD-based (Intel-based) observations only, hence the number of observations is less than 3,280.

Table 2: Upstream Production Capacity Over Time

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

Notes: ICP stands for Integrated Circuit Process. 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 3: Antitrust activity timeline

2001	European Commission opens investigation
2004	Japan Fair trade opens investigation
2005	Japan Fair Trade issues decision that Intel violated rules; Intel complies Korea Fair Trade opens investigation AMD files lawsuits against Intel in Delaware, US, and in Japan
2007	EC brings charges Korea Fair Trade brings charges
2008	State of New York opens investigation FTC opens investigation Korea Fair Trade decision that Intel violated antitrust law Intel appeals Korea Fair Trade Decision
2009	EC decision that Intel violated Article 82; Intel files appeal State of New York files lawsuit AMD cases against Intel end in settlement FTC brings charges
2010	FTC case ends in settlement
2012	State of New York case ends in settlement
2013	Korea Fair Trade upholds ruling against intel appeal
2014	EC decision of 2009 Upheld by the Court; Intel appeals

Notes: Based on information in Intel and AMD annual reports as well as some media references.

Table 4: Antitrust variables

Year	Quarter	Cumulative cases	Pending cases	Year	Quarter	Cumulative cases	Pending cases
2002	1	1	1	2007	1	5	4
2002	2	1	1	2007	2	5	4
2002	3	1	1	2007	3	5	4
2002	4	1	1	2007	4	5	4
2003	1	1	1	2008	1	6	5
2003	2	1	1	2008	2	7	6
2003	3	1	1	2008	3	7	6
2003	4	1	1	2008	4	7	6
2004	1	1	1	2009	1	7	5
2004	2	2	2	2009	2	7	5
2004	3	2	2	2009	3	7	4
2004	4	2	2	2009	4	7	4
2005	1	2	2	2010	1	7	2
2005	2	5	4	2010	2	7	2
2005	3	5	4	2010	3	7	2
2005	4	5	4	2010	4	7	2
2006	1	5	4				
2006	2	5	4				
2006	3	5	4				
2006	4	5	4				

Notes: The numbers of pending and cumulative antitrust cases (lawsuits and investigations) in connection with the Intel Inside program are displayed. See text for complete details.

Table 5: Client policy function: predicted effects

State variable	Predicted effect
1. Intel's restrictions on the client's use of AMD chips	Negative
2. Intel's payments to the client	Negative
3. Intel's restrictions placed on other clients	Negative
4. Intel's payments to other clients	Negative
5. AMD's capacity, cash available for investment	Positive
6. Intel's benchmark-per-dollar	Negative
7. AMD's benchmark-per-dollar	Positive
8. Intensity of legal action against Intel's restraints	Positive

Notes: The table summarizes the predictions implied by the theory regarding the effect of state variables on the client's continuous control: the fraction of chips purchased from AMD (see text).

Table 6: System GMM estimates: downstream clients policy function, focusing on the role of exclusionary restrictions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
w_{it-1}	0.7956***	0.7553***	0.7331***	0.7441***	0.7859***	0.7745***	0.7537***
	(0.0640)	(0.0684)	(0.0715)	(0.0741)	(0.0804)	(0.0834)	(0.0728)
PC Characteristics							
PC price (100$)	-0.0029	-0.0036	-0.0062	-0.0067	-0.0001	-0.0040	-0.0084
	(0.0046)	(0.0046)	(0.0049)	(0.0053)	(0.0054)	(0.0050)	(0.0051)
Brand Advertising	-0.0114**	-0.0151**	-0.0143**	-0.0161**	-0.0153***	-0.0132**	-0.0148**
	(0.0050)	(0.0059)	(0.0063)	(0.0068)	(0.0056)	(0.0060)	(0.0070)
Firm Advertising	-0.0002	0.0020	-0.0009	0.0007	0.0021	-0.0007	0.0002
	(0.0015)	(0.0018)	(0.0018)	(0.0018)	(0.0017)	(0.0015)	(0.0018)
CPU Characteristics							
AMD benchmark/dollar	-0.0021	-0.0024	-0.0026	-0.0023	-0.0023	-0.0024	-0.0038
	(0.0020)	(0.0022)	(0.0023)	(0.0023)	(0.0020)	(0.0021)	(0.0024)
Intel benchmark/dollar	0.0001	0.0002	-0.0006	-0.0007	0.0014	-0.0001	-0.0019
	(0.0021)	(0.0021)	(0.0022)	(0.0022)	(0.0023)	(0.0021)	(0.0022)
Age	0.0001	-0.0004	-0.0008	-0.0006	0.0002	-0.0004	-0.0006
	(0.0008)	(0.0009)	(0.0010)	(0.0009)	(0.0010)	(0.0009)	(0.0009)
Capacity Related							
AMD's Lagged Free Cash	0.0179	0.0251**	0.0216*	0.0176	0.0281**	0.0210*	0.0211**
	(0.0111)	(0.0120)	(0.0113)	(0.0109)	(0.0116)	(0.0115)	(0.0105)
AMD Capacity Index	0.0029**	0.0027**	0.0026**	0.0026**	0.0022*	0.0034**	0.0032***
	(0.0012)	(0.0012)	(0.0013)	(0.0013)	(0.0011)	(0.0014)	(0.0012)
Intel Capacity Index	-0.0013**	-0.0020***	-0.0020***	-0.0012*	-0.0028***	-0.0037***	-0.0014**
	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0009)	(0.0011)	(0.0006)
Restrictions							
Restrictions		-0.0116***			-0.0112***		
		(0.0029)			(0.0030)		
Extreme Restrictions			-0.0278***			-0.0281***	
			(0.0073)			(0.0070)	
Payments from Intel (M$)				-0.0018***			-0.0017***
				(0.0005)			(0.0005)
Payments to Dell (100M$)							-0.0051***
							(0.0019)
Restrictions on others					-0.0014		
					(0.0011)		
Extreme Rest. on others						-0.0057*	
						(0.0029)	

Notes: The dependent variable is the fraction of CPUs purchased from AMD. Robust Clustered Standard Errors in Parentheses. Each regression contains 3236 observations, and includes a constant and a quarterly trend. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. Units: advertising variables are in 10M$. AMD's lagged free cash is in B$. Age is the number of quarters in which the CPU has been offered with the relevant product line (see section 2 for detailed variable descriptions).

Table 7: System GMM estimates: downstream clients policy function, focusing on the role of the legal environment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
w_{it-1}	0.8031***	0.8058***	0.7767***	0.7758***	0.7694***	0.7788***	0.7798***	0.7749***
	(0.0626)	(0.0627)	(0.0681)	(0.0725)	(0.0719)	(0.0676)	(0.0718)	(0.0708)
PC Characteristics								
Price PC (100$)	-0.0014	-0.0017	-0.0019	-0.0014	0.0004	-0.0021	-0.0016	-0.0001
	(0.0043)	(0.0045)	(0.0047)	(0.0042)	(0.0038)	(0.0046)	(0.0042)	(0.0039)
Brand Advertising	-0.0119**	-0.0114**	-0.0156**	-0.0177**	-0.0073**	-0.0151**	-0.0172**	-0.0075**
	(0.0052)	(0.0051)	(0.0063)	(0.0069)	(0.0031)	(0.0062)	(0.0067)	(0.0031)
Firm Advertising	0.0000	-0.0000	0.0007	-0.0003	0.0001	0.0005	-0.0003	0.0003
	(0.0015)	(0.0014)	(0.0016)	(0.0016)	(0.0014)	(0.0016)	(0.0016)	(0.0014)
CPU Characteristics								
AMD benchmark/dollar	-0.0021	-0.0023	-0.0023	-0.0021	-0.0018	-0.0023	-0.0021	-0.0018
	(0.0020)	(0.0020)	(0.0021)	(0.0021)	(0.0018)	(0.0021)	(0.0020)	(0.0018)
Intel benchmark/dollar	0.0007	0.0005	0.0010	0.0013	-0.0004	0.0009	0.0010	-0.0007
	(0.0020)	(0.0021)	(0.0021)	(0.0019)	(0.0020)	(0.0021)	(0.0019)	(0.0020)
Age	0.0004	0.0003	0.0001	-0.0001	-0.0012	0.0001	-0.0001	-0.0013
	(0.0008)	(0.0008)	(0.0008)	(0.0008)	(0.0009)	(0.0008)	(0.0008)	(0.0009)
Capacity Related								
AMD's Lagged Free Cash	0.0184	0.0167	0.0198*	0.0217*	0.0134	0.0206*	0.0234*	0.0154
	(0.0117)	(0.0122)	(0.0114)	(0.0118)	(0.0103)	(0.0117)	(0.0123)	(0.0106)
AMD Capacity Index	0.0032***	0.0032***	0.0032***	0.0034***	0.0036***	0.0031**	0.0032***	0.0034***
	(0.0012)	(0.0012)	(0.0012)	(0.0012)	(0.0011)	(0.0012)	(0.0012)	(0.0011)
Intel Capacity Index	-0.0015**	-0.0014**	-0.0015**	-0.0017***	-0.0015***	-0.0015**	-0.0016***	-0.0014**
	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0006)
Antitrust Cases								
LPC	0.0014							
	(0.0032)							
LCC		-0.0007						
		(0.0023)						
LPC * not DHT			0.0052***					
			(0.0018)					
LPC * not Dell				0.0084***				
				(0.0030)				
LPC * used AMD in $t-1$					0.0187***			
					(0.0060)			
LCC * not DHT						0.0040***		
						(0.0015)		
LCC * not Dell							0.0064***	
							(0.0023)	
LCC * used AMD in $t-1$								0.0146***
								(0.0047)

Notes: The dependent variable is the fraction of CPUs purchased from AMD. Robust Clustered Standard Errors in Parentheses. Each regression contains 3236 observations, and includes a constant and a quarterly trend. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. For units, see notes to Table 6. LPC and LCC are the lagged numbers of pending and cumulative antitrust cases against Intel, repsectively. "Not DHT" is a dummy variable taking the value 1 for clients that are not Dell, HP or Toshiba. See text.

Table 8: System GMM estimates: downstream clients policy function, all variables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
w_{it-1}	0.7949***	0.7834***	0.7750***	0.7876***	0.7760***	0.7617***	0.7604***
	(0.0831)	(0.0771)	(0.0839)	(0.0770)	(0.0841)	(0.0732)	(0.0726)
PC Characteristics							
Price PC (100$)	0.0033	0.0004	-0.0024	-0.0001	-0.0026	-0.0047	-0.0067
	(0.0045)	(0.0039)	(0.0045)	(0.0040)	(0.0046)	(0.0044)	(0.0045)
Brand Advertising	-0.0094***	-0.0080**	-0.0165**	-0.0082**	-0.0162**	-0.0084**	-0.0167**
	(0.0034)	(0.0032)	(0.0069)	(0.0032)	(0.0068)	(0.0039)	(0.0073)
Firm Advertising	0.0013	-0.0001	-0.0006	-0.0001	-0.0006	0.0002	-0.0004
	(0.0014)	(0.0013)	(0.0015)	(0.0013)	(0.0015)	(0.0014)	(0.0017)
CPU Characteristics							
AMD benchmark/dollar	-0.0019	-0.0019	-0.0023	-0.0018	-0.0023	-0.0034*	-0.0036
	(0.0017)	(0.0017)	(0.0021)	(0.0017)	(0.0021)	(0.0020)	(0.0024)
Intel benchmark/dollar	0.0009	-0.0003	0.0008	-0.0007	0.0006	-0.0027	-0.0011
	(0.0023)	(0.0020)	(0.0019)	(0.0020)	(0.0019)	(0.0021)	(0.0020)
Age	-0.0007	-0.0012	-0.0004	-0.0013	-0.0004	-0.0020**	-0.0008
	(0.0012)	(0.0010)	(0.0009)	(0.0010)	(0.0009)	(0.0010)	(0.0009)
Capacity Related							
AMD's Lagged Free Cash	0.0214*	0.0155	0.0227*	0.0174	0.0238*	0.0186*	0.0255**
	(0.0110)	(0.0105)	(0.0118)	(0.0108)	(0.0122)	(0.0102)	(0.0117)
AMD Capacity Index	0.0026**	0.0043***	0.0038***	0.0041***	0.0037***	0.0037***	0.0034***
	(0.0011)	(0.0012)	(0.0013)	(0.0012)	(0.0013)	(0.0011)	(0.0012)
Intel Capacity Index	-0.0028***	-0.0041***	-0.0039***	-0.0040***	-0.0038***	-0.0014**	-0.0016***
	(0.0009)	(0.0010)	(0.0010)	(0.0010)	(0.0010)	(0.0006)	(0.0006)
Restrictions							
Restrictions	-0.0060***						
	(0.0014)						
Extreme Restrictions		-0.0173***	-0.0237***	-0.0178***	-0.0242***		
		(0.0037)	(0.0058)	(0.0038)	(0.0059)		
Payments from Intel (M$)						-0.0011***	-0.0009**
						(0.0003)	(0.0004)
Payments to Dell (100M$)						-0.0051***	-0.0051***
						(0.0019)	(0.0019)
Restrictions on others	-0.0016						
	(0.0011)						
Extreme Rest. on others		-0.0078***	-0.0063**	-0.0077***	-0.0062**		
		(0.0025)	(0.0029)	(0.0025)	(0.0029)		
Antitrust Cases							
LPC * used AMD in $t-1$	0.0161**	0.0167***					
	(0.0064)	(0.0060)					
LPC * not Dell			0.0051**				
			(0.0023)				
LCC * used AMD in $t-1$				0.0131***		0.0142***	
				(0.0048)		(0.0048)	
LCC * not Dell					0.0038**		0.0046*
					(0.0018)		(0.0025)

Notes: LCC and LPC are the lagged numbers of comulative and pending antitrust cases, respectively (see text). "Not DHT" is a dummy variable for clients other than Dell, HP, and Toshiba. See notes for Tables 4 and 5.

Table 9: Pooled OLS, Fixed Effects, and System GMM estimates

	Pooled OLS		Fixed Effects		System GMM	
	(1)	(2)	(3)	(4)	(5)	(6)
w_{it-1}	0.9358***	0.9103***	0.7010***	0.6243***	0.7745***	0.7749***
	(0.0145)	(0.0216)	(0.0419)	(0.0516)	(0.0834)	(0.0708)
PC Characteristics						
Price PC (100$)	-0.0014***	-0.0012***	-0.0053**	-0.0059**	-0.0040	-0.0001
	(0.0004)	(0.0004)	(0.0025)	(0.0025)	(0.0050)	(0.0039)
Brand Advertising	-0.0060***	-0.0051***	0.0004	-0.0016	-0.0132**	-0.0075**
	(0.0022)	(0.0018)	(0.0038)	(0.0040)	(0.0060)	(0.0031)
Firm Advertising	-0.0011**	-0.0007	-0.0043**	-0.0004	-0.0007	0.0003
	(0.0006)	(0.0006)	(0.0021)	(0.0020)	(0.0015)	(0.0014)
CPU Characteristics						
AMD CPU benchmark/dollar	-0.0027**	-0.0027**	-0.0031*	-0.0028	-0.0024	-0.0018
	(0.0012)	(0.0012)	(0.0018)	(0.0018)	(0.0021)	(0.0018)
Intel CPU benchmark/dollar	-0.0011	-0.0015	-0.0002	-0.0007	-0.0001	-0.0007
	(0.0009)	(0.0009)	(0.0013)	(0.0012)	(0.0021)	(0.0020)
Age	-0.0000	-0.0005	-0.0010	-0.0012	-0.0004	-0.0013
	(0.0004)	(0.0004)	(0.0012)	(0.0012)	(0.0009)	(0.0009)
Capacity Related						
AMD's Lagged Free Cash	0.0191**	0.0152*	0.0207**	0.0226**	0.0210*	0.0154
	(0.0091)	(0.0090)	(0.0099)	(0.0101)	(0.0115)	(0.0106)
AMD Capacity Index	0.0012	0.0007	0.0003	0.0003	0.0034**	0.0034***
	(0.0013)	(0.0012)	(0.0018)	(0.0016)	(0.0014)	(0.0011)
Intel Capacity Index	-0.0018**	-0.0001	-0.0015	-0.0003	-0.0037***	-0.0014**
	(0.0009)	(0.0005)	(0.0010)	(0.0007)	(0.0011)	(0.0006)
Exclusionary Restrictions						
Extreme Restrictions	-0.0115***		-0.0146***		-0.0281***	
	(0.0023)		(0.0035)		(0.0070)	
Extreme Rest. on others	-0.0042**		-0.0024		-0.0057*	
	(0.0018)		(0.0019)		(0.0029)	
Antitrust Cases						
LCC * used AMD in $t-1$		0.0054***		0.0122***		0.0146***
		(0.0015)		(0.0027)		(0.0047)

Notes: LCC: Lagged number of cumulative cases (see text). Pooled OLS results are reported in columns 1 and 2. Fixed effects results are reported in columns 3 and 4. Columns 5 and 6 replicate System GMM estimates reported above in Tables 6 and 7, repsectively. For detailed variable definitions, see notes for Tables 6 and 7.

Table 10: System GMM estimates, the Home segment

	(1)	(2)	(3)	(4)	(5)	(6)
w_{it-1}	0.7473***	0.7424***	0.7422***	0.7169***	0.7489***	0.7069***
PC Characteristics						
Price PC (100$)	-0.0017	-0.0040	-0.0037	-0.0027	-0.0020	-0.0063
Brand Advertising	-0.0074	-0.0168*	-0.0197*	-0.0222*	-0.0084	-0.0171
Firm Advertising	-0.0001	0.0002	-0.0008	0.0019	-0.0005	-0.0010
CPU Characteristics						
AMD CPU benchmark/dollar	-0.0014	-0.0021	-0.0019	-0.0022	-0.0015	-0.0024
Intel CPU benchmark/dollar	-0.0013	0.0006	0.0008	0.0012	-0.0013	-0.0002
Age	-0.0017	0.0001	-0.0003	-0.0004	-0.0019	-0.0006
Capacity Related						
AMD's Lagged Free Cash	0.0193	0.0257	0.0292	0.0362**	0.0220	0.0284
AMD Capacity Index	0.0035**	0.0032	0.0033*	0.0029	0.0042**	0.0035
Intel Capacity Index	-0.0014	-0.0016*	-0.0017*	-0.0028**	-0.0040***	-0.0037**
Exclusionary Restrictions						
Restrictions				-0.0126***		
Extreme Restrictions					-0.0205***	-0.0336***
Restrictions on others				-0.0006		
Extreme Rest. on others					-0.0073*	-0.0044
Antitrust Cases						
LCC * used AMD in $t-1$	0.0168*				0.0157*	
LCC * not DHT		0.0042*				0.0012
LCC * not Dell			0.0078**	0.0048		

Notes: LCC: Lagged number of cumulative cases (see text). For space conservation, we only report the coefficients and indicate statistical significance at the 1%, 5% and 10% levels by ***, **, and *, respectively. For detailed variable definitions, see notes for Tables 6 and 7.

Table 11: System GMM estimates, the Business segment

	(1)	(2)	(3)	(4)	(5)	(6)
w_{it-1}	0.8313***	0.8588***	0.8512***	0.8712***	0.8495***	0.8686***
PC Characteristics						
Price PC (100$)	0.0018	0.0005	-0.0001	0.0027	0.0020	0.0004
Brand Advertising	-0.0058*	-0.0098*	-0.0112**	-0.0108*	-0.0063*	-0.0094*
Firm Advertising	0.0003	0.0004	-0.0003	0.0012	0.0002	-0.0000
CPU Characteristics						
AMD CPU benchmark/dollar	-0.0021	-0.0024	-0.0021	-0.0022	-0.0021	-0.0024
Intel CPU benchmark/dollar	-0.0001	0.0010	0.0007	0.0020	0.0001	0.0008
Age	-0.0006	0.0002	0.0000	0.0004	-0.0006	0.0001
Capacity Related						
AMD's Lagged Free Cash	0.0120	0.0150	0.0172	0.0243	0.0133	0.0155
AMD Capacity Index	0.0035**	0.0032**	0.0031**	0.0019	0.0042***	0.0039**
Intel Capacity Index	-0.0014*	-0.0014*	-0.0014*	-0.0029**	-0.0041***	-0.0040***
Exclusionary Restrictions						
Restrictions				-0.0072***		
Extreme Restrictions					-0.0144***	-0.0178***
Restrictions on others				-0.0021		
Extreme Rest. on others					-0.0083***	-0.0078**
Antitrust Cases						
LCC * used AMD in $t-1$	0.0110***				0.0095**	
LCC * not DHT		0.0029**				0.0015
LCC * not Dell			0.0043**	0.0017		

Notes: LCC: Lagged number of cumulative cases (see text). For space conservation, we only report the coefficients and indicate statistical significance at the 1%, 5% and 10% levels by ***, **, and *, respectively. For detailed variable definitions, see notes for Tables 6 and 7.

Table 12: Dynamic Tobit estimates: downstream clients policy function

	(1)	(2)	(3)	(4)	(5)	(6)
w_it-1	0.791***	0.703***	0.825***	0.703***	0.700***	0.704***
	(0.033)	(0.037)	(0.033)	(0.037)	(0.037)	(0.037)
PC Characteristics						
Price PC (100$)	-0.029***	-0.026***	-0.028***	-0.026***	-0.026***	-0.024***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Brand Advertising	-0.036	-0.031	-0.031	-0.031	-0.031	-0.041
	(0.030)	(0.031)	(0.030)	(0.031)	(0.031)	(0.031)
Firm Advertising	0.022***	0.009**	0.008*	0.009**	0.009**	0.008*
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
CPU Characteristics						
AMD benchmark/dollar	-0.006***	-0.006***	-0.007***	-0.006***	-0.006***	-0.007***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Intel benchmark/dollar	-0.001	-0.001	-0.002	-0.001	-0.000	-0.001
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Age	-0.014***	-0.014***	-0.015***	-0.014***	-0.014***	-0.015***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Capacity Related						
AMD's Lagged Free Cash	0.031	0.038	-0.006	0.038	0.030	0.007
	(0.028)	(0.026)	(0.027)	(0.026)	(0.026)	(0.026)
AMD Capacity Index	0.007	0.005	0.009*	0.005	0.006	0.009*
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Intel Capacity Index	-0.005***	-0.004	-0.004**	-0.004	-0.005*	-0.005***
	(0.002)	(0.003)	(0.002)	(0.003)	(0.003)	(0.002)
Exclusionary Restrictions						
Restrictions	-0.047***					
	(0.006)					
Extreme Restrictions		-0.064***		-0.064***	-0.063***	
		(0.014)		(0.014)	(0.014)	
Payments from Intel (M$)			-0.008***			-0.010***
			(0.003)			(0.003)
Payments to Dell (100M$)			-0.004			-0.009**
			(0.004)			(0.004)
Restrictions on others	-0.000					
	(0.002)					
Extreme Rest. on others		0.002		0.002	0.001	
		(0.007)		(0.007)	(0.007)	
Antitrust Cases						
LCC * used AMD in $t-1$		0.028***		0.028***		
		(0.003)		(0.003)		
LPC * used AMD in $t-1$					0.036***	0.039***
					(0.004)	(0.004)

Notes: LCC and LPC are the lagged numbers of comulative and pending antitrust cases, respectively (see text). "Not DHT" is a dummy variable for clients other than Dell, HP, and Toshiba. See notes for Tables 6 and 7.

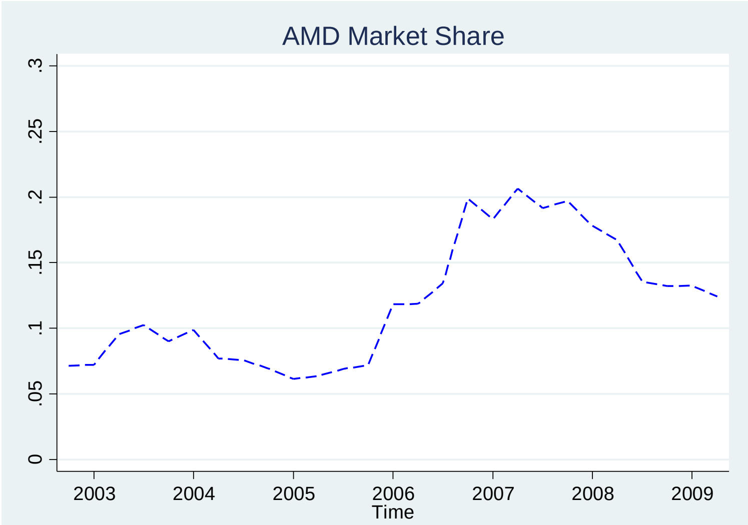

Figure 1: AMD's market share over time

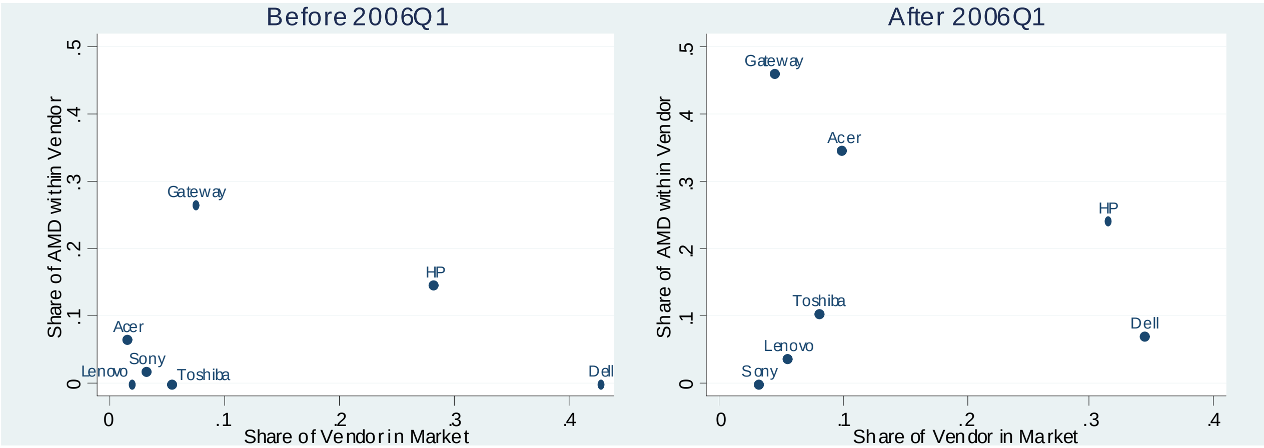

Figure 2: Selected downstream customers' share of the PC market and the rate of their utilization of AMD's chips

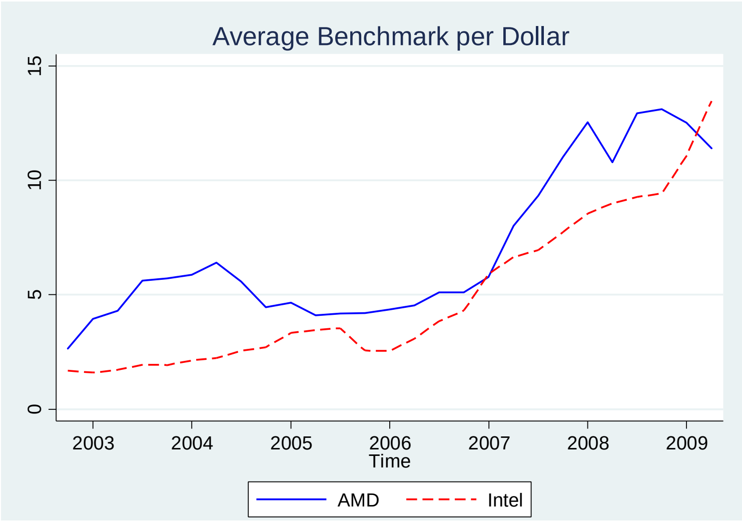

Figure 3: CPU Benchmark Quality Per Dollar, over time, for Intel and AMD

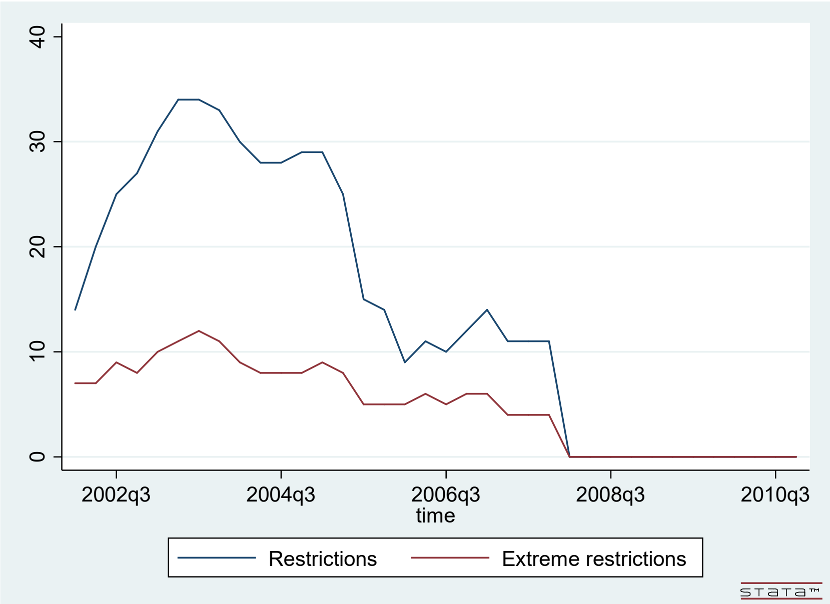

Figure 4: Evolution of the number of restrictions on downstream clients (see text)

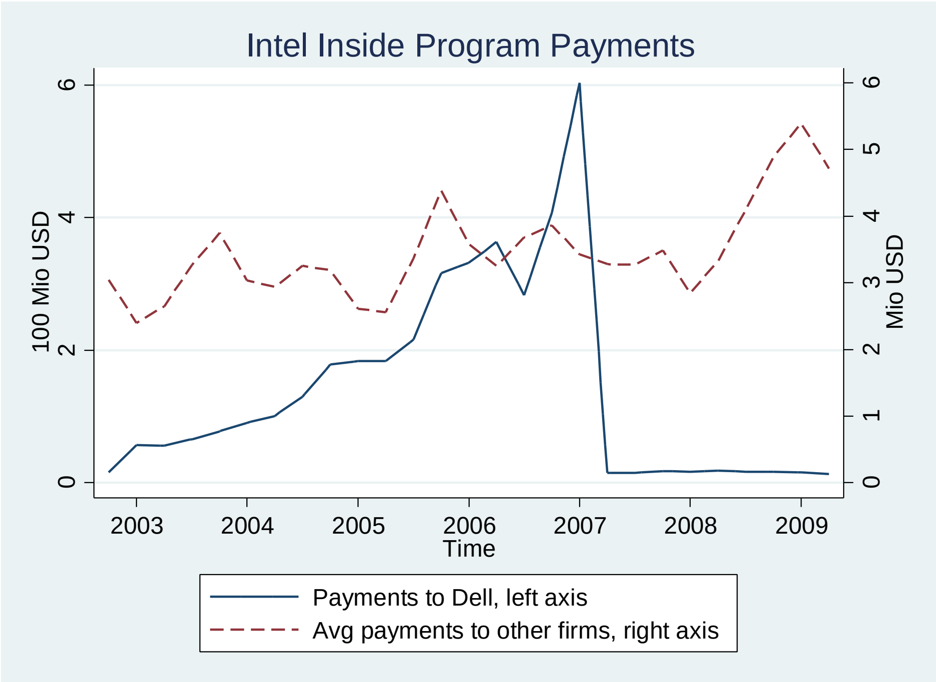

Figure 5: Evolution of Intel's Payments to PC Firms