Incapacitating the Competition: The Impact of Vertical Restraints on Technology Adoption

Laura Grigolon and Michelle Sovinsky*

January 8, 2026

Abstract

Vertical restraints imposed on some downstream buyers can affect non-contracted firms by weakening upstream suppliers' effective viability. We study these contracting externalities using Intel's exclusionary agreements with PC manufacturers in the microprocessor market. Combining litigation-based measures of restraints with PC data, we estimate dynamic models of AMD adoption that allow for cross-buyer spillovers. We find that exclusivity imposed on a given buyer significantly reduces adoption by other, non-contracted buyers, generating sizable and persistent market-wide effects. The paper provides the first empirical quantification of the economic magnitude of contracting externalities and highlights the broader competitive risks posed by exclusionary contracting.

JEL Codes: L42, L63, K21, D22

Keywords: contracting externalities, vertical restraints, semiconductors

* Laura Grigolon: University of Mannheim and CEPR, laura.grigolon@uni-mannheim.de; Michelle Sovinsky: University of Mannheim and CEPR, michelle.sovinsky@uni-mannheim.de. Sovinsky acknowledges support from the European Research Council (ERC) Grant #725081 FORENSICS and from the German Research Foundation (DFG) through CRC TR 224 (Project B04). We are grateful for helpful comments from Luís Cabral, Chiara Fumagalli, Massimo Motta, Nicolas Schutz, Frank Verboven, and seminar participants at APIOS (Asia-Pacific Industrial Organization Conference), Barcelona Applied IO Workshop, Cambridge, CEPR Symposium, CREST Paris, Düsseldorf (DICE), EARIE, Edinburgh, EEA, European University Institute (EUI), European Competition Network Chief Economists Meeting, ETH Zurich, German IO meetings, Gothenburg, Hamburg, Hebrew University of Jerusalem, IIOC, Liverpool, MaCCI Summer Institute, Mannheim, Oxford, Pompeu Fabra, Tel Aviv, and Tilburg. We thank Andres Pechy, Toby Webber, and Yihan Yan for excellent research assistance.

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)

When manufacturers enter into exclusive agreements with their downstream customers, they restrict the ability of these customers to engage with upstream rivals. As a result, vertical restraints may hinder the viability of upstream firms, thereby reducing market competition. However, exclusive agreements can also foster competition by encouraging firms to improve service quality or secure investments, mitigating downstream free-riding problems.[1]

Empirically discerning whether these restraints promote or hinder competition is challenging.[2] Crucially, these effects may extend beyond the immediate parties to an agreement. The resulting *contracting externalities* (Segal, 1999) arise when a restraint imposed on one downstream customer indirectly alters the incentives of others. As noted by Whinston (2006), an upstream supplier may "induce a particular buyer or a subset of buyers to sign *[an exclusive deal]* because by doing this he can monopolize *other* buyers without paying them anything." These forces are particularly salient in markets where a supplier's ability to provide future support or upgrades is shaped by its current sales and scale. In such settings, extensive exclusivity with some buyers can weaken a rival's effective viability market-wide, making its products less attractive even to buyers not directly subject to the restraints.[3]

In this paper, we investigate the understudied contracting externalities and persistent effects of exclusive agreements, focusing on how restraints targeting specific buyers alter competitors' incentives and market outcomes. We examine Intel's alleged agreements with personal computer (PC) firms as a case study. During the period under study (2002–2009), the x86 microprocessor

[1] As a result, in the US, exclusive dealing is not *per se* illegal; rather, each case is evaluated under the "rule of reason" standard, as reaffirmed in *Tampa Elec. Co. v. Nashville Coal Co., 365 U.S. 320 (1961)*. While exclusive dealing may violate Section 3 of the Clayton Act and Section 2 of the Sherman Act, its *per se* illegality was rejected in *Standard Oil Co. v. United States (Standard Stations)*, 337 U.S. 293, 305–06 (1949). For further discussion of the competitive effects and legal treatment of exclusive dealing in the US, see Areeda and Kaplow (1997).

[2] Asker (2016) develops empirical tests to distinguish procompetitive from foreclosure motives in vertical contracts. Fumagalli and Motta (2017) examine the application of price-cost tests to loyalty discounts and exclusive dealing arrangements. They show that contracts referencing rivals, that is, contracts in which the terms depend on how much the buyer sources from competing suppliers, can serve as a powerful exclusionary tool and facilitate buyer-specific discrimination by dominant firms.

[3] Atalay et al. (2011) show 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 similarly incorporates downstream customers who reassess a supplier's future viability.

market was dominated by two suppliers: Intel, with approximately 80% of the market share, and its smaller competitor, Advanced Micro Devices (AMD), accounting for most of the remaining 20%. These suppliers provided essential components, Central Processing Units (CPUs), to a downstream oligopoly of PC manufacturers, including Dell, HP, and Toshiba. From late 2002 until 2007, Intel engaged in various exclusive or near-exclusive deals with PC manufacturers in the CPU market, including, among others, conditional rebates, the exclusion or delayed launch of AMD-based machines, and market-share-based contracts.[4] These practices attracted significant scrutiny from antitrust authorities worldwide, as well as from AMD, resulting in billions of dollars in fines and settlements.

To organize the analysis, we first outline a simple two-period conceptual framework of downstream sourcing in which AMD's future "viability", capturing its effective ability to supply, invest, and be perceived as a credible long-run supplier, depends on current adoption. Intel can offer exclusive or near-exclusive contracts to selected buyers in the first period. Excluding some buyers reduces AMD's current sales, lowers its subsequent viability, and thereby depresses other buyers' incentives to adopt AMD in the next period. This framework formalizes the notion of contracting externalities in our setting: exclusivity with one set of buyers generates negative spillovers on others through the dynamic viability channel. We derive predictions about (i) the direct effect of restraints on targeted product lines or buyers and (ii) cross-buyer spillovers when restraints are imposed on other lines or firms, which we take to the data.

We use information from litigation to construct quantitative indices measuring the extent of exclusionary restraints in Intel's transactions with specific customers. We match these indices with market-level data on PC and CPU unit sales to calculate, for each PC brand over time, the share of PCs equipped with AMD processors. We combine this information with data on each CPU supplier's technological advancement and production capacity. Finally, we develop separate quantitative indices that measure the intensity of antitrust activity against Intel over time.

Using these data, we analyze the relationship between exclusive agreements, downstream technology adoption, and upstream capacity-related factors. Specifically, we estimate dynamic panel models of product line–level AMD adoption that allow for rich unobserved heterogeneity and dynamic persistence, and we exploit both internal (lag-based) instruments and external instruments based on antitrust activity to address the endogeneity of Intel's restraints.

First, we find robust evidence that Intel's contractual practices directly depress AMD adoption on targeted product lines. The dynamic estimates show that own-line restraints have economically meaningful effects in the short run: an additional restraint lowers AMD's share on the treated

[4] While some of the contractual forms used by Intel were not, *de jure*, exclusive – meaning they would not strictly trigger breach if a PC firm sourced outside Intel – they were designed to induce exclusivity or near-exclusivity by threatening the loss of benefits for downstream firms through credible threats. More details are provided in Section 3.

product line by 0.22–0.40 percentage points (roughly 2–4% of the average adoption level). Because sourcing decisions are highly persistent, these short-run effects magnify substantially over time. The implied long-run reductions range from roughly 2.4 to 2.8 percentage points evaluated at the sample means of the restraint indices, corresponding to about 20–30% of the average AMD adoption level across product lines.

Second, we present novel evidence of contracting externalities. We distinguish between *standard* restraints (which count the number of restraints imposed on a downstream PC firm in a given quarter) and *extreme* restraints (those that effectively preclude meaningful AMD adoption). Restraints imposed on *other* firms generate sizable spillovers: in our preferred specifications, an additional standard restraint imposed on any competing firm reduces other firms' AMD purchases by 0.1 percentage points in the short run, while an additional extreme restraint reduces purchases by 0.3 percentage points. At prevailing restraint intensity—where product lines face, on average, more than 10 standard and nearly 4 extreme cross-firm restraints—these estimates imply market-wide short-run reductions of 0.6 to 1.3 percentage points. These effects are amplified over time, with long-run reductions of 0.4 percentage points per additional standard restraint and 2.6 percentage points per additional extreme restraint (roughly 4% and 25% of AMD's average share). At typical observed restraint levels, these per-unit effects imply market-wide reductions of roughly 4 percentage points for standard restraints and 10 percentage points for extreme restraints.

The results are robust to a wide set of checks, including placebo tests on the timing of restraints and alternative specifications such as first-difference IV models. Across all specifications, the evidence consistently shows both direct exclusion and substantial cross-buyer spillovers. Taken together, these patterns provide consistent empirical evidence of contracting externalities: exclusivity directed at some buyers reduces AMD's effective viability and thereby lowers adoption even among buyers not directly subject to the contractual provisions.

While contracting externalities have been recognized in theory, empirical evidence has been largely absent. This paper provides, to our knowledge, the first direct measurement of these spillovers in a major industrial setting. The results demonstrate that exclusionary contracts can meaningfully alter market outcomes not only for targeted buyers but also for firms that never enter into such agreements. These findings underscore that effective antitrust scrutiny of exclusive dealing must look beyond bilateral contracting partners and consider the broader market-wide consequences of dynamic viability effects.

Related Literature Our paper contributes to the literature on vertical restraints by quantifying contracting externalities and the persistent competitive dynamics of exclusive contracts. We show how these restraints distort downstream input choices not only through direct contractual terms but also by altering rivals' capacity investments and triggering self-reinforcing expectations

about supplier viability. In the theoretical literature, such contracting externalities are the primary mechanism through which exclusive dealing becomes anticompetitive: restraints imposed on some buyers weaken a rival's viability, thereby distorting other buyers' adoption incentives. Our empirical results document this mechanism in practice.

The competitive effects of exclusive dealing have been debated since the Chicago School critique, which dismissed the view that these contracts could be used by a firm to exclude a rival (Posner, 1976). Post-Chicago models demonstrate several mechanisms through which exclusive deals may generate anticompetitive effects. A first channel is the *coordination/threshold* mechanism: in Rasmusen et al. (1991), refined by Segal and Whinston (1996, 2000), entry requires the entrant to secure a minimum number of buyers to cover fixed costs.[5] When entry is viable only at a minimum scale, an incumbent can exploit buyers' lack of coordination to profitably exclude rivals, even when overall market foreclosure is incomplete. Consequently, signing an exclusive contract imposes a negative contracting externality on the rest. This mechanism closely aligns with Intel's alleged practices, where partial exclusivity agreements with PC manufacturers may have prevented AMD from achieving the minimum efficient scale needed to compete effectively in the CPU market.

A second, complementary channel is cross-market rent extraction in *noncoincident markets*: Bernheim and Whinston (1998) and Whinston (2006) show how exclusives in one relationship shift bargaining positions and surplus extraction in related markets, creating contracting externalities beyond the directly contracted market.

More recent theoretical contributions have incorporated dynamics and asymmetric information to further explain when exclusivity is likely to harm competition. Calzolari and Denicolò (2013, 2015) develop a framework in which exclusive contracts can be anticompetitive when one firm holds a dominant position. They argue that if the dominant firm's competitive advantage is sufficiently large, it can impose exclusivity without fully compensating buyers, as rivals may not be able to offer competitive alternatives. This is consistent with Intel's ability to maintain its market share despite AMD's technological improvements, a prediction our empirical results substantiate. Dynamic considerations amplify these concerns: reducing a rival's profits today increases its probability of exit tomorrow (Cabral and Riordan, 1994), while buyers' expectations about a supplier's future viability can depress current demand and reinforce foreclosure (Katz and Shapiro, 1985).[6]

Exclusion can also be implemented via practices that do not carry a contractual obligation of exclusivity, such as resale price maintenance, loyalty rebates, and lump-sum payments to transfer industry profits to retailers and induce them not to accommodate efficient upstream entry. Asker and Bar-Isaac (2014) show how such vertical practices can facilitate exclusion. In the Intel litigation,

[5] Fumagalli and Motta (2006) extend this work to settings with buyer–retailers, with downstream competition moderating exclusion.

[6] A related strand of work examines how incumbents sustain dominance through preemptive innovation (Gilbert and Newbery, 1982).

plaintiffs alleged that Intel employed lump-sum rebates conditional on hardware manufacturers' loyalty, amid growing competition from AMD microprocessors. These payments were described as a "bribe" to secure Intel's market dominance, with the implicit threat that greater reliance on AMD chips would lead to the withdrawal of these rebates.[7] In Asker and Bar-Isaac (2014), the threat is plausible as the equilibria indicate that, upon entry by a rival, Intel lacks the incentive to continue such payments, causing downstream firms to lose this revenue stream if they shift to the entrant's products. DeGraba and Simpson (2013) propose a similar theory of harm when discussing the Intel case. Similarly, DeGraba (2013) models "naked exclusion" – exclusionary conduct lacking any efficiency justification – in which dominant firms pay buyers to foreclose smaller rivals, showing that exclusives or near exclusives can lead to anticompetitive outcomes even if the rival remains in the market and is profitable. This model illustrates how a dominant firm (Intel) can maintain its market power by limiting the rival (AMD) access to downstream markets. Complementing this mechanism, Chao et al. (2018) demonstrate that all-units discounts, that is, retroactive rebates that apply once a threshold purchase share is met, can act as a partial foreclosure device: they raise the dominant firm's profits and market share while reducing the rival's, even without explicit exclusivity clauses. These results highlight how pricing-based vertical restraints can produce foreclosure effects akin to those generated by formal exclusive contracts.

In sum, contemporary theories recognize multiple mechanisms, including coordination failures, spillover effects, and asymmetric information, through which exclusive dealing can strategically foreclose competitors. Of course, no theoretical model can account for all the complex features of the Intel-AMD case and the microprocessor industry at once.

Empirically, the effects of exclusivity are mixed and context-dependent, with identification complicated by endogenous contract choice. Efficiency rationales appear in beer distribution (Sass, 2005; Asker, 2016), while anticompetitive effects are documented in automobiles and retail (Nurski and Verboven, 2016; Ater, 2015). Hortaçsu and Syverson (2007) find little foreclosure evidence in cement, emphasizing the importance of productivity differences. Conlon and Mortimer (2021) find that quantity-based "vertical rebates", similar to the loyalty-based payments used by Intel, have foreclosure effects in the snack food market. Our paper complements this literature by documenting contracting externalities in a high-innovation, capacity-constrained technology market.

Finally, Goettler and Gordon (2011) analyze competition between Intel and AMD in the PC microprocessor industry using a dynamic oligopoly model with durable goods and endogenous innovation. They find that AMD's presence constrains Intel's prices, benefiting consumers, but may slow the overall rate of innovation. Their model does not directly incorporate specific contractual restraints, such as exclusivity agreements, and rationalizes AMD shares through strong Intel fixed

[7] State of New York v. Intel Corp., Complaint, §99 (S.D.N.Y. 2009), available at https://www.intel.com/pressroom/legal/docs/NY_AG_v._Intel_COMPLAINT.pdf.

effects. Their counterfactual analyses simulate scenarios with varying degrees of foreclosure. These simulations show that some level of foreclosure increases prices, but also leads to higher innovation rates and potentially higher consumer surplus, as the increase in innovation offsets the negative effects of higher prices. Their work highlights the complex trade-offs involved in antitrust policy regarding vertical restraints that motivate our analysis, while abstracting from the complexity of contracting externalities, which is the focus of our work.

The remainder of the paper is organized as follows. Section 2 develops the conceptual framework. Section 3 describes the data and institutional background. Section 4 presents the empirical strategy. Section 5 reports the main results. Section 6 presents robustness checks. Section 7 concludes.

2 A Framework of Contracting Externalities in the CPU market

We present a simple two-period framework of downstream customers' sourcing choices that captures how exclusivity with some buyers can depress a rival supplier's future viability and, through that channel, other buyers' incentives to adopt the rival. Our motivation follows the idea of contracting externalities in Bernheim and Whinston (1998).

We consider two upstream firms, an incumbent (Intel) and a smaller rival (AMD), serving a set of downstream buyers $i \in \mathcal{I}$. To aggregate buyers' adoption into AMD's overall period-1 demand, we assign each buyer a weight $\theta_i > 0$, normalized so that $\sum_{i \in \mathcal{I}} \theta_i = 1$. These weights scale buyers' sourcing decisions into AMD's total period-1 sales, which in turn determine its next-period viability.

AMD enters the first period with an initial viability state $S_1 \geq 0$, which captures AMD's overall ability to supply and compete in future periods (capacity, financial strength, perceived reliability). Viability evolves endogenously with period-1 AMD sales.

Each individual buyer represents only a small fraction of AMD's global demand. We therefore treat each $i \in \mathcal{I}$ as an "atomistic" buyer whose own sourcing decision has a negligible effect on the aggregate state S_2. Buyers are forward-looking, but because their individual choices do not affect S_2, each buyer's continuation value is independent of its period-1 sourcing choice. This implies that buyers do not internalize how their own sourcing contributes to AMD's future viability, yielding a clean maximal benchmark for uninternalized contracting externalities. By contrast, Intel's contracting strategy makes clear that, in aggregate, these decisions are material for AMD's viability. We return to the implications of this assumption in the model discussion at the end of the section.

Time is $t = 1, 2$. In each period t, AMD's quality or price–performance advantage relative to

Intel is denoted by $\Delta Q_t \in \mathbb{R}$ and is common knowledge. If buyer i sources a share $w \in [0,1]$ from AMD when viability is S_t, its reduced-form flow payoff is:

$$\pi_{it}(w; S_t) = \alpha\, \Delta Q_t\, w - (\kappa_0 - \kappa_1 S_t)\, w - \phi\, w^2,$$

where $\alpha > 0$ scales the value of AMD's advantage; $\kappa_0 > 0$ and $\kappa_1 > 0$ capture how viability reduces switching and operational frictions; and $\phi > 0$ is a convex adjustment-cost (inertia) parameter.

AMD's viability updates once between periods according to:

$$S_2 = S_1 + \eta \sum_{j \in \mathcal{I}} \theta_j w_{j1}, \tag{1}$$

where $\eta > 0$ measures the sensitivity of future viability to AMD's aggregate period-1 sales, $\sum_{j \in \mathcal{I}} \theta_j w_{j1}$.[8] This captures the idea that current adoption builds scale, financial strength, and reputation, reinforcing AMD's competitive position. In the Intel–AMD context, capacity constraints and scale economies make this feedback particularly salient: lower period-1 sales limit AMD's ability and incentive to expand production and R&D, weakening its credibility as a long-term supplier.

Intel's second-period profit $\Pi_{I2}(S)$ is continuously differentiable and strictly decreasing in AMD viability, $\Pi'_{I2}(S) < 0$. Intel's first-period profit depends on its contracts and sales; it is summarized below.

The timing of the game is as follows:

1. The initial state $(S_1, \Delta Q_1, \Delta Q_2)$ is common knowledge.

2. Intel offers each buyer i a take-it-or-leave-it exclusivity contract specifying a transfer T_i (rebate or lump sum). If buyer i accepts, it commits to exclusive sourcing in period 1 (i.e., $w_{i1} = 0$) and receives the transfer T_i. If it rejects, it sources freely and receives no transfer.

3. In period 1, each buyer i who rejects exclusivity chooses $w_{i1} \in [0,1]$; those who accept choose $w_{i1} = 0$. Period-1 payoffs are realized, and transfers T_i are paid to accepting buyers.

4. AMD's viability evolves according to equation (1).

5. In period 2, all buyers simultaneously choose $w_{i2} \in [0,1]$ without contractual restrictions. Industry profits are $\pi_{i2}(\cdot)$ for buyers and $\Pi_{I2}(S_2)$ for Intel.

[8] We choose a linear and deterministic specification for simplicity and tractability. A more general law of motion could incorporate depreciation, concavity (diminishing returns to scale), or a minimum viable scale threshold, the latter making exclusion disproportionately powerful if it keeps AMD below critical mass. Our simple form captures the core feedback mechanism while providing a clear benchmark for the contracting externality.

The solution concept is subgame-perfect equilibrium, solved by backward induction.

In period 2, given S_2, buyer i chooses w_{i2} to maximize $\pi_{i2}(w; S_2)$. Because $\pi_{i2}(\cdot; S_2)$ is strictly concave, the unique interior optimum satisfies:

$$w_{i2}^{\text{int}}(S_2) = \frac{\alpha \Delta Q_2 - (\kappa_0 - \kappa_1 S_2)}{2\phi}.$$

The best response is therefore:

$$w_{i2}^*(S_2) = \begin{cases} 0, & w_{i2}^{\text{int}}(S_2) \le 0, \\ w_{i2}^{\text{int}}(S_2), & 0 < w_{i2}^{\text{int}}(S_2) < 1, \\ 1, & w_{i2}^{\text{int}}(S_2) \ge 1, \end{cases} \qquad \frac{\partial w_{i2}^*}{\partial S_2} = \begin{cases} \frac{\kappa_1}{2\phi}, & 0 < w_{i2}^{\text{int}}(S_2) < 1, \\ 0, & \text{otherwise.} \end{cases}$$

In period 1, for a buyer not bound by exclusivity, the same reasoning yields:

$$w_{i1}^{\text{int}}(S_1) = \frac{\alpha \Delta Q_1 - (\kappa_0 - \kappa_1 S_1)}{2\phi}, \qquad w_{i1}^*(S_1) = \begin{cases} 0, & w_{i1}^{\text{int}}(S_1) \le 0, \\ w_{i1}^{\text{int}}(S_1), & 0 < w_{i1}^{\text{int}}(S_1) < 1, \\ 1, & w_{i1}^{\text{int}}(S_1) \ge 1. \end{cases}$$

If the buyer accepts exclusivity, $w_{i1} = 0$ by contract. Payoffs are normalized so that sourcing exclusively from Intel ($w_{i1} = 0$) yields zero payoff; $\pi_{it}(w; S_t)$ is thus the incremental surplus from sourcing from AMD relative to this baseline.

Let $\pi_{i1}^{\text{uncon}} \equiv \max_{w \in [0,1]} \pi_{i1}(w; S_1)$ and $\pi_{i1}^{\text{excl}} \equiv \pi_{i1}(0; S_1) = 0$. An exclusion offer ($\text{EXCL}_i = 1, T_i$) is accepted if and only if

$$T_i \ge \pi_{i1}^{\text{uncon}}.$$

Intel, offering take-it-or-leave-it contracts, sets $T_i^* = \pi_{i1}^{\text{uncon}}$ when exclusion is profitable.[9]

Let $\mathcal{E} \subseteq \mathcal{I}$ denote the set of excluded buyers. Unconstrained buyers $j \notin \mathcal{E}$ choose w_{j1}^*, while $w_{j1} = 0$ for $j \in \mathcal{E}$. AMD's second-period viability is then

$$S_2(\mathcal{E}) = S_1 + \eta \sum_{j \notin \mathcal{E}} \theta_j w_{j1}^*.$$

[9] If contracts cap $w_{i1} \le \bar{w}_i$ rather than require $w_{i1} = 0$, we can simply replace w_{i1}^* by $\min\{w_{i1}^{\text{int}}(S_1), \bar{w}_i\}$ throughout. The signs of the profitability condition, the externality, and all comparative statics are unchanged, though the effects become weak when the cap binds.

Intel's total profit under the exclusion set $\mathcal{E}$ is:

$$\Pi_I(\mathcal{E}) = \Pi_{I1}(\mathcal{E}) + \Pi_{I2}\big(S_2(\mathcal{E})\big),$$

where $\Pi_{I1}(\mathcal{E})$ denotes Intel's total period-1 profit from sales to all buyers (both excluded and non-excluded) net of any transfers T_i^* paid to excluded buyers. We treat Π_{I2} as a reduced-form profit function capturing Intel's improved future market power when AMD is weaker.

Profitability of excluding buyer i. Should Intel offer an exclusivity contract to buyer i, or let the buyer buy freely? Intel's total profit changes if it excludes buyer i are:

$$\Delta\Pi_I(i) = \underbrace{\Big(\Pi_{I1}^{\text{excl } i} - \Pi_{I1}^{\text{no-excl } i}\Big)}_{\text{change in period-1 profit}} + \underbrace{\Big(\Pi_{I2}\big(S_2^{\text{excl } i}\big) - \Pi_{I2}\big(S_2^{\text{no-excl } i}\big)\Big)}_{\text{change in period-2 profit}},$$

where the first term captures the short-run cost of paying T_i^*, and the second term captures the long-run gain from lower AMD viability.

AMD's viability next period, holding the existing exclusion set $\mathcal{E}$ fixed (and thus $w_{j1} = 0$ for all $j \in \mathcal{E}$), is:

$$S_2^{\text{no-excl } i} = S_1 + \eta(\theta_i w_{i1}^* + \sum_{j \neq i} \theta_j w_{j1}^*), \quad S_2^{\text{excl } i} = S_1 + \eta \sum_{j \neq i} \theta_j w_{j1}^*,$$

Excluding buyer i raises Intel's profit if $\Delta\Pi_I(i) > 0$, with T_i^* satisfying the buyer's participation constraint.[10]

Equilibrium existence and characterization Backward induction yields a subgame-perfect equilibrium. Each buyer's sourcing choice $w_{it}^*(S_t)$ is unique in every period because the payoff function $\pi_{it}(w; S_t)$ is strictly concave in w.

Because the exclusion decision is made over the finite family of subsets $\mathcal{P}(\mathcal{I})$, Intel's problem $\max_{\mathcal{E} \subseteq \mathcal{I}} \Pi_I(\mathcal{E})$ admits an optimal solution $\mathcal{E}^*$.[11]

[10] Formally, Intel chooses an exclusion set of buyers. The marginal profitability $\Delta\Pi_I(i)$ depends on which other buyers are excluded, so it may be positive for some exclusion sets and negative for others.

[11] Corner cases arise when the interior solution $w_{it}^{\text{int}}(S_t)$ lies outside the feasible range $[0, 1]$. In such cases, the buyer's optimal share is truncated to the nearest boundary, that is, $w_{it}^* = 0$ if $w_{it}^{\text{int}}(S_t) \leq 0$ and $w_{it}^* = 1$ if $w_{it}^{\text{int}}(S_t) \geq 1$. The comparative statics derived below remain valid in weak form: changes in parameters can no longer increase w_{it}^* beyond 1 or decrease it below 0, but the direction of effects (higher S_t increasing AMD adoption) continues to hold wherever the solution is interior.

Comparative statics and contracting externalities For interior solutions and a given exclusion set $\mathcal{E}$, period-2 AMD adoption by buyer i responds as follows:

$$\begin{aligned}
&\text{(i)} \quad \frac{\partial w_{i2}^*}{\partial \Delta Q_2} = \frac{\alpha}{2\phi} > 0, \\
&\text{(ii)} \quad \frac{\partial w_{i2}^*}{\partial \phi} = -\frac{\alpha \Delta Q_2 - (\kappa_0 - \kappa_1 S_2)}{2\phi^2} < 0, \\
&\text{(iii)} \quad \frac{\partial w_{i2}^*}{\partial S_2} = \frac{\kappa_1}{2\phi} > 0.
\end{aligned}$$

Thus, (i) better AMD quality/price–performance increases adoption; (ii) greater inertia ϕ dampens all responses and lowers adoption; and (iii) higher AMD viability S_2 makes AMD more attractive to all buyers.

Most importantly, if Intel excludes buyer j in period 1, w_{j1} is forced to zero; AMD sells less, so S_2 falls ($S_2^{\text{excl } j} < S_2^{\text{no-excl } j}$). Others then reduce period-2 adoption purely through this viability channel. The discrete effect on any other buyer $i \neq j$ is therefore:

$$\Delta w_{i2}^* = w_{i2}^*(S_2^{\text{excl } j}) - w_{i2}^*(S_2^{\text{no-excl } j}) = -\frac{\kappa_1}{2\phi}\, \eta\, \theta_j\, w_{j1}^* \ \leq 0,$$

strictly negative whenever $w_{j1}^* > 0$ and w_{i2}^* is interior. With a set $\mathcal{E}$ excluded, the externality adds up linearly:

$$\Delta w_{i2}^* = -\frac{\kappa_1}{2\phi}\, \eta \sum_{j \in \mathcal{E}} \theta_j w_{j1}^* \leq 0.$$

Exclusivity with a large or would-be early adopter (θ_j high, w_{j1}^* large) suppresses AMD's short-run sales, weakens its subsequent viability, and, absent any direct restraint on them, tilts other buyers away from AMD in the future.

Excluding buyer i is therefore profitable for Intel if the future gain from lowering S_2 exceeds the period-1 sacrifice (including the participation transfer):

$$\left(\Pi_{I1}^{\text{excl } i} - \Pi_{I1}^{\text{no-excl } i}\right) + \left(\Pi_{I2}(S_2^{\text{excl } i}) - \Pi_{I2}(S_2^{\text{no-excl } i})\right) > 0, \quad \Pi_{I2}'(S) < 0.$$

This is the dynamic foreclosure logic. Exclusive contracts reduce AMD's period-1 sales, weaken its viability S_2, and, through reduced scale, lower all buyers' incentives to adopt AMD in the next period. The resulting decline in second-period adoption is a *contracting externality*: a negative cross-buyer spillover generated by dynamic feedback rather than direct information exchange. In our reduced-form setup, excluding some buyers depresses S_2, which, in turn, makes AMD less attractive to all other buyers. This captures the central cross-buyer spillover that, in richer models with endogenous reservation transfers, can make exclusion decisions strategic complements and give rise to "divide-and-conquer" dynamics as in the naked-exclusion literature (Rasmusen et al.,

1991; Segal and Whinston, 2000).[12]

Table 1 summarizes the model's comparative statics. These derivatives characterize buyers' period-2 best responses holding the exclusion set $\mathcal{E}$ fixed.[13] In the empirical analysis, we abstract from the specific channels through which viability evolves and focus instead on how exclusivity-induced changes in S_t generate cross-buyer spillovers in future adoption.

Table 1: Comparative Statics of Period-2 AMD Adoption

Parameter / Shift	**Effect on w_{i2}^* (interior)**	**Interpretation**
$\uparrow \Delta Q_2$	Positive	Higher AMD advantage increases adoption
$\uparrow \phi$	Negative	Higher switching frictions dampen responsiveness
Exclude buyer j in $t = 1$ ($w_{j1}^* \to 0$)	Negative	Contracting externality via lower AMD future viability (S_2)

Discussion This framework provides a simple baseline for understanding contracting externalities in vertical markets. The assumption that individual buyers behave atomistically and do not meaningfully affect AMD's future viability S_2 is a useful benchmark. The assumption implies that a buyer's decision has no measurable effect on S_2; when deciding whether to accept Intel's exclusivity offer, a buyer only considers its own direct period-1 product loss, without internalizing the effect of its own adoption on AMD's long-run position. The entire viability externality is therefore uninternalized and operates fully across buyers.

In the PC market, only a very small number of extremely large buyers were sufficiently important for their sourcing decisions to plausibly influence AMD's perceived viability. For example, the European Commission notes that *"Fujitsu-Siemens, which in 2006, although it was the next largest OEM [original equipment manufacturer] after IBM in terms of market share, saw itself as too small to legitimise AMD for enterprise"*.[14] The Commission's notion of "legitimization" directly parallels the viability channel in our model: AMD's ability to compete in the future requires achieving sufficient scale and visibility today.

[12] In particular, our setup is conceptually related to Doganoglu and Wright (2010), who study exclusive dealing in the presence of network effects. In our framework, dynamic "viability" spillovers play a role analogous to network effects: early exclusive contracts reduce AMD's viability, which, in turn, lowers AMD's attractiveness to other buyers.

[13] In principle, sufficiently large parameter changes could also alter the optimal exclusion set $\mathcal{E}^*$, which would then shift S_2 and w_{i2}^*. Such general-equilibrium adjustments are outside the scope of these local comparative statics.

[14] European Commission (2009), Intel Decision, Recital 1588.

Intel's conduct is consistent with this asymmetry. Its most restrictive conditional rebates were targeted precisely at large buyers; those firms may have, at least partially, internalized AMD's future viability and therefore required stronger inducements to remain exclusive to compensate for both the lost period-1 surplus and the lost future benefit from a stronger AMD (see Section 3). This creates a coordination game among large buyers. If other large buyers are expected to stay non-exclusive and support AMD, then it is individually optimal to also support AMD. Conversely, if other large buyers are expected to take exclusivity and abandon AMD, then AMD's future is doomed (S_2 will be low). The large buyer then faces a prisoner's dilemma: supporting AMD alone is costly and futile, so it becomes optimal to take Intel's exclusivity payment. Intel's strategy becomes about breaking this coordination, using high transfers and, for example, sequential offers to convince one pivotal large buyer to defect.[15] Our model abstracts from such internalization by large buyers, providing a clean benchmark for the fringe. Consistent with this benchmark, the empirical results in Section 5 indicate that the contracting externalities were driven significantly by the aggregated actions of this large fringe, as Intel's strategy focused on securing the pivotal, large buyers whose choices could tilt the ecosystem in Intel's favor.

Second, we abstract from the possibility that contracts are renegotiated over time or that buyers interact strategically. In practice, renegotiation could allow Intel to renew or adjust exclusives, and strategic behavior among buyers could influence who accepts exclusivity and when.

Third, in our framework, AMD is not a strategic actor: its quality advantage (ΔQ_t) is exogenous, and it does not adjust pricing, capacity, or R&D in response to Intel's contracts or its own viability. These decisions are obviously central to the competitive dynamic. Making ΔQ_2 a function of S_2, so that greater viability improves AMD's future quality, would strengthen the feedback loop and could create nonlinear thresholds reinforcing our dynamic foreclosure logic.

Finally, the framework abstracts from downstream market growth (which we control for in the empirical analysis). If growth occurs for products outside Intel's exclusive contracts, AMD may scale despite partial foreclosure; if growth occurs inside the contracted products, the dynamic feedback is stronger because the incumbent captures a growing base of demand.

[15] 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.

3 Data

We analyze the US market for CPUs from 2002 until mid-2009, a period during which Intel's relationship with prominent downstream customers was characterized by exclusive deals and vertical restraints. We focus on CPUs installed in personal computers purchased by US consumers in the home and business segments.[16] Sales to the home and business segments account for the majority of PC sales. Our dataset combines multiple sources, including information on PC and CPU sales and attributes, CPU manufacturers' production capacity, restraints characterizing Intel's vertical contracts, and the scope of the global legal action taken in connection with these restraints.

We estimate dynamic panel models to capture the impact of Intel's vertical restraints on downstream PC firms' adoption of AMD technology. Our unit of observation is the combination of a PC firm (e.g., Acer), brand (e.g., Aspire D), market segment (Home or Business), and quarter (from 2002:1 to 2009:2), resulting in around 3,800 observations.

3.1 PC Makers' Adoption of AMD and CPU Quality

We obtained quarterly PC and CPU sales data from Gartner Group. Table 2 reports descriptive statistics for PCs. During this period, there are 30 PC firms and 169 PC brands, yielding 204 unique firm-brand combinations. Sales are split almost evenly between home and business segments. The largest PC producers are Dell (31% market share) and HP (22%). There are no close competitors; however, a few firms have market shares between 5% and 10% (on average), including Acer, Fujitsu Siemens, Gateway, Lenovo, IBM, and Toshiba. PC prices vary widely, ranging from $138 to $3,218.[17] Over 60% of PCs sold are mobile devices (e.g., laptops, notebooks, ultraportables).

We define a product line as a firm–brand–segment (home/business) combination. We observe the CPU firm (e.g., AMD) and model (e.g., Athlon 64 X2) associated with each product line. The dependent variable is the share of sales in a product line-quarter with an AMD chip installed. For example, Acer's Aspire D was sold to home consumers in 2006:2. Some of these were equipped with an AMD Athlon 64 chip and others with an Intel Celeron chip.

About 22% of PCs sold over the period were equipped with an AMD chip. Figure 1 shows the evolution of the market share of AMD over time. A noticeable increase occurs around 2006, when AMD's share rises from around 12% to over 25%. This gain, however, was temporary, and AMD's share declined to 15% by 2009. We discuss several developments that occurred during that time,

[16] We exclude Apple products as those exclusively used IBM's chips during much of the sample period. We also do not include servers, as server sales were not recorded in the dataset before 2005. However, we do record exclusionary restraints on the use of AMD chips in servers and incorporate those into the analysis. As explained below, this is consistent with the institutional background: market participants believed that adopting AMD chips in servers could have important spillovers into other market segments, such as desktops.

[17] We compute the average price by dividing revenue by the number of units.

Table 2: PC Descriptive Statistics

	Mean	Std. Dev.	Min	Max
PC Price (100$)	7.73	4.02	1.38	32.18
Sold in Home Market	0.46	0.50	0.00	1.00
Mobile PC	0.61	0.49	0.00	1.00
No. of PC Firms	30			
No. of PC Brands	169			
No. of Product Groups	204			
No. of CPU Models	26			
PC contains AMD CPU	0.22	0.41	0.00	1.00

Notes: All variables expressed in monetary terms were deflated using the quarterly consumer price index (Bureau of Labor Statistics, base year 2000). A product group is a PC firm–PC brand combination.

including technological trends and strategic changes in the competitive arena, as well as changes in the intensity of exclusive restraints.

Figure 1: AMD Market Shares

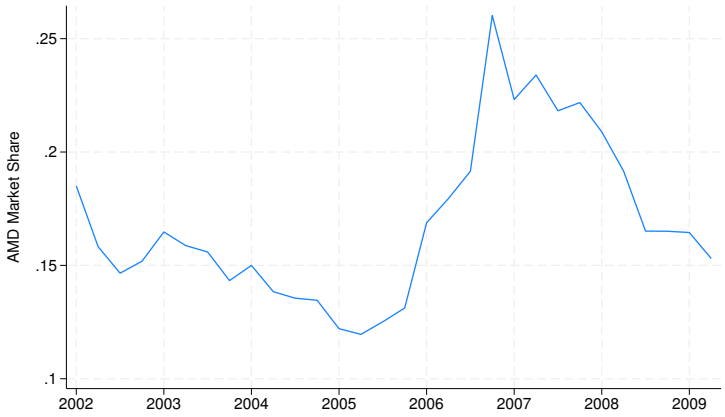

Figure 2 presents AMD's market share for selected PC firms, along with these firms' respective shares of the PC market before and after 2007.[18] While some firms continued with exclusivity into 2007, Dell discontinued its exclusive arrangements with Intel in late 2006. Figure 2 shows an increase in AMD adoption at the beginning of 2007 across *all firms*, where the rate of increase varies substantially. This shift reflects, in part, the removal of exclusivity restrictions. For instance, Dell, Lenovo, and Toshiba offered virtually no AMD-based systems before 2007 due to their exclusive

[18] Several mergers occurred in the industry: Gateway/emachines merged in 2004 and subsequently merged with Acer in 2007. Lenovo and IBM merged in 2005.

agreements with Intel.[19] However, the data also exhibit an increase in AMD adoption among firms *not* bound by exclusivity with Intel in the earlier sample period. Systemax, for example, expanded its AMD offerings post-2007 despite never facing Intel-imposed restraints. This suggests the presence of contracting externalities or other competitive forces unrelated to exclusive restraints. We explore this in detail through our estimation approach, designed to disentangle these contracting externalities from other forces.

Figure 2: AMD Purchases by Selected PC Firms

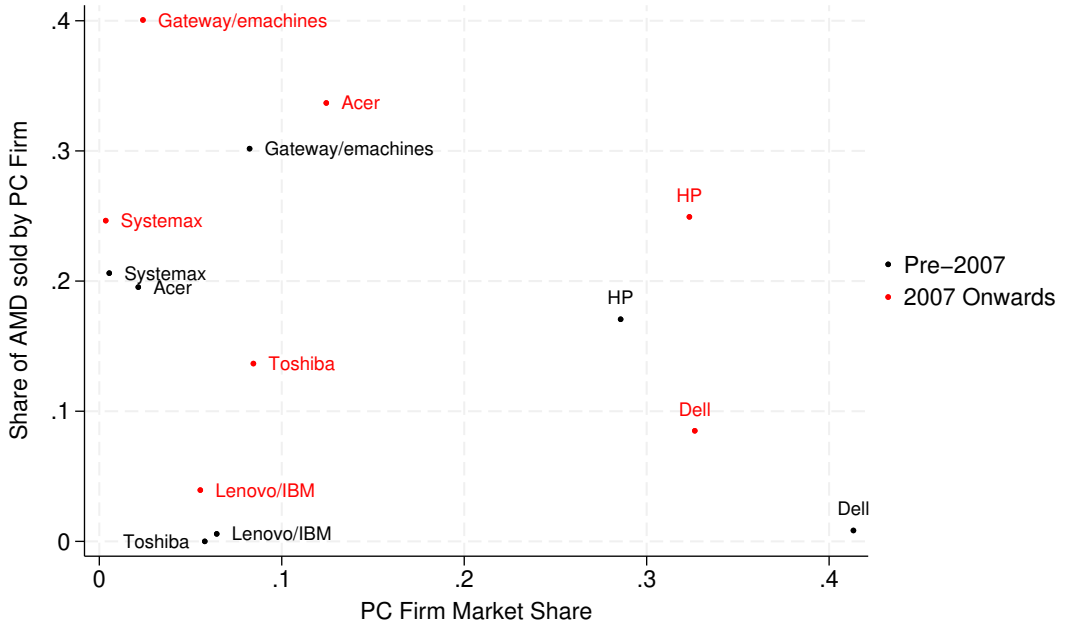

Panel A of Table 3 presents descriptive statistics for the CPUs sold in each PC. We obtained CPU prices from Instat and published list prices.[20] As shown in the table, Intel chips were, on average, 35% more expensive than their AMD counterparts. The subsequent rows demonstrate that this price premium does not reflect superior Intel quality. To measure quality, we use a continuous measure of the CPU model performance (the benchmark) obtained from Passmark.[21] The benchmark for Intel chips was, on average, 13% lower than AMD chips.

[19] Internal documents revealed in litigation showed that Intel's financial support was often conditioned on exclusivity or near-exclusivity. A 2002 Dell document stated that the "...original basis for the fund is...Dell's loyalty to Intel," meaning "no AMD processors." Similarly, an HP executive wrote in a 2002 email: "PLEASE DO NOT...communicate to the regions, your team members, or AMD that we are constrained to 5% AMD by pursuing the Intel agreement."

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

[21] Accessed from www.cpubenchmark.net. Details are provided in Appendix A1.

Table 3: Descriptive Statistics

	Obs	Mean	Std. Dev.	Min	Max
A. CPU Characteristics					
Price AMD CPU (100$)	3600	1.04	0.64	0.19	5.43
Price Intel CPU (100$)	3858	1.40	0.52	0.20	3.77
Benchmark AMD CPU (in 1000)	3600	0.78	0.47	0.27	3.50
Benchmark Intel CPU (in 1000)	3858	0.69	0.58	0.17	6.45
CPU Benchmark/Dollar AMD	3600	9.88	8.03	1.27	39.33
CPU Benchmark/Dollar Intel	3858	6.11	6.51	1.09	45.10
Age AMD Brand	3600	6.59	4.29	1.00	23.00
Age Intel Brand	3858	7.09	5.42	1.00	30.00
B. Capacity and Cashflow					
AMD Capacity Index	3858	7.61	3.72	3.00	13.00
Intel Capacity Index	3858	31.56	6.47	23.00	44.00
AMD Free Cash (100M$)	3858	7.89	3.25	2.82	19.05
C. Antitrust Variables & Exclusionary Restraints					
Num. Pending Antitrust Cases Against Intel	3858	3.22	1.69	1.00	6.00
Exclusionary Restraints Index	1540	3.32	1.37	1.00	6.00
Extreme Restraints Index	1540	1.30	0.68	0.00	3.00
Intel Payments to Dell (M$)	744	579	396	0.00	1070
Intel Payments to non-Dell PC Firms (M$)	3114	4.60	25.98	0.00	193

Notes: An observation is a product line-quarter combination. AMD (Intel) statistics are reported only over AMD-based (Intel-based) CPU models, hence the counts differ. The underlying CPU characteristics are aggregated to the product line–quarter level as described in Appendix A1.

The value delivered by chip manufacturers is captured by the benchmark-per-dollar metric (CPU benchmark score divided by CPU price). This indicates the superior value of AMD chips relative to Intel, which were 61% better in terms of this indicator. AMD's technological success is reflected, for example, in a New York Times article covering the State of New York antitrust case against Intel:

> "In 2005, Michael S. Dell's namesake company was getting pounded. His competitors were selling personal computers and servers built on cheap, popular and powerful chips from Advanced Micro Devices, while Mr. Dell had stuck loyally with slower chips from Intel. In an email note to Intel's chief executive, Paul S. Otellini, Mr. Dell threatened to switch to A.M.D."[22]

[22] "State Accuses Intel in an Antitrust Suit," Ashlee Vance, New York Times, November 5 2009.

Likewise, a 2002 internal HP presentation noted that AMD's Athlon desktop processor "had a unique architecture" and was "more efficient on many tasks," adding that AMD offers "no-compromise performance at superior value." [23]

However, in 2006, Intel introduced a new generation of Intel chips, specifically the Intel Core product family, which offered substantial improvements over incumbent generations. Indeed, the benchmark-per-dollar metric increases more rapidly in the latter part of the sample period for Intel, even surpassing AMD's in some cases (e.g., the highest values of the benchmark-per-dollar metric for Intel are associated with core processors). The European Commission 2009 decision also notes this trend

> "... 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." [24]

Considering this information together with Figure 1, we obtain a preliminary view of a 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, but increased around the time that Intel started to compete along a technological edge. However, this pattern aligns with the timeline of Intel's exclusivity arrangements, which we analyze next.

3.2 CPU Production Capacity

The production process in the microprocessor industry is multidimensional, and crucial inputs of productive capacity are disclosed in firms' annual reports. These include the number of fabrication facilities (FABs), the silicon wafer size used at each FAB (larger wafers enable more simultaneous CPU production), and the integrated circuit (IC) process node (measured in nanometers, where smaller nodes allow both higher CPU yields per wafer and improved CPU power efficiency). Using this information, we construct a firm-level *capacity index*.

A firm's capacity index is the yearly sum of rank points assigned to each FAB, where each FAB's points reflect its wafer size and IC process node. Specifically, process nodes are ranked from largest to smallest in nanometers (with smaller nodes reflecting more advanced manufacturing capabilities), and wafer sizes are ranked from smallest to largest (larger wafers implying greater production capacity). Each FAB receives the sum of its two rank points, and the firm's capacity index is the total of these points across all of its active FABs in that year. All rankings are performed within-year to reflect the evolving technology frontier, so the index increases with both the number of FABs a firm operates and the technological sophistication of those facilities.

[23] European Commission 2009 decision (COMP/C-3 /37.990 - Intel)
[24] European Commission 2009 decision (COMP/C-3 /37.990 - Intel)

Table 4: 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 4 shows the evolution of the processor makers' production technology and capacity over time.[25] AMD consistently lags behind Intel in both IC process technology and wafer size. In addition, capacity expansion required substantial capital expenditures; in 2007, constructing a single FAB required approximately $5 billion (Brown and Linden, 2009). Using AMD's quarterly financial reports, we calculate AMD's quarterly free cash flow.[26] Panel B reveals that AMD's free cash in each quarter was, on average, $789 million, significantly below the investment needed for new FAB construction. Industry analysts emphasize the significance of Intel's production capacity advantage. A blogger following the industry remarked 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."[27]

The blogger's prediction may, in fact, have materialized. While AMD offered better price-to-performance value in 2002, Intel ultimately regained technological leadership in subsequent years. AMD's chronic production capacity constraints and insufficient capital reserves for expansion may have contributed to this reversal.

[25] We exclude pilot FABs and smaller FABs. Including these would further amplify Intel's capacity advantage.

[26] Data retrieved from http://ir.amd.com/ on September 18, 2014.

[27] 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 9, 2017). The original URL is no longer active.

3.3 Antitrust Activity and Exclusive Restraints

The primary explanatory variables measure the degree of contractual restraints associated with the buyer's terms of trade with Intel, as well as those imposed on other buyers. To quantify these relationships, we construct variables capturing the nature of Intel's vertical contracts. The main empirical challenge is the absence of a unified, authoritative source detailing all these contractual arrangements, which are not directly observed. Instead, we rely on information extracted from legal proceedings pertaining to Intel's practices.

These cases include a private lawsuit by AMD and antitrust cases brought by the European Commission (EC), the Japanese Fair Trade Commission (FTC), the Korean FTC, the State of New York, and the United States FTC. These cases address Intel's behavior over various time frames spanning 2002:4 to 2007:4. Additionally, we draw on information from the U.S. Securities and Exchange Commission (SEC) case against Dell. Although not an antitrust case, it centered on the allegation that Dell failed to disclose to investors that a substantial share of its operating income was derived from payments received from Intel in return for exclusivity. As such, the case provides insight into Dell's exclusive relationship with Intel. According to the SEC's complaint against Dell,

> "Intel paid the computer maker rebates as part of a deal in which Dell agreed not to use microchips manufactured by Intel's rival AMD. We're not talking small change: The payments totaled $4.3 billion between 2003 and 2006."

Table 5 summarizes the timeline of legal actions against Intel. The cases differed in both scope and legal outcome. In 2005, the Japanese FTC issued a decision finding that Intel had obtained exclusivity or near-exclusivity from major Japanese PC manufacturers in violation of antitrust laws. Intel chose to comply with the decision by agreeing to refrain from certain practices, such as contractually requiring customers to use Intel chips exclusively. In 2009, Intel paid AMD $1.25 billion to settle a lawsuit filed by AMD in 2005. The Korean FTC ruled against Intel in 2008 and rejected Intel's appeal in 2013. The European Commission fined Intel €1.06 billion in 2009. This decision was upheld by the General Court in 2014, but, following Intel's appeal, was partially annulled. In 2023, the Commission narrowed the scope of the case to include only naked restrictions and imposed a revised fine of €376 million. The naked restrictions involved Intel paying PC firms to halt or delay the launch of AMD products and to limit the sales channels available for AMD models. In the US, antitrust cases brought by the State of New York and the FTC were settled with minimal or no financial penalties.

Antitrust Index We use the case files to construct an index capturing the scope and magnitude of antitrust activity. The index records the number of pending antitrust cases brought against Intel

Table 5: Antitrust activity timeline

2001	AMD files complaint with European Commission (EC)
2004	Japan Fair Trade (FT) opens investigation
2005	Japan FT issues decision that Intel violated rules; Intel complies
	Korea FT opens investigation
	AMD files lawsuits in Germany, Japan, and US
2007	EC brings charges
	Korea FT brings charges
2008	State of New York opens investigation
	US Federal Trade Commission (FTC) opens investigation
	Korea FT decision that Intel violated antitrust law
	Intel appeals Korea FT Decision
2009	EC decision that Intel violated rules; Intel appeals
	State of New York files lawsuit
	AMD cases against Intel end in settlement
	US FTC brings charges
2010	US FTC case ends in settlement
2012	State of New York case ends in settlement
2013	Korea FT upholds ruling against Intel appeal
2014	EC decision of 2009 upheld by the court; Intel appeals
2022	EC decision of 2009 partially annulled by the court
2023	EC re-imposes fine based on narrowed scope

by AMD and regulators worldwide in the relevant quarter.[28] Panel C of Table 3 shows that, on average, there were 3.22 pending cases per quarter, with a maximum of 6 pending cases in 2008, quarters 2, 3, and 4. These case-count variables are used to characterize the legal environment and document antitrust pressure over time. In the estimation, we use underlying legal events, such as the opening of investigations, formal decisions, and statements of objection, as instruments for the restraint variables.

Exclusionary Restraints Index The case files reveal a variety of instruments through which Intel may have affected the adoption of AMD's technology by downstream customers. We define,

[28] Information on the antitrust activity is obtained from Intel's annual reports. For regulatory cases, we use the date on which a formal investigation was initiated as the case's starting date. For AMD's lawsuits, we used the filing date. The ending date is defined as either the date on which a decision was rendered by the relevant authority or court (regardless of whether the decision was later appealed) or the date of settlement, where applicable.

at the product-line-quarter level, binary indicators that equal one if an instrument was employed.[29]

Specifically, our analysis focuses on the following types of restraints: caps on the amount sold of AMD-based PCs; rebates conditional on the extent of purchases from AMD (i.e., conditional rebates); exclusion of AMD from certain product lines or delayed launches of specific AMD-based machines (i.e., a type of naked restriction); commitments to increase Intel's market share (i.e., the share of CPUs the client buys from Intel); restraints on the distribution channels that could be used to sell AMD-based products (i.e., another form of naked restriction); limitations on the marketing for AMD-based products; restraints imposed on bidding on contracts using AMD-based products; threats to withdraw or divert funding to rivals, or other forms of retaliation for selling AMD-based products; and guarantees of preferred access to Intel CPUs.[30]

Given the complex nature of the restraints, we again use an index approach to quantify their presence and track their temporal evolution.[31] Not all restraints are equally restrictive. Some practices limited AMD adoption at the margin, for example, by capping volumes or conditioning rebates on relative purchases, while firms could still offer AMD-based products. Other practices effectively eliminated AMD from a firm's relevant product space for a sustained period, either by imposing zero caps, excluding AMD from entire product lines, or threatening retaliation in response to AMD adoption. Because these latter practices are likely to have qualitatively different implications for both short-run adoption and longer-run competitive dynamics, we distinguish between *standard* and *extreme* restraints in our empirical analysis. Our measure of standard restraints counts the number of all restraints imposed on a downstream PC firm in a given quarter.

The second measure counts only those restraints we classify as *extreme*, defined as practices that effectively preclude meaningful AMD adoption within the affected product lines. These include: a zero-cap on AMD-based machines; exclusion from certain product lines or delayed launches of specific AMD-based machines; threats or retaliation; and commitments to increase Intel's share of the client's CPU purchases. We provide details on the construction of the index in Appendix A1.

[29] While the public documents contain rich information about Intel's restraints, they are not comprehensive. The text of the lawsuits often states that the listed practices are illustrative rather than exhaustive. The variables we construct to capture the restraints are, therefore, subject to measurement error. That said, the plaintiffs had an incentive to provide detailed information on what they believed to be the important aspects of these restraints, based on detailed internal documents. Moreover, to the extent that measurement error induces attenuation bias, our estimates can be interpreted as lower bounds on the true effects of the restraints. The fact that we nonetheless find substantial effects reinforces our conclusions.

[30] The latter benefit was particularly valuable to downstream customers. As one industry insider noted, "survival practically depends on being able to get allocations of the newest chips, which are always in short supply coming out of the gate." See ZDNet (Dell–AMD coverage). Furthermore, the State of New York case states: "Access to adequate and timely supply of products from Intel was a major concern for all OEMs, whose business was extremely time-sensitive."

[31] We should note that Intel has consistently denied these claims. Determining the legal status of such restraints is a matter for the courts. Our analysis uses these alleged restraints as an empirical context to explore their potential effects, conditional on their existence.

Panel C of Table 3 reports summary statistics. On average, there were 3.32 restraints in place per brand-segment-quarter and 1.30 extreme restraints. The maximum number observed in a single product line-quarter was six restraints, of which up to three were classified as extreme. Figure 3 shows the total count of product line data cells that were affected by different numbers of restraints over the sample period. This count was high early in the sample period, and later declined; no restraints are observed after 2007. This pattern provides sharp temporal variation in the deployment of restraints. The cross-sectional variation at the firm level is also substantial: the case files identify the start and end dates of restraints for each PC firm, and the number and intensity of restraints vary significantly across firms, as well as in their timing.

Figure 3: Evolution of restraints Imposed by Intel

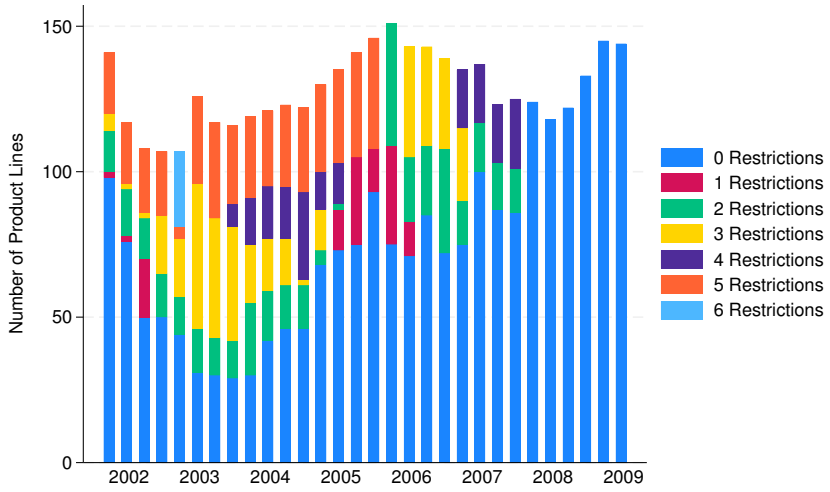

Examples of this variation are abundant. As previously noted, restraints affecting Japanese customers began in 2002 but abruptly ended in 2005, when Intel decided to comply with the Japanese authorities' ruling. Additionally, Intel applied different restraints across customers and market segments. For example, HP faced limits on the number of AMD-based machines it could sell in the business sector. The distribution of these restraints across customers and segments is not random; however, our empirical strategy addresses this endogeneity concern.

The decline in the prevalence of restraints in the later part of the sample, particularly the absence of any restraints after 2007, may be traced to increased antitrust activity. This interpretation aligns with contemporaneous accounts. For example, the SEC case against Dell states:

> "(Dell executives) ...also understood that the Intel...payments were at risk because of the near continuous scrutiny directed at Intel by various competition authorities around the world and, to the degree that...payments were deemed anticompetitive, such payments could abruptly end."

In other words, market participants appear to have believed that the rise in antitrust scrutiny had a chilling effect on Intel's restraints on AMD adoption. The escalation in regulatory pressure thus created a quasi-natural experiment that led to a sharp decline in the deployment of restraints, helping identify their potential impact.

One possible concern is that the observed decline in restraints may partly reflect truncation error: while our product market data extends through mid-2009, some of the legal documents we rely on are dated earlier, raising the possibility that restraints may have persisted after 2007 but are simply not observed in our data. This concern can, however, be rather easily dismissed. The documents underlying our analysis cover the entire relevant period. Although the AMD lawsuit was filed in 2005, other legal proceedings provide documentation extending into later years. For example, the FTC case against Intel, settled in 2010, includes a May 2010 memorandum describing Intel's actions in 2008 and 2009.[32] However, those actions pertain to products such as GPUs, chipsets, and compilers, rather than to CPUs. The document does not specify CPU-related restraints after 2007, strengthening the credibility of our index, which falls to zero beginning in 2008. The European Commission's 2009 decision explicitly covers the period 1997-2007 and draws on evidence collected during site inspections at Intel and several European PC retailers in February 2008. This decision further supports the notion that the restraints were concentrated earlier in the sample period: "Most of the individual abuses concerned 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." The State of New York case, filed in November 2009, also provides corroborating information. Many of these legal proceedings contain overlapping evidence, and none report CPU-related restraints persisting beyond 2007.

We also observe rebates offered by Intel to PC firms via the "Intel Inside" program. The exact amounts paid to Dell are available from the SEC's case *SEC v. Dell Inc.* (July 22, 2010), which covers the period from 2003 to 2006. These payments were substantial and deviated markedly from Intel's public description of the program, which claimed that 3% of the CPU costs would be rebated to PC manufacturers to finance advertising for PC models equipped with Intel CPUs. For the period after 2007, we compute Intel's payments to Dell as 3% of Dell's CPU costs, computed using Gartner sales data and the price dataset described previously. We apply the same 3% methodology to estimate payments to other PC firms throughout the sample period. This variable is defined at the firm level and aggregated over all brands and segments.

Figure 4 provides an overview of these payments. Payments to Dell are displayed on the left axis, while the maximum payments to other PC firms appear on the right axis. Payments to Dell were nearly 100 times the amount the firm should have received under the advertised 3% rebate. For other PC firms, the average per-quarter per-firm payment varied over time between $2.3 and

[32] See FTC (2010) Memorandum.

$5.5 million, with maximum payments ranging between $50 and $200 million.

Figure 4: Evolution of Intel's Payments to PC Firms

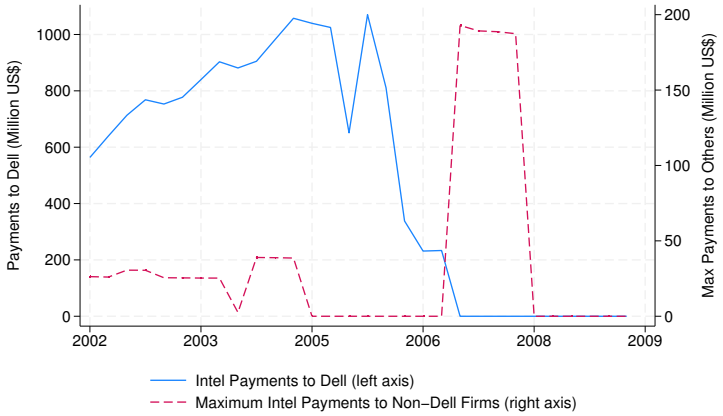

3.4 Data Patterns

A few patterns emerge from the data during the period of heightened antitrust concern (2002:4 – 2007:4). First, the rate of adoption of AMD's technology increased significantly, though it declined again toward the end of the sample period (see Figure 1). Second, AMD's benchmark-per-dollar advantage began to erode as Intel regained a technological edge. Finally, around the same time, Intel's use of exclusionary restraints began to decline (see Figure 3). This joint data variation supports the hypothesis that Intel's restraints had previously slowed the adoption of AMD's technology: as the restraints were gradually phased out, AMD gained market share despite Intel's increasing technological advantage.

Another notable pattern is Intel's advantage in production capacity, which made it an unavoidable trading partner for downstream customers. This capacity advantage enabled Intel to engage in exclusive dealing, an option that was not available to AMD, given its inability to meet 100% of a customer's demand on its own. 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."[33]

Furthermore, as discussed in Section 2, Intel did not need to engage in exclusive deals with all clients as

[33] Source: Intel NYAG Complaint

> "The flip-side of large [PC firms]' importance in legitimising a product is that smaller [PC firms] are not able to do so in the same way. This is explicitly recognised by one such [PC firm], Fujitsu Siemens, which in 2006, although it was the next largest [PC firm] after IBM in terms of market share ... expressed concerns vis-à-vis AMD that it saw itself 'as too small to legitimize AMD for enterprise." [34]

This suggests that contracting externalities may have played a role in the adoption of AMD technology.

In addition, the fact that antitrust scrutiny expanded concurrently with the phasing out of exclusive restraints raises the possibility that increasing regulatory pressure was an important factor in curbing Intel's exclusionary practices. Anecdotal evidence cited above also indicates that downstream customers may have interpreted the heightened antitrust activity as a predictor of future reductions in Intel's exclusionary restraints, thereby increasing their current willingness to adopt AMD technology.

4 Econometric Framework

4.1 Empirical design

Our empirical setting exploits a panel of PC firm–brand–segment combinations (e.g., Acer's Aspire D line for the business segment), defined as product lines and indexed by i; we observe them across quarterly time periods $t = 1, ..., T$. We estimate the following dynamic specification:

$$w_{it} = \rho\, w_{i,t-1} + \beta_p\, p_{it} + \beta_c\, c_t + \beta_x\, x_{it} + \beta_r^{\text{own}}\, r_{it}^{\text{own}} + \beta_r^{\text{riv}}\, r_{it}^{\text{riv}} + \mu_i + \tau_t + \varepsilon_{it}, \tag{2}$$

where the dependent variable w_{it} denotes the fraction of product line i's sales with an AMD chip installed at time t.[35] The term μ_i captures time-invariant unobserved heterogeneity at the PC firm–brand–segment level, τ_t is a quarterly time trend, and ε_{it} is an idiosyncratic error term.[36]

The right-hand-side variables capture the observable drivers of downstream firm adoption decisions at the product line level:

Lagged AMD share, $w_{i,t-1}$: The inclusion of the lagged share term captures the path-dependent dynamics of product i adoption. Firms with prior AMD integration benefit from learning-by-doing and leverage existing production-line investments, thereby reducing the marginal

[34] European Commission 2009 decision (COMP/C-3 /37.990 - Intel)

[35] When c_t and x_{it} are vectors, $\beta_c c_t$ and $\beta_x x_{it}$ denote the corresponding inner products.

[36] We estimate a dynamic panel model rather than a structural model. While this limits our ability to simulate counterfactual scenarios, it allows us to flexibly control for rich product-level heterogeneity without incurring the computational cost associated with a large state space.

cost of subsequent AMD adoption. Switching between Intel and AMD is costly as chips are not "pin compatible," and platform transitions involve sunk costs. Consequently, higher values of $w_{i,t-1}$ are expected to increase the probability of continued AMD utilization.

PC prices, p_{it}: For each product i in quarter t, we compute the sales-weighted average prices across PC models. Any other time-invariant PC characteristics are absorbed by the fixed effect μ_i. AMD technology may be especially attractive in certain lines, for example, in low-margin "value" products using cost-efficient AMD processors.

CPU manufacturing capacity and liquidity, c_t: The vector c_t includes AMD's and Intel's global manufacturing capacity indices and the lagged amount of free cash available to AMD. Capacity reflects long-run FAB investment decisions, such as facility count, wafer size, and process technology, that evolve slowly and do not respond to quarter-level fluctuations in U.S. product line adoption. These indices capture medium-run supply availability and thus proxy for the constraints relevant to PC manufacturing, which operates with minimal inventories and requires reliable within-quarter chip deliveries. Because capacity is global while our dependent variable reflects U.S.-specific adoption, any feedback from w_{it} to capacity would operate only through longer-run profitability and is unlikely to occur within the quarter. We include AMD liquidity because AMD faced tighter financing constraints than Intel, which affected its ability to sustain production and reliably supply customers.

Technology variables, x_{it}: The vector x_{it} contains CPU-level attributes, including the extent of technological progress, measured by the benchmark-per-dollar indices for Intel and AMD, and the age of the CPUs used by the product line, measured by the number of quarters in which the segment-brand-CPU family combination has been available.[37] Benchmark-per-dollar indices for Intel and AMD quantify each supplier's quality-adjusted price competitiveness. The age of CPU families proxies for obsolescence risk.

Own Vertical Restraints, $r_{it}^{\mathbf{own}}$: The vector r_{it}^{own} captures upstream vertical restraints on the focal firm. We include measures capturing the intensity of exclusionary restraints imposed on the relevant downstream firm, as well as counts of extreme restraints (see Section 3), total payments Intel made to the relevant PC maker, and payments to Dell.

Vertical restraints on rivals, $r_{it}^{\mathbf{riv}}$: These indices capture exclusivity imposed on other downstream products or buyers. Whereas r_{it}^{own} records restraints directly imposed on line i, the rival-restraint term r_{it}^{riv} aggregates restraints imposed on market rivals and thus reflects the extent to which AMD is restricted outside the focal line. Higher values of r_{it}^{riv} therefore proxy for a more exclusionary environment facing AMD. In the conceptual framework, such restraints reduce rivals' period-1 AMD shares w_{j1} toward zero, thereby lowering AMD's aggregate sales and, in turn, its subsequent viability (S_2 in our conceptual framework). A lower S_2 weakens AMD's ability

[37] The number of quarters available is counted starting from the first quarter in 2002.

to finance future R&D and capacity expansion investments, making its technology less attractive to all buyers in the next period. The case files support the presence of such a channel.[38] In addition, reputational or informational spillovers may reinforce this effect: widespread adoption of AMD technology by PC makers could "legitimize" it as a credible substitute for Intel's platform, amplifying the impact of rival restraints on overall adoption.[39] In sum, such restraints on AMD adoption by specific PC makers may thus diminish the perceived benefits for others using AMD chips.[40]

The econometric specification operationalizes the theoretical structure in reduced form. Equation (2) is a reduced-form analogue of the two-period model: the vector c_t captures slow-moving supply- and finance-related conditions that affect firms' ability to supply and serve buyers, r_{it}^{own} and r_{it}^{riv} implement the direct exclusion and cross-buyer externality channels, and ρ captures dynamic persistence consistent with convex adjustment costs. By controlling for observable supply-side and financial constraints, the estimated spillover effects reflect dynamic cross-buyer responses that operate beyond capacity, pricing, and liquidity channels.

4.2 Identification

Our primary objective is to identify the causal effect of upstream vertical restraints on downstream input choices at the product line level. A key empirical challenge is the potential endogeneity of these restraints: one would expect them to be correlated with unobserved product line-specific factors (μ_i) and time-varying shocks (ε_{it}). Specifically, Intel may have set these restraints in response to unobserved factors affecting the downstream firms' demand for AMD's chips, such as anticipated shifts in the relative attractiveness of AMD processors or expectations about future demand.

Sources of identifying variation Our setting offers several sources of variation that we exploit for identification. First, we observe sharp variation in the deployment of exclusive restraints. Temporal variation is predominantly driven by the sharp decline in the presence of the restraints

[38] The European Commission 2009 decision 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."

[39] The European Commission 2009 decision states that "Intel itself expressed concern that success for AMD with HP corporate desktops would lead to a 'spillover' possibility of ... products into corporate space 'legitimizing' AMD platforms."

[40] Operationally, in the empirical analysis we focus on cross-firm restraints (restraints imposed on other firms' product lines in the same quarter and segment). This ensures the spillover effect is identified from variation external to the focal firm's own contractual relationship with Intel.

towards the end of the sample period, plausibly driven by increased antitrust scrutiny. In addition, the case files provide firm-specific details on the timing and nature of restraints set by Intel, revealing substantial cross-sectional variation in exclusionary restraints across PC firms.

Second, as noted above, we observe variation in the rate of AMD adoption, both over time and across PC manufacturers. The rise in AMD's market share occurred at different rates among downstream customers, depending on the extent to which their transactions with Intel were subject to vertical restraints (Figure 2).

Third, the richness of the data allows us to control for critical confounding factors, including the evolving technological landscape, as captured by measures of Intel and AMD's CPU performance and value, as well as capacity and liquidity variables. Finally, we explicitly address the endogeneity of exclusionary restraints in our econometric methodology.

Dynamic panel endogeneity and internal instruments Our identification strategy exploits the panel structure of the data to address unobserved heterogeneity in the incentives to adopt AMD technology. Unobserved heterogeneity may arise if some firms are inherently better positioned to benefit from AMD chips due to differences in their demand, product mix, or production flexibility. A standard fixed-effects estimator would difference out time-invariant heterogeneity (μ_i), but it relies on a strict exogeneity assumption for the regressors $z_{it} \equiv \{w_{i,t-1}, p_{it}, c_t, x_{it}, r_{it}^{\text{own}}, r_{it}^{\text{riv}}\}$.[41] This assumption is violated in our context for two reasons.

First, the presence of the lagged dependent variable $w_{i,t-1}$ implies that, after eliminating μ_i, the transformed error term is mechanically correlated with the transformed lagged dependent variable (i.e., Nickell bias). Second, Intel's restraints are likely to respond to past demand shocks: if a positive shock to AMD adoption in brand–segment i at time t induces Intel to impose stricter restraints in period $t+1$, then ε_{it} is correlated with $r_{i,t+1}^{\text{own}}$ and $r_{i,t+1}^{\text{riv}}$, violating strict exogeneity.[42]

We therefore treat $w_{i,t-1}$, p_{it}, r_{it}^{own}, and r_{it}^{riv} as potentially endogenous. The technology and capacity variables (x_{it} and c_t) are treated as predetermined. To address endogeneity, we employ the Arellano–Bover/Blundell–Bond system GMM estimator as our leading specification.

First-differencing Equation (2) eliminates the time-invariant component μ_i:

$$\Delta w_{it} = \rho\, \Delta w_{i,t-1} + \beta_p\, \Delta p_{it} + \beta_c\, \Delta c_t + \beta_x\, \Delta x_{it} + \beta_r^{\text{own}}\, \Delta r_{it}^{\text{own}} + \beta_r^{\text{riv}}\, \Delta r_{it}^{\text{riv}} + \Delta \tau_t + \Delta \varepsilon_{it}. \tag{3}$$

Following Arellano and Bond (1991), we use suitable lags of the levels of the endogenous and predetermined variables ($w_{i,t-3}$) as internal instruments for the corresponding first differences. Arellano and Bover (1995) note that when variables are highly persistent, such lagged levels may be

[41] Strict exogeneity requires $\mathbb{E}[\varepsilon_{it} \mid z_{i1}, \ldots, z_{iT}, \mu_i] = 0$ for all $t = 1, \ldots, T$; see, for example, Chapter 10 of Wooldridge (2002).

[42] See, for example, Chapter 10 of Wooldridge (2002).

weak instruments in the differenced equations. To improve efficiency and alleviate weak-instrument concerns, we adopt the system GMM estimator of Blundell and Bond (1998), which combines the equations in first differences with the equations in levels and augments the instrument set with lagged differences as instruments for the levels.[43] All system GMM specifications are estimated using the two-step estimator with Windmeijer-corrected standard errors (Windmeijer (2005)). To limit instrument proliferation, we restrict all GMM-style instruments to a single lag depth (the third lag). For the lagged dependent variable and PC prices, we additionally collapse the instrument matrix. As a result, the total number of instruments remains always below the number of firm–brand–segment groups.

External instruments We complement these internal (lag-based) instruments with external instruments for selected regressors. For PC prices p_{it}, we use the price of 50-percent ferrosilicon as a cost shifter in the level equation. Silicon is the primary material used in semiconductor chips, so movements in ferrosilicon prices affect manufacturers' marginal costs and, in turn, PC prices.

Most importantly, to strengthen the identification of the effects of vertical restraints, we also rely on legal-exposure measures as external instruments for Intel's restraint variables in the system GMM estimator. These measures, constructed from the timing of antitrust investigations and lawsuits brought against Intel by competition authorities around the world, capture changes in Intel's expected legal risk from pending investigations and formal proceedings.

Specifically, we construct the following lagged legal-exposure variables from case files and contemporaneous regulatory documents: (i) the total number of pending antitrust proceedings against Intel in the previous quarter; (ii) a measure of legal exposure interacting pending proceedings with quarters in which AMD was actively offering competing products, capturing legal pressure when AMD posed a meaningful competitive threat; (iii) a measure removing legal actions centered around Dell, whose contractual relationships with Intel were heavily litigated, to separate general regulatory pressure from customer-specific disputes; and (iv) a measure removing legal actions involving the PC firms (Dell, HP, Toshiba) most directly implicated in exclusive-dealing allegations, isolating legal shocks that shifted Intel's contracting incentives more broadly rather than those tied to particular firms.

Because these legal events are initiated and adjudicated by public agencies and courts, their exact timing is partly shaped by procedural constraints that are plausibly exogenous. Nevertheless, we treat their use as external instruments as a maintained assumption: enforcement intensity may respond to broader market conditions (including the degree of foreclosure or AMD's competitive strength) that could also affect adoption. Our identifying restriction is therefore that, conditional on product line fixed effects, time effects, and detailed controls for technology, capacity, prices,

[43] See Blundell and Bond (2023) for an overview.

and other market-wide factors, quarter-to-quarter variation in lagged legal exposure is orthogonal to residual product line–specific shocks to AMD adoption, and affects adoption primarily through Intel's contracting behavior. In Section 6, we assess the robustness of our results to this assumption by re-estimating the model without legal instruments.

Regarding their relevance, increases in Intel's legal exposure raise the expected cost of maintaining or expanding exclusive-dealing arrangements, thereby weakening Intel's ability or willingness to impose such restraints on PC manufacturers. Conditional on product line fixed effects, time effects, and detailed controls for technology, capacity, and prices, these regulatory and legal milestones should not directly alter AMD's technological attractiveness or production costs. A remaining concern is that enforcement activity could directly affect customers' perceptions of AMD and Intel (e.g., through publicity). We mitigate this concern by (i) using lagged exposure measures, (ii) conditioning on rich measures of AMD and Intel's performance, capacity, and prices that capture the main channels through which changing competitive conditions might influence both adoption and enforcement, and (iii) absorbing market-wide shocks through time effects. Under these assumptions, the primary channel through which legal exposure affects AMD adoption is via Intel's contractual conduct.

This logic is consistent with the empirical vertical restraints literature, which treats legal or regulatory shocks as natural experiments shifting the incidence or strength of vertical contracts while being orthogonal to underlying demand conditions (Lafontaine and Slade, 2008; Babina et al., 2023).[44]

Appendix Table B1.1 shows that the legal instruments are strongly relevant. The Kleibergen–Paap Wald F-statistics exceed the Stock–Yogo critical values commonly used to rule out weak instruments, and first-stage tests reject underidentification with p-values near zero. The lagged legal variables enter the first-stage equations with intuitive signs and substantial explanatory power, even after controlling for product line fixed effects and rich controls.

5 Results

We present the results for our leading specification, the system GMM estimates of Equation (2). Table 6 reports the baseline specifications, including legal payments and the various restraint indices. Table 7 examines contracting externalities across firms, while table 8 examines buyer heterogeneity by re-estimating the model after excluding large, pivotal buyers. Table 9 presents

[44] All lag operators refer to quarterly lags. In the system GMM specifications, we use the third lag of the lagged dependent variable, PC prices, and restraint variables as GMM-style instruments in differences and levels. We collapse the instrument matrix to limit instrument proliferation. The legal-exposure measures enter as additional external instruments lagged by one quarter. Results are robust to using second or fourth lags instead of third lags.

robustness checks based on alternative timing assumptions and placebo specifications.

At the bottom of each table, we report short- and long-run economic effects of the restraint variables. Let $r_{it,k}$ denote the k-th restraint measure included in the specification. Because the model is linear in levels, the short-run marginal effect of $r_{it,k}$ on AMD adoption is given directly by $\hat{\beta}_{r,k}$, the estimated coefficient on $r_{it,k}$.

The long-run effects correspond to a permanent one-unit increase in $r_{it,k}$ from period t onward. Under the autoregressive structure, the cumulative impact on $w_{i,t+j}$ is $\hat{\beta}_{r,k} + \rho\hat{\beta}_{r,k} + \rho^2\hat{\beta}_{r,k} + \cdots$. As $j \to \infty$, this geometric series converges to the long-run multiplier (Hamilton, 1994):

$$\frac{\hat{\beta}_{r,k}}{1-\rho}.$$

Finally, the tables report the Arellano–Bond tests for first-order and second-order serial correlation in the differenced residuals. Absence of second-order correlation (that is, a non-rejected AR(2) test) is required for the validity of the lagged instruments and is empirically verified for each specification.

Direct Effects In Table 6, we report the core estimates of the dynamic model of AMD adoption. The coefficient on lagged AMD share, $w_{i,t-1}$, is large and highly significant across all four models, ranging from 0.87 to 0.96 (values below one in absolute value are consistent with a stable autoregressive process). This substantial persistence reflects the strong state dependence inherent in CPU sourcing: once a product line adopts AMD, switching back is costly due to redesign, validation, and supply-chain adjustments. The high degree of inertia also implies that even modest short-run effects of Intel's restraints translate into quantitatively meaningful long-run impacts via the multiplier $1/(1-\rho)$.

Turning to the PC characteristics, the effect of PC price is consistently negative but generally small and statistically weak, indicating that AMD adoption is marginally lower for higher-priced systems. AMD's benchmark-per-dollar metric is positively associated with AMD adoption (significant at the 10% level in models with restraints). Intel's benchmark-per-dollar metric exerts a negative effect (significant across all models). These patterns confirm that our quality measures capture relevant relative performance considerations in the sourcing decision.

Capacity variables also align with the model: higher AMD capacity increases adoption, whereas Intel's capacity depresses it, consistent with capacity signaling future supply reliability. AMD's lagged free cash flow is generally positive but insignificant.

In column 2, legal payments from Intel ("Intel Inside" rebates) have a negative and highly significant effect. The coefficient implies a short-run effect of –0.09 percentage points (pp) per $1 million of payments and a long-run effect of –0.74 pp. Relative to the average AMD share in the

estimation sample, this corresponds to a long-run reduction of roughly 7%.

In column 3, the restraints index is negative and significant. A one-unit increase reduces AMD share by 0.22 pp in the short run and 2.09 pp in the long run ($\approx$ 20 percent relative to the mean). Given the sample mean of the restraint index of 1.35, the implied long-run impact at the sample mean is approximately –2.8 pp, or a 27% reduction in AMD adoption relative to its average level.

In column 4, extreme restraints have, as expected, larger effects. The short-run marginal impact is –0.40 pp, and the long-run impact is –4.43 pp. Evaluated at the sample mean of the extreme index (0.53), the implied total long-run reduction is about 2.4 pp, or roughly 23% of the average AMD adoption rate.

In sum, the results indicate that Intel's contractual practices meaningfully depress AMD adoption even in the short run and, due to the strong dynamic persistence in sourcing, generate substantial long-run reductions, amounting to 20–30% of the adoption AMD would otherwise achieve on average.

Contracting Externalities We now turn from direct effects of restraints on the treated product line to their spillover effects on other lines. To capture these indirect effects, we introduce two exposure measures. The first is the standard rival-exposure index, which counts the number of restraints imposed on *other firms'* operating in the same segment and quarter (excluding the focal firm itself). The second is the extreme exposure index, which counts only extreme restraints. These measures allow us to test if Intel's agreements with a buyer's competitors affect that buyer's own adoption decisions.

Table 7 reports the results. The estimates confirm the mechanism developed in Section 2: restraints imposed elsewhere in the segment reduce AMD's perceived viability and, in turn, depress AMD adoption on untreated lines.

Columns (1) and (3) show the direct effect of Intel's restraints on the treated product line, already documented in Table 6. Columns (2) and (4) introduce the key tests of contracting externalities. Column (2) provides evidence of cross-firm contracting externalities using the variable r^{riv}. In the short run, a one-unit increase in the cross-firm restraint index reduces AMD adoption by 0.06 percentage points. Standard restraints reduce AMD's share by 0.39 pp per unit in the long run; evaluated at the sample mean of the cross-firm restraint index (10.1), this implies a long-run reduction in AMD adoption of nearly 4 pp.

Column (4) provides analogous evidence for extreme restraints. In the short run, a one-unit increase in the extreme cross-firm restraint index reduces AMD adoption by 0.32 percentage points. Extreme restraints reduce AMD's share by 2.61 pp per unit in the long run, and given that product lines face, on average, nearly 4 extreme cross-firm restraints, this implies a reduction of roughly 10 pp.

It is notable that, once we control for r^{riv}, the coefficient of the direct effects (standard and extreme restraints) becomes larger, suggesting that failing to control for the exclusionary environment can mask the true severity of bilateral restraints.

Taken together, the results offer clear empirical support for contracting externalities. The identified pattern is consistent with the mechanism outlined in Section 2: Intel's use of exclusionary restrictions reduced AMD's perceived viability, lowering adoption even among buyers not directly subject to these contractual provisions.

Buyers' heterogeneity The theoretical framework in Section 2 predicts that contracting externalities operate most clearly through smaller buyers who are not in the position to internalize the effect of their sourcing decisions on AMD's future viability. In contrast, Intel's most restrictive and costly inducements were directed at a small set of large, pivotal buyers whose adoption choices could materially affect AMD's viability. The model therefore implies that, among smaller buyers, cross-buyer contracting externalities should be prominent, while direct exclusionary effects should be weaker.

To assess this prediction, we re-estimate the main specification excluding Dell and HP, thereby isolating the fringe of atomistic product lines. As shown in Table 8, removing large buyers substantially attenuates the estimated direct effects of exclusion: the short-run effect of the own restraint index becomes statistically insignificant. This attenuation reflects a decline in the prevalence of direct exclusion among smaller buyers. In the full sample, the mean value of the own restraint index is 1.35 (0.53 for the extreme measure), whereas in the subsample excluding large buyers, it falls to 0.64 (0.23 for the extreme measure). By contrast, exposure to cross-firm restraints is essentially unchanged across samples. Thus, excluding large buyers disproportionately removes the most intensive forms of direct exclusion while leaving the scope for contracting externalities largely intact.

Intuitively, the cross-firm externality effects remain negative, statistically significant, and economically meaningful. Excluding large buyers increases the magnitude of the cross-firm index effect from -0.06 percentage points in the short run (-0.39 in the long run) to -0.21 percentage points (-1.35 in the long run), and the cross-firm extreme effect remains large, with a short-run impact strengthening from -0.32 to -0.54 percentage points and the long-run effect increasing from -2.61 to -4.10 percentage points. These results support the model's prediction that, especially for smaller buyers, exclusion operates primarily through contracting externalities rather than through direct bilateral restraints.

Table 6: Direct Own-Restraint Effects on AMD Share

	(1) Baseline	(2) Legal Payments	(3) Restraints Index	(4) Extreme Index
Lagged Share	0.9611***	0.8750***	0.8937***	0.9104***
	(0.1090)	(0.0415)	(0.0359)	(0.0383)
PC Price (100$)	-0.0018	-0.0069**	-0.0021	-0.0028
	(0.0055)	(0.0031)	(0.0020)	(0.0024)
CPU Age	0.0000	-0.0004	0.0001	-0.0001
	(0.0012)	(0.0006)	(0.0005)	(0.0006)
AMD Benchmark/Dollar	0.0003	0.0012*	0.0011*	0.0008
	(0.0013)	(0.0007)	(0.0006)	(0.0007)
Intel Benchmark/Dollar	-0.0020*	-0.0021***	-0.0016*	-0.0018***
	(0.0011)	(0.0007)	(0.0009)	(0.0007)
AMD Capacity Index	0.0022	0.0018**	0.0022***	0.0020**
	(0.0022)	(0.0009)	(0.0007)	(0.0008)
Intel Capacity Index	-0.0013*	-0.0006	-0.0009***	-0.0009**
	(0.0008)	(0.0004)	(0.0003)	(0.0004)
AMD's Free Cash $t-1$ (B$)	-0.0031	0.0023	0.0039	0.0044
	(0.0118)	(0.0066)	(0.0058)	(0.0054)
Legal Payments from Intel (M$)		-0.0009***		
		(0.0004)		
Restraints Index			-0.0022***	
			(0.0007)	
Extreme Restraints Index				-0.0040**
				(0.0018)
Observations	2895	2895	2895	2895
Groups	295	295	295	295
AR(1) p-value	0.085	0.280	0.092	0.064
AR(2) p-value	0.195	0.881	0.337	0.238
SR: Own restraint		-0.09	-0.22	-0.40
Std. Err.		0.035	0.069	0.176
LR: Own restraint		-0.74	-2.09	-4.43
Std. Err.		0.350	0.796	2.573

Notes: This table reports system-GMM estimates of the dynamic specification in Equation (2). The dependent variable is the share of a product line's sales using an AMD processor. Each column adds a measure of Intel's contractual restraints: direct legal payments (2), the overall restraints index (3), or the extreme-restraints index (4). All specifications include product-line fixed effects and quarter time trend. Additional controls include lagged AMD share, PC characteristics, performance measures, AMD and Intel capacity indices, and AMD's lagged free cash. AR(1) and AR(2) rows report $p-$values from Arellano–Bond autocorrelation tests. Short-run (SR) effects report the impact of a one-unit increase in the corresponding restraint variable on AMD share, expressed in percentage points. Long-run (LR) effects apply the dynamic multiplier $1/(1-\hat{\rho})$ implied by the estimated coefficient on lagged AMD share; standard errors for LR effects are computed using the delta method. For legal payments, effects are per additional $1 million; for the restraints indices, they are per additional (extreme) restraint. Standard errors in parentheses are Windmeijer-corrected and clustered at the product-line level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Contracting Externalities in AMD Adoption

	(1) Own (Index)	(2) Cross-firm (Index)	(3) Own (Extreme)	(4) Cross-firm (Extreme)
Lagged Share	0.8937***	0.8418***	0.9104***	0.8791***
	(0.0359)	(0.0376)	(0.0383)	(0.0428)
PC Price (100$)	-0.0021	-0.0040**	-0.0028	-0.0029
	(0.0020)	(0.0018)	(0.0024)	(0.0025)
CPU Age	0.0001	0.0003	-0.0001	0.0005
	(0.0005)	(0.0006)	(0.0006)	(0.0007)
AMD Benchmark/Dollar	0.0011*	0.0017**	0.0008	0.0011
	(0.0006)	(0.0007)	(0.0007)	(0.0009)
Intel Benchmark/Dollar	-0.0016*	-0.0018***	-0.0018***	-0.0014**
	(0.0009)	(0.0006)	(0.0007)	(0.0007)
AMD Capacity Index	0.0022***	0.0024***	0.0020**	0.0032***
	(0.0007)	(0.0008)	(0.0008)	(0.0010)
Intel Capacity Index	-0.0009***	-0.0012***	-0.0009**	-0.0018**
	(0.0003)	(0.0004)	(0.0004)	(0.0007)
AMD's Free Cash $t-1$ (B$)	0.0039	0.0099	0.0044	0.0069
	(0.0058)	(0.0075)	(0.0054)	(0.0072)
Restraints Index	-0.0022***	-0.0035***		
	(0.0007)	(0.0008)		
Restraints Index Others		-0.0006*		
		(0.0003)		
Extreme Restraints Index			-0.0040**	-0.0085***
			(0.0018)	(0.0022)
Extreme Index Others				-0.0032**
				(0.0015)
Observations	2895	2895	2895	2895
Groups	295	295	295	295
AR(1) p-value	0.092	0.336	0.064	0.139
AR(2) p-value	0.337	0.985	0.238	0.466
SR: Own restraint	-0.22	-0.35	-0.40	-0.85
Std. Err.	0.069	0.077	0.176	0.217
SR: Rival restraints		-0.06		-0.32
Std. Err.		0.033		0.149
LR: Own restraint	-2.09	-2.23	-4.43	-7.05
Std. Err.	0.796	0.629	2.573	3.262
LR: Rival restraints		-0.39		-2.61
Std. Err.		0.223		1.809

Notes: This table reports system-GMM estimates of the dynamic specification in Equation (2). The dependent variable is the share of a product line's sales using an AMD processor. Columns introduce measures of Intel's contractual restraints at the product-line level: own-line restraints (1 and 3), and cross-firm restraints (2 and 4), each in standard and extreme forms. All specifications include product-line fixed effects and a quarter time trend. Additional controls include lagged AMD share, PC characteristics, performance measures, AMD and Intel capacity indices, and AMD's lagged free cash. AR(1) and AR(2) rows report $p-$values from Arellano–Bond tests for serial correlation. Short-run (SR) effects report the impact of a one-unit increase in the corresponding restraint measure on AMD adoption, expressed in percentage points. Long-run (LR) effects apply the multiplier $1/(1-\hat{\rho})$ implied by the lagged-share coefficient; standard errors for LR effects are computed using the delta method. Standard errors in parentheses are Windmeijer-corrected and clustered at the product-line level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Contracting Externalities Excluding Large Buyers

	(1) Index	(2) Extreme
Lagged Share	0.8440***	0.8693***
	(0.0399)	(0.0418)
PC Price (100$)	-0.0073**	-0.0067*
	(0.0035)	(0.0038)
CPU Age	-0.0007	-0.0007
	(0.0009)	(0.0009)
AMD Benchmark/Dollar	0.0022**	0.0017**
	(0.0009)	(0.0008)
Intel Benchmark/Dollar	-0.0020*	-0.0022**
	(0.0010)	(0.0010)
Restraints Index	-0.0009	
	(0.0011)	
Restraints Index Others	-0.0021***	
	(0.0005)	
Extreme Restraints Index		0.0065
		(0.0048)
Extreme Index Others		-0.0054***
		(0.0013)
Observations	2023	2023
Groups	188	188
AR(1) p-value	0.345	0.220
AR(2) p-value	0.932	0.640
SR: Own restraints	-0.09	0.65
Std. Err.	0.112	0.484
SR: Rival restraints	-0.21	-0.54
Std. Err.	0.046	0.129
LR: Own restraints	-0.57	4.96
Std. Err.	0.712	4.529
LR: Rival restraints	-1.35	-4.10
Std. Err.	0.389	1.492

Notes: This table reports system-GMM estimates of the dynamic specification in Equation (2) estimated on the subsample that excludes large buyers (Big OEMs). The dependent variable is the share of a product line's sales using an AMD processor. Column (1) includes the standard restraints index and the corresponding cross-firm index (restraints imposed on other firms in the same quarter × segment). Column (2) replaces these measures with their extreme counterparts. All specifications include product-line fixed effects and a quarter time trend, and control for PC characteristics and performance measures. AR(1) and AR(2) rows report p-values from Arellano–Bond tests for serial correlation. Short-run (SR) effects report the immediate percentage-point impact of a one-unit increase in the relevant restraint measure. Long-run (LR) effects apply the multiplier $1/(1-\hat{\rho})$ implied by the estimated lagged-share coefficient; standard errors for LR effects are computed using the delta method. Standard errors in parentheses are Windmeijer-corrected and clustered at the product-line level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6 Robustness

Restraints' Timing and Placebo Table 9 reports two placebo tests for our main specification. In columns (1) and (2), we replace current restraints with their one-period leads. Future values of the restraint and extreme-restraint variables have small, statistically insignificant coefficients, indicating that the estimated effects are not driven by anticipation and that the timing of the restraint variables is correctly identified in our data. In column (3), we use a randomly generated "placebo" restraint index that mirrors the empirical range of the true index (0–6). The coefficient on this placebo index is also statistically insignificant and close to zero, suggesting that our findings are not an artifact of spurious correlation or overfitting.

Excluding legal instruments As discussed above, the use of legal-exposure measures as external instruments relies on a maintained exclusion restriction, since enforcement activity may respond to broader market conditions that also affect AMD adoption. Although we mitigate this concern using product line fixed effects, time effects, and detailed controls for technology, capacity, prices, and other market-wide factors, we assess the sensitivity of our results to this assumption. Specifically, we re-estimate the main specifications excluding all legal instruments and relying exclusively on internal lagged instruments within the system-GMM framework.[45]

The results, reported in Appendix Table B1.2, support the main findings in Table 7. While the coefficient on the standard cross-firm restraint index becomes imprecisely estimated in the absence of legal instruments, the effects associated with extreme restraints remain negative, statistically significant, and economically meaningful. In particular, extreme cross-firm restraints reduce AMD adoption by about 0.34 percentage points per unit in the short run (compared to 0.32 in the baseline) and by about 1.38 percentage points in the long run (2.61 in the baseline), where the attenuation in the long-run effect is primarily due to a lower estimate of state dependence (ρ) rather than a change in the immediate impact of the restraints themselves. Thus, the core evidence of contracting externalities does not hinge on the inclusion of legal instruments and is also supported by identification from internal dynamics alone, with the short-run impact of extreme rival restraints remaining virtually identical to the baseline estimate.

Alternative specification Appendix Table B1.3 reports IV estimates in first differences that flexibly absorb product-line and firm-by-year heterogeneity. The objective is to test whether changes in restraints are associated with changes in adoption once all level persistence is removed.[46]

[45] We use deeper lags of the endogenous variables to ensure that the AR(2) test does not reject.

[46] In this robustness check, we estimate a static first-difference model and do not include the first-differenced lagged dependent variable. Including $\Delta w_{i,t-1}$ would reintroduce the dynamic structure and require instrumentation for $\Delta w_{i,t-1}$, whereas the goal here is to evaluate short-run co-movement in differences after eliminating level persistence.

Despite the demanding specification, which removes most cross-sectional and low-frequency variation, the evidence of exclusion remains. Intel's legal payments significantly reduce AMD adoption in differences (column 1), and the cross-firm restraint measures (columns 4–5) continue to exhibit negative and precisely estimated effects. Own-line restraint variables become statistically weaker, as expected. All models pass first-stage diagnostics, confirming that the legal instruments remain informative.

Year Fixed Effects As a robustness check, Appendix Table B1.4 re-estimates the specifications using year fixed effects instead of the quarterly time trend in our preferred specification to capture broad technological cycles. Because year dummies absorb much of the year-to-year persistence, they leave less identifying variation for slowly moving covariates such as AMD and Intel capacity; these variables, therefore, are removed under year controls. The results for all restraint measures remain qualitatively unchanged. Legal payments, own-line restraints, and cross-firm spillovers all retain the expected signs and broadly similar magnitudes. The long-run effects become slightly larger in absolute value, particularly for cross-firm spillovers, strengthening the evidence of exclusionary effects and their externalities.

7 Conclusions

In this study, we empirically document contracting externalities of exclusionary restraints in the semiconductor industry. Our results show that when the relationship between Intel and a downstream client becomes more exclusive, other clients also reduce their purchases from AMD, Intel's main rival. While such contracting externalities have been explored theoretically, they have not, to the best of our knowledge, been documented empirically prior to this study.

Our results suggest that AMD's ability to innovate and offer competitive chips was important for expanding its market share. However, technological leadership was muted by Intel's ability to engage in a broad set of exclusionary vertical restraints, reflecting a fundamental incumbency advantage. These results bear implications for understanding competition in the microprocessor market. More generally, they underscore the need for antitrust scrutiny of exclusionary practices beyond direct contracting partners.

Table 9: Robustness and Placebo Checks

	(1) Lead Restraints Index	(2) Lead Extreme Index	(3) Random Placebo Index
Lagged Share	0.8809***	0.9077***	0.9151***
	(0.0306)	(0.0366)	(0.0311)
PC Price (100$)	0.0010	0.0006	-0.0010
	(0.0029)	(0.0020)	(0.0012)
CPU Age	0.0011	0.0011*	0.0005
	(0.0007)	(0.0006)	(0.0006)
AMD Benchmark/Dollar	0.0016**	0.0008	0.0008*
	(0.0006)	(0.0005)	(0.0005)
Intel Benchmark/Dollar	-0.0009	-0.0007	-0.0013**
	(0.0006)	(0.0005)	(0.0006)
AMD Capacity Index	0.0038***	0.0035***	0.0018*
	(0.0010)	(0.0010)	(0.0009)
Intel Capacity Index	-0.0016***	-0.0014***	-0.0006
	(0.0004)	(0.0004)	(0.0005)
AMD's Free Cash $t-1$ (B$)	0.0036	0.0034	0.0037
	(0.0085)	(0.0072)	(0.0058)
Restraints Index $t+1$	-0.0012		
	(0.0009)		
Extreme Restraints Index $t+1$		-0.0025	
		(0.0015)	
Placebo Restraints Index			0.0006
			(0.0036)
Observations	2638	2638	2895
Groups	278	278	295
AR(1) p-value	0.148	0.078	0.052
AR(2) p-value	0.774	0.490	0.203

Notes: This table reports robustness and placebo tests for the system-GMM specification in Equation (2). The dependent variable is the share of a product line's sales using an AMD processor. Column (1) replaces contemporaneous restraints with their *future* values to test for anticipation effects. Column (2) repeats the same test using the extreme–restraints index. Column (3) uses a randomly generated placebo index. All specifications include product-line fixed effects and a quarterly time trend, and control for lagged AMD share, PC characteristics, performance measures, AMD and Intel capacity indices, and AMD's lagged free cash. AR(1) and AR(2) rows report p-values from Arellano–Bond tests for serial correlation in first differences. Standard errors in parentheses are Windmeijer-corrected and clustered at the product-line level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

References

Areeda, Phillip and Louis Kaplow, "Antitrust Analysis," 1997.

Arellano, Manuel and Olympia Bover, "Another look at the instrumental variable estimation of error-components models," *Journal of Econometrics*, 1995, *68* (1), 29–51.

_ **and Stephen Bond**, "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *The Review of Economic Studies*, April 1991, *58* (2), 277.

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

_ **and Heski Bar-Isaac**, "Raising Retailers' Profits: On Vertical Practices and the Exclusion of Rivals," *American Economic Review*, February 2014, *104* (2), 672–86.

Atalay, Enghin, Ali Hortaçsu, James Roberts, and Chad Syverson, "Network structure of production," *Proc. Natl. Acad. Sci. U. S. A.*, March 2011, *108* (13), 5199–5202.

Ater, Itai, "Vertical Foreclosure Using Exclusivity Clauses: Evidence from Shopping Malls," *Journal of Economics & Management Strategy*, September 2015, *24* (3), 620–642.

Babina, Tania, Simcha Barkai, Jessica Jeffers, Ezra Karger, and Ekaterina Volkova, "Antitrust Enforcement Increases Economic Activity," Technical Report, National Bureau of Economic Research August 2023.

Bernheim, B. Douglas and Michael D. Whinston, "Exclusive Dealing," *Journal of Political Economy*, 1998, *106* (1), 64–103.

Blundell, Richard and Stephen Bond, "Initial conditions and moment restrictions in dynamic panel data models," *Journal of Econometrics*, 1998, *87* (1), 115–143.

_ **and Steve Bond**, "Initial conditions and Blundell–Bond estimators," *Journal of Econometrics*, 2023, *234*, 101–110. Jubilee Issue - Celebrating our Fiftieth Anniversary: 1973–2023.

Brown, Clair and Greg Linden, *Chips and Change: How Crisis Reshapes the Semiconductor Industry*, The MIT Press, August 2009.

Cabral, Luis M. B. and Michael H. Riordan, "The Learning Curve, Market Dominance, and Predatory Pricing," *Econometrica*, September 1994, *62* (5), 1115.

Calzolari, Giacomo and Vincenzo Denicolò, "Competition with Exclusive Contracts and Market-Share Discounts," *American Economic Review*, October 2013, *103* (6), 2384–2411.

_ **and** _ , "Exclusive Contracts and Market Dominance," *American Economic Review*, November 2015, *105* (11), 3321–51.

Chao, Yong, Guofu Tan, and Adam Chi Leung Wong, "All-units discounts as a partial foreclosure device," *The RAND Journal of Economics*, 2018, *49* (1), 155–180.

Conlon, Christopher T. and Julie Holland Mortimer, "Efficiency and Foreclosure Effects of Vertical Rebates: Empirical Evidence," *Journal of Political Economy*, December 2021, *129* (12), 3357–3404.

DeGraba, Patrick, "Naked exclusion by a dominant input supplier: Exclusive contracting and loyalty discounts," *International Journal of Industrial Organization*, September 2013, *31* (5), 516–526.

_ **and John Simpson**, "Loyalty discounts and theories of harm in the Intel investigations," *Journal of Antitrust Enforcement*, 10 2013, *2* (1), 170–202.

Doganoglu, Toker and Julian Wright, "Exclusive dealing with network effects," *International Journal of Industrial Organization*, March 2010, *28* (2), 145–154.

Fumagalli, Chiara and Massimo Motta, "Exclusive Dealing and Entry, when Buyers Compete," *American Economic Review*, May 2006, *96* (3), 785–795.

_ **and** _ , "On the Use of Price-Cost Tests in Loyalty Discounts and Exclusive Dealing Arrangements: Which Implications from Economic Theory Should be Drawn?," *Antitrust Law Journal*, 2017, *81* (2), 537–586.

Gilbert, Richard and David Newbery, "Preemptive Patenting and the Persistence of Monopoly," *American Economic Review*, 1982, *72* (3), 514–526.

Goettler, Ronald L. and Brett R. Gordon, "Does AMD Spur Intel to Innovate More?," *Journal of Political Economy*, 2011, *119* (6), 1141–1200.

Hamilton, James D., *Time Series Analysis*, Princeton, NJ: Princeton University Press, 1994.

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

Katz, Michael L. and Carl Shapiro, "Network Externalities, Competition, and Compatibility," *The American Economic Review*, 1985, *75* (3), 424–440.

Lafontaine, Francine and Margaret Slade, "10 Exclusive Contracts and Vertical Restraints: Empirical Evidence and Public Policy," *Handbook of antitrust economics*, 2008, p. 391.

Nurski, Laura and Frank Verboven, "Exclusive Dealing as a Barrier to Entry? Evidence from Automobiles," *The Review of Economic Studies*, 01 2016, *83* (3), 1156–1188.

Posner, Richard A., *Antitrust Law: An Economic Perspective*, Chicago: University of Chicago Press, 1976.

Rasmusen, Eric B., J. Mark Ramseyer, and Jr. Wiley John S., "Naked Exclusion," *American Economic Review*, 1991, *81* (5), 1137–1145.

Sass, Tim R., "The Competitive Effects of Exclusive Dealing: Evidence from the U.S. Beer Industry," *International Journal of Industrial Organization*, April 2005, *23* (3-4), 203–225.

Segal, I., "Contracting with Externalities," *The Quarterly Journal of Economics*, May 1999, *114* (2), 337–388.

Segal, Ilya R. and Michael D. Whinston, "Naked Exclusion and Buyer Coordination," Working Paper 1780, Harvard University, Institute of Economic Research, Cambridge, MA 1996. Harvard Institute of Economic Research Working Paper No. 1780.

Segal, Ilya R and Michael D Whinston, "Naked Exclusion: Comment," *American Economic Review*, March 2000, *90* (1), 296–309.

Whinston, Michael D, *Lectures on antitrust economics*, The MIT Press, June 2006.

Windmeijer, Frank, "A finite sample correction for the variance of linear efficient two-step GMM estimators," *Journal of Econometrics*, 2005, *126* (1), 25–51.

Wooldridge, Jeffrey M., *Econometric Analysis of Cross Section and Panel Data*, Cambridge, MA: MIT Press, 2002.

A1 Appendix A: Data Details

AMD prices AMD list prices (per 1,000 units) were collected from archived versions of AMD's corporate price lists using the Wayback Machine. These list prices are reported at the CPU *model* level (e.g., "Athlon 64 2800+") with variable frequency. We aggregate them to the CPU *family* level (e.g., "Athlon 64"), matching the Gartner market-share data available at quarterly frequency.

Intel prices Intel list prices were obtained from Intel's historical price catalogues collected from multiple archival sources. For the period 2002–2007, we supplement these with Instat's Rosetta Stone dataset, which provides quarterly CPU core-level prices. A CPU core may appear in multiple CPU families (e.g., "Northwood" appears in both the Pentium 4 and Mobile Celeron lines), and cores within a family may change over time.

We therefore match cores to the Gartner CPU families at the platform group (desktop/mobile), market segment (mainstream/value/ultraportable), family, speed, and quarter levels.[47]

CPU Benchmark Calculation We obtain CPU performance benchmarks from Passmark Software Pty Ltd, which provides standardized performance scores for microprocessors across multiple computing tasks. For each CPU model (e.g., "Athlon 64 2800+"), Passmark assigns a single numerical benchmark score that reflects its relative performance. These scores are updated periodically; in our sample, benchmark values remain constant until 2005 and vary quarterly thereafter as new CPU models are introduced and tested.

We match Passmark benchmarks to CPU models in the Gartner sales data using CPU vendor, model name, and clock speed (where available). For CPU models with multiple speed grades, we assign the benchmark corresponding to the specific speed when available; otherwise, we use the average benchmark across speed grades within the model family.

To construct the benchmark-per-dollar measures used in the analysis, we proceed as follows. For each CPU model m sold in the U.S. home or business market in quarter t, we compute:

$$\text{benchmark-per-dollar}_{mt} = \frac{\text{Passmark benchmark}_m}{\text{CPU price}_{mt} \times \text{CPI adjustment}},$$

where CPU prices are obtained as described in the preceding sections and adjusted to real 2000

[47] For CPUs not matched at the full platform/type/family/speed/quarter level, we proceed with sequential relaxations: (1) drop the market-segment criterion; (2) match using family/marketing name, speed, year, and quarter only; (3) if no time-based match is available, match using platform group, family/marketing name, and speed while ignoring time; and (4) for remaining observations, use the mean price across CPUs with the same marketing name, year, and quarter. The sequence is non-monotonic by design: Step (2) prioritises a time-based match within a family even if platform group is dropped; Step (3) restores platform comparability (laptop vs. desktop vs. server) when no time-based match exists.

dollars using the Consumer Price Index.

We then aggregate at the product line (PC firm × brand × segment × quarter). For each product line, we calculate AMD and Intel benchmark-per-dollar as the sales-weighted average of benchmark-per-dollar over all AMD and Intel CPU models sold in that product line during the quarter. If a product line sells zero AMD or Intel CPUs in a given quarter, the corresponding benchmark-per-dollar is set to zero for that observation.

Exclusionary Restraints Indices We draw on information revealed in legal proceedings to construct our indices of exclusionary restraints. Table A1.1 summarizes the sources. We constructed indicators for the following restraints put into place by Intel for the use of AMD's technology: (caps) caps on the amount sold of AMD-based PCs; (exclusion) exclusion of AMD from certain product lines or delayed launch of specific AMD-based machines; (distribution) restraints on the distribution channels that could be used to sell AMD-based products; (rebates) provision of rebates in exchange for selling certain amounts of Intel-based machines; (marketing) limitations on the marketing PC firms could undertake for AMD-based products; (bidding) restraints imposed on bidding on contracts using AMD-based products; (threats) threats to remove funding, divert funding to rivals, or other retaliation, as a consequence of selling AMD-based PCs; (shares) promises to increase market shares, (supply) and guarantees of preferred supply of Intel CPUs.

We provide a few examples of how we translated text from case files into values for our exclusionary variables for particular firms in particular time periods.

- From the Japan Fair Trade Commission (2005): *"IJKK [Intel Kabushiki Kaisha], since May 2002, has made the five major Japanese OEMs refrain from adopting competitors' CPUs for all or most of the PCs manufactured and sold by them or all of the PCs that belong to specific groups of PCs referred to as 'series', by making commitments to provide the five OEMs with rebates and/or certain funds referred as 'MDF' (Market Development Fund) in order to maximize their MSS, respectively, on condition that (a) the Japanese OEMs make MSS at 100% and refrain from adopting competitors' CPUs. (b) the Japanese OEMs make MSS at 90%, and put the ratio of competitors' CPUs in the volume of CPUs to be incorporated into the PCs manufactured and sold by them down to 10%; or (c) the Japanese OEMs refrain from adopting competitors' CPUs to be incorporated into PCs in more than one series with comparatively large amount of production volume to others."* This text results in the indicator variable for "caps" and "exclusion" to take the value of one for Hitachi, Sony, Fujitsu, Toshiba and NEC starting from 2002 quarter 2.

- From the European Commission (2009): *"In a presentation of 10 January 2003 on Dell rebates, [Intel Executive] outlined a list of objectives to be achieved by Intel in a high-level*

executive meeting with Dell. This includes the following objective: 'Get [Dell executives] to clearly understand our meet-comp process and how it applies to DELL- i.e. if they have AMD in their arsenal they'll have less meet comp exposure hence less meet comp dollars avail to them—even the possibility that meet-comp dollars that we're applied [sic] to DELL go somewhere else...' " This text results in the indicator variable for "threats" to take the value of one for Dell for 2003 quarter 1.

- From Federal Trade Commission (2010): In May 2006, *"Intel worked to limit the scope and pace of Dell's adoption of AMD. Intel offered $120 million in additional funds to fix the profitability of Dell's first quarter and agreed on several other issues. First, Intel agreed to increase its payments to Dell for the second quarter by another $150 million. Second, Intel agreed to allow Dell to announce a limited AMD-based server, but in return, Dell had to agree to deliver two messages: (1) a full-fledged endorsement of Intel's new products and (2) no wiggle room in Dell's statements for anyone to 'construe that there would be additional offerings beyond (an AMD Opteron multi-processor) server.' Dell announced on May 18th that it would introduce a single AMD server in the fall 2006. Intel felt Dell had breached their agreement despite the fact that the reference to AMD was limited to a single line in a quarterly financial call. Art Roehm told Dell that Intel considered 'the deal off.' "* This text results in the indicator variable for "exclusion" and "marketing" to take the value of one for Dell in 2006 quarter 2.

- From United States District Court for State of Delaware (2005a): *"After Gateways 2004 merger with eMachines AMD attempted to revive the relationship it had enjoyed with Gateway until 2001 but experienced extremely limited success. While Gateway built one AMD-powered desktop model at the request of Circuit City AMD remained locked out entirely of Gateways direct internet sales its commercial offerings and its server line. According to Gateway executives their Company has paid high price for even its limited AMD dealings. They claim that Intel has beaten them into guacamole in retaliation."* This text results in the indicator variable for "exclusion," "distribution" and "threats" to take the value of one for Gateway in 2005 quarter 1 and 2.

- From Federal Trade Commission (2010): *"Intel and Toshiba reaffirmed their exclusive arrangement in early 2001 and the arrangement lasted through 2007. Intel gave Toshiba hundreds of millions of dollars, priority CPU supply, engineering support, early product samples, supply line management support, marketing support and other support in return for Toshiba's commitment to purchase CPUs exclusively from Intel during that time. Intel also gave Toshiba money to help it meet financial earnings targets in 2003 as a token of appreciation for Toshiba's loyalty to Intel."* This text results in the indicator variable for "caps",

"rebate," and "supply" to take the value of one for Toshiba starting in 2001 until 2007 quarter 4.

Table A1.1: Legal Sources Used to Construct Exclusion Indices

Case	Reference
Japanese Fair Trade Commission	Japanese Fair Trade Commission (2005), *"The JFTC rendered a recommendation to Intel K.K."* https://www.jftc.go.jp/en/pressreleases/yearly_2005/mar/2005_mar_8_files/2005-Mar-8.pdf
Japanese Fair Trade Commission	Okumura, T. (2006), *"Case Study: Regulation on Exclusive Dealing in Japan."* https://www.slideserve.com/dante/case-study-regulation-on-exclusive-dealing-in-japan-powerpoint-ppt-presentation?
AMD vs. Intel, Delaware	Advanced Micro Devices (2005), *"In the United States District Court for the District of Delaware – Complaint."* https://web.archive.org/web/20140405162439/http://www.amd.com/Documents/AMD-Intel_Full_Complaint.pdf
Korean Fair Trade Commission	Corrective measures against Intel's abuse of market dominance
Korean Fair Trade Commission	Intel's abuse of market dominance in Korea. American Antitrust Institute (2009), English translation of the case.
European Commission	European Commission (2009a), *"COMP/37.990 Intel, Commission Decision."* https://web.archive.org/web/20140809055643/http://ec.europa.eu/competition/sectors/ICT/intel_provisional_decision.pdf
European Commission	European Commission (2009b), *"Summary of Commission Decision COMP/37.990 Intel."* https://publications.europa.eu/en/publication-detail/-/publication/c7a4517a-39cf-4bfd-a3d3-8464d8ee37b7
European Commission	European Commission (2009c), *"Antitrust: Commission publishes decision concerning Intel's abuse of dominant position."*
State of New York	State of New York (2009), *"In the United States District Court for the District of Delaware – Complaint."* https://www.intel.com/pressroom/legal/docs/NY_AG_v._Intel_COMPLAINT.pdf
Federal Trade Commission (US)	Federal Trade Commission (2009), *"In the Matter of Intel Corporation – Complaint,"* Docket No. 9341. https://www.ftc.gov/sites/default/files/documents/cases/091216intelcmpt.pdf
Federal Trade Commission (US)	Federal Trade Commission (2010), *"In the Matter of Intel Corporation – Memorandum in Opposition to Hewlett-Packard Company's Motion to Quash Intel's Subpoena Duces Tecum,"* Docket No. 9341. https://www.ftc.gov/sites/default/files/documents/cases/2010/05/100517intelmemoopphpmoquash.pdf
SEC vs. Dell	Securities and Exchange Commission (2010), *"Securities and Exchange Commission vs. Dell Inc., et al."* United States District Court, District of Columbia. https://www.sec.gov/litigation/complaints/2010/comp21599.pdf

B1 Appendix B: Additional Tables

Table B1.1: First-Stage Regressions and Instrument Strength

	(1) Legal Payments from Intel (M)	(2) Restraints Index	(3) Extreme Restraints Index
Legal Payments from Intel (M$) $t-2$	0.0426		
	(0.0370)		
Legal Payments from Intel (M$) $t-3$	0.1689***		
	(0.0363)		
Pending antitrust cases/investigations against Intel $t-1$	-0.0689	-0.4005***	-0.1208***
	(0.1170)	(0.0412)	(0.0169)
Pending antitrust × AMD-offering line $t-1$	-0.0211	-0.1063***	-0.0242**
	(0.0256)	(0.0224)	(0.0114)
Pending antitrust × non-Dell line $t-1$	0.4224***	0.2463***	0.1740***
	(0.1422)	(0.0699)	(0.0223)
Pending antitrust × non-Dell/HP/Toshiba $t-1$	-0.1799***	0.1603**	-0.0639***
	(0.0460)	(0.0638)	(0.0200)
Restraints Index $t-2$		0.5226***	
		(0.0281)	
Restraints Index $t-3$		-0.1456***	
		(0.0380)	
Extreme Restraints Index $t-2$			0.3928***
			(0.0294)
Extreme Restraints Index $t-3$			-0.1061**
			(0.0430)
First-stage $F-$stat (KP Wald)	20.13	245.96	53.52

Notes: This table reports first-stage regressions corresponding to the baseline system-GMM specifications. Each column regresses one endogenous restraint variable, Intel Inside legal payments, the standard restraints index, or the extreme restraints index, on its own lags and on the set of legal instruments capturing the timing and intensity of antitrust investigations and related legal events. Controls include lagged AMD share, PC characteristics, capacity measures, financial variables, and a linear quarterly trend, with product line fixed effects absorbed. Standard errors in parentheses are Windmeijer-corrected and clustered at the product-line level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B1.2: Contracting Externalities without Legal IVs

	(1) Index (no legal IVs)	(2) Extreme (no legal IVs)
Lagged Share	0.7786***	0.7530***
	(0.0555)	(0.0729)
PC Price (100$)	-0.0067***	-0.0070**
	(0.0022)	(0.0028)
CPU Age	0.0001	0.0001
	(0.0008)	(0.0007)
AMD Benchmark/Dollar	0.0023***	0.0026***
	(0.0008)	(0.0009)
Intel Benchmark/Dollar	-0.0023***	-0.0022***
	(0.0007)	(0.0006)
AMD Capacity Index	0.0021**	0.0028***
	(0.0008)	(0.0010)
Intel Capacity Index	-0.0014**	-0.0017***
	(0.0006)	(0.0006)
AMD's Free Cash $t-1$ (B$)	0.0095	0.0045
	(0.0095)	(0.0074)
Restraints Index	-0.0039***	
	(0.0014)	
Restraints Index Others	-0.0009	
	(0.0005)	
Extreme Restraints Index		-0.0088***
		(0.0028)
Extreme Index Others		-0.0034***
		(0.0013)
Observations	2895	2895
Groups	295	295
AR(1) p-value	0.767	0.511
AR(2) p-value	0.151	0.117
SR: Own restraints	-0.39	-0.88
Std. Err.	0.136	0.284
SR: Rival restraints	-0.09	-0.34
Std. Err.	0.054	0.130
LR: Own restraints	-1.76	-3.58
Std. Err.	0.648	1.197
LR: Rival restraints	-0.40	-1.38
Std. Err.	0.211	0.654

Notes: This table reports system-GMM estimates of Equation (2) re-estimated without using legal-exposure variables as external instruments. The dependent variable is the share of a product line's sales using an AMD processor. Columns report specifications with standard restraint indices (Column 1) and extreme restraint indices (Column 2), together with the corresponding cross-firm measures. Identification relies exclusively on internal lagged instruments. All specifications include product-line fixed effects, a quarter time trend, and controls for PC characteristics, performance measures, capacity indices, and AMD's lagged free cash. AR(1) and AR(2) rows report p-values from Arellano–Bond tests. Short-run (SR) effects report immediate percentage-point impacts. Long-run (LR) effects apply the multiplier $1/(1-\hat{\rho})$. Standard errors in parentheses are Windmeijer-corrected and clustered at the product-line level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B1.3: First-Difference Robustness

	(1) Intel payments	(2) Own + Cross Index	(3) Own + Cross Extreme
ΔCPU Age	0.0018	0.0025	0.0024
	(0.0020)	(0.0021)	(0.0021)
ΔAMD Benchmark/Dollar	0.0034***	0.0032***	0.0030***
	(0.0011)	(0.0011)	(0.0011)
ΔIntel Benchmark/Dollar	-0.0012	-0.0008	-0.0009
	(0.0012)	(0.0012)	(0.0012)
ΔAMD Capacity Index	0.0039***	0.0035**	0.0061***
	(0.0015)	(0.0014)	(0.0016)
ΔIntel Capacity Index	-0.0003	-0.0019**	-0.0028**
	(0.0007)	(0.0009)	(0.0013)
ΔAMD's Free Cash $t-1$ (B$)	0.0013	0.0028	-0.0048
	(0.0102)	(0.0104)	(0.0108)
ΔPC Price (100)	-0.0076**	-0.0074**	-0.0070*
	(0.0036)	(0.0038)	(0.0038)
ΔLegal Payments from Intel (M)	-0.0059**		
	(0.0023)		
ΔRestraints Index		-0.0047	
		(0.0037)	
ΔExtreme Restraints Index			-0.0225*
			(0.0126)
ΔRestraints Index Others		-0.0052***	
		(0.0017)	
ΔExtreme Index Others			-0.0130**
			(0.0052)
First-stage $F-$stat (KP Wald)	587.01	117.05	64.84
Observations	2583	2583	2583

Notes: This table reports first-difference instrumental-variables estimates of the effect of Intel's contractual practices on AMD adoption. The dependent variable is the change in the AMD share of a product line's sales. All specifications difference out product-line fixed effects and include firm by year dummies in levels. Columns instrument the differenced restraint measures using lagged levels and legal variables. Additional controls include the first differences of PC characteristics, performance measures, AMD and Intel capacity indices, and AMD's lagged free cash. Standard errors are clustered at the product-line level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B1.4: Direct and Cross-firm Effects with Year Controls

	(1) Baseline	(2) Legal Payments	(3) Own Index	(4) Own + Cross	(5) Own Extreme	(6) Own + Cross Extreme
Lagged Share	0.7384***	0.9024***	0.9102***	0.8814***	0.9249***	0.9213***
	(0.0911)	(0.0319)	(0.0330)	(0.0318)	(0.0307)	(0.0316)
PC Price (100$)	0.0049	-0.0043**	-0.0016	0.0020	-0.0017	0.0041*
	(0.0031)	(0.0020)	(0.0013)	(0.0023)	(0.0013)	(0.0021)
CPU Age	0.0009	0.0000	0.0002	0.0016***	0.0000	0.0021**
	(0.0014)	(0.0006)	(0.0005)	(0.0006)	(0.0006)	(0.0009)
AMD Benchmark/Dollar	0.0005	0.0008	0.0009	0.0010*	0.0007	0.0004
	(0.0009)	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0005)
Intel Benchmark/Dollar	-0.0015	-0.0018**	-0.0011	-0.0011	-0.0014**	-0.0008
	(0.0013)	(0.0008)	(0.0008)	(0.0008)	(0.0006)	(0.0008)
Legal Payments from Intel (M$)		-0.0008***				
		(0.0003)				
Restraints Index			-0.0019**	-0.0021**		
			(0.0008)	(0.0010)		
Restraints Index Others				-0.0010**		
				(0.0005)		
Extreme Restraints Index					-0.0028*	-0.0071***
					(0.0016)	(0.0027)
Extreme Index Others						-0.0056***
						(0.0018)
Observations	2895	2895	2895	2895	2895	2895
Groups	295	295	295	295	295	295
AR(1) p-value	0.176	0.089	0.065	0.194	0.043	0.141
AR(2) p-value	0.053	0.350	0.245	0.700	0.167	0.531
SR: Own restraint		-0.08	-0.19	-0.21	-0.28	-0.71
Std. Err.		0.030	0.075	0.104	0.156	0.266
SR: Rival restraints				-0.10		-0.56
Std. Err.				0.047		0.182
LR: Own restraint		-0.85	-2.09	-1.75	-3.72	-9.02
Std. Err.		0.393	0.948	0.795	2.444	4.839
LR: Rival restraints				-0.86		-7.06
Std. Err.				0.432		4.390

Notes: This table reports system-GMM estimates of the dynamic specification in Equation (2), augmented with year fixed effects. The dependent variable is the share of a product line's sales using an AMD processor. Columns introduce alternative measures of Intel's contractual restraints at the product-line level: (i) legal payments (Column 2), (ii) own-line restraint indices (Columns 3 and 5), and (iii) own plus cross-firm restraint indices (Columns 4 and 6), each in both standard and extreme forms. All specifications include product-line fixed effects, year dummies, and the full set of PC characteristics, performance measures. We exclude AMD and Intel capacity indices and AMD's lagged free cash. AR(1) and AR(2) rows report p-values from Arellano–Bond tests for serial correlation. Short-run (SR) effects report the immediate percentage-point impact of a one-unit increase in the corresponding restraint measure. Long-run (LR) effects apply the multiplier $1/(1-\hat{\rho})$ implied by the estimated lagged-share coefficient; standard errors for LR effects are computed using the delta method. Standard errors in parentheses are Windmeijer-corrected and clustered at the product-line level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.