Innovation Policy in Product Line Markets: Evidence from GPUs*

Jiahao Chen†

Job Market Paper

August 2026

Abstract

I study the effects of innovation policies in vertically differentiated multi-product markets, using a dynamic oligopoly model with endogenous pricing, product-line adjustment, and R&D spending. I estimate the model in the desktop GPU market, a duopoly between NVIDIA and AMD. In this market, a dollar of the industry's R&D generates about $8.8 of within-market surplus net of that dollar, the firms capturing $3.0 and consumers the rest. Despite the wide wedge, the policy effects are mixed and largely modest. I simulate three stylized interventions: an R&D credit and two product rebates, one at the cutting edge, one on the mass market. Even the best-performing instrument, the R&D credit, returns only $0.82 of surplus per public dollar, because additionality is low: most of its outlay reimburses research the firms would have conducted anyway. The rebate on cutting-edge products generates essentially zero net surplus and sets back the technology frontier. Subsidizing both firms, it ignites a price war at the top of the line, benefiting consumers but undercutting the very reason firms advance the frontier. The mass-market rebate sets off the same war lower down the line but expands adoption enough that R&D survives and the consumer gain outweighs the margin loss, generating $0.53 of surplus per public dollar. Innovation policy in product line markets is thus a more intricate design problem than in markets with simpler structure. Each instrument induces a distinct equilibrium response in prices, qualities, and R&D. The static and dynamic components of that response can offset or reinforce each other, attenuating or even reversing the intended effect.

*I am grateful to Ying Fan, Zach Brown, Uday Rajan, and Jingyuan Wang for their invaluable guidance and support. All errors are my own.

†University of Michigan. Email: cjiahao@umich.edu.

1 Introduction

When the GPU (Graphics Processing Unit) was invented in the mid-1990s, it was a piece of computer hardware sold for one purpose: rendering better graphics in PC games. Three decades later, what was then a circuit board slotted into a gamer's desktop is now the building block of modern data centers, stacked rack upon rack in warehouses, training frontier AI models. This leap did not happen overnight. It accumulated, one chip generation at a time, through the continual R&D of two firms, NVIDIA and AMD. Success stories like this, where private R&D yielded sizable social value, provided the rationale for long-standing public policies designed to incentivize innovation.

Economists have long recognized the case for public support of innovation: a firm cannot capture all the value its research creates, so the social value of research exceeds the private value (Arrow 1962). But two questions remain contested: how large that gap is, and how the government should get involved. On the first question, the leading macro calculation (Jones and Summers 2022) puts the economy-wide social return at roughly \$5 to \$20, a fourfold span.[1] At the micro level, direct empirical evidence is scarce outside the pharmaceutical market (Bryan and Williams 2021). On the second question, the public dollars at stake are not small: governments now commit tens of billions to spur innovation.[2] They do so through two different levers: *push* policies, such as R&D tax credits and NIH grants, lower the cost of doing the research; *pull* policies, such as advance purchase commitments and electric-vehicle subsidies, raise the reward for research whose output reaches the market. The heart of the matter is whether a public dollar buys anything at all or pays for what would have happened anyway, and which lever delivers the best bang for the buck.

I study both questions in vertically differentiated multi-product markets[3] using a dynamic oligopoly model, which I apply to the desktop GPU market. I first estimate the primitives of this market. Based on those estimates, I derive the private and the within-market social returns to the firms' R&D by running a counterfactual that shuts down their R&D programs. I then simulate and compare three stylized policy interventions: a credit to the firms' R&D costs, a rebate on the cutting-edge products, and a rebate on the mass-market products. I report what a public dollar buys under each instrument and demonstrate how the interplay between product-line positioning and dynamic R&D incentives drives their effects apart.

[1] This wide range comes from an accounting exercise disciplined by growth theory, not from an econometric design: it prices the economy's productivity gains against its innovation spending, and the span reflects how broadly that spending is measured, how quickly benefits are assumed to arrive, and which benefits are counted.

[2] The United States CHIPS and Science Act of 2022 alone appropriated \$52 billion for semiconductor manufacturing and R&D. Similar programs are in place in the European Union and across Asia.

[3] I refer to such markets as product line markets throughout the paper.

Estimating the returns to R&D and comparing push against pull on equal terms need three things. The first is a cardinal measure of R&D output. In most empirical settings, the direct output of research is not observed; output is instead read off a proxy (productivity growth, patent counts, citation counts), and the estimated return to R&D is only as good as the proxy.[4] One exception is the pharmaceutical market, where drug approvals and clinical trials give a direct count of the output of R&D, and researchers exploit exogenous variation in expected market size to identify how that output responds.[5] The second requirement is a dollar-valued measure of welfare. Even in the pharmaceutical case, a count of new drugs is not welfare. Translating that output into dollars requires a model of consumers. A policy cost-benefit analysis further requires dividing the resulting surplus by the public dollars spent. The IO literature routinely measures surplus through demand estimation and cost recovery. Yet to my knowledge, no paper measures the surplus a support instrument buys per public dollar, consumer surplus included, while accounting for how the policy endogenously moves firms' R&D spending and the technology frontier.[6] Finally, the comparison needs a technique for constructing the policy counterfactual. The reduced-form literature exploits exogenous variation to identify a policy's effect on measured outcomes, such as R&D spending, new products, or the movement of scientists across regions, but falls short where market power is first-order,[7] and where natural experiments are lacking.[8] The structural approach instead estimates primitives off other variation and re-solves the model to recover the counterfactual outcomes; it carries its own costs, chiefly stronger assumptions on firm conduct and a heavier data appetite.

[4]The productivity branch of this tradition is laid out in Griliches (1979).

[5]Acemoglu and Linn (2004) use demographic-driven variation in pharmaceutical market size to identify its effect on new-drug entry; Finkelstein (2004) uses changes in vaccine policy to identify their effect on investment in clinical trials. This line of work is commonly referred to as the demand-pull literature.

[6]On one side, Petrin (2002) and Goettler and Gordon (2011) measure the surplus from new products and innovation, but neither evaluates a public support instrument, and neither normalizes surplus by public spending. On the other, Peters et al. (2017) and Maican et al. (2023) evaluate how R&D policy moves firm value, productivity, and the return to R&D, but neither reports consumer surplus. Maican et al. (2023) do weigh the benefits of an R&D subsidy against its cost to the government, with R&D responding endogenously, but count only firm profit on the benefit side and note that they "cannot quantify gains to consumers." The demand-side subsidy literature evaluates instruments per public dollar, but for the adoption of a fixed product set: Springel (2021) and De Groote and Verboven (2019) evaluate demand subsidies per public dollar with the technology frontier exogenous to policy and without allowing firms to adjust their R&D. Banares-Sanchez et al. (2026) compare demand, production, and innovation subsidies for Chinese solar manufacturing at a common fiscal cost, with the innovation margin responding endogenously. Their firms are monopolistically competitive and each supplies a single variety, innovation is a discrete, deterministic cost reduction of fixed size, and the benefit priced per subsidy dollar is the change in solar and coal manufacturing revenue, with consumer welfare measured as a representative downstream consumer's exogenous budget divided by the electricity price index.

[7]The reason is SUTVA: subsidizing one firm induces its rival to re-optimize prices and qualities, so any treated-control contrast bundles the instrument's direct effect with the rival's strategic response.

[8]To my knowledge, no historical policy event provides a clean natural experiment for a common-budget comparison of push and pull; even one that did would identify outcome responses, not surplus. The closest work takes other routes: Hsu and Schwartz (2008) value push and pull vaccine-development contracts in a calibrated single-firm real-options model, and Peters et al. (2012) relate solar-PV patenting to push and pull policy measures across countries. For the broader instrument-by-instrument evidence, see the survey of that literature in Goolsbee and Jones (2022).

Because welfare is hard to measure, policy debate often leans heavily on a more tractable statistic: additionality, the amount of private R&D a public dollar crowds in. R&D spending is observed directly, in financial reports, tax records, and R&D surveys, so additionality is easier to measure and far better documented: for the R&D tax credit, the most-studied instrument, leading reviews find that a public dollar induces roughly a dollar of additional private R&D.[9] Governments use this statistic to certify programs, gauge value for money, rank instruments, and drive reform.[10] That additionality is not a welfare measure is recognized in the policy evaluations themselves and in the literature, but how well the indicator tracks ultimate market outcomes remains unclear.

The desktop GPU market provides three advantages for studying these questions. First, the output of R&D is measured directly rather than proxied: the quality of a chip is mainly reflected in its compute speed, which is tracked by standardized software. The same benchmark rates every GPU in the sample, so the cutting edge of 2020 can be compared with what was released in 2005 on a single metric. Second, I observe both firms' product lines over 15 years. Each firm maintains a low-end, a midrange, and a premium tier, and the panel records each tier's price, sales, and quality quarter by quarter, providing the variation needed to estimate demand and recover marginal costs. Third, the market structure held still while the technology raced. The same two firms compete throughout, neither entering nor exiting, and the market was left laissez-faire: no major industrial policy targeted GPUs over the sample. That stillness makes the market an ideal wind tunnel for studying the effects of a hypothetical policy intervention.

I build a model tailored to the institutional details of the GPU market: a dynamic oligopoly with endogenous pricing, quality positioning, and R&D. On the dynamic side, the firms track each other's technology frontier (the highest quality each can currently produce) and choose how much to spend on R&D. When R&D succeeds, the firm's frontier jumps to a higher level, and, through a within-firm cost spillover channel, the advance also lowers the production cost of the tiers below it. On the static side, the firms compete by choosing quality and price across the product line, taking into account both business-stealing and cannibalization effects. The policy comparison re-solves the estimated environment under each instrument.

[9] Behind that figure are regressions of firms' reported R&D spending on the tax price of R&D, identified off policy-driven variation in that price, with the estimated response converted into induced spending per dollar of forgone tax revenue (Hall and Van Reenen 2000; Hall 2022).

[10] HM Treasury, proposing to merge the United Kingdom's two R&D reliefs into a single scheme, cited the RDEC-style credit's up-front certainty as "leading to higher additionality" (HM Treasury 2023). The Dutch cabinet certified the WBSO, the Dutch R&D tax credit, as effective on the strength of an evaluation that leads with additionality (Ministerie van Economische Zaken en Klimaat 2019), even as the evaluators cautioned that bang-for-the-buck analysis is too narrow a lens (Dialogic et al. 2019). The OECD presents the incrementality ratio as "in principle, comparable across countries and across policy instruments" (OECD 2020), the comparative use that the ordering result of Section 6 speaks to directly.

I present two sets of results. The first is a pair of surplus-based private and social returns to firm R&D. The within-market social return is about $8.8 per R&D dollar, against an industry net private return of about $3.0.[11] The wedge captures the surplus that accrues to consumers and that the firms cannot appropriate. Compared with the literature, the $8.8 sits at the lower end of the $5 to $20 that economy-wide growth accounting delivers: that range is built to capture the whole economy's gains, while my measure counts only the surplus created inside this one market. The pattern matches the proxy tradition as well, which finds social returns of roughly 1.5 to 2.7 times private (Bloom et al. 2013). It is worth noting that neither match is built in; my structural approach shares almost no machinery with either literature, reading output directly off the quality index and pricing it into surplus through estimated demand inside an equilibrium model of a single market.

The second result is a comparison of stylized push and pull policies under a common ex ante budget of $5 billion over ten years. On the push side, I apply an R&D credit that lowers the firms' cost of research, in effect a more aggressive and simplified version of the federal R&E tax credit the United States has offered in some form for over four decades.[12] On the pull side, I consider two versions of a product rebate, both implemented as per-unit payments that lower a product's effective marginal cost. The premium rebate is paid at the cutting edge, but on a small share of units sold. The non-premium rebate is paid on the mass-market models, aiming to expand overall adoption; it is the analog of an electric-vehicle purchase subsidy that restricts eligibility to models priced below a cap.[13] The fundamental difference between push and pull is who bears the risk: a push dollar goes out when the research money is spent, whether or not the research pans out; a pull dollar, in contrast, goes out only when the research output makes it to the market.

The policy effects are mixed and largely modest. Despite the large social-private wedge, no intervention converts that wedge into surplus at anything close to the return to the firms' own R&D. The R&D credit is the overall best among the three: it crowds in about $0.18 of private R&D per public dollar, returns about $0.82 of within-market surplus per public dollar, and lifts the end-of-horizon technology frontier by about one percent for every $1 billion of outlay. The premium rebate backfires: it crowds out R&D on net, returns near zero in surplus, and pushes the frontier back by about three-tenths of a percent per $1 billion spent. The non-premium rebate sits

[11] Both returns are reported net of the R&D dollar itself, so zero is the break-even point: a dollar of R&D brings the industry about $4.0 of gross profit and $9.8 of gross surplus. The growth-accounting ratios cited in this section are gross, so the $9.8 is the directly comparable figure.

[12] The enacted credit is small in size and more intricate than a flat rate; for more details, see Hall (2022).

[13] Electric-vehicle purchase subsidies are usually filed under environmental policy, but they are also innovation policy: they reward carmakers for improving the quality of electric vehicles, on which long-run adoption depends.

in between: its net crowd-in is near zero, yet it returns about $0.53 of surplus per public dollar and nudges the frontier up due to its differential impact on the leading firm and the follower. Across all three, the crowd-in falls far short of the roughly dollar-for-dollar figure the tax-credit literature reports.[14] And even the highest return among the three instruments is small next to the return to firm R&D: the thin crowd-in limits what a public dollar can set in motion, and the equilibrium price and quality adjustments bundled with each instrument attenuate the effect further.

At first glance, these numbers look like a contradiction. If a dollar of the firms' R&D generates $8.8 of social value, how can instruments that subsidize R&D and its products return so little, and how can additionality be so low? Part of the answer is in the definition: the wedge measures how the surplus is split, not how strong the firms' incentive to invest is. A dollar of R&D already returns the firms about $4.0 of gross profit. With that incentive, they have pushed research spending to where the marginal return equals the marginal cost, and the $8.8 is the average value of the firms' entire R&D program, not the value of the marginal dollar. Most of a credit's outlay therefore reimburses research the firms would have conducted anyway. The rest of the answer plays out in the product market. A rebate, paid on both firms' products, sets off a price war: fighting for the subsidized buyers, the firms often pass on more than the rebate itself, so consumers pay less and buy more while margins shrink across nearly the whole line, leaving the industry poorer after netting out the transfer. Whether the war also kills innovation depends on where the rebate lands, because advancing the frontier is rewarded in two forms. At the top of the line, a better product commands a premium, so the reward is a margin. Below the top, a frontier advance saves production cost on every low-end and midrange unit sold, so the reward is volume. A price war trades margin for volume. The non-premium rebate therefore leaves the frontier's reward intact and simply widens adoption, while the premium rebate destroys the very margin the race is run for: R&D falls, mostly at the firm holding the technology frontier, hence the frontier retreats, and the static gain is offset by the dynamic response. The R&D credit, in contrast, never reshapes the product market and works on the frontier directly: it lowers the cost of research, both firms advance faster along the path they were already on, and the within-firm cost spillover carries the frontier gain down the whole product line.

It is extremely important to emphasize that the magnitudes of my empirical findings carry

[14]The gap is informative rather than a contradiction. The regressions behind the dollar-for-dollar figure estimate the response of firms' reported R&D to the tax price of R&D, averaged across firms pooled from many industries; what I estimate is the equilibrium response of one concentrated, innovation-driven industry to a credit stacked on the federal credit already in force. Since additionality likely varies across industries, a market of this kind can sit well below the cross-industry average, a form of heterogeneity that the pooled designs are not built to recover.

no external validity. The returns, the welfare signs, and the instrument ranking are specific to this market and the 2005–2020 window. The experiments evaluate stylized instruments, chosen as metaphors for real-world policies; they do not measure the effects of the CHIPS Act, DoD procurement programs, patent boxes, export controls, or any other enacted program. I also focus on within-market effects: the surplus measure excludes downstream spillovers, such as those to the AI industry, as well as distributional, environmental, and general-equilibrium channels. Folding them in would make the returns and the policy effects hard even to define, let alone to measure empirically: what is the return to calculus, or to fire?

Nonetheless, at a conceptual level, some lessons may travel. First, the gap between the social and the private return is large in this industry; but a large gap does not mean public support generates returns of the same size. Second, in product line markets, an instrument's welfare effect runs through the equilibrium response, not the R&D response alone: strategic adjustment of R&D, prices, and quality positioning can attenuate or even reverse a policy's intended effect, and the price war that delivers the consumer gain can destroy the reward for innovation. Third, additionality is an incomplete yardstick for choosing among instruments. It tracks the R&D response and misses the equilibrium response, so it can misread the surplus a public dollar buys: here it understates the R&D credit and scores the two rebates nearly alike, missing the real divergence in what they buy. Lessons of this kind suggest that innovation policy in product line markets is a more intricate design problem than it is in markets where the policy margin is a discrete project's arrival. For example, pharmaceutical markets, where the evidence on innovation policy is most direct, differ in exactly this sense: therapeutic classes segment demand, the patent system awards each compound a period of exclusivity, and the margin the demand-pull designs identify is whether a new drug arrives at all, rather than where in a product line an advance lands.

The rest of the paper is organized as follows. Section 2 provides institutional background on the market and describes the dataset. Section 3 builds the model. Section 4 presents the identification assumptions, the estimation procedure, and the parameter estimates. Section 5 measures the returns to firm R&D. Section 6 runs the instrument comparison and discusses the policy implications of my results. Section 7 concludes.

2 Background and Data

A modern personal computer runs on two kinds of processors. The CPU (Central Processing Unit) is the computer's brain: a general-purpose processor that works through tasks largely one after another and handles most routine computing, from filling in a spreadsheet to sending an email. The GPU is built instead to run many simple calculations at once, and handles a computer's image and video rendering.[15] Nearly every computer ships with integrated graphics, a basic GPU built into the CPU or the motherboard, enough for routine work but not for heavy computation. Buyers who need the heavy computation, gamers above all, add a discrete GPU, a standalone graphics card. I study this market in its pre-AI era: 2005–2020. During this period, the discrete GPU was mainly a consumer product, bought by gamers and by a small group of computer scientists who discovered that machine learning runs on the same parallel arithmetic. The AI boom that followed, most visibly after ChatGPT launched in late 2022, has shifted the GPU market toward business-facing demand concentrated among large enterprise buyers; that era falls outside the scope of this paper.

The desktop consumer discrete-GPU market over my sample period is a stable duopoly: NVIDIA and AMD, whose GPUs jointly account for essentially all units sold.[16] Both firms are so-called fabless designers, meaning they design their chips and contract the fabrication to outside foundries such as TSMC. Each firm maintains a product line that targets buyers across a wide range of willingness to pay. I observe unit sales, prices, and product quality for each firm at three tiers: low-end, midrange, and premium. Products arrive in successive generations that refresh the line, and within a generation each tier holds only a handful of GPUs.

These data form a balanced firm-tier-quarter panel: both firms offer all three tiers in each of the sample's 64 quarters (2005Q1–2020Q4). Prices are inflation-adjusted, and quality is the PassMark G3D benchmark, a published performance test score. I observe each firm's research spending in its quarterly reports and allocate it to the GPU segment in proportion to the segment's share of revenue. A potential buyer's outside option is to add no discrete GPU and rely on the system's built-in graphics. Every desktop that could take a discrete GPU has a CPU, so desktop CPU sales count the potential buyers, and market size is proxied by that count.

Table A.1 summarizes the product line, the market aggregates, and the measurement and sources of each series. Over the sample, benchmark quality at the frontier rose to nearly 100 times its 2005 level, while the premium-tier price roughly doubled. The tiers separate cleanly on both

[15] If this page is being read on a screen, a GPU drew it.

[16] The Radeon GPU line entered AMD through its 2006 acquisition of ATI; I treat the line as one continuing firm.

price and measured quality: the premium tier sets the performance frontier, and the low-end and midrange tiers deliver a fraction of premium benchmark performance at a fraction of the price. The sustained rise at the frontier is why the model in Section 3 treats the quality frontier as moving and R&D-driven rather than as a static dimension. Behind the moving frontier sits research spending: the two firms' combined R&D runs at roughly a tenth of market revenue, $1.1 billion against $12.9 billion in an average year; Figure A.2 plots the by-firm series. Figure A.1 plots the premium-tier price and frontier benchmark series year by year.

The sharp contrast between quality growth and price growth implies that any measure of the return to innovation must account for quality change. If buyers value quality, a hundredfold quality gain sold at about twice the old price leaves most of the gain with them, and that share never appears in the firms' revenue. The demand estimates of Section 4 turn measured quality into willingness to pay, which lets Section 5 price that gain in surplus terms.

3 Model

The model is a dynamic multi-product oligopoly in the Ericson–Pakes tradition (Ericson and Pakes 1995), with endogenous product-line choice, Bertrand–Nash pricing, and endogenous R&D. Time t runs in quarters. There are two firms, NVIDIA and AMD, indexed by f. Each firm maintains three tiers: low-end, midrange, and premium. A product is a firm-tier pair $j = (f, \text{tier})$, six inside products in all; $j = 0$ denotes the outside option.[17]

The dynamic state is a five-dimensional vector, $x_t = (q_{1t}^{\max}, q_{2t}^{\max}, u_{0t}, z_t, M_t)$. The firm frontier $q_{ft}^{\max}$ is the highest quality firm f can currently produce. I refer to the firm with the higher frontier as the period's frontier leader and to the other as the follower; the roles switch endogenously as the state evolves. The difference between the two frontiers is the firms' technology gap, to which Section 6 returns. The outside-option level u_{0t} is the value of not buying a discrete GPU (using integrated graphics). The process-innovation state z_t is the industry-wide technical-progress index that lowers the cost of a given quality over time, reflecting the steady improvement of the foundries' manufacturing technology. Market size M_t counts the potential buyers. Of the five states, only the firm frontiers are endogenous, moved by the firms' R&D; the other three evolve exogenously.

The timing of the game is as follows. Each period opens with firms observing whether last

[17] I do not model entry, exit, or adjustment in the number of tiers. The data come at the three-tier level, so the model is cast at the same level: both firms offer the same three tiers in every sample quarter. Endogenizing the menu size would add substantial computational burden without corresponding insight into the research question.

period's R&D advanced their frontiers. The within-period game then begins: firms simultaneously choose the qualities of their low-end and midrange tiers, and then compete in prices à la Bertrand–Nash. I assume each firm sets the quality of its premium tier at its frontier; Section 3.2.1 specifies the two steps and justifies this assumption. Consumers choose among the six inside products and the outside option; demand is a random-coefficients model following Berry et al. (1995). Firms close the period by simultaneously choosing this period's R&D, which resolves at the next period's opening; this one-period lag is the standard time-to-build assumption of the dynamic-oligopoly literature (Ericson and Pakes 1995; Pakes and McGuire 1994). Each firm maximizes the expected present value of its profits, discounted at the per-quarter factor β. Equilibrium is Markov-perfect. I now specify each block of the model in turn and discuss the underlying assumptions. Table G.1 collects the recurring symbols.

3.1 Demand

Demand is a standard BLP model.[18] Each quarter, every potential buyer picks one of the six inside products or the outside option. Consumer i's utilities are

$$u_{ijt} = \gamma\, q_{jt} + \alpha_i\, p_{jt} + \psi_{f(j)} + \xi_{jt} + \varepsilon_{ijt}, \qquad j \in \mathcal{J}_t, \tag{3.1}$$

$$u_{i0t} = u_{0t} + \varepsilon_{i0t}, \tag{3.2}$$

where $\mathcal{J}_t$ collects the quarter's six inside products, q_{jt} is quality, p_{jt} is price, ξ_{jt} is unobserved quality, u_{0t} is the outside-option level, and ε_{ijt} is an i.i.d. type-I extreme-value taste shock. The brand fixed effect $\psi_{f(j)}$ captures brand loyalty and the software ecosystem specific to each firm. The single random coefficient is on price, $\alpha_i = \bar{\alpha} + \sigma_{\text{price}}\, \nu_i$ with $\nu_i \sim \mathcal{N}(0,1)$. Integrating the choice probabilities over the price coefficient gives product j's market share,

$$s_{jt} = \int \frac{\exp\left(\gamma\, q_{jt} + \alpha_i\, p_{jt} + \psi_{f(j)} + \xi_{jt}\right)}{\exp(u_{0t}) + \sum_{k \in \mathcal{J}_t} \exp\left(\gamma\, q_{kt} + \alpha_i\, p_{kt} + \psi_{f(k)} + \xi_{kt}\right)} \, \mathrm{d}\Phi(\nu_i), \tag{3.3}$$

with Φ the standard-normal distribution; the outside share replaces the numerator with $\exp(u_{0t})$.

[18] The demand system is static and abstracts from forward-looking replacement timing. For the modeling of dynamic consumers, see Goettler and Gordon (2011).

3.2 Supply

3.2.1 Stage Game

Within each quarter, the stage game is a two-step simultaneous-move game: firms first commit the qualities of their interior tiers (the low-end and midrange), then set prices for all products. The game is solved by backward induction. This timing assumption fits the industry: a fabless designer must contract with a foundry and commit in advance which products to manufacture, whereas retail prices can be adjusted flexibly within the quarter.[19] After qualities are committed, the demand and marginal-cost shocks, ξ_{jt} and $\varepsilon_{jt}^{\mathrm{MC}}$, realize, and the pricing step begins.

I describe the two steps in backward order. Pricing is a differentiated Bertrand–Nash game over all six products. The first-order condition for product j of firm f is

$$s_{jt} + \sum_{k \in \mathcal{J}_f} \left(p_{kt} - \mathrm{MC}_{kt}\right) \frac{\partial s_{kt}}{\partial p_{jt}} = 0, \qquad j \in \mathcal{J}_f, \tag{3.4}$$

where $\mathcal{J}_f$ is firm f's menu of three tiers and the marginal cost MC_{kt} is specified at the end of this subsection; solving the six conditions jointly gives the equilibrium pricing map p^*. Each firm internalizes that a price cut at one tier draws demand from its own other two, so each product carries the standard inverse-elasticity markup, widened where substitution runs within the firm. A frontier advance thus monetizes through the menu, not through one cell of it; it earns at the premium tier, shifts same-tier price competition, and cannibalizes the firm's own interior tiers, so the rent from advancing is itself a product-line object.

In the quality step, each firm offers its premium tier at the frontier quality and holds its interior tiers at or below it. The premium tier is identified with the frontier as a modeling simplification.[20] This makes the frontier an observed state, read off the premium benchmark, so the innovation process can be estimated from observed frontier advances rather than from a filtered latent state. The firm chooses its interior qualities to extract surplus from consumers with lower willingness to pay while protecting the rent its premium tier earns.[21]

[19] Quality-before-price timing is standard in the endogenous product choice literature (Fan and Yang 2020), and it leaves qualities predetermined relative to the period's demand and cost shocks. Section 4 justifies this assumption further in my empirical context and discusses the implications for identification and estimation.

[20] Intuitively, the assumption emerges from the firm's dynamic incentive. Maintaining a positive wedge between the true frontier and premium quality is costly: the slack is capability paid for with R&D but never sold. In the quality step, a frontier above every chosen quality leaves the constraint slack, and a slack constraint carries no shadow value, so the spending that holds the wedge open buys nothing at the margin. The firm therefore shrinks the wedge, and in equilibrium the wedge is zero: the commercially relevant frontier and the best product sold move together.

[21] The interior-quality problem parallels the classic quality-discrimination menu of Mussa and Rosen (1978), extended

Formally, each firm chooses its two interior qualities before the shocks realize, foreseeing the price equilibrium that will follow. It solves

$$\max_{\{q_{j't}\}_{j' \in \mathcal{J}_f^{\text{int}}}} \sum_{j \in \mathcal{J}_f} \mathbb{E}\big[\pi_{jt}\big(p^*(q), q\big)\big] \qquad \text{subject to} \quad q_{j't} \le q_{ft}^{\max} \ \text{ for } j' \in \mathcal{J}_f^{\text{int}}, \tag{3.5}$$

where $\mathcal{J}_f^{\text{int}}$ collects firm f's two interior tiers, $\pi_{jt} = M_t\, s_{jt}\,(p_{jt} - \text{MC}_{jt})$ is product j's variable profit, and the expectation is taken over the demand and marginal-cost shocks, neither yet realized when qualities are chosen. Following Fan (2013), the first-order condition for interior product j' is

$$\sum_{j \in \mathcal{J}_f} \left[\frac{\partial\, \mathbb{E}[\pi_{jt}]}{\partial q_{j't}} + \sum_{k \in \mathcal{J}_t} \frac{\partial\, \mathbb{E}[\pi_{jt}]}{\partial p_{kt}} \frac{\partial p_{kt}^*}{\partial q_{j't}} \right] = 0, \qquad j' \in \mathcal{J}_f^{\text{int}}. \tag{3.6}$$

The bracket adds a direct term, the profit response with prices held fixed, to an indirect term that works through the equilibrium price response of every product in the market, obtained from the pricing map p^* by the implicit function theorem. Raising a lower tier's quality wins some demand outright, but it raises the product's marginal cost, cannibalizes the firm's higher-priced tiers, and drags down equilibrium prices, including the rival's. The cannibalization is the own-product force distinctive to product-line competition (Fan and Yang 2020). The cost margin and the cannibalization together are why a firm holds its interior tiers strictly below the frontier.

The marginal cost of product j in quarter t is

$$\text{MC}_{jt} = c_0 + z_t\, q_{jt} + c_3\, q_{jt}^2 + c_4 \left[\frac{q_{jt}}{q_{f(j)t}^{\max}} \right] + c_5 \left[\frac{q_{jt}}{q_{f(j)t}^{\max}} \right]^2 + \varepsilon_{jt}^{\text{MC}}. \tag{3.7}$$

The first three terms price the chip on its own: a constant, a linear quality term, and a quadratic quality term. The coefficient on the linear term is the process-innovation state z_t, equal to $c_1 + c_2\, t$ with $c_2 < 0$, so the cost of delivering a given quality falls over time. The decline captures advances in upstream foundry technology, exogenous to both firms, such as finer process nodes and rising yields as each node matures. The law of motion and long-run floor of z_t are specified in Section 3.2.2. The last two terms depend on $q_{jt}/q_{f(j)t}^{\max}$, the product's quality as a share of its firm's own frontier. This frontier-relative block captures design-technique reuse, a within-firm cost spillover: techniques

to competing multi-product sellers by Johnson and Myatt (2003), with two further departures: the menu is fixed at three tiers, and the top quality is pinned by the technology frontier rather than chosen freely. The parallel is interpretive, not formal: demand is discrete choice with idiosyncratic tastes, so consumers do not sort through self-selection constraints.

developed at the cutting edge carry down the firm's own lineup.[22] It gives the firm an additional incentive to advance its frontier: an advance lowers the cost of its interior tiers. The channel captures the intertwined process innovation and product innovation in this industry. Fixed costs are assumed constant in quality and so drop out of the quality step; Section 4.2 carries the justification and the confirming evidence.

3.2.2 Dynamic Game

Over time, firms compete by choosing R&D spending. After the stage game clears, firm f chooses an R&D intensity $\iota_{ft} \in [0, 1]$, its R&D spending as a share of expected per-period variable profit. Its value function on the state x_t solves

$$V_f(x_t) = \max_{\iota_{ft} \in [0,1]} \mathbb{E}\big[\pi_f^*(x_t)\big]\big(1 - \chi\, \iota_{ft}\big) + \beta\, \mathbb{E}\big[V_f(x_{t+1}) \,\big|\, I_{ft},\, x_t\big], \tag{3.8}$$

where $\mathbb{E}[\pi_f^*(x_t)]$ is expected equilibrium variable profit from the stage game. The firm's R&D spending is $I_{ft} = \iota_{ft}\, \mathbb{E}[\pi_f^*(x_t)]$, the dollars it books as R&D. The economic cost of that spending is $\chi\, I_{ft}$, where the cost scale χ converts each booked dollar into economic cost. A scale above one captures the economic costs that do not surface on the accountant's books: the managerial attention and organizational capacity a program absorbs, and the adjustment costs of assembling and holding a research team.

R&D carries an intrinsic risk of failure. The probability that firm f's frontier advances at the next period's opening, conditional on its current frontier and its accounting spending, is

$$h_f\big(I_{ft},\, q_{ft}^{\max}\big) = \frac{r_{ft}}{1 + r_{ft}}, \qquad r_{ft} = \exp\big(a_f - \varphi\, q_{ft}^{\max}\big)\, g\big(I_{ft}\big), \tag{3.9}$$

where a_f is a firm-specific success-rate parameter, $\varphi > 0$ is the fishing-out parameter, and g is an increasing, concave function of R&D spending with $g(0) = 0$;[23] the saturating map $r/(1 + r)$ turns the innovation rate into a probability, following Pakes and McGuire (1994). The hazard function h_f has three economic properties: zero R&D gives a zero advance probability; additional own R&D raises the probability at a diminishing rate; and, other things fixed, a higher current frontier

[22] The practice is documented on both sides of the duopoly: NVIDIA, for example, described its Pascal generation's GPUs as different configurations of the same architectural building blocks, from the flagship down to the entry level, and AMD wrote in its white paper that its Graphics Core Next architecture was designed to span a family of products.

[23] The baseline exposure is log-concave, $g(I) = \kappa^{-1} \log\big(1 + \kappa I\big)$, with curvature $\kappa = 0.5$; I report the sensitivity of the dynamic estimates to the hazard family and this curvature in Appendix F.2.

lowers the advance probability, the restriction φ imposes. Conditional on success, the size of the advance is an i.i.d. draw from a log-normal jump distribution. The hazard depends only on the firm's own R&D and its own frontier, with no cross-firm knowledge spillover. This exclusion fits the industry: a chip design is among a firm's most closely guarded secrets, and an advance by one firm does not readily travel to the other. The technical progress the two firms do share, the foundries' manufacturing improvement, is already in the model as the process-innovation state z_t.

The optimal intensity balances the cost and reward of research. Differentiating equation 3.8 with respect to ι_{ft} and dividing through by expected profit gives, at an interior optimum,

$$\chi = \beta \frac{\partial h_f}{\partial I_{ft}} \Big(\mathbb{E}\big[V_f(x_{t+1}) \,\big|\, x_t,\ \text{advance}\big] - \mathbb{E}\big[V_f(x_{t+1}) \,\big|\, x_t,\ \text{no advance}\big] \Big), \tag{3.10}$$

where the two expectations condition on whether firm f's own frontier advances and integrate over everything else: the size of the jump and the rival's outcome. The condition prices the firm's trade-off. The left side is the marginal economic cost of an R&D dollar; the right side is its discounted marginal benefit, the extra advance probability the dollar buys times the value of advancing over standing still. Because g is concave and the saturating map compresses high rates, the marginal benefit falls as spending rises, while the value gap does not depend on the spending itself. A higher cost scale therefore meets the falling benefit schedule at lower spending: given the continuation values, optimal intensity is weakly decreasing in the cost scale. That monotonicity lets a single intensity moment identify χ (Section 4.3). The scale is also the parameter the supply counterfactual acts on: a constant R&D tax credit is, in effect, a downward shift in χ.

The model does not restrict whether the industry underinvests or overinvests in R&D. Two forces pull the private reward away from the social value of an advance in opposite directions: the firm does not capture the consumer surplus its advances create, which depresses the private reward below the social value, and the business-stealing effect can lead to overinvestment, as a transfer between the firms creates no new value. The comparison is further complicated by the potential under- or over-provision of quality. Whether the equilibrium innovation rate sits above or below the socially optimal level is hence left to the data.

The remaining three states evolve exogenously and deterministically, each disciplined by its own in-sample dynamics. The outside option improves over time but not without limit: its level

drifts upward along a linear trend toward a long-run ceiling,

$$u_{0,t+1} = \min\{u_{0t} + \Delta_u,\ u_{0,\max}\}, \tag{3.11}$$

where the per-period drift Δ_u is estimated from the recovered demand time effects and the ceiling $u_{0,\max}$ is a closure object beyond the sample. The process-innovation state, the technical-progress index in marginal cost (equation 3.7), falls as manufacturing technology advances on a publicly telegraphed industry roadmap, again not without limit: it drifts downward toward a long-run floor,

$$z_{t+1} = \max\{z_t + c_2,\ z_{\min}\}, \tag{3.12}$$

where $c_2 < 0$ is the in-sample quality-by-time cost trend and the floor $z_{\min}$ is again a closure object.[24]

Market size mean-reverts in logs toward a long-run level,

$$\log M_{t+1} = \log M_\infty + \rho_M(\log M_t - \log M_\infty), \tag{3.13}$$

with the long-run level M_∞ and the persistence $\rho_M \in (0, 1)$ calibrated to the late-sample level and persistence of the observed series.

I focus on Markov-perfect equilibria of the game. Because all state transitions are time-homogeneous, the infinite-horizon equilibrium is stationary: policies and value functions depend on the state alone, not on calendar time. The state also stays effectively bounded. The exogenous laws of motion impose their long-run bounds directly, while the frontiers are bounded in equilibrium: advances from a high frontier become both harder to draw and worth too little additional product-market profit to justify the spending, so R&D tapers off. For the Ericson–Pakes class in general, existence of a pure-strategy Markov-perfect equilibrium is not guaranteed, and equilibria need not be unique (Doraszelski and Pakes 2007; Doraszelski and Satterthwaite 2010). As is standard practice in this literature, I assume the data are generated by a single equilibrium and solve for it. Counterfactuals re-solve the equilibrium under changed primitives.

[24] A closure bound is needed for two reasons: realism, since neither the outside option nor manufacturing technology improves without limit, and tractability, since bounding the exogenous states keeps the stationary dynamic game solvable. In the baseline both trends flatten at their bounds about five years past the estimation sample, well beyond the horizon over which I score the simulations (Section 4.2), so the closure does not bind on the simulated paths. Appendix D.1 makes this precise, and I report sensitivity to later plateau dates in Appendix F.2.

4 Identification and Estimation

There are four blocks of primitives to be estimated: demand, marginal cost, the transition process, and the R&D-cost scale. For each block, I state the identification assumptions and estimation procedure, then report the estimate and model fit. One noticeable feature of the methodology is the neural-network surrogate that approximates the within-period equilibrium (Section 4.3.2). The appendices carry additional detail, such as robustness checks (Appendix F). Readers primarily interested in the policy results can skip ahead to Sections 5 and 6 without loss of continuity.

4.1 Demand

The timing of the stage game makes price endogenous in the demand system, a familiar problem in BLP-style estimation: qualities are committed before the period's demand shocks realize, while prices are set afterward. Price therefore co-moves with the demand unobservable ξ_{jt}.

I address the endogeneity issue with price instruments. The lead instrument is a shift-share built on the semiconductor-wafer producer-price index,[25] interacted with the premium and midrange tier indicators, with the low-end tier omitted as the base. Because the wafer index is a single quarterly time series, it is absorbed by the quarter fixed effects. The identifying variation therefore comes from the tiers' differential exposure to that shock: high-end chips use larger dies, which have lower fabrication yields, so their marginal cost is more responsive to wafer prices than the low end's. The wafer price is set upstream in the global silicon market, orthogonal to the period's GPU demand shocks.

I supplement the wafer cost shifter with two families of instruments: the Gandhi–Houde quadratic differentiation instruments in own and rival quality (Gandhi and Houde 2019), constructed with the routines of Conlon and Gortmaker (2020) as sums of squared quality distances to the firm's other products and to rival products; and a supplementary moment in the square root of own quality. Both families are functions of product characteristics predetermined before the period's demand shock realizes.[26] The demand first stage is strong, with a joint F statistic of about 46 (Appendix B.1).

[25] The silicon wafer is the single most important physical input to fabricating a chip.

[26] Because the demand shocks are serially uncorrelated and qualities are fixed before ξ_{jt} is drawn, no part of the current shock is forecastable when the menu is set, so these functions are orthogonal to ξ_{jt}, unlike the price they instrument for; this is the standard basis for instrumenting price with functions of predetermined characteristics (Eizenberg 2014). The square-root own-quality term carries one further requirement: as a nonlinear function of a characteristic that itself enters utility, it is a valid excluded instrument only under the linear-in-quality index, which keeps it out of the demand equation. It enters as a supplementary moment. Conlon and Gortmaker (2020) use analogous moments.

Table 1 reports the estimates: the parameters in Panel A, with a positive quality coefficient and a negative price coefficient, both estimated precisely; the implied own-price elasticities in Panel B; and the markups those elasticities imply through pricing optimality in Panel C. The elasticities rise across tiers, about 1.91 at the low-end tier, 3.29 at the midrange, and 4.58 at the premium. Because utility is linear in price, the own-price elasticity rises with the price level: on a more expensive product a given percentage price change is a larger dollar amount, hence a larger utility loss, so the premium tier, priced highest, faces the most elastic demand. It does so even though premium buyers are the least price-sensitive: the price-level effect outweighs their lower sensitivity. The estimate of the random coefficient on price (σ_{price}) shows a high level of consumer heterogeneity, which sorts the least price-sensitive, highest-willingness-to-pay buyers toward the premium tier and the most price-sensitive toward the low-end tier or the outside option.

Table B.2 reports the full elasticity matrix. The dollar markup and the Lerner index then measure market power in two ways. In dollars the markup rises across tiers, from about \$80 to \$100 to \$160. The markup ratio falls across tiers, from 0.62 to 0.37 to 0.28, because price climbs faster than the markup. The premium tier thus carries the largest absolute markup but the thinnest rate: it reverses the presumption that luxury goods carry the fattest markup rates, but is a reasonable outcome in this market. These demand-stage estimates mostly stay close to the reported specification across the moment-set, estimation-sample, market-size, and aggregation checks (Appendix F.1, Table F.1).

The outside-option level is backed out from the quarter fixed effects; it rises roughly linearly, by about \$6 per quarter, as integrated graphics improve (Figure D.1), and enters the dynamic model as an evolving exogenous state (Section 4.3.1). Together, the included variables explain about 70 percent of the variation in the mean utilities recovered from the standard Berry inversion procedure (Berry et al. 1995).

4.2 Static Cost Primitives

Marginal costs are recovered from the standard pricing first-order conditions. I then estimate the coefficients of the marginal-cost function by OLS on the recovered series. The stage game's timing means no instrument is needed at this step: qualities are chosen and the frontier is realized before the period's cost shock is drawn,[27] and the cost shocks are serially uncorrelated, so the regressors are predetermined with respect to that shock and least squares is consistent. Table 2 reports the

[27]The timing fits the industry. Committing a quality binds the design to a foundry process and to a supply chain that must be coordinated far in advance. The period's cost shock, by contrast, reflects input prices and fabrication yields, neither of which is settled until well after the design is frozen.

Table 1. Demand estimates

	Estimate (SE)
Panel A. Parameters	
Price	−17.50 (5.21)
Quality	60.73 (11.76)
Random coefficient on price (σ_{price})	5.35 (2.86)
Firm fixed effects	Yes
Quarter fixed effects	Yes
Observations (firm-tier-quarter)	384
Panel B. Own-price elasticities	
Low-end	1.91
Midrange	3.29
Premium	4.58
Panel C. Markups and marginal cost (low-end / midrange / premium)	
Average price	\$132 / \$267 / \$576
Recovered marginal cost	\$51 / \$167 / \$416
Markup	\$81 / \$100 / \$160
Lerner index (markup ratio)	0.615 / 0.373 / 0.278

Notes: Two-step GMM estimates of the demand system of Section 3.1 with a single random coefficient on price; instruments are the tier-interacted wafer cost shifter, the quadratic Gandhi–Houde differentiation instruments, and the square-root own-quality moment. Heteroskedasticity-robust, unclustered standard errors are in parentheses; sample period 2005–2020. Panel B averages the absolute own-price elasticities across markets within each tier; the all-product mean is 3.26. Panel C markups follow from the estimated elasticities through the pricing first-order conditions. Price, marginal cost, and markup are tier means; the Lerner index is tier-mean markup over tier-mean price. Recovered marginal cost is positive for all but eight of 384 observations. Appendix B.3 reports prices, marginal costs, and markups by firm and tier.

Table 2. Marginal-cost estimates

	Estimate (SE)
Constant (c_0)	0.048 (0.017)
Quality (c_1)	1.289 (0.421)
Quality × time, process trend (c_2)	−0.046 (0.009)
Quality squared (c_3)	17.759 (1.387)
Relative quality, $q/q^{\max}$ (c_4)	−0.103 (0.067)
Relative quality squared (c_5)	0.335 (0.054)
R^2	0.87
Observations (firm-tier-quarter)	384

Notes: Ordinary least squares on the marginal-cost series recovered in Table 1, Panel C; standard errors in parentheses. All parameters except the linear relative-quality term (c_4) are significant at the 0.01 level.

coefficients. Marginal cost rises and is convex in own quality, consistent with the falling fabrication yields of the larger dies that higher quality requires. The process-innovation trend enters with a negative slope, so delivering a given quality gets cheaper as upstream manufacturing technology improves. The frontier-relative terms explain a nontrivial share of the variation in marginal cost: dropping them lowers the R^2 by 0.15. The data thus support the within-firm cost spillover of Section 3.2.1. Overall, this rich specification fits the recovered series closely, with an R^2 of 0.87.

The spillover coefficients carry an incentive implication: a given interior quality becomes cheaper to build as the firm's frontier advances. What a firm learns in pushing its frontier is not specific to the chip that first embodies it. Design techniques first developed for cutting-edge models are reused across the firm's own lineup, so an advance lowers the marginal cost of the interior tiers rather than upgrading the flagship alone.[28] The interior tiers therefore give the firm a stake in advancing its frontier.

Manufacturing high-end chips requires expensive capital investment. I nonetheless treat fixed cost as constant in quality, so it drops out of the firm's quality-choice problem. The fixed cost is largely one-time: chips of the same generation share the same basic architecture, and the expense of designing that architecture is incurred before the generation's tier qualities are set, so it is invariant

[28] In the model, the frontier-relative block of equation 3.7 is $c_4\left(q_{jt}/q^{\max}_{f(j)t}\right) + c_5\left(q_{jt}/q^{\max}_{f(j)t}\right)^2$. Differentiating marginal cost with respect to the firm's frontier $q^{\max}_{f(j)t}$, holding q_{jt} fixed, gives $\partial \mathrm{MC}_{jt}/\partial q^{\max}_{f(j)t} = -\left(q_{jt}/(q^{\max}_{f(j)t})^2\right)\left[c_4 + 2c_5\, q_{jt}/q^{\max}_{f(j)t}\right]$, which is negative whenever $q_{jt}/q^{\max}_{f(j)t} > -c_4/(2c_5)$: an advance lowers the marginal cost of every product built above that ratio. At the estimates the threshold is roughly 15 percent, computed from the unrounded coefficient estimates. On the estimation panel every midrange observation lies above it, as do most low-end observations.

to them.[29] The larger platform costs, such as the move to extreme-ultraviolet lithography, are borne further upstream: by the foundry, TSMC, and by its equipment suppliers, chief among them ASML, rather than by the fabless designer. They also arrive on an update cycle of their own. I therefore assume that repositioning an interior tier's quality reuses the same design and the same manufacturing process, so the right-hand side of the quality first-order condition is zero rather than a function of quality, as is more typical in the endogenous product choice literature (Fan 2013).

The data support the assumption. At the observed qualities, the left-hand side of equation 3.6 can be computed from the estimates: the direct profit derivatives come from the demand system and the cost function, and the equilibrium price response from the implicit function theorem applied to the pricing first-order conditions. I evaluate it at every firm-quarter for the two interior tiers, integrating over the demand and cost shocks, following the total-differentiation approach of Fan (2013). Interpreted as an implied marginal fixed cost of quality, the object is large and swings widely in the early years; it falls by nearly two orders of magnitude across the sample and sits near zero over the post-2010 window (Figure B.1). I read the early years as an emerging industry not yet settled into stable quality-setting, and 2011 onward as the stationary regime the dynamic model describes. The static primitives, demand and marginal cost, are therefore recovered on the full sample, while the R&D-cost parameter and the counterfactuals, which both rely on the quality first-order condition, use the post-2010 window, 2011Q1–2020Q4. That window also sets the paper's scoring horizon: every counterfactual object is computed over ten years, using the 2011Q1 values as the initial state.

4.3 Dynamic Estimation

The dynamic block has two kinds of primitives: the laws of motion for the state variables, and the R&D-cost scale χ, the economic cost of the research activity. As is standard in dynamic-games estimation, I estimate the transition parameters offline, outside the game: those parameters are identified off the functional-form assumptions and the observed state transitions, with no equilibrium computation needed. The cost scale is identified by the monotonicity argument of Section 3.2.2: the R&D-cost margin enters the firm's payoff linearly in intensity (equation 3.8), the marginal benefit of intensity falls as intensity rises, through the concave hazard and the saturating map, and the value gap between advancing and standing still does not depend on the spending itself. A higher cost scale therefore meets the falling benefit schedule at a lower intensity (equation 3.10),

[29] The industry calls this one-time cost non-recurring engineering (NRE).

so a higher χ implies a lower mean R&D intensity. I estimate the scale by full solution: I solve the dynamic game, compute the model's mean R&D intensity, and match it to the data. With one moment for one parameter, the estimator is exactly identified, and the monotonicity makes the match unique: exactly one value of the scale is consistent with the observed intensity. With two firms and one market, full solution is the only practical way to estimate the dynamic game; the data are not rich enough for two-step methods such as Bajari et al. (2007), whose first stage recovers the firms' policy functions from observed play, a step for which the short single-market time series offers too little support.

4.3.1 The Transition Process: Frontier Dynamics and Exogenous States

The innovation hazard (equation 3.9) is estimated by maximum likelihood on the innovation-success panel: each firm's record of frontier advances over the sample period, paired with the R&D spending and the current frontier level. The estimation recovers the two firm-specific success-rate parameters and the common fishing-out parameter. The conditional frontier-jump distribution has its log-mean and log-variance fixed at their sample moments. The hazard is estimated using the full history rather than the shorter post-2010 window.[30]

The three exogenous states are estimated as deterministic processes, each disciplined by its own in-sample dynamics. The outside-option level is recovered by inverting the quarter fixed effects (Section 4.1) and converted to dollars with the price coefficient. Because the benchmark measures product quality directly, the quarter effects load onto the outside option rather than onto unmeasured quality drift in the inside goods. The fitted trend runs from roughly −\$120 in 2005 to roughly +\$265 in 2020. A noticeable feature of the recovered series is that it shows a dip in 2017–2018 and then returns quickly to its overall path (Figure D.1). This dip likely picks up the crypto-mining boom of those quarters: mining then ran on GPUs, and the surge in prices and volumes is rationalized by the demand system as an outside option that is temporarily particularly bad. I exclude those quarters when fitting the trend.

The process-innovation state needs no separate estimation: z_t is the time trend already estimated in the marginal-cost function, $z_t = c_1 + c_2\, t$ (Table 2). Market size over my sample period shows the pattern of a maturing technology: desktop CPU sales fall through the middle years and level off toward the end as adoption saturates. I therefore model market size as mean-reverting in logs

[30] The post-2010 window of Section 4.2 guards the quality first-order condition, whose implied fixed cost settles only after 2010; the transition estimates never use that condition, so the full history is available to them.

toward a long-run level, calibrated on the post-2010 window; like the other states, it enters the dynamic model as a deterministic fitted path rather than as stochastic draws.

4.3.2 The Static-Equilibrium Surrogate

The full-solution estimator re-solves the entire dynamic game at every trial value of χ. The quality-then-price stage game inside each period compounds the burden: every update of the value functions needs that game's equilibrium at a large number of states, and solving it exactly at each one is computationally prohibitive. An exact solve is a nested pair of fixed points: a quality equilibrium computed by best-response iteration, each step of which re-solves the six-product pricing equilibrium at every draw of the period's demand and cost shocks.

I approximate the stage game with two nested neural networks, a pricing surrogate and a quality surrogate. The pricing surrogate maps the six product qualities, the outside-option level, and the process-innovation index to the shock-integrated outcomes of the pricing equilibrium: expected prices, shares, variable profits, and consumer surplus. The quality surrogate maps the two frontiers and the same two exogenous states to the equilibrium interior qualities. Market size enters neither network: it scales profits and consumer surplus without moving equilibrium prices or qualities, so it multiplies the surrogate outputs directly. Training mirrors the nesting, in two passes: I solve the pricing game exactly on a large sample of quality menus and states and fit the pricing surrogate to those solutions; then, with the fitted pricing surrogate as the inner map, I solve the quality game across sampled states and fit the quality surrogate to the result. For each surrogate I train 20 independent networks, identical except for the random seed, and use their equal-weight ensemble as the final model.[31]

The training draws span the historical range and a wider region beyond it, and include the observed post-2010 states and small perturbations around each. For the policy exercises of Section 6 that change the stage game, I re-simulate and re-fit the surrogates, appending to the training set the states visited by the estimated equilibrium's simulated paths, so each re-fit trains densely where the policy paths travel. Appendix C details the construction and training. Approximating equilibrium objects and decision rules with neural networks is a standard tool in high-dimensional dynamic programming and reinforcement learning, but remains less common in empirical dynamic games.[32]

[31] Unlike a regression, training a neural network carries intrinsic randomness; averaging independent runs is insurance against an unrepresentative draw.

[32] Bodéré (2026) similarly uses a neural network to approximate stage-game equilibrium profits and to construct basis functions for the value function in a full-solution empirical dynamic game.

The surrogates fit well. Held-out R^2 is 0.99 for the pricing and 0.98 for the quality game, and the networks improve on polynomial and LASSO benchmarks (Table C.1), with the largest gains on shares and profits, the objects the surplus measures are built from. I choose the training bounds so that state visitation in the dynamic estimation and the counterfactuals is confined to the training region, with no extrapolation beyond the bounds.

4.3.3 Solving the Game and Estimating the Cost Scale

At each trial value of the R&D-cost scale, I solve the dynamic game by iterative best response: I initialize both firms' policies at zero investment, compute each firm's value function and R&D policy as a best response to the rival's, and iterate the pair to a fixed point. The dynamic game has a five-dimensional continuous state space and a continuous control, which rules out exact solution; I approximate the value function on a Chebyshev basis.[33] At each state the solver visits, the surrogates supply the period payoff, the expected equilibrium variable profit $\mathbb{E}[\pi_f^*(x_t)]$. Appendix D.2 details the solution method and the simulation protocol.

The cost scale is estimated by the method of simulated moments: with the game solved at a trial value, I simulate the model over the post-2010 window, 2011Q1–2020Q4, compute its mean R&D intensity, pooled across firms and quarters, and adjust the scale until the simulated mean matches the sample value. The observed intensity mirrors the model's definition, each firm's R&D spending divided by the expected variable profit that the estimated static primitives imply at the observed states, and averages about 20 percent. Because the parameter is exactly identified, I match this moment essentially exactly. Appendix D.3 states the estimator formally.

4.3.4 Estimation Results and Model Fit

Table D.1 reports all the parameters of the dynamic block. Panel A carries the main parameter, and the only one estimated inside the dynamic game: the R&D-cost scale χ is 4.62. The scale reads as an economic cost of about 4.6 times the firm's accounting R&D cost. The excess over the accounting outlay is a familiar object in the user-cost-of-R&D literature, where it appears as the marginal adjustment cost of R&D (Hall and Lerner 2010), a term posited but rarely measured directly. I read the parameter estimate as a self-consistent, reduced-form measure of that user cost, implied by the observed intensity of the firms. It should not be read as a direct measure of the cost of research. Appendix F.2 reports its sensitivity (Table F.2): across the hazard-specification variants estimated

[33] The parametric value-function approximation follows Sweeting (2013).

there, the scale ranges from about 4.4 to 4.6, the window and closure-horizon comparisons reported alongside them run from about 3.0 to 4.9, and the reading of an economic cost well above the accounting cost holds in every case. Figure D.2 plots the estimation objective against χ: it falls to a single interior minimum at the estimate.

Panels B and C collect the offline blocks of Section 4.3.1: Panel B the frontier transition (the hazard's by-firm success-rate parameters, the common fishing-out parameter, and the frontier-jump distribution matched to its sample log-moments); Panel C the fitted exogenous-state processes (the outside-option trend, the process-innovation state, and the market-size law of motion). Across both panels the estimates take economically sensible signs and magnitudes. Panel D carries the discount factor, fixed at $\beta = 0.99$ per quarter (about 4 percent annually).

Despite having only one free dynamic parameter, matched to a single pooled moment (the mean R&D intensity), the model fits the data fairly well. I simulate the estimated equilibrium and plot the 10th-to-90th percentile bands of its by-firm paths of market share, frontier quality, and R&D intensity, layering the observed paths on top (Figure D.3). None of these paths is targeted in the estimation, yet the observed paths sit largely within the bands. The exceptions concentrate in the final years, when both firms' R&D intensities climb above the bands, most persistently NVIDIA's, whose frontier runs above as well.[34] With the primitives in hand, I simulate the market under counterfactual restrictions and policy schemes, and compare the expected outcomes with the estimated equilibrium.

5 Returns to Firm R&D

I open the counterfactual analysis with a measurement: how much surplus a dollar of a firm's own R&D generates, how much profit it creates, and the wedge between the two.

The return to a firm's R&D is measured against the world in which that R&D does not happen. I set each firm's R&D to zero over the ten-year horizon (2011–2020) and re-simulate the equilibrium path; the expected difference from the estimated equilibrium is what the firm's R&D produces, for the firm in operating profit and for the market in within-market surplus. The shutdown runs under two conventions. In the by-firm experiments, the rival keeps the R&D policy rule it follows in the estimated equilibrium (its realized spending still changes because the state path differs), so

[34] One reading of NVIDIA's late-sample surge is R&D aimed at demand outside this market: if the firm was already positioning for the boom in AI accelerators, which arrived after the sample ends, an R&D program sized to that opportunity would rationally exceed what desktop-GPU profits alone justify, and the model prices only the latter.

each measures the effect of a unilateral deviation from the equilibrium strategy to zero R&D. In the full-industry experiment, both firms' R&D shuts down at once, so neither firm falls behind and the competitive pressure to innovate vanishes. In both conventions, the exogenous state variables evolve as in the estimated environment. Throughout, I focus on within-market surplus as the conservative policy-evaluation object: it captures the benefits to discrete-GPU consumers, and it excludes the downstream and user spillovers (to AI, cloud computing, autonomous vehicles, national security) and the diffuse channels (health, international diffusion) whose scope has no settled convention.

I report each return as a ratio of present values: the net private return is the change in the firm's own operating profit per dollar of its own R&D, net of that dollar; the social return is the change in within-market surplus[35] per that same dollar; and the social-private wedge is the difference between the two. Formally, for the shutdown of firm f's R&D, write Δx_{t+h} for the equilibrium-minus-shutdown change in a value x in quarter $t+h$ of the horizon. With sums discounted at the per-quarter factor β over $h = 0, \ldots, 39$, the two returns are

$$\text{net private return} = \frac{\sum_{h=0}^{39} \beta^h \, \Delta\Pi_{f,t+h}}{\sum_{h=0}^{39} \beta^h \, I_{f,t+h}} - 1, \qquad \text{social return} = \frac{\sum_{h=0}^{39} \beta^h \, \Delta W_{t+h}}{\sum_{h=0}^{39} \beta^h \, I_{f,t+h}},$$

where the minus one nets out the R&D dollar itself, because operating profit is measured before R&D spending; the social return carries none, because the outlay already sits inside $W = \text{CS} + \Pi - I$. Each denominator is the present value of the firm's own R&D program: its spending $I_{f,t+h}$ along the estimated-equilibrium path (the shutdown path spends nothing). The full-industry variant applies the same construction with both firms shut down and the industry's discounted R&D in the denominator. These returns are program-level averages over the firm's entire R&D, reported as dollar-denominated present values: the average return to the whole program, not the marginal effect of an additional R&D dollar.[36]

Table 3 reports the net private and social returns to firm R&D, and the wedge between them, by firm and for the full industry. The social return exceeds the net private return in every row. A dollar of AMD's R&D returns \$3.49 to the firm and \$6.07 to the market, a wedge of \$2.58; a dollar

[35] Within-market surplus $W = \text{CS} + \Pi - I$ deducts the firms' accounting R&D spending I from consumer surplus plus operating profit measured before R&D. This convention treats R&D outlays as a resource cost rather than a transfer. How much of the full user cost χ (Section 4.3.4) is real resource cost rather than rents in the labor market is not separately identified. Under the polar alternative that treats the full excess over the accounting outlay as real resource cost, both returns fall in level, the industry wedge is exactly unchanged, and the by-firm wedges widen slightly.

[36] The marginal return is the harder object to define and interpret in a dynamic framework: the effect of an additional R&D dollar depends on when the increment happens and how spending adjusts around it, so each timing convention defines a different experiment. The program average requires no such convention.

Table 3. Returns to firm R&D

	Present-value dollars per dollar of own R&D		
R&D shut down	Net private return	Social return	Social-private wedge
AMD	3.49	6.07	2.58
NVIDIA	3.71	6.97	3.27
Full industry	3.00	8.82	5.82

Notes: The net private return is the change in the firm's operating profit, net of the R&D dollar; the social return is the change in within-market surplus, downstream spillovers excluded; the social-private wedge is the social minus the net private return. Each by-firm row shuts down the named firm's R&D while the rival keeps the policy rule it follows in the estimated equilibrium; the full-industry row shuts down both firms' R&D at once and divides by the industry's discounted R&D. Present values run over the ten-year horizon 2011–2020, 40 quarters, discounted at the quarterly factor of Section 4.3.4. For scale, the full-industry shutdown removes about $8 billion of discounted R&D and with it about $67 billion of within-market surplus.

of NVIDIA's returns $3.71 and $6.97, a wedge of $3.27.[37] The full-industry row carries the lowest net private return in the table, $3.00, the highest social return, $8.82, and the largest wedge, $5.82 per R&D dollar. In present-value terms, the full-industry shutdown removes roughly $8 billion of discounted R&D and, with it, roughly $67 billion of within-market surplus, about half the revenue the market realized over the decade.

The full-industry net private return, about $3.0 per R&D dollar, lies about $0.5 to $0.7 below the by-firm private values. The gap captures business stealing: in the by-firm rows, part of what a firm's R&D earns is the shares and margins defended against a rival that keeps innovating. When both firms stop at once, neither falls behind and this component nets out of the industry aggregate. Conceptually, the effect is similar to the rent dissipation of a patent race (Loury 1979). The social side runs the other way, and more strongly: the full-industry social return, $8.82, exceeds the returns in the by-firm rows by roughly $1.8 to $2.7 per R&D dollar. The simultaneous shutdown removes about $17.1 billion more within-market surplus than the two unilateral shutdowns combined. More broadly, when market power is first-order, the return to R&D is an equilibrium object: its definition requires an assumption about rival behavior, and firm-level measures can mismeasure the return to the industry's R&D.

It is informative to put my numbers beside the estimates from the literature. To the best of my

[37]The by-firm rows hold the rival to its estimated-equilibrium policy rule. In a companion exercise, the rival instead re-optimizes its R&D policy against the permanent shutdown. The private-return rows change by less than $0.03; the social-return responses are larger and asymmetric: AMD's rises by about $0.11, while NVIDIA's falls by about $0.33 (Appendix E.2).

knowledge, there is no structural paper that delivers surplus-based private and social returns to R&D from an estimated market equilibrium, so no like-for-like benchmark exists. But the broader economics of innovation literature offers magnitudes close enough to compare. The first comparison is the level of the social return. Jones and Summers (2022) put the social value of innovation investment, economy-wide and inclusive of cross-market spillover channels, at roughly \$5 to \$20 per dollar invested; against the narrower base of R&D spending alone, their estimate is about \$13. The \$8.82 here is within-market surplus per dollar of the industry's own R&D, net of that dollar (gross, \$9.82), and it omits every channel beyond this one market. Provided those omitted spillovers are positive, it is a lower bound on the economy-wide value of the same R&D dollar. Consistent with that reading, it sits comfortably at the lower end of their band. The second comparison is the ratio of social to private returns. Bloom et al. (2013) use a firm panel to separate knowledge spillovers to technological neighbors from business stealing among product-market rivals, converting the estimates into private and social returns: their gross social returns run roughly 1.5 times private under IV and 2.7 under OLS. Their wedge prices the cross-firm knowledge-spillover channel, which my model excludes; mine comes instead from inappropriable within-market surplus and product-market interaction. Jones and Williams (1998) likewise find social returns well above private.

Overall, by an alternative route, my estimates line up with the literature's prior: the return to innovation is high, and the gap between private and social returns is wide. A gap this wide puts a question on the table: can public support convert any of it into surplus? I take up this question in Section 6 by simulating alternative support instruments matched on a common ex ante budget. The comparison shows how the market's product-line structure and the firms' dynamic interaction generate equilibrium effects beyond the public dollar's direct impact.

It is important to highlight that while a wedge is the canonical economic rationale for public support, in reality it is seldom the only motivation for large-scale intervention, and often not the main one. Governments also press industrial policy on grounds of geopolitical tension, supply-chain resilience, and externalities. Sometimes technological leadership itself is the policy target. My analysis remains positive and sits beneath these motivations: it measures what the public dollar buys, an input to any of these cases rather than a verdict on them. I discuss the framework's broader policy implications, beyond my numerical results, in Section 6.4.

6 Effects of Public Support

I now turn to the analysis of public support for innovation. Governments have recently begun committing large packages of public money to such support; the CHIPS and Science Act of 2022, for example, provides $52 billion for the semiconductor industry, including $39 billion in manufacturing incentives and $13 billion for R&D and workforce development.[38] To study what such an intervention buys, I simulate three stylized support instruments in the estimated GPU market, each committing about $5 billion over the ten-year window 2011–2020,[39] and compare market outcomes. I first define each policy counterfactual and the conventions under which outcomes are reported, then measure the returns in welfare and in frontier advancement per public dollar. Section 6.4 draws out the policy implications.

6.1 Policy Design and Reporting

The three instruments are chosen to offer a stylized push-versus-pull comparison: on the push side, an aggressive constant-rate R&D tax credit that lowers the effective cost of research; on the pull side, two alternatives that subsidize demand for a targeted part of the product line. One pays a rebate per unit sold on the premium tier of both firms; the other pays a rebate per unit sold on the non-premium tiers (the low-end and midrange tiers of both firms). The premium rebate is the model analog of demand pull concentrated on cutting-edge products, such as public procurement of high-performance compute; the non-premium rebate is the analog of adoption subsidies that cover products priced below a cutoff, such as electric-vehicle purchase credits. Each instrument is evaluated by re-solving the estimated environment under it and comparing outcomes to the estimated equilibrium. The design holds the ex ante budget fixed and varies only its form.

Formally, the R&D credit is a reduction in the R&D-cost scale χ: under a credit at rate τ, a booked R&D dollar costs the firm $\chi - \tau$ rather than χ, and the government pays the difference, so the credit's outlay in a quarter is τ times the firms' booked R&D spending that quarter. The estimated scale already embeds the public support in force over the sample period, principally

[38] The act was signed in August 2022, under the Biden administration, and attacked under the Trump administration on the basis of low additionality and low returns: in his March 2025 address to Congress, President Trump called it "a horrible, horrible thing" ("We give hundreds of billions of dollars and it doesn't mean a thing") and urged that it be scrapped. The alternative settled on was not non-intervention but a different instrument: in August 2025 the administration converted Intel's remaining CHIPS funding, about $8.9 billion, into a 9.9 percent government equity stake. Public equity in private firms is a subject of its own, and I do not study that instrument here.

[39] About $5 billion over ten years is a policy-relevant scale: roughly seven times the combined federal R&E tax credits (Section 41) the two firms disclosed in their 10-K filings over 2011–2020 (roughly $0.7 billion).

the economy-wide federal R&E credit, so the R&D credit stacks on the existing credit rather than landing on a blank market. A rebate is a per-unit payment b on every covered unit sold; in the stage game it is equivalent to lowering the covered products' marginal cost by b, and its outlay in a quarter is b times covered units sold. Each instrument's rate is a constant, chosen so that planned outlay, evaluated at estimated-equilibrium behavior, equals the same present-value budget B:

$$\tau \sum_{h=0}^{39} \beta^h \, \mathbb{E}\big[I_{t+h}\big] \;=\; b \sum_{h=0}^{39} \beta^h \, \mathbb{E}\big[Q^{\text{cov}}_{t+h}\big] \;=\; B,$$

where I_{t+h} is the two firms' combined R&D spending and Q^{cov}_{t+h} the covered units sold, both along the estimated-equilibrium path (expectations over simulated paths), and B is the present value of the common budget, spread evenly across the 40 quarters. The R&D credit leaves the stage game unchanged, so it runs on the surrogates of Section 4.3.2 as estimated; each rebate changes the stage game itself, so prices, shares, qualities, and profits under it come from surrogates re-fit with the rebate applied to its covered tiers (Appendix C.5).

At the $5 billion budget, the R&D credit is a substantially more aggressive version of the federal credit the firms already receive: about a 12 percent reduction in the estimated R&D-cost scale, equivalently about $0.55 per dollar of the firms' accounting R&D spending (Table F.3 reports the credit's returns at nearby rates). The premium rebate comes to about $86 per premium unit (about 17 percent of the average premium price), and the non-premium rebate to about $10 per covered unit (about 4 percent of the average price of the units it covers). The applied rate triggers an equilibrium response in price, quality, and R&D. Realized outlay is therefore an equilibrium outcome, not a constant: each instrument stimulates the very base its outlay is proportional to (the credit raises R&D spending, a rebate raises covered sales), so realized spending drifts from plan. I match the budget ex ante (the plan comes to about $4.1 billion in present value) because this keeps the instrument a primitive of the experiment. It is also simpler than an ex post match, which requires solving a computationally intensive fixed-point loop. Returns are computed on the realized spending.

I report three per-dollar objects: additionality, within-market surplus per public dollar, and frontier advancement per public dollar. The first two are

$$\text{additionality} = \frac{\sum_{h=0}^{39} \beta^h \, \Delta I_{t+h}}{G}, \qquad \text{surplus per public dollar} = \frac{\sum_{h=0}^{39} \beta^h \, \Delta W_{t+h}}{G},$$

where ΔI_{t+h} is the change in the two firms' combined R&D spending in quarter $t + h$ relative to the estimated equilibrium, averaged over simulated paths, ΔW_{t+h} the same change in the within-market surplus W of Section 5, and G the instrument's realized discounted outlay: the quarterly outlays defined above, discounted and summed over the horizon.[40] Frontier advancement per public dollar is the percent change in the industry frontier at the end of the horizon, per billion dollars of realized outlay. The public outlay itself is excluded from the surplus measure W (a rebate dollar is a transfer, not surplus) and enters the accounting only as the denominator G; a public dollar is treated as a pure transfer, with a marginal cost of public funds of one, so the per-dollar surplus objects are read at face value. The object definitions are collected in Appendix E.3. The welfare returns follow in Section 6.2 and the returns in frontier advancement in Section 6.3.

6.2 Additionality and Welfare per Public Dollar

Table 4 reports additionality, within-market surplus and consumer surplus per public dollar, the decomposition of the welfare change, and budget drift under each instrument. Additionality is +0.18 for the R&D credit, −0.03 for the premium rebate, and 0.00 for the non-premium rebate: only the credit crowds in private R&D on net; the premium rebate crowds it out, and the non-premium rebate leaves it nearly unchanged. Within-market surplus per public dollar is +0.82 for the R&D credit, +0.06 for the premium rebate, and +0.53 for the non-premium rebate. The decomposition behind these figures shows how the welfare components move, with induced R&D carried in its own column. Under the credit, consumer surplus rises +\$2.8 billion and profit +\$1.8 billion against +\$0.8 billion of induced R&D, for ΔW = +\$3.8 billion. Under each rebate, consumer surplus rises while the producer operating-profit component, measured net of rebate transfers, falls: +\$5.9 billion against −\$5.6 billion under the premium rebate; +\$3.3 billion against −\$0.9 billion under the non-premium rebate. These negative producer-profit changes are entries in the welfare decomposition, not transfer-inclusive private payoffs; adding the rebate receipts back raises private payoffs at both firms under either rebate. Induced R&D is small under both rebates, −\$0.2 billion under the premium rebate and near zero under the non-premium rebate; the totals come to ΔW = +\$0.5 billion (premium rebate) and +\$2.4 billion (non-premium rebate).

Three patterns in the results are worth highlighting. The instrument ranking is consistent across both headline metrics: the R&D credit first, the non-premium rebate second, and the premium

[40] Direct expenditure is the whole of the denominator: a policy-induced change in the cost of support already embedded in the estimated equilibrium is an indirect fiscal footprint, not part of it.

rebate last. But the agreement ends there. The surplus column spans a far wider range than the additionality column, +0.82 to +0.06 against +0.18 to −0.03, and the two statistics can part ways instrument by instrument: the non-premium rebate induces essentially no private R&D yet clears the second-highest surplus per public dollar in the table. The decomposition columns show why: in every row, induced R&D is the smallest entry, so an instrument's return in surplus is decided by the far larger moves it sets off in consumer surplus and ex-transfer producer profit, not by the R&D it adds or displaces. Indeed, the consumer side barely separates the instruments, each delivering roughly 0.6 to 0.7 of consumer surplus per public dollar; what ranks them is the producer column, where only the R&D credit leaves the firms better off once rebate transfers are netted out. Read together, the columns give push and pull distinct signatures: only the push instrument, the R&D credit, buys additional R&D and within-market surplus together; the two pull instruments buy no additional R&D on net and instead reallocate surplus toward consumers through their product-market responses.

The table also shows that additionality is an incomplete policy indicator. It misreads the comparison at the level and is silent at the margin that decides it. At the level it understates the credit: most of the credit's outlay reimburses R&D the firms would have booked anyway, a transfer rather than a resource cost, so its denominator carries spending that was never in question. Between the two rebates it offers only a near-tie, −0.03 against 0.00, atop a real divergence in their returns: the non-premium rebate clears +0.53 of surplus per public dollar and edges the frontier up, while the premium rebate clears neither.

These welfare changes come from how each instrument moves prices, qualities, and shares along the product line. Table 5 traces the push-versus-pull contrast at the tier level. Under the R&D credit the whole menu lifts (Panel A): quality rises on every tier, +3.2 percent at the low-end, +3.3 at the midrange, and +4.7 at the premium. Share rises on the lower tiers, +0.66 and +0.68 percentage points, with the premium tier roughly flat. The lift is even because the credit works through the frontier, which every tier shares: the premium tier carries it directly as its own quality, and the lower tiers draw on it through the within-firm cost spillover of Section 3.2.1, so no tier's gain is bought by repositioning another. Panel A's one sizable price move is at the top: the pooled premium price rises by about +$30 per unit. The rise is quality-led: the markup column shows about +$8, so the higher price mainly prices a genuinely better product. The menu-wide lift is where the credit's row in Table 4 comes from: consumers get a better menu at little-changed prices below the premium, the +$2.8 billion of consumer surplus, and the firms sell that menu at higher

Table 4. Additionality and welfare per public dollar

	Per public dollar			Discounted totals ($B)				Budget
Instrument	Additionality	Welfare	CS	ΔCS	$\Delta\Pi$	ΔI	ΔW	drift
R&D credit	+0.18	+0.82	+0.62	+2.8	+1.8	+0.8	+3.8	+11%
Premium rebate	−0.03	+0.06	+0.64	+5.9	−5.6	−0.2	+0.5	+122%
Non-premium rebate	0.00	+0.53	+0.73	+3.3	−0.9	0.0	+2.4	+9%

Notes: Additionality is discounted induced private R&D per dollar of realized public outlay (the crowd-in statistic); welfare and CS per public dollar are the discounted changes in the within-market surplus W and in consumer surplus per dollar of the same outlay. Realized outlay is the public outlay actually spent over the ten-year horizon, discounted, and budget drift is that outlay relative to the common planned budget, about $4.1 billion discounted: realized outlays span about $4.5 billion (non-premium rebate) to $9.2 billion (premium rebate). Each instrument's rate is in Section 6.1. Each dollar column is a discounted 40-quarter total, an expected outcome averaged over simulated equilibrium paths, relative to the estimated equilibrium and measured within the market, with downstream spillovers excluded. The relation $\Delta W = \Delta\text{CS} + \Delta\Pi - \Delta I$ is an accounting identity rather than a decomposition into independent channels; $\Delta\Pi$ is producer operating profit net of rebate transfers, and ΔI is the firms' R&D spending. Entries are rounded independently, so printed columns may fail to close by up to 0.1, and the per-dollar columns are computed from unrounded values, so ratios of printed entries can differ in the second decimal. A public dollar is treated as a pure transfer; a reader who applies a marginal cost of public funds above one subtracts the excess from the surplus columns uniformly, leaving the instrument ranking unchanged.

markups, the +$1.8 billion of profit. Each rebate instead reshapes the menu around its covered tiers (Panels B and C) rather than lifting it: under the non-premium rebate the uncovered premium tier is dragged along only modestly, and under the premium rebate the tier's markup collapses.

Panels B and C of Table 5 carry the mechanism behind the two rebates' rows in Table 4. Two features organize it: a rebate paid on both firms' covered products sets off a price war, and advancing the frontier is rewarded in two forms, a margin at the top of the line and, through the within-firm cost spillover, saved production cost on the volume sold below it. Take the premium rebate first (Panel B). On its covered tier, the premium, the consumer price falls $131.2, well more than the rebate itself; the markup falls with it, $124.2; quality slips 1.3 percent; and buyers pile in, +5.98 percentage points of share. The uncovered tiers are dragged along the menu, the midrange repriced down $27.8 and the low-end up $8.8, both shedding share to the cheapened top. The drag fades with distance: the midrange loses $28.2 of markup to the war, the low-end none at all. Consumers gain on price and on adoption, and that is where the rebate's +$5.9 billion of consumer surplus comes from. But the same price war that hands buyers the transfer compresses the covered tier's displayed ex-transfer markup by $124.2; even after adding back the roughly $86 rebate receipt, the private margin falls by about $38. That margin is the reward for holding the frontier. Producer operating profit net of the rebate transfer falls $5.6 billion, and R&D falls with the reward that

Table 5. Tier-level price, quality, and share response under public support

Tier	Δ price ($)	Δ markup ($)	Δ quality (%)	Δ share (pp)
Panel A. R&D credit				
Low-end	+3.4	+1.5	+3.2	+0.66
Midrange	+3.8	+1.7	+3.3	+0.68
Premium	+29.7	+8.2	+4.7	−0.04
Panel B. Premium rebate				
Low-end	+8.8	+1.2	+1.8	−1.30
Midrange	−27.8	−28.2	−0.7	−0.46
Premium	−131.2	−124.2	−1.3	+5.98
Panel C. Non-premium rebate				
Low-end	−4.9	−13.5	+3.1	+1.69
Midrange	−43.3	−46.3	+1.1	+2.25
Premium	−19.5	−20.8	+0.3	−0.24

Notes: Entries are 40-period averages over simulated paths, relative to the estimated equilibrium. The price, markup, and quality columns are tier-pooled averages across both firms; the share column is the per-tier total, both firms' share changes summed. Under the two rebates the price column is the price consumers pay, with the rebate applied. The displayed markup is consumer price minus predicted marginal cost, the operating margin net of the rebate transfer; the per-unit rebate is not added to the displayed margin. The product-level percentage changes behind the markup column under the two rebates are in Table E.2.

drives it, −$0.2 billion on net. The fall is NVIDIA's: it cuts its own R&D $0.25 billion as that return thins, while AMD's is essentially unchanged. The tier's quality slip is that withdrawal made visible: the premium product's quality is the frontier itself, so the R&D cut lands on the covered product.

The non-premium rebate works the same way one tier down (Panel C). It covers the low-end and midrange, and on both the consumer price falls, $4.9 at the low-end and $43.3 at the midrange, the midrange by well more than the roughly $10 rebate. Here quality rises rather than slips, +3.1 and +1.1 percent, and adoption rises on both covered tiers, +1.69 and +2.25 percentage points. Those share gains are mostly new buyers rather than trade-down from the top: the tier shares are shares of the whole market, and the total inside share rises about 3.7 percentage points while the uncovered premium tier gives up only 0.24. The premium tier is dragged along too, $19.5 down on price and $20.8 on markup, with quality nearly flat. Consumers gain $3.3 billion. Producer profit net of the transfer falls $0.9 billion, less than under the premium rebate because the compression lands on thinner low-end and midrange margins. Those thin margins are not the frontier's reward: below the premium tier the frontier pays in volume, through the production-cost savings its advance generates over covered sales, and covered volume expands under the rebate. R&D therefore barely moves on net, the two firms pulling against each other, AMD up $0.19 billion and NVIDIA down $0.17 billion. Within-market surplus per public dollar still comes to +0.53.

The aggregate profit declines are not organized by the rebates' coverage sets. Under the non-premium rebate, the largest single-product ex-transfer operating-profit gain arrives at a covered low-end product while the deepest loss lands on another covered product; under the premium rebate every product loses, and an uncovered midrange product loses the most. The driver is the equilibrium repricing across the menu (the per-product breakdown is in Appendix E.6). On the producer side too, the menu, not the rebate's covered set, decides where the ex-transfer component goes.

The premium rebate comes with one scale caveat. Its realized outlay drifts far above plan: +122 percent (the budget-drift column of Table 4), about $9.2 billion realized against the roughly $4.1 billion plan. The drift is directionally consistent with the premium tier facing the most elastic demand of the three tiers (Section 4.1). The R&D credit and the non-premium rebate, the two instruments the surplus column ranks first and second, realize nearly identical outlays close to plan, $4.6 billion (credit) and $4.5 billion (non-premium rebate). That comparison is therefore read at a common realized scale, while comparisons involving the premium rebate run across different realized scales.

Table 6. Static and dynamic components of welfare per public dollar under the two rebates

Instrument	Static welfare effect	Dynamic state response	Welfare per public dollar
Premium rebate	+0.19	−0.13	+0.06
Non-premium rebate	+0.54	−0.01	+0.53

Notes: The static welfare effect holds R&D and the industry frontier at the levels of the estimated equilibrium, replays only the rebate's within-period pricing and quality responses, and is expressed per realized public dollar of that replay. The dynamic state response is the remainder once R&D, the frontier, and market positions respond; the two sum to the total, the within-market surplus per public dollar of Table 4. The split is a replay diagnostic computed with the rebate-refit surrogates, not a structural partition into independent channels. The R&D credit changes no stage-game object and has no static component, so only the two rebates appear. Conventions are those of Table 4. Construction detail is in Appendix E.4.

Both rebates mix a within-period force and a dynamic one: the price war reprices the menu inside the period, while the change in the frontier's reward matters only as R&D, the frontier, and market positions respond. To separate the two, I hold R&D and the industry frontier at the levels of the estimated equilibrium and replay only each rebate's within-period pricing and quality responses. The welfare effect per public dollar of that replay is the static welfare effect; the rest of the total welfare per public dollar in Table 4 is the dynamic state response. Table 6 reports the split.

The R&D credit changes no stage-game object, so it has no static layer to separate: it works entirely through the dynamic state response, and its absence from the table follows from the construction, not from a result. Both rebates, by contrast, are welfare-positive as static policies. The premium rebate's static welfare effect is +0.19 per public dollar: with R&D and the frontier held at the estimated equilibrium, repricing and reallocating the menu raises surplus. The dynamic state response then takes back most of that gain, −0.13, for a total of +0.06, essentially nothing per public dollar. The non-premium rebate's static effect is +0.54 and its dynamic state response is negligible, so the total holds at +0.53, essentially the static number.

6.3 Frontier Advancement per Public Dollar

Table 7 reports additionality, frontier advancement per public dollar, and the frontier gain per dollar of induced R&D. Frontier advancement per public dollar is the percent change in the end-of-horizon industry frontier per $1 billion of realized outlay. The R&D credit leads at 1.11: a billion dollars of credit outlay lifts the terminal frontier about one percent. The non-premium rebate manages a

Table 7. Frontier advancement under public support

Instrument	Additionality	Frontier advancement	
		% per $1B public outlay	% per $1B induced R&D
R&D credit	+0.18	1.11	6.23
Premium rebate	−0.03	−0.29	–
Non-premium rebate	0.00	0.11	–

Notes: Additionality is as in Table 4. Frontier advancement is measured on the industry frontier $q_t^{\max}$, the higher of the two firms' frontier qualities. The two frontier columns share that numerator, the percent change in the end-of-horizon frontier, and differ only in the denominator: realized outlay in the first, induced private R&D in the second, both in billions. The induced-R&D column is reported only for the R&D credit: neither rebate induces private R&D on net, so the ratio would put the frontier change over an R&D withdrawal (premium rebate) or over a near-zero net change (non-premium rebate), in neither case a return, and their cells print "–". For the R&D credit, frontier advancement per public dollar equals additionality times frontier advancement per dollar of induced R&D, computed from unrounded values, so the printed product can differ in the second decimal. Realized outlays span about $4.5 billion (non-premium rebate) to $9.2 billion (premium rebate); budget drift relative to the common planned budget is in Table 4, and each instrument's rate is in Section 6.1. Conventions are those of Table 4. For scale, the estimated equilibrium's industry frontier grows about 371 percent over the same 40-quarter horizon.

small positive 0.11, and the premium rebate is negative, −0.29; the one instrument aimed squarely at the frontier tier pushes the frontier back. The frontier column thus ranks the three instruments in the same order as the surplus column, the R&D credit ahead of the non-premium rebate ahead of the premium rebate.

How much frontier advancement an instrument buys depends not only on how much R&D it induces but on where the induced R&D sits. The industry frontier is the higher of the two firms' frontier qualities, so R&D induced at the firm holding the frontier advances it directly, while R&D induced at the follower narrows the gap between the firms and moves the industry frontier little. Table 8 locates each instrument's R&D response: it attributes the response to the leader and follower roles on the estimated-equilibrium path. The R&D credit raises R&D at both roles by nearly the same amount, +11.1 percent at the leader and +10.4 at the follower. The premium rebate cuts R&D at both roles, and the cut is larger at the leader, −5.2 percent against −1.5. The non-premium rebate barely moves either role, −0.2 and +0.7 percent. Held against the frontier column of Table 7, the pattern lines up instrument by instrument: the one instrument that adds R&D at the leader role is the one that advances the frontier, the instrument that cuts leader R&D most is the one that moves

Table 8. Percent change in R&D by frontier role under public support

	Percent change in role R&D	
Instrument	Frontier leader	Frontier follower
R&D credit	+11.1	+10.4
Premium rebate	−5.2	−1.5
Non-premium rebate	−0.2	+0.7

Notes: Entries are the percent change in each frontier role's discounted R&D relative to the estimated equilibrium. Frontier leader and follower are state-contingent roles, not fixed firm identities: on each path-period the leader is the firm with the higher frontier quality on the estimated-equilibrium path. Each firm's induced R&D is attributed to its role at that path-period, discounted, and summed over the simulated paths and the 40-quarter horizon; the denominator is that role's discounted R&D on the estimated-equilibrium path, about $3.5 billion for the leader and $4.1 billion for the follower in present value. Construction detail and the discounted dollar levels are in Appendix E.7.

it back, and the instrument that leaves both roles nearly untouched buys only a small frontier gain.

The role patterns can be explained by the margin each instrument targets. The R&D credit lowers the cost of research itself, the same margin at either role, so it lifts both roles evenly and the R&D it induces includes the leader's. The premium rebate thins the reward for holding the frontier, the margin at the top of the line (Section 6.2), and the withdrawal is deepest at the role that earns that reward; the attribution reflects who holds the frontier along the path, not a firm investing differently because it leads. The non-premium rebate's near-zero role entries hide an offsetting move across the firms, one raising its R&D while the other lowers its own (Appendix E.7). The shift runs toward the firm that trails more often than it leads, so most of what it buys is a narrower gap between the firms; but the roles are state-contingent, the gaining firm's frontier rises while the other's falls, and the industry frontier is the within-path maximum, so the joint frontier still edges up even as total R&D barely moves.

The frontier gain per dollar of induced R&D, the second frontier column of Table 7, completes the accounting. It is defined only for the R&D credit, at about 6.2 percent per $1 billion of induced private R&D, and it is what turns the credit's +0.18 of additionality into 1.11 of frontier advancement per public dollar: the credit's dollar moves the frontier both because it induces private R&D on net and because the R&D it induces sits at both roles, the leader's included. For the two rebates the ratio has no meaning as a return: neither induces private R&D on net. Under the premium rebate

the division would set a frontier retreat over an R&D withdrawal, two declines whose quotient prints as a positive number; under the non-premium rebate the near-zero net change leaves nothing meaningful to divide by. Table 7 therefore marks those cells with "–".

6.4 Discussion and Policy Implications

The instrument comparison delivers two per-dollar rankings, by welfare and by frontier advancement, together with the net R&D response each instrument induces. The R&D credit crowds in +0.18 of private R&D per public dollar, returns +0.82 of within-market surplus per public dollar, and advances the end-of-horizon frontier by 1.11 percent per $1 billion of outlay; the premium rebate crowds R&D out, returns essentially nothing per public dollar, and pushes the frontier back; the non-premium rebate induces near-zero net R&D yet returns +0.53 and edges the frontier up. Read as a design guide, the pattern echoes the principle of targeting: pursue an objective with the instrument aimed most directly at it.

But the targeting principle holds here only under a particular reading of directness: distance to the decision margin, not to the technological object. The premium rebate is the instrument aimed most squarely at the frontier product, yet the worst for the frontier: -0.29 percent per $1 billion, with the R&D cut concentrated at whichever firm leads the frontier. The reason for the backfire is that demand for the frontier good and the reward for advancing the frontier are different objects. The reward comes in two forms, a margin at the top of the line and production cost saved on the volume sold below it. A rebate paid at the top lifts demand for the good while setting off the price war that compresses the margin. Buying cutting-edge output or subsidizing a flagship product targets the frontier directly in product space, but demand for that product need not move with the frontier's reward, and how closely they do is itself an equilibrium outcome.

The distinction bears on a wider argument over how directive public support should be. The classic objection to directed support is that the government cannot pick winners, and that even a well-chosen target can be spent badly through administrative frictions. My experiment grants the government both: a well-chosen target and frictionless administration. A pull policy can still trigger an equilibrium response in the product market that runs against what it intends. Here it manifests as the price war, which hands consumers lower prices at the frontier's expense: the two rebates differ only in covered tiers yet end with opposite frontier signs. In another market the surprise could run the other way: a subsidy that softens price competition might raise markups or lower the quality of non-premium products, leaving consumers worse off while the firms gain. A

push instrument carries less of this risk. Lowering the cost of research whether or not the research succeeds changes no stage-game object, so the direction of the R&D response is signed in advance; only its size, and with it the additionality a public dollar buys, is an empirical question.

The natural place to look for that size is the social-private wedge in R&D. There is now near consensus among economists that the wedge is large. I confirm it in my setting, even when cross-firm knowledge spillovers are absent. But a wide wedge need not mean that R&D sits far below its efficient level, nor that a well-designed instrument will convert much of the gap. The wedge measures how the surplus the R&D program created is divided, while the marginal response to a cost subsidy is set by how fast returns diminish where the firms already operate. Under incremental technological advance, how much the frontier can improve in a period is often concave in R&D spending, and the difficulty rises as the frontier moves up. The firms' own returns could have already pulled spending into the flat region of that curve, so the efficient level may sit close above the observed one and a further reduction in the cost of research moves spending little. In contrast, where innovation is discrete, as in drug development, a public dollar can tip whether a project happens at all, and direct support can be far more effective.

Conceptually, this heterogeneity in policy response must be there: industries differ in their innovation technologies and in how far private returns have pushed research spending along them. It is, however, hard to measure. The most direct evidence on additionality comes from cross-industry panel regressions, which suggest a roughly dollar-for-dollar figure. This averaging is hard to undo. The more finely a causal design tries to isolate one industry's response, the more its identifying comparisons run between that industry's firms, and the more strongly it violates the stable unit treatment value assumption it rests on. The heterogeneity creates a selection issue for the pooled figure: governments select strategic-support targets on the perceived value of the technology, a trait often positively correlated with private returns, so realized additionality in those sectors should sit below the pooled benchmark, as it does here at +0.18. The right tail carries the opposite warning: the sectors where a public dollar would induce the most R&D are, by the same logic, those where private incentives to innovate are weakest. The weak private incentives may reflect a large surplus the firms cannot capture, but they could also arise in an industry where innovation is secondary and there is little surplus to create, in which case, absent some other externality, the intervention is hard to justify.

Aggregate additionality has a further issue in industries with strong innovation rivalry: it tracks the sum of induced R&D spending across firms, and the sum is invariant to reallocation between

them. The industry's cutting edge, however, depends largely on how well the front-running firm is doing. A policy intended to accelerate technological growth could induce more R&D in aggregate yet tilt the allocation toward the firms that trail, or even weaken the leader's incentive; such a policy narrows the gap between the firms while doing little to advance the industry frontier, and nothing in the aggregate figure would show it. In my experiments, the premium rebate's aggregate R&D response is only modestly negative, but the cut concentrates at the leader role, which turns it into a frontier retreat.

The policy comparison also surfaces a trade-off between consumer surplus and innovation. All three instruments deliver roughly 0.6 to 0.7 of consumer surplus per public dollar, a modest but non-trivial gain over the ten-year window. Beneath that flat column, the welfare and frontier columns diverge sharply, and the divergence has one source: a price war is a gain on the consumer ledger and a withdrawal of the frontier's reward on the producer ledger. Policies aimed at improving consumer welfare in technology-driven markets, such as adoption subsidies, can therefore work well in the short term yet backfire after years: the consumer gain is real, while the cost sits outside the evaluation window, in the frontier the future inherits. That long-run cost is hard to quantify from data alone, because the frontier path the market would otherwise have taken is never observed. Recovering it requires a model, not just a causal design.

Finally, my model provides a framework for positive economic analysis of policies whose motivation is not mainly economic. Large industrial policies today are often motivated by national security, geopolitical tension, supply-chain resilience, or technological leadership in its own right, and the value of achieving those objectives is hard to quantify in welfare terms. What the framework can measure is instead the price of pursuing them. Whatever the objective, the instrument works through the product market in the same way, and the analysis proceeds unchanged; what changes is the reading of its output, from a verdict on the policy to the cost of the objective. For example, requiring that strategic goods such as advanced chips be produced domestically can raise their marginal cost; the increase enters the model as, in effect, a tax on production. The surplus consumers and firms surrender, and any innovation forgone over the longer run, are then not failures of the policy but its price. The framework can therefore be used to ask which instrument achieves a given non-economic goal at the least economic cost, or with the least disruption to the market.

7 Conclusion

This paper develops a framework for analyzing the private incentive to innovate and the effects of public support: a dynamic oligopoly model of product line markets with endogenous pricing, quality positioning, and R&D. I apply it to the desktop GPU market, quantifying the returns to R&D and simulating three stylized support instruments on a common budget. The return is large: a dollar of the industry's R&D generates about $8.8 of within-market surplus; the firms' net return is about $3.0, and the rest is consumer surplus the firms cannot appropriate. Despite this wide wedge, the policy effects are mixed and modest at best: a public dollar routed through any of the three instruments buys far less surplus than a dollar the firms spend on R&D themselves.

The contrast has two sources. First, the marginal effect of support is small. The firms' own returns have already carried R&D spending to where the next dollar just breaks even for them: the $8.8 is what the program earns on average, not at the margin. A credit therefore pays largely for research that needed no inducing. Second, the product market responds. A rebate sets off a price war that compresses margins across the product line, and at the top of the line the margin lost to the war is the reward for holding the frontier. The rebate can therefore kill the very innovation it is meant to buy. The broader lesson is that pricing, positioning, and R&D move together in product line markets: an instrument aimed at any one margin lands on all three, and the R&D it induces is an input response, not a welfare verdict.

Three future directions follow. The first is to take the framework to other markets: nothing in the model is specific to GPUs, to two firms, or to three tiers per product line. A few other innovation-driven markets fit naturally into the framework. Electric vehicles, for example, can be viewed as a product line market, where driving range serves as the vertically differentiating quality dimension, and carmakers invest heavily to improve battery density. Another recent example is tiered access to large language models, where benchmark performance prices the menu and the within-firm cost spillover has a literal analog in the so-called distillation technique: work on a firm's frontier model lowers the cost of its smaller ones. The payoff extends beyond any single market. Growth accounting prices the economy's entire innovation effort at one average return inside a wide band; returns estimated market by market would build the distribution beneath that average, each carrying its social-private split and policy counterfactual. Such disaggregation can help uncover the sources of heterogeneity in private innovation incentives and policy effects.

The second direction is to measure what the within-market boundary leaves out. I exclude

downstream spillovers deliberately: the exclusion keeps the surplus object well defined and, provided the omitted spillovers are positive, makes the $8.8 a lower bound on the economy-wide value of the same dollar. The omitted channel is concrete: the chips sold to render games became the hardware that trains modern AI. Quantifying the surplus that GPU frontier advances created downstream would show how far the broader gains exceed the within-market measure.

The third direction concerns the yardstick itself. Additionality tracks the R&D response a public dollar induces, not the product-market response that decides what the dollar ultimately buys. Yet it is the most readily available indicator: R&D spending appears in financial reports and tax records, while the equilibrium response took the full apparatus of this paper to see, machinery few agencies can build. The practical questions are how to build better yardsticks where detailed data are limited: when the input response alone is a reliable guide to the outcomes policy aims at, and what is worth collecting beside it when it is not. Tier-level prices, qualities, and market shares can be observed for each program separately, and the product-market response shows up in them. An estimated framework of this kind offers a testbed: candidate indicators built from such data can be scored against the model's welfare and frontier outcomes before any is trusted in the field.

References

Acemoglu, D. and J. Linn (2004): "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry," *The Quarterly Journal of Economics*, 119, 1049–1090.

Arrow, K. J. (1962): "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, ed. by Universities-National Bureau Committee for Economic Research and Committee on Economic Growth of the Social Science Research Council, Princeton, NJ: Princeton University Press, 609–626.

Bajari, P., C. L. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75, 1331–1370.

Banares-Sanchez, I., R. Burgess, D. László, P. Simpson, J. Van Reenen, and Y. Wang (2026): "Ray of Hope? China and the Rise of Solar Energy," NBER Working Paper 34893, National Bureau of Economic Research.

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

Bloom, N., M. Schankerman, and J. Van Reenen (2013): "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 81, 1347–1393.

Bodéré, P. (2026): "Dynamic Spatial Competition in Early Education: An Equilibrium Analysis of the Preschool Market in Pennsylvania," Working paper, Yale School of Management, July 2026.

Bryan, K. A. and H. L. Williams (2021): "Innovation: market failures and public policies," in *Handbook of Industrial Organization*, ed. by K. Ho, A. Hortaçsu, and A. Lizzeri, Amsterdam: North-Holland, vol. 5, 281–388.

Conlon, C. and J. Gortmaker (2020): "Best Practices for Differentiated Products Demand Estimation with PyBLP," *The RAND Journal of Economics*, 51, 1108–1161.

De Groote, O. and F. Verboven (2019): "Subsidies and Time Discounting in New Technology Adoption: Evidence from Solar Photovoltaic Systems," *American Economic Review*, 109, 2137–2172.

Dialogic, APE, and UNU-MERIT (2019): "Evaluatie WBSO 2011–2017," Evaluation report, Dialogic, APE, and UNU-MERIT, Utrecht, commissioned by the Dutch Ministry of Economic Affairs and Climate Policy; transmitted to the Tweede Kamer, April 2019.

Doraszelski, U. and A. Pakes (2007): "A Framework for Applied Dynamic Analysis in IO," in *Handbook of Industrial Organization*, ed. by M. Armstrong and R. H. Porter, Amsterdam: North-Holland, vol. 3, 1887–1966.

Doraszelski, U. and M. Satterthwaite (2010): "Computable Markov-Perfect Industry Dynamics," *The RAND Journal of Economics*, 41, 215–243.

Eizenberg, A. (2014): "Upstream Innovation and Product Variety in the U.S. Home PC Market," *The Review of Economic Studies*, 81, 1003–1045.

Ericson, R. and A. Pakes (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *The Review of Economic Studies*, 62, 53–82.

Fan, Y. (2013): "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market," *American Economic Review*, 103, 1598–1628.

Fan, Y. and C. Yang (2020): "Competition, Product Proliferation, and Welfare: A Study of the US Smartphone Market," *American Economic Journal: Microeconomics*, 12, 99–134.

Finkelstein, A. (2004): "Static and Dynamic Effects of Health Policy: Evidence from the Vaccine Industry," *The Quarterly Journal of Economics*, 119, 527–564.

Gandhi, A. and J.-F. Houde (2019): "Measuring Substitution Patterns in Differentiated-Products Industries," NBER Working Paper 26375, National Bureau of Economic Research.

Genz, A. and B. D. Keister (1996): "Fully Symmetric Interpolatory Rules for Multiple Integrals over Infinite Regions with Gaussian Weight," *Journal of Computational and Applied Mathematics*, 71, 299–309.

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

Goolsbee, A. and B. F. Jones (2022): "Introduction to 'Innovation and Public Policy'," in *Innovation and Public Policy*, ed. by A. Goolsbee and B. Jones, Chicago: University of Chicago Press, 1–11.

Griliches, Z. (1979): "Issues in Assessing the Contribution of Research and Development to Productivity Growth," *The Bell Journal of Economics*, 10, 92–116.

Hall, B. H. (2022): "Tax Policy for Innovation," in *Innovation and Public Policy*, ed. by A. Goolsbee and B. Jones, Chicago: University of Chicago Press, 151–188.

Hall, B. H. and J. Lerner (2010): "The Financing of R&D and Innovation," in *Handbook of the Economics of Innovation*, ed. by B. H. Hall and N. Rosenberg, Amsterdam: North-Holland, vol. 1, 609–639.

Hall, B. H. and J. Van Reenen (2000): "How effective are fiscal incentives for R&D? A review of the evidence," *Research Policy*, 29, 449–469.

HM Treasury (2023): "R&D Tax Reliefs Review: Consultation on a Single Scheme," Consultation document, HM Treasury, London.

Hsu, J. C. and E. S. Schwartz (2008): "A model of R&D valuation and the design of research incentives," *Insurance: Mathematics and Economics*, 43, 350–367.

Johnson, J. P. and D. P. Myatt (2003): "Multiproduct Quality Competition: Fighting Brands and Product Line Pruning," *American Economic Review*, 93, 748–774.

Jones, B. F. and L. H. Summers (2022): "A Calculation of the Social Returns to Innovation," in *Innovation and Public Policy*, ed. by A. Goolsbee and B. Jones, Chicago: University of Chicago Press, 13–59.

Jones, C. I. and J. C. Williams (1998): "Measuring the Social Return to R&D," *The Quarterly Journal of Economics*, 113, 1119–1135.

Krizhevsky, A., I. Sutskever, and G. E. Hinton (2012): "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances in Neural Information Processing Systems 25*, ed. by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Red Hook, NY: Curran Associates, 1097–1105.

Loury, G. C. (1979): "Market Structure and Innovation," *The Quarterly Journal of Economics*, 93, 395–410.

Maican, F. G., M. Orth, M. J. Roberts, and V. A. Vuong (2023): "The Dynamic Impact of Exporting on Firm R&D Investment," *Journal of the European Economic Association*, 21, 1318–1362.

Ministerie van Economische Zaken en Klimaat (2019): "Kamerbrief over de evaluatie WBSO 2011–2017," Parliamentary letter, Ministry of Economic Affairs and Climate Policy, The Hague, cabinet letter to the Tweede Kamer transmitting the Dialogic/APE/UNU-MERIT evaluation of the WBSO; issued jointly with the Ministry of Finance.

Mussa, M. and S. Rosen (1978): "Monopoly and Product Quality," *Journal of Economic Theory*, 18, 301–317.

OECD (2020): "The Effects of R&D Tax Incentives and Their Role in the Innovation Policy Mix: Findings from the OECD microBeRD Project, 2016–19," OECD Science, Technology and Industry Policy Paper 92, OECD Publishing, Paris.

Pakes, A. and P. McGuire (1994): "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model," *The RAND Journal of Economics*, 25, 555–589.

Peters, B., M. J. Roberts, V. A. Vuong, and H. Fryges (2017): "Estimating dynamic R&D choice: an analysis of costs and long-run benefits," *The RAND Journal of Economics*, 48, 409–437.

Peters, M., M. Schneider, T. Griesshaber, and V. H. Hoffmann (2012): "The impact of technology-push and demand-pull policies on technical change – Does the locus of policies matter?" *Research Policy*, 41, 1296–1308.

Petrin, A. (2002): "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110, 705–729.

Springel, K. (2021): "Network Externality and Subsidy Structure in Two-Sided Markets: Evidence from Electric Vehicle Incentives," *American Economic Journal: Economic Policy*, 13, 393–432.

Sweeting, A. (2013): "Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry," *Econometrica*, 81, 1763–1803.

Table A.1. The GPU market: product line and market aggregates, 2005Q1–2020Q4

Panel A. Product line by tier

Tier	Mean price ($2018)	Relative performance	Unit share (%)
Low-end	123	11.2	62.5
Midrange	297	53.5	30.0
Premium	726	100.0	7.6

Panel B. Market aggregates

Aggregate	Value	Basis
Market revenue	12.9	$B per year, average 2005–2020
Discrete GPUs sold	14.6	million per quarter, average 2005–2020
Attach rate	46.2	% of desktop CPU sales, pooled sample
Combined GPU-segment R&D	1.1	$B per year, average 2005–2020
R&D-to-revenue ratio	0.10	pooled sample
Frontier benchmark growth	98.7	multiple, 2005–2020
Premium-tier price growth	1.9	multiple, 2005–2020

Notes: Panel A reports unit-weighted means over the cleaned panel. Each firm–tier–quarter cell pools the products active in it, weighted by a popularity proxy. Relative performance is measured over the pooled sample, with the premium tier normalized to 100. The printed shares sum to 100.1 because of rounding. Panel B reports sample-wide magnitudes: average annual market revenue, average quarterly discrete-GPU units, the pooled attach rate, average annual combined GPU-segment R&D over fiscal years, the pooled R&D-to-revenue ratio, and endpoint growth in the premium series. The attach rate is discrete-GPU units divided by desktop CPU units, which is also the pooled inside share of the addressable market. AMD's R&D series begins with the ATI acquisition in fiscal 2006, so the first year of the R&D average reflects NVIDIA alone. The panel is balanced: NVIDIA and AMD each offer low-end, midrange, and premium tiers in every one of the sample's 64 quarters. Quality is the PassMark G3D benchmark index; prices combine launch MSRP and tracked retail prices; unit sales are JPR desktop discrete-GPU volumes, with desktop CPU sales as the market size; R&D is firm-reported spending allocated to the GPU segment in proportion to segment revenue. Over the sample, that allocation share averages roughly 60 percent for NVIDIA and roughly a quarter for AMD. Construction from product-level records is documented in the replication package.
Sources: Jon Peddie Research (JPR) for discrete-GPU units; firm filings and Compustat for financials; NVIDIA platform disclosures and AMD segment disclosures for the allocation bridge; PassMark for benchmark performance; and launch and retail-price records for price.

A The Desktop GPU Market, 2005–2020

Table A.1 summarizes the product line and the market aggregates; its notes record what each series measures and where it comes from. Panel A gives the tier ladder: low-end cards account for 62.5 percent of units at a mean price of $123, midrange cards for 30 percent at $297, and premium cards for 7.6 percent at $726, with the two lower tiers delivering 11 and 53 percent of premium benchmark performance.

Section 2 reports the endpoint growth of quality and price; Figure A.1 plots the path between the endpoints. The frontier in panel (b), the best premium-tier benchmark score in each year, climbs as a staircase: flat stretches broken by generational jumps, the largest of which, in 2008 and 2010, more than double the frontier within a year. The premium price in panel (a) moves

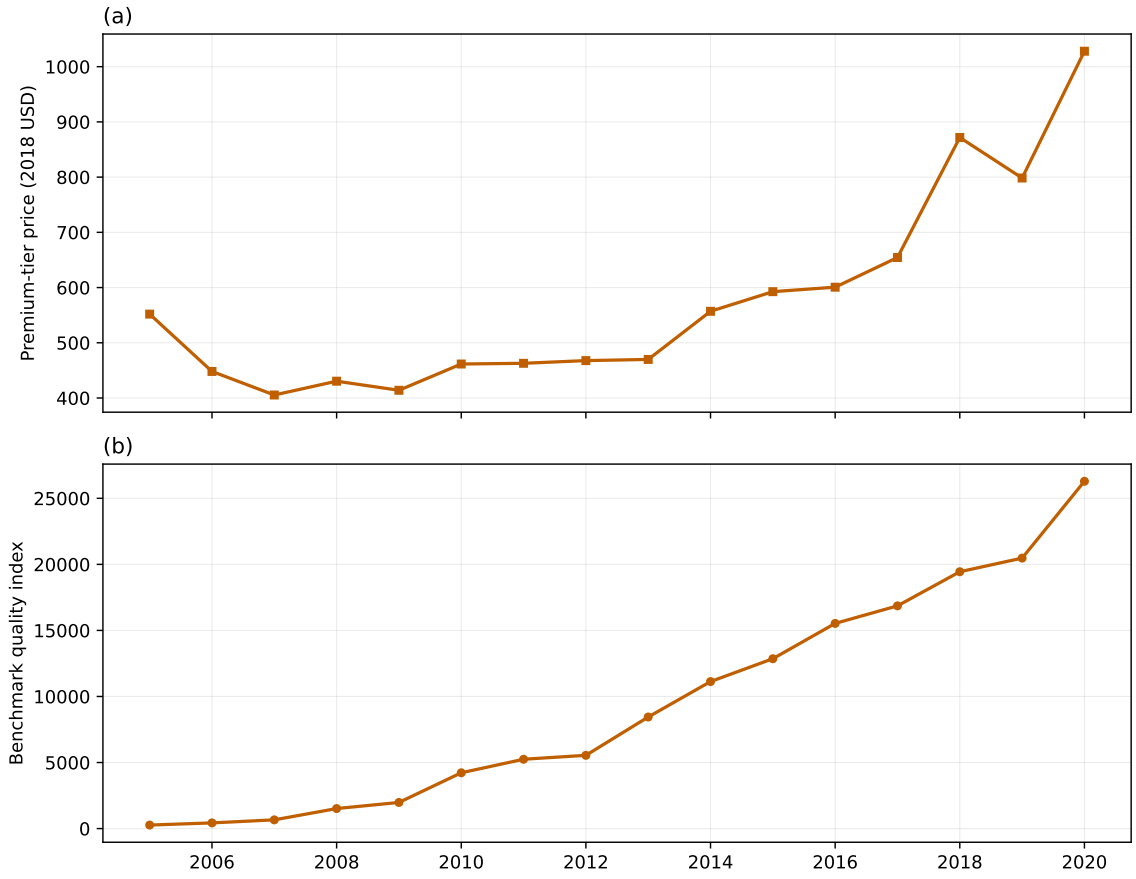


Figure A.1. Premium-tier prices and observed frontier quality

Notes: Panel (a) plots the premium-tier price, pooled across the two firms and averaged within year, in 2018 dollars. Panel (b) plots the observed frontier: the highest premium-tier benchmark index across the two firms in each year, on the same premium series as Table A.1. The 2005–2020 endpoint ratios of the two series equal the growth multiples in Table A.1, Panel B.

far less, holding between roughly $400 and $470 from 2006 through 2013 and climbing only late in the sample. Midway through the sample, these gaming cards found a second constituency: NVIDIA released CUDA, its platform for general-purpose GPU computing, in 2007, and the deep-learning breakthrough of 2012 was trained on two consumer gaming cards (Krizhevsky et al. 2012). Before the large-language-model boom of 2023–2024, however, the GPU market remained largely consumer-facing. It has since turned corporate-facing, with data-center accelerators that sell for tens of thousands of dollars apiece and are bought tens of thousands of units at a time by large AI developers; the market studied here is the consumer-facing one that preceded that shift.

The R&D behind the race appears in Figure A.2, firm by firm. NVIDIA's GPU-segment R&D grows from roughly $0.25 billion in fiscal 2005 to $1.8 billion in fiscal 2020. AMD's runs lower throughout: it falls to roughly $0.2 billion in fiscal 2015 and recovers to $0.56 billion by fiscal 2020.

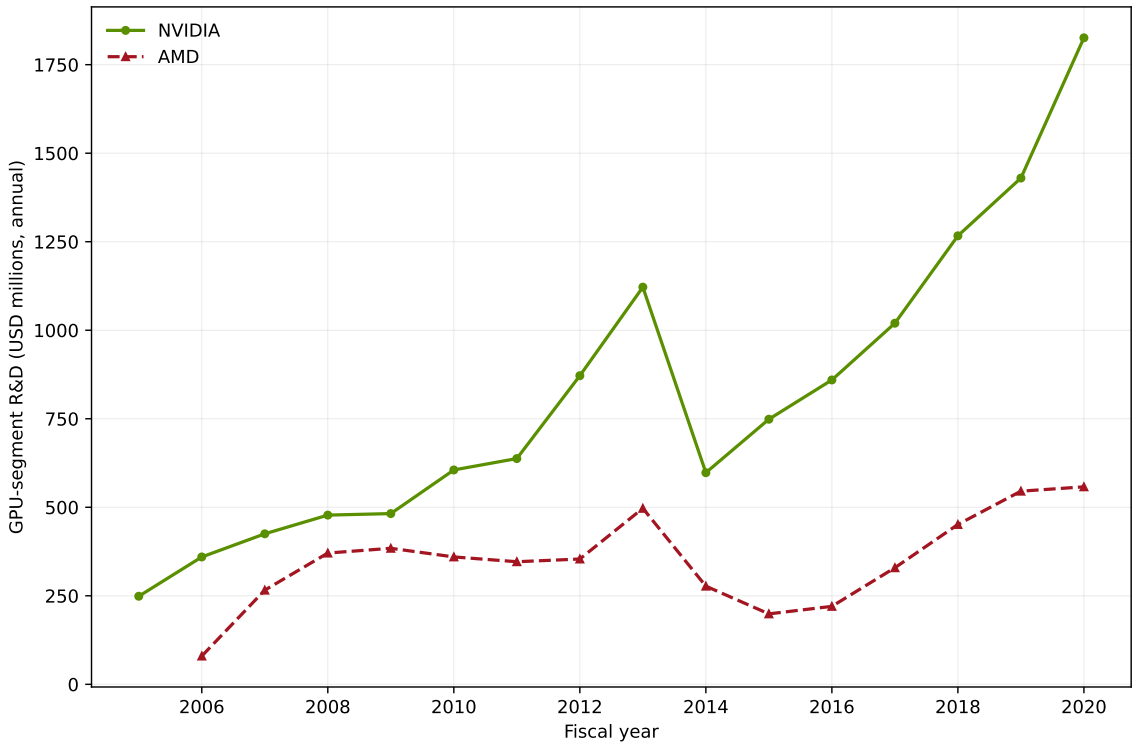


Figure A.2. Allocated GPU-segment R&D

Notes: Each series is the firm's reported R&D allocated to the GPU segment by the segment's revenue share (Section 2), summed within fiscal year, in millions of dollars; NVIDIA's fiscal years are labeled by the calendar year in which they begin. AMD's series begins in fiscal 2006 with the acquisition of ATI, and its first year covers only the partial post-acquisition period.

The two series are the accounting inputs behind the R&D moment of Section 4.

Competitive positions shifted inside the stable duopoly. AMD shipped roughly half of discrete-GPU units in 2005, slid to about a fifth by 2015, the trough of its fiscal-year R&D series in Figure A.2, and recovered to roughly 30 percent in the late 2010s. That history is the concrete face of the leader/follower asymmetry the model of Section 3 tracks: the policy results of Section 6 turn on which firm's R&D a public dollar reaches, and on whether the induced innovation advances the frontier or narrows the gap behind it.

Figure A.3 turns to adoption and market size. Panel (a) plots the attach rate, discrete-GPU units as a share of desktop CPU units in each quarter, around its pooled 46.2 percent (Table A.1, Panel B). Panel (b) plots the denominator: desktop CPU shipments fall from about 37 million units a quarter in 2005 to between roughly 21 and 27 million late in the sample and level off, the maturing-market pattern that Section 4 models as mean reversion toward a long-run level.

The shaded band in panel (a) marks the cryptocurrency-mining episode, the event behind the

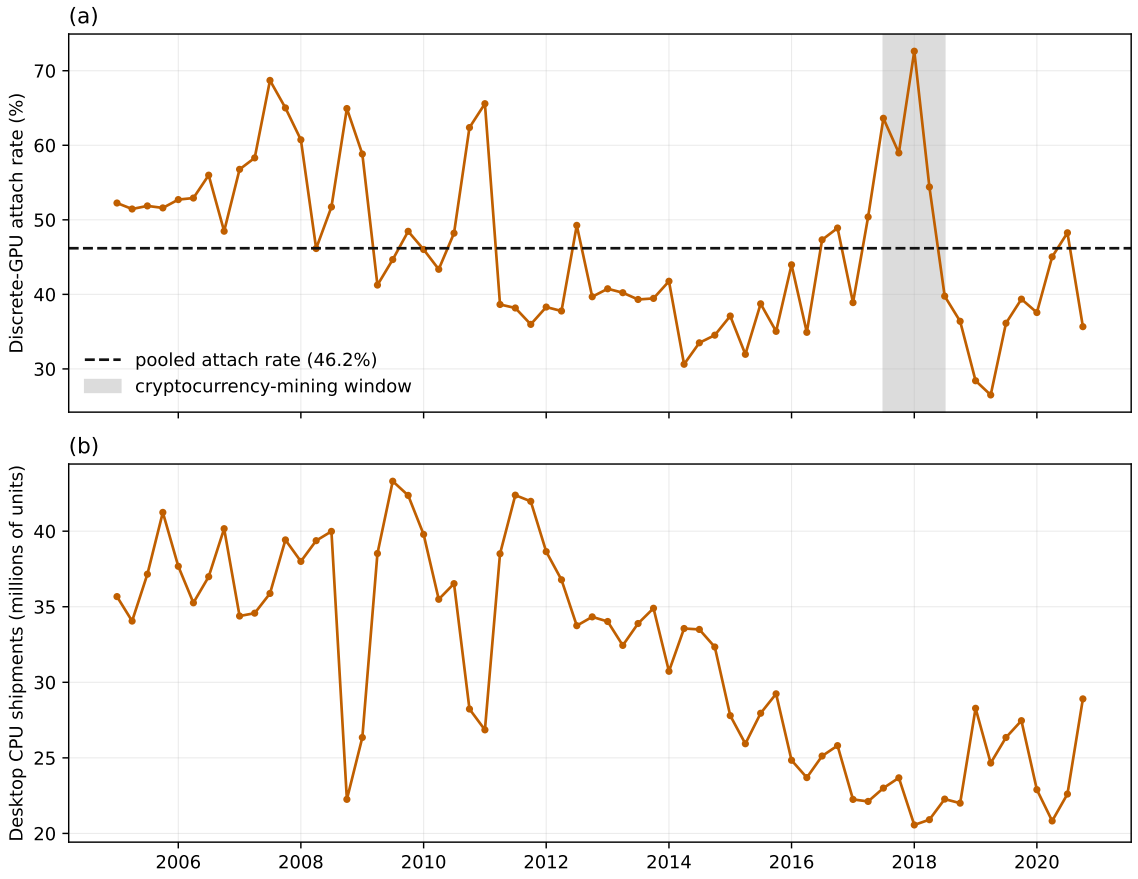


Figure A.3. Discrete-GPU adoption and market size

Notes: Panel (a) plots the quarterly attach rate, discrete-GPU units sold by the two firms divided by desktop CPU units, over 2005Q1–2020Q4; the dashed line marks the pooled attach rate of 46.2 percent (Table A.1, Panel B), and the shaded band marks 2017Q3–2018Q2, the cryptocurrency-mining quarters excluded from the outside-option trend fit (Section 4). Panel (b) plots quarterly desktop CPU shipments in millions of units, the market-size measure.

outside-option exclusion in Section 4. Over 2017Q3–2018Q2 miners bought gaming cards. The attach rate, which had run between roughly 30 and 50 percent in the three preceding years, jumped to 64 percent in 2017Q3 and peaked at 73 percent in 2018Q1; the surge concentrated up the ladder, as midrange units ran roughly 70 percent above their early-2017 pace, premium units roughly doubled, and low-end units fell about 10 percent. Prices peaked in the same quarter: in 2018Q1 NVIDIA's midrange and premium prices stood 72 and 54 percent above their early-2017 averages, and AMD's more than doubled; the bump is visible even in the annual premium series of Figure A.1. The boom then unwound: by the first half of 2019 the attach rate sat below 30 percent, beneath its pre-boom level. In the demand system the surge registers as an outside option that is temporarily much worse: the recovered outside-option series (Figure D.1) drops in exactly these four quarters, and the trend fit excludes them.

B Details for Demand and Cost Estimation

This appendix collects the detail behind the demand and marginal-cost estimates of Section 4. Appendix B.1 reports the estimated first stage that produces the joint F statistic of Section 4.1; Appendix B.2 reports the full elasticity matrix the policy analysis works through; Appendix B.3 spells out the inversion that recovers marginal costs and reports the by-firm split of Table 1's markup panel; and Appendix B.4 constructs the implied marginal fixed cost of quality that motivates the post-2010 window of Section 4.2. Appendix C assembles these primitives into exact static equilibria and their surrogate approximations.

B.1 The Demand First Stage

With one endogenous variable, price, the demand first stage is a single equation. Five excluded instruments enter it together with the fixed effects and the exogenous characteristics of the demand system: the two tier-interacted wafer cost shifters, the differentiation instruments in own and rival quality, and the square-root own-quality moment (Section 4.1). The sample is the complete $6 \times 64 = 384$ product–market panel of the demand estimation.

Table B.1 reports the estimated first stage: the coefficients on the five excluded instruments with their standard errors, the included effects, and the strength diagnostics. All five instruments enter positively, and the wafer cost shifter loads about two and a half times as strongly on the premium tier as on the midrange tier, the die-size gradient that motivates the tier interaction (Section 4.1). The joint F statistic on the five excluded instruments is 46.5, the statistic Section 4.1 rounds to about 46.

B.2 Elasticities

Table B.2 reports the average elasticity matrix. The matrix is not symmetric, because the quantity response to a price change depends on the share and price of the product whose price moves. Reading down the NVIDIA midrange column, every other product's demand rises by between 0.51 and 0.71 percent when that price rises by one percent; in the AMD premium column, no product's demand rises by more than 0.21 percent. In every other product's row, the largest cross-price entry sits in the NVIDIA midrange column: a price move at that product moves demand everywhere in the menu.

Table B.1. Demand first-stage estimates

	Estimate (SE)
Excluded instruments	
Wafer cost shifter × premium tier	1.640 (0.217)
Wafer cost shifter × midrange tier	0.644 (0.123)
Differentiation instrument, own quality	7.198 (2.429)
Differentiation instrument, rival quality	6.462 (1.236)
Square root of own quality	1.303 (0.445)
Included effects	
Quality	Yes
Firm fixed effects	Yes
Quarter fixed effects	Yes
Diagnostics	
Joint excluded-instrument F statistic	46.5
Partial R^2 of excluded instruments	0.42
Observations	384

Notes: The table reports the price first stage of the maintained demand specification: a least-squares regression of price, in thousands of 2018 dollars, on the five excluded instruments and the included effects shown. Standard errors, in parentheses, are heteroskedasticity-robust and unclustered. The joint F statistic tests the five excluded instruments; the partial R^2 measures the excluded set's explanatory contribution after the included effects. The sample matches the demand estimation of Section 4.1.

Table B.2. Mean own- and cross-price elasticities

	AMD low-end	AMD midrange	AMD premium	NVIDIA low-end	NVIDIA midrange	NVIDIA premium
AMD low-end	−1.86	0.20	0.06	0.37	0.51	0.12
AMD midrange	0.20	−3.22	0.11	0.35	0.65	0.25
AMD premium	0.14	0.20	−4.79	0.25	0.71	0.64
NVIDIA low-end	0.21	0.21	0.07	−1.96	0.55	0.15
NVIDIA midrange	0.19	0.23	0.13	0.34	−3.35	0.39
NVIDIA premium	0.11	0.17	0.21	0.19	0.60	−4.37

Notes: Entries average product–market elasticities over the 64-quarter sample. An entry is the percent change in the row product's demand when the column product's price rises by one percent. Diagonal entries are own-price elasticities, whose tier means are in Table 1, Panel B, in absolute value; off-diagonal entries are cross-price elasticities, positive because the six products are substitutes.

B.3 Marginal-Cost Recovery

Marginal costs come out of the pricing first-order conditions. For a quarter's menu, the six conditions of equation (3.4) stack into ownership-matrix form,

$$s_t + (\Omega_t \odot \Delta_t)(p_t - \mathrm{MC}_t) = 0, \tag{B.1}$$

where s_t, p_t, and MC_t stack the six market shares, prices, and marginal costs, Δ_t is the share Jacobian, whose (j,k) entry $\partial s_{kt}/\partial p_{jt}$ is the derivative of product k's share with respect to product j's price, Ω_t is the ownership indicator matrix for the six products, whose (j,k) entry is one when products j and k are sold by the same firm and zero otherwise, and $\odot$ is the elementwise product. I invert equation (B.1) market by market at the estimated demand parameters:

$$\widehat{\mathrm{MC}}_t = p_t + \left(\Omega_t \odot \widehat{\Delta}_t\right)^{-1} s_t. \tag{B.2}$$

The recovery uses no cost-side data and no instrument: markups follow from the estimated elasticities through pricing optimality, and marginal cost is the residual.

Table B.3 reports the recovered levels by firm and tier: for each of the six products, the mean price, marginal cost, dollar markup, and Lerner index, the firm-level split behind the tier means of Table 1. NVIDIA carries the larger dollar markup in every tier, \$90 against \$72 at the low end, \$116 against \$83 at the midrange, and \$200 against \$120 at the premium tier, and the Lerner index orders the two firms the same way. The gap widens as the tier rises, so it is not a level shift common to the menu. The pattern coincides with NVIDIA's positive estimated brand effect and its higher average quality within each tier.

B.4 The Implied Marginal Fixed Cost of Quality

The quality game needs the derivative of fixed cost with respect to quality, not its level. The implied series is the left-hand side of the quality first-order condition, equation (3.6), evaluated at the observed menu: the total derivative of the firm's expected variable profit (equation (3.5)) with respect to the quality of one interior tier, computed for the four interior products, each firm's low-end and midrange tiers. If fixed cost varied with quality, optimality would equate this derivative to the marginal fixed cost, so the series is the marginal fixed cost the observed menu implies. The derivative is total because prices respond to every quality change: the direct profit

Table B.3. Prices, marginal costs, and markups by firm and tier

Product	Price ($)	Marginal cost ($)	Markup ($)	Lerner index
AMD low-end	123	51	72	0.588
AMD midrange	234	151	83	0.354
AMD premium	493	374	120	0.243
NVIDIA low-end	142	51	90	0.639
NVIDIA midrange	299	183	116	0.388
NVIDIA premium	658	458	200	0.304

Notes: Entries are means by firm and tier over the estimation sample: the firm-level split behind the tier means of Table 1, Panel C. Price, marginal cost, and markup are means in dollars; the Lerner index is the ratio of the mean markup to the mean price. Each column is rounded from its own unrounded series, so a printed markup can differ by $1 from the difference of the printed price and marginal cost, and the printed Lerner index from the ratio of the printed markup to the printed price.

effect is combined with the induced response of all six equilibrium prices, obtained by applying the implicit function theorem to the pricing first-order conditions. Expectations over the demand and marginal-cost shocks are computed by sparse-grid quadrature for the normal weight (Genz and Keister 1996), with ξ_{jt} and $\varepsilon_{jt}^{\mathrm{MC}}$ taken as mean-zero and normally distributed, independent across products, each scaled by the standard deviation of its estimated residuals.

Figure B.1 plots the four implied series over the full 2005Q1–2020Q4 sample and marks 2011Q1, the start of the maintained window. The early observations make the case for the window visible at a glance: the pre-2011 swings dwarf everything that follows, and from 2011 onward the series settle near zero at the scale of the plot. The decline is about 80-fold, and nearly all of it happens before 2011. A single stationary R&D-cost environment cannot rationalize static marginal benefits this far apart in scale, and the post-2010 window of Section 4.2 is what removes the early scale from the estimation.

C Details for the Static-Equilibrium Approximation

The dynamic model needs the stage game's outcomes at every state the Bellman solver and the simulated paths visit, far too many for exact computation (Section 4.3.2). I therefore solve a large set of exact games, train constrained neural networks on the solutions, and validate the approximation where the dynamic model travels. The construction comes first, in Appendices C.1–C.5, in the order the pipeline runs. The validation follows in two layers: held-out accuracy (Appendix C.6) and coverage of the visited support (Appendix C.7).

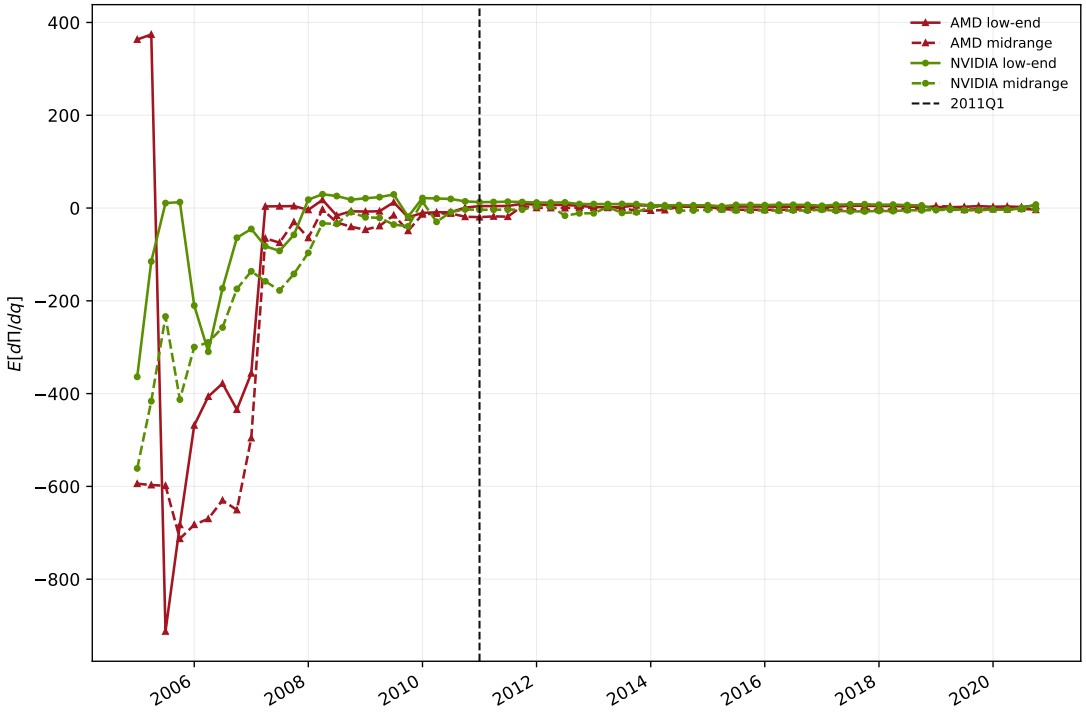


Figure B.1. Implied marginal fixed costs of quality

Notes: The four series are the implied marginal fixed costs of the interior products, each firm's low-end and midrange tiers, over 2005Q1–2020Q4. The vertical marker at 2011Q1 records the start of the maintained window, a maintained sample choice rather than an estimated break date.

C.1 The Exact Stage Game

The stage game takes four static inputs: the two firm frontiers ($q_1^{\max}$, $q_2^{\max}$), the outside-option level (u_0), and the process-innovation index (z). Market size scales quantities, profits, and consumer surplus after the unit-market game (the game at unit market size) is solved. Each firm's premium quality equals its frontier, so the choice variables are the four interior qualities, each firm's low-end and midrange positions. The game is nested. Conditional on a six-product menu and a draw of the demand and marginal-cost shocks, prices solve the ownership-adjusted first-order conditions of equation (B.1). Anticipating that pricing equilibrium, qualities solve the expected-profit conditions of equation (3.6), which include the induced price response of every product in the market; fixed costs are assumed constant in quality (Section 4.2) and drop out. Interior qualities are constrained to lie between zero and the firm's frontier and to preserve the low-end, midrange, premium ordering. The quality layer therefore needs expected profits over the shock distribution.

C.2 Computing One Exact Equilibrium

Shock integration comes first. The exact solver integrates over the product-level demand and marginal-cost shocks, six-vectors with independent components, using the same sparse Kronrod–Patterson rule as the fixed-cost inversion of Appendix B.4, so the training data and the recovered cost primitives share one shock convention.

Equilibrium selection comes second. Both the pricing and the quality games are solved from multiple starts. A start is admissible only if prices are positive, product shares are feasible, qualities are ordered and inside their frontiers, and the first-order conditions hold up to the solver's tolerance. A converged quality-game candidate that carries a detectable profitable one-firm deviation is re-solved from that deviation until none remains, and a candidate that cannot be cleared is dropped. When several admissible equilibria remain, the selection rule keeps the one with the highest total industry variable profit, and the same rule is applied when generating training data, validating held-out states, and fitting the rebate-specific surrogates. The common rule makes equilibrium selection an explicit part of the model rather than an unrecorded numerical accident. Exact solves are what the networks train on, and the design question is where to run them.

C.3 Designing the Training Sample

The base design covers a broad rectangular domain of states and quality menus, deliberately wider than the observed path, with the observed post-2010 states appended as anchors and additional draws concentrated near them. The design also appends a deterministic union of states drawn from the dynamic solvers' sampling geometry, the states on which the value functions are fitted and the boundary regions of the sampled state box, so the static mapping is trained where the dynamic layer reads it. The same design serves the maintained specification and the sensitivity re-estimations of Appendix F.

The sample is generated in two passes. The first draws quality menus and states across the domain, subject to the tier-ordering constraints, solves the exact pricing game at each draw, and fits the pricing network to the shock-integrated solutions; the second solves the quality game across sampled states with the fitted pricing network as the inner map and fits the quality network to the selected equilibria. The split concentrates exact computation where a dense sample is feasible, the pricing surface, rather than inside every quality best-response iterate; its cost is that the quality targets embed the pricing approximation.

For the rebate-specific re-fits of Appendix C.5, after the dynamic model is first estimated the design is augmented: the states visited by the estimated equilibrium's simulated paths are appended, with support extensions near the frontier-ordering constraints and over the policy domains. Throughout, the training bounds are set wide enough that state visitation in the dynamic estimation and the counterfactuals stays inside them, the design fact behind the no-extrapolation statement of Section 4.3.2; Appendix C.7 records the realized rates.

C.4 The Surrogate Networks and Their Economic Constraints

Two neural networks are fitted to these training sets. The quality network Q maps the two frontiers and the two exogenous static states into the four interior qualities:

$$\widehat{q}^{\,\text{int}} = Q\left(q_1^{\max}, q_2^{\max}, u_0, z\right). \tag{C.1}$$

The pricing network $\mathcal{P}$ takes the resulting six-product menu q together with the same two exogenous states and returns the shock-integrated equilibrium outcomes, expected prices, shares, variable profits, and consumer surplus:

$$\left(\widehat{p}, \widehat{s}, \widehat{\pi}, \widehat{CS}\right) = \mathcal{P}(q, u_0, z). \tag{C.2}$$

Market size enters neither mapping: it multiplies the unit-market profits and surplus directly, which makes a fitted pair portable across the market-size process.

Economic constraints are built into the output layers. The quality network returns interior qualities as ordered sigmoid fractions of the firm frontier, so each firm's low-end and midrange qualities respect the tier ordering and stay inside the frontier by construction. The pricing network returns the six inside shares and the outside share through a softmax over seven logits, so shares are nonnegative and sum to one, and its prices pass through a positive output activation.

C.5 Rebate-Specific Re-Fits

Each static environment gets its own fitted pricing-and-quality pair. A rebate changes the stage game itself: it enters the pricing first-order conditions on the covered tiers (Section 6). For each rebate I therefore re-solve the exact games with the rebate applied and fit rebate-specific surrogates on the same state domain, including the visited-state augmentation of Appendix C.3. The re-fit changes the static mapping only; the dynamic state variables, their transitions, and the R&D technology are those of the estimated environment. The R&D credit and the R&D shutdowns act on the dynamic layer and leave the stage game unchanged, so they run on the estimated equilibrium's surrogates. With the re-fits in place, the construction is complete; what remains is its accuracy, checked in two layers.

C.6 Held-Out Accuracy

Accuracy is checked first on held-out fit. The validation sample is a single 20 percent split held out from training and never used for tuning. Table C.1 reports held-out R^2 by output block beside two benchmarks fitted on the same split, the best cubic polynomial and the best LASSO. The four interior qualities are predicted with R^2 between 0.969 and 0.985; on the pricing side, R^2 is about 1.00 for prices, 0.99 for shares, 0.99 for variable profits, and about 1.00 for consumer surplus. The block averages put the networks ahead of both benchmarks, 0.98 against 0.96 and 0.89 on the quality game and 0.99 against 0.81 and 0.74 on the pricing game, with the largest gains on shares and profits, roughly 28 and 29 percentage points of R^2 over the cubic benchmark.

Table C.1. Held-out surrogate accuracy

Output block	Neural-network R^2	Best cubic R^2	Best LASSO R^2
Interior quality, AMD low-end	0.975	0.954	0.891
Interior quality, AMD midrange	0.976	0.957	0.894
Interior quality, NVIDIA low-end	0.985	0.970	0.890
Interior quality, NVIDIA midrange	0.969	0.943	0.876
Interior quality, average	0.977	0.956	0.888
Prices, average	0.998	0.998	0.995
Shares, average	0.994	0.712	0.603
Variable profits, average	0.986	0.691	0.575
Consumer surplus	0.998	0.967	0.949
All pricing outputs, average	0.993	0.809	0.736

Notes: The validation split is a single 20 percent test split held out from training. Group averages are unweighted means across the corresponding network outputs: six products for prices, shares, and variable profits, four for the interior qualities. The cubic benchmark is the best degree-three polynomial fitted on the same split; the LASSO benchmark is the ℓ_1-penalized polynomial whose degree and penalty give the highest mean R^2 within each displayed output block.

C.7 Coverage of the Visited Support

Held-out fit is measured on the training design; the dynamic model decides where the mapping is consumed. For every simulated state, the code therefore records whether the quality or the pricing mapping extrapolates and how close the realized R&D actions and interior qualities sit to their bounds. I choose the bounds of Appendix C.3 to contain visitation, and the recorded rates confirm that they do: on the estimated equilibrium's simulated paths, state visitation stays well inside the sampled state box, and in the estimated equilibrium and every policy counterfactual, neither mapping is ever evaluated outside its training domain. The R&D policy is interior throughout: no simulated intensity reaches either action bound. About 1.3 percent of simulated firm-quarters come within one ten-thousandth of a bound, roughly 0.6 percent near zero investment and 0.7 percent near full reinvestment, but the policy passes smoothly through those points rather than being clipped at them. The realized interior qualities never sit against the feasibility bounds of the firm's problem; the only boundary proximity on record is a handful of evaluations, about 0.01 percent, in which a chosen quality happens to coincide with the rival's frontier, a point the menu crosses freely rather than a constraint. Together these records complete this appendix's approximation audit; Appendix F varies the economic environment, not the numerical layer.

D Details for Dynamic Estimation

This appendix documents the dynamic estimation of Section 4.3. Table D.1 collects the parameters of the dynamic block, which Section 4.3.4 discusses panel by panel.

D.1 Exogenous-State Processes and Long-Run Closure

The three exogenous states, the outside option, the process-innovation index, and market size, evolve deterministically by the laws of motion of Section 3.2.2 (equations 3.11, 3.12, and 3.13); this subsection records the fitted parameters and the closure bounds. The outside-option drift is recovered from the demand system's quarter fixed effects, excluding the 2017Q3–2018Q2 cryptocurrency-mining window (Section 4.3.1). In monetary units, the fitted trend has a slope of \$24.59 per year, with $R^2 = 0.95$ over 60 quarters; converted to utility units, it gives the quarterly drift of 0.11 in Panel C. Figure D.1 plots the recovered series and the fitted trend, including the excluded window.

The closure bounds cap both trends. The process-innovation drift, the time component of the marginal-cost projection, extends the in-sample decline of about 0.05 per quarter until the index reaches the floor of -2.5; the outside option rises by its quarterly drift until it reaches the ceiling of 7.0. On the maintained laws of motion, the ceiling binds from 2026Q2 onward and the floor from 2025Q4 onward, each about five years beyond the 2020Q4 sample end, so neither is reached on the 40-quarter reported paths. Because the policies are permanent in the Bellman problem, the bounds still enter through continuation values (Appendix E.1). Neither bound is a free level: each equals the fitted 2020Q4 value advanced by 20 quarters of the estimated drift, rounded to the nearest 0.5. The in-sample trend disciplines the drift, but the date at which the trend stops is not identified by the sample. Appendix F.2 treats the closure as a sensitivity dimension.

The market-size calibration uses 2011Q1–2020Q4: a log-space AR(1) fit gives a quarterly persistence of 0.85 (standard error 0.09) and a long-run level of 28.43 million desktop units, with $R^2 = 0.72$ over 39 transitions. The initial market size is 36 million units, the rounded 2011–2013 mean.

D.2 Solution Method and Simulation

Firm f's dynamic choice is the R&D intensity $\iota_f(x_t) \in [0, 1]$; the static surrogate supplies the expected variable profit $\mathbb{E}[\pi^*_f(x_t)]$ in millions of dollars, and accounting-equivalent R&D spending

Table D.1. Dynamic parameters and calibrated processes

Panel A. Estimated in the solved dynamic game			
Object	Parameter	Estimate	Standard error
R&D-cost scale	χ	4.616	0.567
Panel B. Innovation arrival and jump distribution, estimated offline			
Object	Parameter	Estimate	Standard error
AMD success-rate parameter	a_{AMD}	−1.056	0.419
NVIDIA success-rate parameter	a_{NVIDIA}	−1.487	0.460
Fishing-out coefficient	φ	0.991	3.278
Mean log jump	m_J	−4.244	0.166
Variance of log jump	v_J	0.660	0.146
Panel C. Fitted exogenous-state processes			
Object	Parameter	Value	
Outside-option quarterly drift	Δ_u	0.11	
Outside-option ceiling	$u_{0,\max}$	7.0	
Process-innovation quarterly drift	c_2	−0.0457	
Process-innovation floor	$z_{\min}$	−2.5	
Market-size persistence	ρ_M	0.85	
Long-run market size	M_∞	28.0 million	
Panel D. Calibrated discount factor			
Object	Parameter	Value	
Quarterly discount factor	β	0.99	

Notes: The R&D-cost scale χ is estimated from a single moment, the average firm-period R&D intensity over the post-2010 sample (Appendix D.3). Its standard error is asymptotic, computed by the delta method from the long-run variance of the targeted moment and the derivative of the simulated moment with respect to the cost scale. A parametric bootstrap that redraws the innovation panel at the point estimate and re-estimates the scale in 256 replications yields a larger standard error of 3.208, reflecting the finite-sample variance of the single targeted moment and the nonlinearity of the mapping from moment to cost scale. Both constructions are conditional on the upstream blocks, holding the demand, marginal-cost, frontier-transition, and surrogate objects at their point estimates; counterfactual quantities in Sections 5 and 6 are likewise reported as point estimates conditional on the same blocks. Success-rate and fishing-out standard errors are from the offline arrival-rate fit to the 2006–2020 innovation history, in which 24 of 120 firm-quarter observations are frontier advances above a small positive threshold; they are computed from the observed-information Hessian of the arrival-rate likelihood, with the fishing-out parameter estimated on the log scale and its natural-scale standard error obtained by the delta method. The log-jump moments are fixed at the sample moments of those advances, and the jump-moment standard errors use sample-moment asymptotics. In Panel C the drifts, the persistence, and the long-run market size are fitted values; the ceiling and the floor are maintained closure bounds (Appendix D.1). The market-size regression gives a persistence of 0.85 and a long-run level of 28.43 million. Panel C reports the rounded values used in the model; for the long-run market size, the model is solved at the rounded value of 28.0 million, about 1.5 percent below the fitted level.

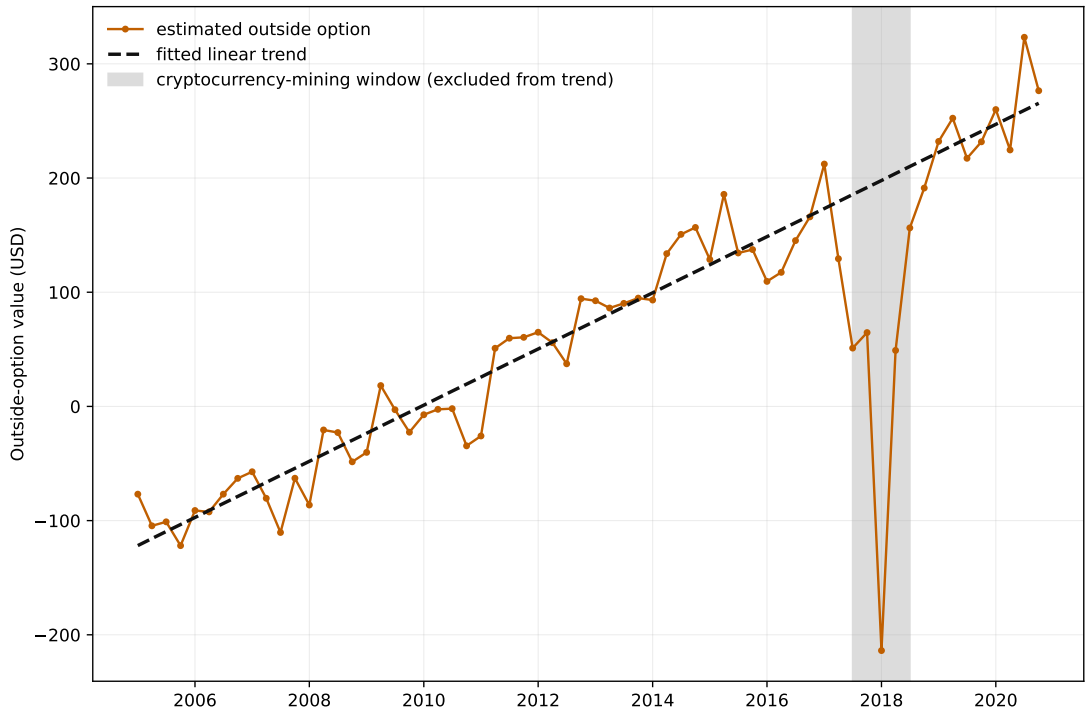


Figure D.1. Recovered outside option and fitted transition

Notes: The recovered outside-option value, plotted in dollars, is obtained from the demand system's quarter fixed effects. The linear trend is fitted after excluding the shaded 2017Q3–2018Q2 cryptocurrency-mining window.

is

$$I_f(x_t) = \iota_f(x_t)\, \max\{\mathbb{E}[\pi_f^*(x_t)],\, \underline{\pi}\}, \tag{D.1}$$

where $\underline{\pi}$ is a numerical profit floor. The floored value enters both the flow cost, whose R&D-cost term in equation (3.8) is $\chi\, \iota_{ft} \max\{\mathbb{E}[\pi_f^*(x_t)], \underline{\pi}\}$, and the innovation hazard, whose argument is $I_f(x_t)$. The floor never binds on the simulated paths. The same conversion is used when comparing simulated R&D with segment-allocated accounting R&D.

The game is solved by damped best-response iteration on the Chebyshev value approximation of Section 4.3.3: each firm's Bellman problem is solved holding the rival's current policy fixed, and iteration continues until both firms' value functions and policies converge. The Bellman equation is enforced on the observed 2011–2020 trajectory states together with a space-filling quasi-random sample covering the state region beyond the observed path.

For each state and action pair, the Bellman expectation enumerates the four advance configurations, neither firm, AMD only, NVIDIA only, or both, and integrates the conditional jump sizes by quadrature, so the expectation is an exact finite sum rather than a simulated average. Forward simulation draws the advance events and jump sizes from the same distributions the solver integrates over, so the solver and the simulator share one transition law. Across values of the cost scale, a given path and period reuse the same success and jump draws, so differences in the simulated moment reflect the parameter rather than the draws.

D.3 The Estimation Moment and the Objective

The cost scale is estimated from one moment, the average of firm-period R&D intensity over the post-2010 sample; the scalar objective is the squared gap between the sample moment and its simulated counterpart:

$$\widehat{m}_I = \frac{1}{2T} \sum_{f=1}^{2} \sum_{t=1}^{T} \iota_{ft}^{\text{data}}, \qquad Q(\chi) = \left[\widehat{m}_I - E_\chi\left(\frac{1}{2T} \sum_{f,t} \iota_{ft} \right) \right]^2. \tag{D.2}$$

The targeted mean intensity is 0.20, and the estimator matches it to numerical precision. The exact match is mechanical for an exactly identified estimator, so the fit evidence rests on the by-firm paths of Appendix D.4, dimensions the estimation does not use, rather than on this moment.

Figure D.2 plots the scalar objective against the cost scale: it falls to a single interior minimum at $\widehat{\chi} = 4.62$, the shape implied by a simulated mean intensity that declines in the scale (the

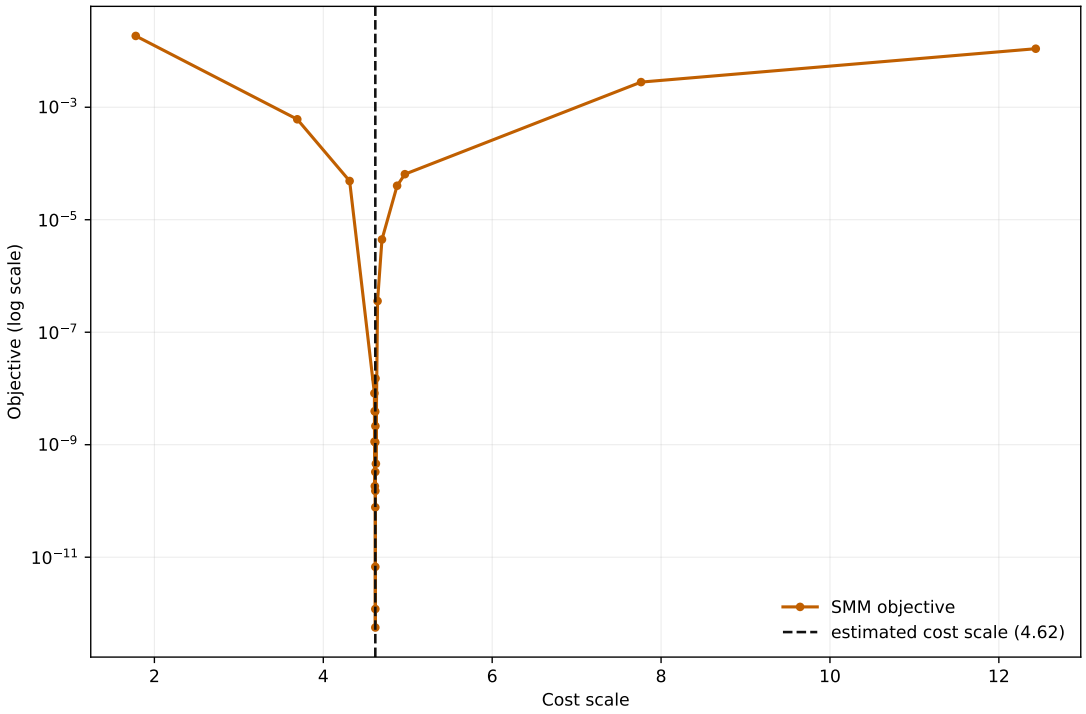


Figure D.2. The estimation objective for the dynamic R&D-cost scale

Notes: The curve is the estimation objective of equation (D.2), the squared gap between the simulated and the observed mean R&D intensity, plotted on a log scale against the cost scale. The dashed vertical line marks the estimate, $\widehat{\chi} = 4.62$.

comparative-statics argument of Section 3.2.2).

D.4 Model Fit Beyond the Estimation Moment

The dynamic estimation matches a single moment, the pooled mean R&D intensity, so the by-firm paths are untargeted. They nonetheless track the data: the observed series for market share, frontier quality, and R&D intensity lie mostly inside the simulated 10th–90th percentile bands.

Figure D.3 overlays the observed paths on the simulated mean and median and the 10th–90th percentile band, by firm and object. The bands represent innovation and jump randomness at the point estimate rather than parameter uncertainty. Across the six panels, the observed series lie inside the band in 88 percent of firm-quarter observations. The largest deviations are in R&D intensity and are mirror-imaged: the single-scale model overstates AMD's average intensity and understates NVIDIA's. That pattern motivates the hazard-window, hazard-family, and closure sensitivities in Appendix F.2.

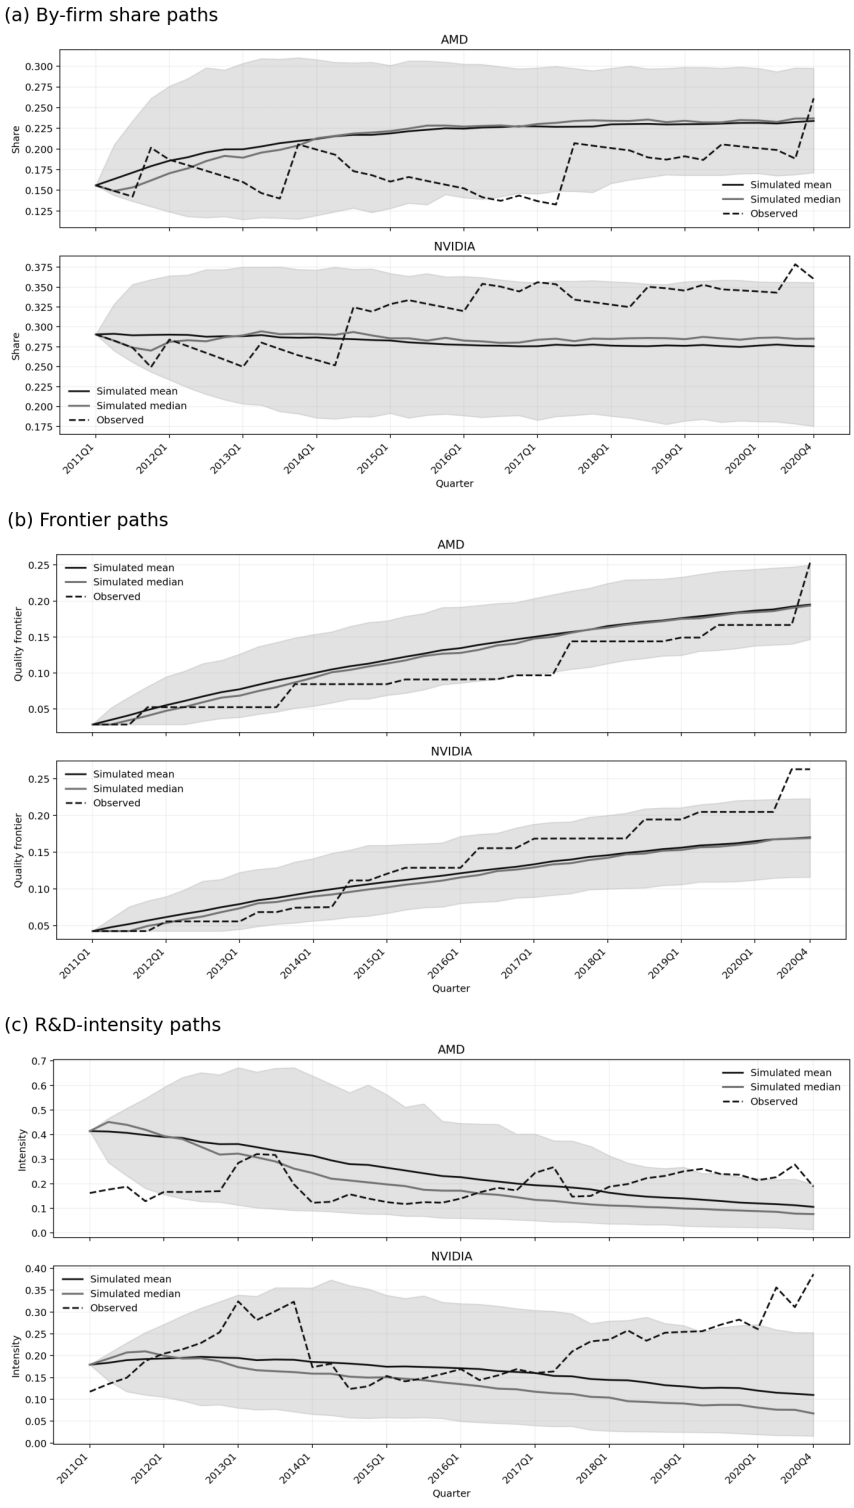


Figure D.3. Observed paths and simulated point-estimate bands

Notes: Dashed paths are observed. Solid paths show the simulated mean and median, with the shaded 10th–90th percentile band across 256 simulations at the point estimate. The bands therefore reflect innovation and jump randomness rather than parameter uncertainty. Across the six panels, 88 percent of observed firm-quarter values lie inside the bands.

E Details for Counterfactuals and Additional Results

This appendix collects the mechanics behind the counterfactuals of Sections 5 and 6 and the product- and firm-level evidence behind their mechanism claims. Every experiment changes one economic primitive at a time, re-solves the relevant equilibrium, and compares the resulting paths against the estimated equilibrium under a common initial state and common random numbers. The shutdown experiments ask what a firm's own R&D buys; the policy experiments ask what an additional public dollar buys. The estimated equilibrium throughout is the model solved at the estimated primitives, whose R&D-cost environment already embeds the historical support.

E.1 Common Counterfactual Protocol

Every maintained counterfactual is permanent in the solved game: agents form continuation values under the counterfactual law of motion rather than treating the intervention as a ten-year surprise, so the ten-year window is a reporting window, and a flow measured in the first 40 quarters can respond to continuation values beyond quarter 40. Each regime also starts every path from the observed 2011Q1 state and reuses the innovation and jump-size draws of the estimated equilibrium across 256 simulated paths, so common random numbers damp simulation noise in the difference without changing the stochastic model. Beyond the horizon, every regime keeps the stationary exogenous-state closure of estimation, with no terminal scrap-value adjustment.

The fixed initial state prevents a policy from receiving credit for a favorable draw before it begins, and the shared closure prevents comparisons from mixing different beliefs about the outside option, process innovation, or market size. Appendix F.2 varies this closure directly.

E.2 Strategic Adaptation of the Rival

The fixed-rival convention is useful for measurement, but the return measured under it need not equal the long-run value of an active innovator when the rival understands that the shutdown is permanent. The constrained re-equilibrium fixes the shutdown firm's R&D at zero and lets the rival re-optimize its dynamic policy against a non-innovating competitor. Writing r_f for a fixed-rival return of Section 5 and $\widetilde{r}_f$ for its rival-reoptimized counterpart, the strategic-adaptation component is their difference, $r_f - \widetilde{r}_f$. The sign alone does not reveal whether the rival raises or lowers its R&D, because re-optimization also changes the state and product-market paths. Table E.1 carries the by-firm comparison.

Table E.1. Strategic adaptation under a unilateral R&D shutdown

Firm	Return object	Fixed-rival policy	Rival re-optimized	Strategic gap
AMD	Private	3.49	3.48	+0.01
AMD	Social	6.07	6.18	−0.11
NVIDIA	Private	3.71	3.74	−0.03
NVIDIA	Social	6.97	6.64	+0.33

Notes: The rival-reoptimized game holds the shutdown firm's R&D at zero permanently and re-solves the rival's dynamic problem. The strategic gap is the fixed-rival return minus the rival-reoptimized return.

E.3 Policy Implementation and Return Accounting

The policy experiments add a new layer of support on top of the estimated equilibrium; Section 6.1 defines the three instruments and their rates. For a demand rebate of b_d dollars per unit on the covered set $\mathcal{K}_d$, the consumer pays

$$p_{kt}^{\text{net}} = p_{kt} - b_d \, \mathbf{1}\{k \in \mathcal{K}_d\}, \tag{E.1}$$

while the producer receives the posted price p_{kt}. A covered product's consumer-net-price change is therefore the posted-price change minus b_d, so pass-through is not the posted-price response alone; for uncovered products, net and posted price changes coincide.

The realized outlays the three rates generate are equilibrium outcomes, not design constants, and they are the one fiscal denominator the per-public-dollar objects share: realized discounted direct outlay on the new layer. For the R&D credit, the outlay is the rate times the discounted accounting R&D generated under the policy; for a rebate, it is the dollar amount times discounted covered units. Their range and budget drift appear in Table 4.

The shutdown and policy families answer different questions: a shutdown removes a firm's own R&D program, while a policy changes the equilibrium environment and lets both firms respond. Frontier advancement per dollar of induced R&D is reported only for the R&D credit, the one instrument that induces private R&D on net; dividing the rebates' frontier changes by an R&D withdrawal or a near-zero net change would not produce a return.

E.4 Static and Dynamic Components: Construction

The static welfare effect and the dynamic state response of Table 6 come from a full counterfactual run and a static replay along the estimated-equilibrium state path, under the same rebate schedule,

Table E.2. Percent change in 40-period average markup by product

Product	Premium rebate (%)	Non-premium rebate (%)
NVIDIA low-end	+0.6	−18.4
NVIDIA midrange	−15.1	−28.9
NVIDIA premium	−62.9	−12.1
AMD low-end	+2.1	−9.7
AMD midrange	−25.7	−38.0
AMD premium	−72.2	−10.6

Notes: Entries are percent changes in the 40-period average markup relative to the estimated equilibrium; the average is taken over paths within each period and then over periods. Markup is consumer price minus predicted marginal cost, the operating margin net of the rebate transfer.

initial state, and random draws. The static replay changes consumer net prices, quantities, posted prices, quality choices within the static menu, marginal costs, and markups, but those changes never accumulate into a different frontier; the full run lets R&D and the frontier states respond. Because static quality is chosen as a function of the frontier state, holding the frontier does not fix every product attribute numerically; it holds the dynamic state transition at the estimated-equilibrium path. Per public dollar, the static effect divides by the replay's own realized outlay and the two components sum to the total by construction. Appendix E.5 reports the product-level markup responses under the two rebates.

E.5 Product-Level Mechanisms Under the Rebates

The two rebates reach welfare through the product menu, and this subsection reports the product-level objects behind the Section 6 mechanism discussion. The exhibits use two aggregations: 40-period averages for the typical menu response, and discounted present values for profit incidence, reported separately in Appendix E.6. Menu responses average each product's per-period difference from the estimated equilibrium over the simulated paths within each period and then over the 40 periods.

Table E.2 reports the percent change in each product's 40-period average markup, the markup table that the mechanism discussion of Section 6 draws on.

Table E.3. Discounted operating-profit incidence by firm

Instrument	AMD (USDb)	NVIDIA (USDb)	Total (USDb)
R&D credit	+0.77	+0.99	+1.76
Premium rebate	−2.22	−3.41	−5.63
Non-premium rebate	+0.47	−1.36	−0.89

Notes: Entries are 40-quarter discounted operating profit net of demand-rebate transfers, relative to the estimated equilibrium, in 2018 USD billions. Entries are the producer-profit components of the welfare identity of Table 4, not transfer-inclusive private payoffs and not changes after subtracting R&D.

E.6 Profit Incidence Across Firms and Products

Operating-profit incidence is measured net of demand-rebate transfers: the fiscal payment is removed from producer receipts before comparing profit with the estimated equilibrium, so the rebate dollar never sits in both the profit numerator and the fiscal denominator. Table E.3 reports the firm-level present values. The R&D credit raises operating profit at both firms, +$0.77 billion at AMD and +$0.99 billion at NVIDIA. The premium rebate lowers the ex-transfer component at both firms, by $2.22 billion at AMD and $3.41 billion at NVIDIA; the non-premium rebate lowers NVIDIA's by $1.36 billion while AMD's rises by $0.47 billion. These entries are not the transfer-inclusive payoffs the firms optimize. Adding the subsidy receipts back makes private payoffs rise at both firms under both rebates: by $1.54 billion at AMD and $2.02 billion at NVIDIA under the premium rebate, and by $2.42 billion and $1.22 billion under the non-premium rebate. The negative entries isolate compression of market-generated producer surplus, the discounted counterpart of Table E.2.

The firm totals conceal within-firm reallocation. Table E.4 carries the product-level breakdown. Under the premium rebate every product loses ex-transfer operating profit, but the coverage set does not order the losses: NVIDIA's uncovered midrange loses the most, $1.5 billion, more than either covered premium product, while AMD's uncovered low-end loses the least, $0.4 billion. Under the non-premium rebate the only material product-level gain, +$0.6 billion, arrives at AMD's covered low-end, yet NVIDIA's equally covered midrange takes the deepest loss, $1.1 billion, and its covered low-end comes out approximately flat. The pattern is a direct expression of equilibrium repricing across the menu: the repricing a rebate sets off decides each product's ex-transfer profit, gain or loss, with little regard for the coverage set.

Table E.4. Discounted product-level operating-profit incidence under the two rebates

Product	Premium rebate (USDb, net of transfer)	Non-premium rebate (USDb, net of transfer)
NVIDIA low-end	−0.8	0.0
NVIDIA midrange	−1.5	−1.1
NVIDIA premium	−1.1	−0.2
AMD low-end	−0.4	+0.6
AMD midrange	−0.6	0.0
AMD premium	−1.1	−0.1

Notes: Entries are 40-quarter discounted operating profit net of demand-rebate transfers, relative to the estimated equilibrium, in 2018 USD billions. Rows split the firm totals of Table E.3 by product; conventions are those of that table. Firm totals may differ from the sum of rounded product entries.

Table E.5. Induced R&D and terminal frontier by firm identity

Instrument	ΔI_{AMD} (USDb PV)	ΔI_{NV} (USDb PV)	ΔI (USDb PV)	$\Delta q^{\max}_{AMD}$ (%)	$\Delta q^{\max}_{NV}$ (%)	$\Delta q^{\max}_{joint}$ (%)
R&D credit	+0.36	+0.46	+0.81	+4.9	+4.8	+5.1
Premium rebate	0.00	−0.25	−0.24	−0.9	−3.0	−2.7
Non-premium rebate	+0.19	−0.17	+0.02	+2.8	−2.2	+0.5

Notes: R&D entries are 40-quarter discounted changes in accounting R&D relative to the estimated equilibrium, in 2018 USD billions; totals may differ from the sum of rounded firm entries. Frontier entries are the percent change in the expected period-39 frontier relative to that equilibrium, with the expectation across simulated paths. The joint frontier is the AMD/NVIDIA maximum, taken within a simulation path before averaging.

E.7 Induced R&D and Frontier Allocation

Table E.5 reports induced R&D and the terminal frontier by firm identity. The table is an identity audit rather than a frontier-role allocation: the frontier leader is a state-contingent role, switching across paths and periods, whereas AMD and NVIDIA are fixed firm names. The joint-frontier column is the percent change in the expected within-path maximum, so it is neither the sum nor the average of the two firm columns. The credit's induced R&D splits nearly evenly across firms, +$0.36 billion at AMD and +$0.46 billion at NVIDIA. The two rebates land on the firms asymmetrically: NVIDIA cuts its own R&D under both, −$0.25 billion under the premium rebate and −$0.17 billion under the non-premium rebate, while AMD's is essentially unchanged under the premium rebate and rises +$0.19 billion under the non-premium rebate. Under the premium rebate NVIDIA's cut is nearly the whole net change, so total R&D and the expected joint frontier both fall; under the non-premium rebate the two firms' moves nearly cancel, leaving total R&D essentially unchanged while the expected joint frontier still edges up.

The role allocation is a different accounting partition: when a policy changes R&D at a given

Table E.6. Allocation of induced R&D across frontier roles

Instrument	Leader ΔI (USDb PV)	Follower ΔI (USDb PV)	Total (USDb PV)
R&D credit	+0.39	+0.42	+0.81
Premium rebate	−0.18	−0.06	−0.24
Non-premium rebate	−0.01	+0.03	+0.02

Notes: Entries are signed 40-quarter discounted changes in role-attributed accounting R&D relative to the estimated equilibrium, in 2018 USD billions. Roles are assigned on the estimated-equilibrium path as in equation (E.2): within each path and period the firm holding the higher frontier quality is the leader and the other the follower, so a policy-induced role switch cannot relabel a dollar. Each instrument's total is the induced-R&D column of Table E.5, which partitions the same total by firm identity; Table 8 reports the same role allocation as a percent change in role-attributed R&D.

path-period state, is that change attributed to the firm classified as the leader or the follower on the estimated-equilibrium path? For role $r \in \{L, F\}$,

$$\Delta I_r = \mathrm{E} \sum_{t=0}^{39} \beta^t \sum_f \mathbf{1}\{f = r_t\} \, \Delta I_{ft}, \tag{E.2}$$

where ΔI_{ft} is the change in firm f's accounting R&D relative to the estimated equilibrium and r_t assigns the firm with the higher frontier on the estimated-equilibrium path as leader and the other as follower. Assigning roles on that path prevents a policy-induced role switch from mechanically relabeling the same dollar.

Table E.6 carries the resulting signed dollar levels by role, and Table 8 in Section 6 reports the same allocation as a percent change in role-attributed R&D. The credit raises R&D almost equally across the two roles, +\$0.39 billion at the leader and +\$0.42 billion at the follower. The premium rebate's cut lands more heavily on the leader role, \$0.18 billion of the \$0.24 billion withdrawn against \$0.06 billion on the follower, an incidence of frontier position rather than a difference in how leaders and followers respond. Under the non-premium rebate the role-level changes are small in dollars, −\$0.01 billion at the leader and +\$0.03 billion at the follower, and those small role-level changes mask a larger, offsetting firm-level reallocation, AMD raising its R&D and NVIDIA cutting its own, shown in Table E.5. These entries describe incidence by role in the estimated equilibrium, not a behavioral effect of becoming a leader or follower.

F Robustness and Sensitivity Analysis

Each check changes one named primitive of the maintained specification and re-estimates the block that owns it: Appendix F.1 carries the demand variants through the recovered marginal costs and tier markups, Appendix F.2 re-estimates the R&D-cost scale under dynamic variants, and Appendix F.3 re-scores the R&D credit at nearby rates. The standard throughout is transparent stability rather than universal invariance: a result is robust when reasonable changes in a primitive preserve its sign and its broad ranking.

F.1 Demand and Static-Supply Specifications

The maintained demand specification combines the excluded-instrument set of Section 4.1, quarter fixed effects, one random coefficient on price, and a time-varying market size. On the static-supply side, prices and qualities are the weighted pooled series, and the recovered marginal costs are projected on the rich cost specification of Section 4.2.

Table F.1 collects the variants. Alternative 1 re-estimates the maintained specification without the square-root own-quality moment, on the identical sample of 384 observations, so it isolates the sensitivity of the estimates to the moment set. Alternative 2 removes the cryptocurrency-demand episode from the estimation sample. Alternative 3 replaces the time-varying market-size path with a constant, leaving the outside option untouched. Alternative 4 changes the static-supply environment as well: it rebuilds the product panel, and with it both demand and supply.

Variants are judged on the downstream objects they feed rather than on coefficient stability alone. The table therefore carries each variant through the recovered marginal costs to the tier markups and reports the tier Lerner indices.

F.2 Dynamic Specification and Estimation

The maintained dynamic block pairs the arrival-rate hazard and log-concave R&D exposure with the 2011Q1–2020Q4 estimation window, a quarterly discount factor of 0.99, and the exogenous-state closure of Appendix D.1; hazard and jump parameters are estimated offline, and the R&D-cost scale χ then matches average R&D intensity.

Table F.2 is organized by the primitive varied. The curvature and hazard-family rows change the mapping from R&D and frontier position to innovation arrival. The hazard-window row changes which episodes identify the cost scale: the maintained hazard is fit to the full 2006–2020 innovation

Table F.1. Demand and static-supply robustness

Specification	Price coef. (SE)	Quality coef. (SE)	Mean own elasticity $\lvert\varepsilon_{jj}\rvert$	Lerner index low-end / midrange / premium
Maintained specification	−17.50 (5.21)	60.73 (11.76)	3.26	0.615 / 0.373 / 0.278
Alternative 1	−25.08 (7.18)	77.68 (16.85)	3.29	0.477 / 0.341 / 0.408
Alternative 2	−17.11 (4.59)	52.83 (8.79)	3.00	0.641 / 0.402 / 0.306
Alternative 3	−18.41 (4.72)	56.30 (8.32)	3.04	0.603 / 0.383 / 0.322
Alternative 4	−15.23 (5.13)	52.94 (10.60)	2.88	0.738 / 0.429 / 0.293

Notes: Each alternative re-estimates the demand system of Table 1 after one change to the maintained specification: Alternative 1 drops the square-root own-quality moment from the moment set, on the identical sample of 384 observations; Alternative 2 excludes the 2017Q3–2018Q2 cryptocurrency-demand quarters; Alternative 3 replaces the time-varying market size with a constant; Alternative 4 rebuilds the product panel from median rather than weighted prices and qualities, and with it both demand and supply. Standard errors, in parentheses, are heteroskedasticity-robust and unclustered. Mean own elasticity is the mean absolute own-price elasticity across the six products; Lerner indices are tier-mean markup over tier-mean price, from posted price and recovered marginal cost.

history, and this row restricts the fit to the post-2010 history (80 firm-quarters, 19 events). The plateau rows change continuation values through the levels at which the exogenous states settle.

The in-sample drifts of the outside option and the process-innovation state are data inputs; the levels at which those drifts plateau are continuation-value assumptions. The maintained closure provides about five years of headroom after the sample. The horizon is anchored on the cost side: reported wafer prices step up sharply at the 2-nanometer generation, and cost per transistor, the object the process-innovation state proxies, is widely reported as flattening around the end of 2025, so the sample-era improvement trend stops being a defensible expectation past that date. The sensitivity is one-sided, toward later plateaus only: realized 2021–2023 improvement rules out an earlier one.

F.3 The R&D Credit at Nearby Rates

The R&D credit enters the policy comparison at a budget-matched rate: the credit rate is set so that the credit's planned outlay equals the common program budget. Table F.3 asks whether the credit's returns are special to that rate, reporting the maintained rate alongside two nearby rounded rates. Additionality is about 0.18 at all three rates. Surplus per public dollar is 0.82 at a rate of 0.450, 0.80 at 0.500, and 0.82 at the maintained rate of 0.545, and frontier advancement holds between 1.11 and 1.14 percent per billion dollars throughout. The credit's positive verdict is locally stable in the rate, not a knife-edge confined to the maintained rate.

Table F.2. Sensitivity of the R&D-cost scale

Specification	$\widehat{\chi}$ (SE)
Maintained specification	4.616 (0.567)
Exposure curvature 0.25	4.567 (0.629)
Exposure curvature 1.00	4.418 (0.471)
Poisson arrival hazard	4.504 (0.417)
Post-2010 hazard window	2.986 (0.351)
Plateau seven years after the sample	4.879 (1.514)
Plateau ten years after the sample	4.875 (0.588)

Notes: Each row re-estimates the R&D-cost scale χ from the average R&D-intensity moment after changing one primitive of the maintained specification, which pairs the arrival-rate hazard with log-concave R&D exposure at curvature 0.50, fits the hazard to the 2006–2020 innovation history, and lets the exogenous states plateau five years after the 2020Q4 sample end. The curvature rows vary the exposure curvature within the maintained hazard family; the Poisson row changes the hazard family at the matched curvature; the hazard-window row fits the hazard to the post-2010 innovation history alone (80 firm-quarters, 19 events); and the plateau rows defer the exogenous-state plateau to seven and ten years after the sample end. Parentheses report asymptotic standard errors, computed by the delta method from the long-run variance of the targeted moment and the derivative of the simulated moment with respect to the scale; the parametric bootstrap standard error of the maintained specification, 3.208, is reported in the notes to Table D.1.

Table F.3. The R&D credit at nearby rates

Credit rate	Realized outlay ($B)	Additionality	Surplus per public dollar	Frontier % per $1B
0.450	3.712	0.178	0.818	1.128
0.500	4.167	0.179	0.803	1.144
0.545†	4.581	0.178	0.824	1.107

Notes: The credit rate τ is the reduction in the R&D-cost scale: a booked R&D dollar costs the firm $\chi - \tau$. Realized outlay is the discounted direct public outlay over the 40-quarter horizon; frontier % per $1B is the percent change in the expected terminal joint frontier relative to the estimated equilibrium, per $1B of realized outlay. The rate marked † is the maintained budget-matched rate; the other two are nearby rounded rates.

G Notation

This appendix collects the symbols used in the model and the appendices, with the unit conventions behind the reported tables. Table G.1 lists the static demand and supply objects first, then the dynamic states and R&D objects, and last the policy and welfare objects; a product j is a firm–tier alternative, and the market index t is quarterly.

Table G.1. Notation for the model and the appendices

Object	Definition
Static demand and supply	
p_{jt}	Posted price of product j in thousands of 2018 dollars
q_{jt}	PassMark G3D benchmark index
s_{jt}	Inside-product market share
$\mathcal{J}_t$, $\mathcal{J}_f$, $\mathcal{J}_f^{\text{int}}$	The quarter's six inside products; firm f's three-tier menu; its two interior tiers
γ; $\bar{\alpha}$, σ_{price}; $\psi_{f(j)}$	Demand parameters: quality coefficient; mean price coefficient and price-sensitivity dispersion; firm fixed effect
ξ_{jt}, $\varepsilon_{jt}^{\text{MC}}$	Unobserved demand and marginal-cost shocks
Δ_t	Share Jacobian, (j,k) entry $\partial s_{kt}/\partial p_{jt}$ (Δ_u is the outside-option drift; a bare Δ marks a change from the estimated equilibrium)
Ω_t	Ownership matrix, (j,k) entry one when j and k share a firm
MC_{jt}, $\widehat{\text{MC}}_t$	Marginal production cost, in thousands of 2018 dollars; the hat marks the costs recovered from the pricing first-order conditions
c_0–c_5	Marginal-cost projection coefficients (equation (3.7), Table 2); $c_1 + c_2\,t$ forms z_t
π_{jt}	Product variable profit, $M_t\, s_{jt}\, (p_{jt} - \text{MC}_{jt})$, in 2018 USD millions
Dynamic states and R&D	
x_t	Dynamic state vector, $x_t = (q_{1t}^{\max}, q_{2t}^{\max}, u_{0t}, z_t, M_t)$
$q_{ft}^{\max}$	Firm f's quality frontier
z_t, $z_{\min}$	Process-innovation state, $z_t = c_1 + c_2 t$, and its long-run floor
u_{0t}, $u_{0,\max}$	Outside-option level and its ceiling
M_t; ρ_M, M_∞	Addressable desktop market size, in millions of units; its persistence and long-run level
χ	R&D-cost scale
ι_{ft}	Firm f's R&D intensity: spending relative to expected equilibrium stage-game profit $\mathbb{E}[\pi_f^*(x_t)]$
I_{ft}	Accounting GPU-segment R&D, in 2018 USD millions; the hazard exposure uses $I_{ft}/100$, in hundreds of millions
h_f	Innovation arrival probability, $r_{ft}/(1 + r_{ft})$ with arrival rate r_{ft}; distinct from the horizon index h, the role index r, and the return r_f of Appendix E.2
a_f, φ	Firm-specific success-rate intercept and common fishing-out parameter
κ	Curvature of the log-concave R&D exposure in the innovation hazard
m_J, v_J	Log-mean and log-variance of the frontier jump-size distribution
β	Quarterly discount factor
Policy and welfare	
W	Within-market surplus, $W = \text{CS} + \Pi - I$: consumer surplus plus producer operating profit net of rebate transfers, less R&D outlay
B	Common planned present-value budget of the three instruments
G	Realized discounted public outlay of a policy
τ; b_d	R&D-credit rate; rebate per covered unit, with d indexing the rebate design

Notes: Prices and marginal costs enter the demand system and the pricing game in thousands of 2018 dollars. Converting a margin back to dollars and multiplying by unit sales in millions of units yields profit in 2018 USD millions, the unit carried by the dynamic model.