# 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 product line markets, using a dynamic oligopoly model with endogenous pricing, product-line adjustment, and R&D. I apply the model to the desktop GPU market, a duopoly between NVIDIA and AMD. In this market, a dollar of the industry's R&D generates about $9.0 of within-market surplus net of that dollar, the firms capturing $3.0 and consumers the rest. Despite the wide wedge, the return to public support is far more modest than the return to the firms' own R&D. I simulate three stylized interventions: an R&D credit and two product rebates, one at the cutting edge, one on the mass market. The R&D credit is the only instrument that induces a material increase in private R&D, yet it returns only $0.51 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 returns $0.23 per public dollar 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, returning $1.09 per public dollar. These results suggest that innovation policy in product line markets is 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 the return 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 pricing, product-line adjustment, 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 the violation of the stable unit treatment value assumption (SUTVA): subsidizing one firm induces its rival to re-optimize prices and qualities, so the rival is not an untreated control.

[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 16 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 competed throughout, with no entry or exit, and the market was left laissez-faire: no major industrial policy targeted GPUs over the sample. This 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 $9.0 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 $9.0 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 translating it into surplus through 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 research and experimentation (R&E) tax credit the United States has offered in some form since 1981.[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 no single instrument dominates. Despite the large social-private wedge, the return to public support is far more modest than the return to the firms' own R&D. The R&D credit is the only instrument that induces a material increase in private R&D: it crowds in about $0.13 of private R&D per public dollar, returns about $0.51 of within-market surplus per public dollar, and lifts the end-of-horizon technology frontier by about eight-tenths of a percent for every $1 billion of outlay. The premium rebate backfires on innovation: it crowds out R&D on net,

[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 $10.0 of gross surplus. The growth-accounting ratios cited in this section are gross, so the $10.0 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.

returns about $0.23 of surplus per public dollar, and pushes the frontier back by about three-tenths of a percent per $1 billion spent. The non-premium rebate induces essentially no net R&D, yet it returns the most surplus of the three, about $1.09 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.

At first glance, these numbers look like a contradiction. If a dollar of the firms' R&D generates $9.0 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 on average, not how strong the firms' incentive to invest is at the margin. A dollar of R&D already returns the firms about $4.0 of gross profit. With that strong private incentive, they have pushed research spending to where the marginal return is low. Since an R&D credit pays on the whole program, most of its outlay 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 almost the whole line. Whether the war also kills innovation depends on where it is fought, 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, knowledge gained during frontier research carries down to the production of the low-end and midrange tiers, saving cost on the units 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, the added volume repaying the margins given up, while the premium rebate destroys the very prize the race is run for: R&D falls, mostly at the firm holding the technology frontier, hence the frontier retreats, and the dynamic response offsets much of the static gain. The R&D credit, in contrast, never reshapes the product market 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 product line.

It is extremely important to emphasize that the magnitudes of my empirical findings carry no external validity. The returns, the welfare signs, and the instrument ranking are specific to

[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. 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.

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 social return and the social-private gap are both large in this industry, but neither means a public dollar returns what a dollar of firm R&D does. 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 but misses the equilibrium response along other strategic dimensions, so it can misread the surplus a public dollar buys. In my policy experiments, it scores the two rebates nearly alike despite the wide divergence in what they deliver, and it puts the R&D credit first while the top instrument on surplus is the non-premium rebate, which additionality reads as an innovation nonevent. 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 abundant, 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] Almost 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: they design their chips and contract with outside foundries such as TSMC to fabricate them. 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.

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. Each tier's price and quality average over the products the firm offers in that tier in the quarter. 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.[17] 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 each desktop CPU sale counts as a potential buyer. I use those sales as the proxy for market size.

Table A.1 summarizes the product line, the market aggregates, and the measurement and

[15] If this page is 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.

[17] Both firms report research spending as a single company-wide line; a segment split is not required by law, and neither firm provides one. Modeling competition in the firms' other markets, where each faces different rivals, is beyond the scope of this paper (see Rhodes et al. 2026), so I adopt the proportional allocation as a simplifying assumption. The GPU segment averages roughly 60 percent of NVIDIA's revenue over the sample and roughly a fifth of AMD's.

sources of each series. Over the sample, benchmark quality at the frontier rose to about 100 times its 2005 level, while the premium-tier price roughly doubled. Figure A.1 plots the two series year by year. The tiers separate cleanly on both 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, $1.2 billion against $12.9 billion in an average year; Figure A.2 plots the by-firm series. Across the full sample, that spending totals about $19 billion. Appendix A provides a fuller account of the market and these series.

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. Section 5 then prices 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.[18]

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'

[18] 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.

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 period's R&D advanced their frontiers. Firms then simultaneously choose 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). The within-period game follows: 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 in detail 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). Each firm maximizes the expected present value of its profits, discounted at the per-quarter factor $\beta$. 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.[19] 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, $\xi_{jt}$ is unobserved quality, $u_{0t}$ is the outside-option level, and $\varepsilon_{ijt}$ 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 $\Phi$ the standard-normal distribution; the outside share replaces the numerator with $\exp(u_{0t})$.

[19] The demand system is static and abstracts from forward-looking behavior. A full dynamic-consumer model is beyond the scope of this paper's research question. For the modeling of durable-goods demand, 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. 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.[20] After qualities are committed, the demand and marginal-cost shocks, $\xi_{jt}$ and $\varepsilon_{jt}^{\mathrm{MC}}$, are realized, and the pricing step begins. The game is solved by backward induction.

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 $\mathrm{MC}_{kt}$ 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 the frontier 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. Equating the premium tier with the frontier is a simplifying assumption.[21] 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.

[20] 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.

[21] 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.

Formally, each firm chooses its two interior qualities, 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. 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 developed at the cutting edge carry down the firm's own lineup.[22] It gives the firm an additional

[22] Each new GPU generation is typically built on a new architecture designed for its flagship products; lower-end GPUs of the same generation share that architecture and inherit the new design techniques as a matter of course. The practice is documented on both sides of the duopoly, in the firms' architecture white papers.

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 first-order condition for quality; Section 4.2 carries the justification and the confirming evidence.

### 3.2.2 Dynamic Game

Over time, firms compete by choosing R&D spending. At each period's opening, 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 $x_t = (q_{1t}^{\max}, q_{2t}^{\max}, u_{0t}, z_t, M_t)$ 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 $\chi$ 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 lowers the advance probability, the restriction $\varphi$ 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

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

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 $\iota_{ft}$ 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. This monotonicity lets a single intensity moment identify $\chi$ (Section 4.3). The scale is also the parameter the R&D-credit counterfactual acts on: a constant R&D tax credit is, in effect, a downward shift in $\chi$.

The private optimality condition imposes no restriction on 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 lie largely outside the firms' control and evolve exogenously, each on a deterministic path 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,

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

where the per-period drift $\Delta_u$ 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 $M_\infty$ 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.

# 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. Readers primarily interested in the policy results can skip ahead to Sections 5 and 6 without loss of continuity.

[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.

### 4.1 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 are realized, while prices are set afterward. Price therefore co-moves with the demand unobservable $\xi_{jt}$.

I address the endogeneity issue with price instruments. The lead instrument is a shift-share built on the semiconductor-wafer producer-price index, lagged one quarter,[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 with 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 is realized.[26] The demand first stage is strong, with a joint $F$ statistic of about 49. Appendix B.1 reports the full first-stage regression table.

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 1.94 at the low-end tier, 3.29 at the midrange, and 4.41 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. The premium tier, priced highest, faces the most elastic demand even though premium buyers are the least price-sensitive: the price-level effect outweighs their lower sensitivity. Table B.2 reports the

[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 $\xi_{jt}$ is drawn, no part of the current shock is forecastable when the menu is set, so these functions are orthogonal to $\xi_{jt}$, 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.

full elasticity matrix. The estimate of the random coefficient on price ($\sigma_{\text{price}}$) 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.

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 \$165. The markup ratio falls across tiers, from 0.61 to 0.37 to 0.29, 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. I report sensitivity analysis across moment choice, sample range, and market-size definition in Table F.1.

The outside-option value 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 73 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. Table 2 reports the 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: 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 coefficients for the frontier-relative terms 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

[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, in 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 | −17.83 (5.30) |
| Quality | 60.51 (12.01) |
| Random coefficient on price ($\sigma_{\text{price}}$) | 5.63 (2.87) |
| Firm fixed effects | Yes |
| Quarter fixed effects | Yes |
| Observations (firm-tier-quarter) | 384 |
| *Panel B. Own-price elasticities* | |
| Low-end | 1.94 |
| Midrange | 3.29 |
| Premium | 4.41 |
| *Panel C. Markups and marginal cost (low-end / midrange / premium)* | |
| Average price | \$132 / \$267 / \$570 |
| Recovered marginal cost | \$52 / \$167 / \$405 |
| Markup | \$80 / \$100 / \$165 |
| Lerner index (markup ratio) | 0.608 / 0.374 / 0.289 |

*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 3.21. 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 four of 384 observations. Appendix B.3 reports prices, marginal costs, and markups by firm and tier.

**Table 2.** Marginal-cost estimates

| | Estimate (SE) |
|---|---|
| Constant ($c_0$) | 0.046 (0.017) |
| Quality ($c_1$) | 1.066 (0.411) |
| Quality × time, process trend ($c_2$) | −0.043 (0.009) |
| Quality squared ($c_3$) | 18.179 (1.352) |
| Relative quality, $q/q^{\max}$ ($c_4$) | −0.077 (0.067) |
| Relative quality squared ($c_5$) | 0.304 (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.

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 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 thus assume that repositioning an interior tier's quality reuses the same design and the same manufacturing process, so the right-hand side of equation 3.6 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 recovered on the full sample. The R&D-cost parameter and the counterfactuals, which both rely on the quality first-order condition, use the post-2010 window, 2011Q1–2020Q4. This window also sets the paper's scoring horizon: every counterfactual object is computed over ten years, using the 2011Q1 values as the initial state.

[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 13 percent, computed from the unrounded coefficient estimates. On the estimation panel every midrange observation lies above it, as do most low-end observations.

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

### 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 $\chi$, 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), so a higher $\chi$ 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 simulated moment's monotonicity in the scale makes the match unique. With two firms and one market, full solution is the only practical way to estimate the dynamic game, as the data are not rich enough for two-step methods such as Bajari et al. (2007). Overall, the dynamic block is less tightly disciplined by the data than the static primitives, an intrinsic limitation of the single-market setting. I therefore keep the dynamic model and its estimation parsimonious, and report sensitivity analysis in Appendix F.2.

#### 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

[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.

unmeasured quality drift in the inside goods. The fitted trend runs from roughly −$120 in 2005 to roughly +$260 in 2020. The recovered series 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 model market size as mean-reverting in logs 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 $\chi$. 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 using different random seeds, 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]

The surrogates fit well. Held-out $R^2$ is 0.99 for the pricing and 0.97 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 thus 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 21 percent. Because the parameter is exactly identified, I match this moment essentially exactly. Appendix D.3 states the estimator formally.

[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.

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

#### 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 $\chi$ is 5.13. The scale reads as an economic cost of about 5.1 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. Figure D.2 plots the estimation objective against $\chi$: it falls to a clean single interior minimum at the estimate, confirming that the matched moment identifies the scale.

Panels B and D 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 D 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 are of economically sensible magnitude. Within Panel B, the firm-specific success-rate parameters are estimated precisely, while the fishing-out parameter is not, a consequence of the fact that frontier advances are rare events.[34] Panel C 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 holds up well beyond that moment. 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 on top the model-implied paths at the observed state path (Figure D.3). None of these paths is targeted in the estimation, yet they sit largely within the bands. The exceptions mostly fall in the final years, when both firms' R&D intensities climb above the bands, most persistently NVIDIA's, whose frontier runs above as well.[35] With the primitives in hand, I simulate the market under counterfactual restrictions and policy schemes, and compare the expected outcomes with the estimated equilibrium.

[34]With a small sample, theoretically motivated parametric assumptions naturally play a larger role in disciplining the estimation. I therefore re-estimate the R&D-cost scale under alternative assumptions about the hazard and the other primitives of the dynamic specification; Table F.2 reports the results. Across the variants estimated, the scale ranges from about 3.2 to 5.5, and the reading of an economic cost well above the accounting cost holds in every case.

[35]One potential 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 ended, an R&D program sized to that opportunity would rationally exceed what desktop-GPU profits alone justify, and the model prices only the latter.

# 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 permanently, re-simulate the equilibrium path, and evaluate the effects over the ten-year window (2011–2020) on which the dynamic estimation is based; 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 its estimated-equilibrium R&D policy rule, although its realized spending still changes because the state path differs. Each by-firm experiment therefore measures the effect of a unilateral deviation in which the shutdown firm sets R&D to zero. 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. Appendix E.1 carries the full counterfactual protocol. 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).

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[36] 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 $\Delta x_{t+h}$ 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 $\beta$ over $h = 0, \ldots, 39$, the two returns are

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

[36] 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. I keep the accounting convention as the baseline because it is more comparable to the broader innovation literature, which focuses largely on reported R&D. The deduction is deliberately smaller than the cost $\chi I$ the firm bears in its dynamic problem (equation 3.8): the excess over the booked outlay is not a recorded expenditure, and how much of it is real resource cost rather than rents in the labor market is not separately identified, so how much of it to deduct from surplus has no unique answer. Under the polar alternative that treats that full excess as real resource cost, both returns fall in level, the industry wedge is exactly unchanged, and the by-firm wedges widen slightly.

where the minus one nets out the R&D dollar itself, because operating profit is measured before R&D spending; the social return needs no such adjustment, because the outlay is already deducted 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 ratios of present values: the average return to the whole program, not the marginal effect of an additional R&D dollar.[37]

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 \$3.33 to the firm and \$6.00 to the market, a wedge of \$2.67; a dollar of NVIDIA's returns \$3.81 and \$7.57, a wedge of \$3.76.[38] The full-industry row carries the lowest net private return in the table, \$3.00, the highest social return, \$8.95, and the largest wedge, \$5.95 per R&D dollar. In present-value terms, the full-industry shutdown removes roughly \$7 billion of discounted R&D and, with it, roughly \$64 billion of within-market surplus, about half the undiscounted revenue the market realized over the decade.

The full-industry net private return, about \$3.0 per R&D dollar, lies about \$0.3 to \$0.8 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, \$8.95, exceeds the returns in the by-firm rows by roughly \$1.4 to \$3.0 per R&D dollar. The simultaneous shutdown removes about \$14.6 billion more within-market surplus than the two unilateral shutdowns combined. More broadly, with market power, 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 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.

[38]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 \$0.68, while NVIDIA's falls by about \$0.17 (Appendix E.2).

**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 | 3.33 | 6.00 | 2.67 |
| NVIDIA | 3.81 | 7.57 | 3.76 |
| Full industry | 3.00 | 8.95 | 5.95 |

*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 $7 billion of discounted R&D and with it about $64 billion of within-market surplus. Appendix F.2 repeats the full-industry experiment under alternative dynamic specifications.

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 $8.95 here is within-market surplus per dollar of the industry's own R&D, net of that dollar (gross, $9.95), 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, while my model focuses on inappropriable within-market surplus.

Overall, by an alternative route, my estimates confirm 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.[39] To study what such an intervention buys, I simulate three stylized support instruments in the estimated GPU market, each committing about $5 billion over 2011–2020,[40] 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

[39]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.

[40]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 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 $\chi$: under a credit at rate $\tau$, a booked R&D dollar costs the firm $\chi - \tau$ rather than $\chi$, and the government pays the difference. The credit's outlay in a quarter is $\tau$ times the firms' booked R&D spending that quarter. The estimated scale already embeds the public support in force over the sample period, principally 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 $b$. 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 $Q^{\text{cov}}_{t+h}$ 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 and 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 an 11 percent reduction in the estimated R&D-cost scale, equivalently about $0.58 per dollar of the firms' accounting R&D spending. The premium rebate comes to about $81 per premium unit, and the non-premium rebate to about $10.30 per covered unit. The applied rate triggers an equilibrium response in price, quality, and R&D. Realized outlay is an equilibrium outcome rather than 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 and is conceptually less transparent. 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 $\Delta I_{t+h}$ is the change in the two firms' combined R&D spending in quarter $t + h$ relative to the estimated equilibrium, averaged over simulated paths, $\Delta W_{t+h}$ the corresponding 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.[41] 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. Implementation details are in Appendix E.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 +0.13 for the R&D credit, −0.05 for the premium rebate, and +0.02 for the non-premium rebate: only the credit induces a material net increase in private R&D; the premium rebate crowds out private R&D, and the non-premium rebate adds essentially none. Within-market surplus per public dollar is +0.51 for the R&D credit, +0.23 for the premium rebate, and +1.09 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 +\$1.8 billion and profit +\$1.1 billion against +\$0.6 billion of induced R&D, for $\Delta W$ = +\$2.3 billion. Under each rebate, consumer surplus rises, +\$6.0 billion under the premium rebate and +\$4.8 billion under the non-premium rebate, while the producer operating-profit component, measured net of

[41] 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 the denominator $G$.

rebate transfers, splits the two: it falls $4.4 billion under the premium rebate and rises $0.3 billion under the non-premium rebate. This decline is not a claim that the premium rebate hurts the firms: rebate receipts are a transfer, so they are stripped from the profit column. Both firms are in fact better off under either rebate (Appendix E.6); what the negative entry measures is the price war's compression of market-generated producer surplus. Induced R&D is −$0.5 billion under the premium rebate and a small +$0.1 billion under the non-premium rebate; the totals come to $\Delta W$ = +$2.0 billion (premium rebate) and +$5.0 billion (non-premium rebate).

Three patterns in the table stand out. First, the leading instruments under the two criteria are different: additionality puts the R&D credit first, while surplus per public dollar puts the non-premium rebate first. Second, in every row induced R&D is the smallest of the three components, never more than $0.6 billion in magnitude against consumer and profit moves that reach $6.0 billion and $4.4 billion: the crowd-in statistic tracks the smallest line in the ledger while the two larger lines decide the totals, which is why the two headline metrics disagree. Third, the three instruments take different shapes across consumer surplus and profit. The R&D credit's row is balanced: consumer surplus, profit, and R&D rise together. The premium rebate's is offsetting: its consumer gain, the table's largest, is mostly undone by the decline in ex-transfer profit, leaving the smallest total. The non-premium rebate's is concentrated: its consumer gain stands essentially unopposed and makes up almost the whole of the table's largest total. The shapes distinguish push from pull: only the push instrument buys a material increase in private R&D along with its surplus; the two pull instruments buy surplus through the product market instead, delivering more consumer surplus per public dollar than the credit, +0.68 and +1.05 against its +0.40. They differ on the producer side: the premium rebate's consumer gain comes largely at the firms' expense, while the non-premium rebate's consumer gain arrives alongside a small producer gain.

Those patterns suggest that additionality is an incomplete policy indicator because an instrument's return is largely decided by responses it does not track. Across instruments it picks the wrong winner: the credit it puts first sits second in surplus, behind a rebate it scores as adding essentially no R&D. Between the two rebates it offers only a near-tie, −0.05 against +0.02, atop a real divergence in their returns: the non-premium rebate clears more than a dollar of surplus per public dollar and edges the frontier up, while the premium rebate returns +0.23 and sets the frontier back. The blind spot extends to innovation itself: additionality is an industry-wide statistic, silent on which firm's R&D a public dollar induces, and Section 6.3 shows that this allocation decides whether an instrument advances the industry frontier or merely narrows the gap between the firms.

**Table 4.** Additionality and welfare per public dollar

| | Per public dollar | | | Discounted totals ($B) | | | | Budget |
|---|---|---|---|---|---|---|---|---|
| Instrument | Additionality | Welfare | CS | $\Delta$CS | $\Delta\Pi$ | $\Delta I$ | $\Delta W$ | drift |
| R&D credit | +0.13 | +0.51 | +0.40 | +1.8 | +1.1 | +0.6 | +2.3 | +8% |
| Premium rebate | −0.05 | +0.23 | +0.68 | +6.0 | −4.4 | −0.5 | +2.0 | +112% |
| Non-premium rebate | +0.02 | +1.09 | +1.05 | +4.8 | +0.3 | +0.1 | +5.0 | +11% |

*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 (R&D credit) to $8.8 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 $\Delta I$ 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.

The aggregate 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 and shows how the effects are distributed across buyers at different tiers. Under the R&D credit the whole menu lifts (Panel A): quality rises on every tier, +2.1 percent at the low-end, +2.1 at the midrange, and +3.1 at the premium. Share rises on the lower tiers, +0.41 and +0.44 percentage points, with the premium tier roughly flat. The lift is relatively even because the R&D credit works like fast-forwarding technological development: it lowers the cost of research, so firms raise R&D. The added R&D shows up twice: directly, as higher premium quality; and through the within-firm cost spillover of Section 3.2.1, which carries the gain down to the lower tiers. Panel A's one sizable price move is at the top: the pooled premium price rises by about +$20 per unit. The rise is quality-led: the markup column shows about +$6, 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 +$1.8 billion of consumer surplus, and the firms sell that menu at higher markups, the +$1.1 billion of profit. Overall, the credit does not reshape the product market directly: demand and the stage game's marginal-cost schedule are unchanged. Each rebate, in contrast, reshapes the product-line equilibrium (Panels B and C): its effects run beyond its covered tiers, through the pricing and product-line adjustments it sets off.

**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 | +1.5 | +0.8 | +2.1 | +0.41 |
| Midrange | +1.9 | +1.0 | +2.1 | +0.44 |
| Premium | +19.8 | +6.4 | +3.1 | −0.03 |
| *Panel B. Premium rebate* | | | | |
| Low-end | +14.0 | +0.9 | +3.9 | −1.64 |
| Midrange | −15.3 | −20.8 | +0.8 | −0.79 |
| Premium | −126.3 | −117.4 | −1.9 | +5.84 |
| *Panel C. Non-premium rebate* | | | | |
| Low-end | −7.9 | −19.2 | +5.0 | +1.92 |
| Midrange | −45.7 | −50.3 | +2.2 | +2.43 |
| Premium | −8.5 | −11.4 | +0.6 | −0.48 |

*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.

Panels B and C of Table 5 carry the mechanism behind the two rebates' rows in Table 4. A rebate paid on both firms' covered products is competed away rather than pocketed: it triggers a price war on the covered tiers, and the war trades margin for volume. What that trade does to innovation depends on where the war is fought, because the frontier's reward takes a different form along the product line. At the top the reward is a margin. Below the top it is volume: through the within-firm cost spillover, a frontier advance saves production cost on the units sold there. The premium rebate sends the war to the tier where the reward is the margin (Panel B). On its covered tier, the premium, the consumer price falls \$126.3, well more than the rebate itself; the markup falls with it, \$117.4; and buyers pile in, +5.84 percentage points of share. The war fades down the menu: the midrange is repriced down \$15.3 and loses \$20.8 of markup; the low-end loses none, its price \$14.0 higher carrying a 3.9 percent quality gain; and buyers from both trade up to the cheapened top. Consumers gain on price and on adoption, and that is where the premium rebate's +\$6.0 billion of consumer surplus comes from. But the war that hands buyers the transfer compresses the market-generated margins: producer operating profit net of the rebate transfer falls \$4.4 billion. The premium tier's compression is not merely the excluded transfer showing up in the ledger: even with the roughly \$81 receipt added back, the private margin falls by about \$36. Because that margin is the reward for holding the frontier, R&D moves with it: NVIDIA cuts R&D \$0.48 billion while AMD's is essentially unchanged, for a net −\$0.5 billion of induced R&D. The covered tier's quality slips 1.9 percent: the premium product's quality is the frontier itself, so the R&D cut lands on the very product the rebate covers.

The non-premium rebate sends the same war to the tiers where the reward is volume (Panel C). It covers the low-end and midrange, and on both the consumer price falls, \$7.9 at the low-end and \$45.7 at the midrange, the latter well more than the \$10.30 rebate; the war costs them \$19.2 and \$50.3 of markup. Quality rises rather than slips, +5.0 and +2.2 percent, and adoption rises on both covered tiers, +1.92 and +2.43 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 3.9 percentage points while the uncovered premium tier gives up only 0.48. The premium tier is dragged along too, \$8.5 down on price and \$11.4 on markup, with quality roughly flat. Consumers gain on price, on quality, and on adoption, and the non-premium rebate's consumer surplus comes to +\$4.8 billion. Producer profit net of the transfer rises just \$0.3 billion: the compression hits thin low-end and midrange margins, but the adoption the rebate sets off expands the volume those margins are earned on, and the two all but cancel. The gain is NVIDIA's,

$0.47 billion, while AMD gives up $0.18 billion. R&D rises as well, a small +$0.1 billion on net, with both firms adding, AMD $0.09 billion and NVIDIA $0.03 billion. Overall, the thin margins do not drive the rebate's welfare effect; the expanding volume does. Within-market surplus per public dollar comes to +1.09.

More broadly, the ex-transfer profit changes need not line up with the rebates' coverage sets: the equilibrium adjustment of the whole menu picks each product's gain or loss. Under the non-premium rebate, the largest single-product gain arrives at a covered low-end product, yet the deepest loss lands on an uncovered premium product and one covered midrange product loses as well. Under the premium rebate almost every product loses: the one gain, at an uncovered low-end product, rounds to zero, and an uncovered midrange product loses the most. The per-product breakdown is in Appendix E.6.

The premium rebate comes with one scale caveat. Its realized outlay drifts far above plan: +112 percent (the budget-drift column of Table 4), about $8.8 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 surplus column's top two instruments, realize roughly equal outlays close to plan, $4.5 billion (credit) and $4.6 billion (non-premium rebate). The comparison between them is thus read at a common realized scale, while comparisons involving the premium rebate run across different realized scales.

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 it is absent from the table by construction. Both rebates, in contrast, are welfare-positive as static policies. The premium rebate's static welfare effect is +0.43 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 about half of that gain, −0.20, for a total of +0.23. The non-premium rebate's static effect is +1.00, and its dynamic state response adds to it, +0.09, for a total of +1.09: for this rebate the state response works with the static gain rather than against it.

**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 | +0.43 | −0.20 | +0.23 |
| Non-premium rebate | +1.00 | +0.09 | +1.09 |

*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 re-fit 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.

### 6.3 Frontier Advancement per Public Dollar

Table 7 reports additionality, frontier advancement per public dollar, and the percent change in R&D at the frontier-leader and follower roles under each instrument. The R&D credit leads on frontier advancement at +0.81: a billion dollars of credit outlay lifts the terminal frontier about eight-tenths of a percent. The non-premium rebate manages a small positive +0.19, and the premium rebate is negative, −0.32; the one instrument paid on the frontier tier pushes the frontier back. Per $1 billion of induced private R&D, the credit raises the frontier by 6.05 percent; no other instrument induces a material net increase in private R&D, so the ratio is meaningful for the credit alone. Together, the table's two per-dollar columns show that industry-wide additionality is not an ideal indicator for the advancement of the industry frontier: the frontier outcomes are far more dispersed than the additionality reading would suggest, and the two rebates are nearly alike (near zero) by that metric, yet they move the frontier in opposite directions.

The misalignment has a deeper source: 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; R&D induced at the follower closes the distance behind the leader but moves the industry frontier little. The leader and follower columns of Table 7 locate each instrument's R&D response. The R&D credit raises R&D at both roles by about the same amount, +8.9 percent at the leader and +8.0 at the follower. The premium rebate cuts R&D at both roles, and the cut lands more than twice as heavily on the leader, −9.3 percent against −3.8. The non-premium

**Table 7.** Frontier advancement and R&D by frontier role under public support

| | | Frontier advancement | Percent change in role R&D | |
|---|---|---|---|---|
| Instrument | Additionality | % per $1B outlay | Frontier leader | Frontier follower |
| R&D credit | +0.13 | +0.81 | +8.9 | +8.0 |
| Premium rebate | −0.05 | −0.32 | −9.3 | −3.8 |
| Non-premium rebate | +0.02 | +0.19 | +2.5 | +0.8 |

*Notes:* Additionality is as in Table 4. Frontier advancement is the percent change in the end-of-horizon industry frontier $q_t^{\max}$, the higher of the two firms' frontier qualities, per $1 billion of realized public outlay. The role columns are the percent change in each role's discounted R&D relative to its estimated-equilibrium level, about $3.4 billion at the leader and $3.7 billion at the follower in present value. Frontier leader and follower are state-contingent roles, not fixed firm identities: at each path-period the leader is the firm with the higher frontier quality on the estimated-equilibrium path, and a firm's induced R&D is attributed to its role at that path-period. Construction detail and the dollar levels are in Appendix E.7. Conventions and realized outlays are those of Table 4; each instrument's rate is in Section 6.1.

rebate raises both roles modestly, +2.5 percent at the leader and +0.8 at the follower. The leader column orders the instruments exactly as the frontier column does: the R&D credit, which raises leader R&D most, advances the frontier most; the premium rebate, the only instrument that cuts leader R&D, is the only one that moves the frontier back; and the non-premium rebate's modest leader-role rise buys a modest 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 raises R&D at both roles because it works on the frontier's other reward, the production cost its advance saves on covered volume (Section 6.2); the rebate expands that volume, so the return to advancing the frontier thickens at both firms. The rise tilts toward the leader role, and at the firm level toward the firm that trails more often than it leads. The two tilts are consistent because the roles are state-contingent: the more-often-trailing firm still leads on about half the path-periods, and its added R&D counts at the leader role when it does. The allocation between the roles, not the size of the aggregate response, decides an instrument's frontier return: the leader-role share turns the non-premium rebate's small net addition into a frontier gain, and the leader-role concentration of the premium rebate's cut turns a modest withdrawal into a frontier retreat.

### 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 +0.13 of private R&D per public dollar, returns +0.51 of within-market surplus per public dollar, and advances the end-of-horizon frontier by +0.81 percent per $1 billion of outlay; the premium rebate crowds R&D out, returns +0.23 per public dollar, and pushes the frontier back; the non-premium rebate induces essentially no net R&D yet returns +1.09 and edges the frontier up. As a design guide, the pattern echoes the principle of targeting: pursue an objective with the instrument aimed most directly at it. The principle applies objective by objective: additionality and frontier advancement are strongest under the R&D credit, the instrument aimed at the cost of research itself, while within-market surplus per public dollar is highest under the non-premium rebate, the instrument aimed at mass-market adoption.

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: $-0.32$ 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 funds directed to 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 shows up as the price war at the top of the line, 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 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 just 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 +0.13. 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, which is invariant to reallocation between

them. The industry's cutting edge 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. Tilting a policy deliberately toward the leader carries its own risk. Too large a lead can blunt the follower's incentive to close the gap, or even push it out of the market if the gap grows too wide.[42]

The policy comparison also surfaces a trade-off between consumer surplus and innovation. Consumer surplus per public dollar runs from +0.40 under the R&D credit through +0.68 under the premium rebate to +1.05 under the non-premium rebate, and the column pulls against frontier advancement: the non-premium rebate delivers the most consumer surplus per public dollar yet buys only a small frontier gain, while the R&D credit, the instrument that advances the frontier most, delivers the least. A price war is a win for consumers, but it can reduce firms' incentive to innovate. A policy aimed at consumer welfare in a technology-driven market can therefore work well in the short term yet backfire after years, as the premium rebate does here: the consumer gain is real, while the cost outlasts the evaluation window, sitting in the frontier technology the future inherits. This 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 an equilibrium 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.

[42]This is the inverted-U argument as in Aghion et al. (2005).

# 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, quantify the surplus-based private and social returns to R&D, and simulate three stylized support instruments on a common ex ante planned budget. The returns are large, and most of the value a dollar of R&D creates is consumer surplus the firms cannot appropriate. Yet 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 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, extended beyond two firms and three tiers, to other innovation-driven markets. A recent example is tiered access to large language models, where the within-firm cost spillover has a literal analog in distillation: 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. Market-by-market estimates would recover the distribution beneath that average, each market carrying its own social-private split and policy counterfactual.

The second direction is to measure what the within-market boundary leaves out: the chips sold to render games became the hardware that trains modern AI. Quantifying the downstream surplus created by advances in the GPU frontier would show how far the economy-wide value of an R&D dollar exceeds the within-market measure.

The third direction concerns the yardstick itself. Additionality, the R&D response a public dollar induces, is the most readily available indicator: R&D spending appears in financial reports and tax records. But the equilibrium response that decides what it ultimately buys took the full apparatus of this paper to see, machinery few agencies can build. The practical question is how to build better yardsticks where detailed data are limited: when the input response alone is a reliable guide, and what else agencies should collect when it is not. Tier-level prices, qualities, and market shares are observable 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 its 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.

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005): "Competition and Innovation: An Inverted-U Relationship," *The Quarterly Journal of Economics*, 120, 701–728.

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.

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.

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.

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.

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.

Rhodes, A., J. Zhou, and J. Zhou (2026): "Digital Ecosystems and Data Regulation," TSE Working Paper 25-1621, Toulouse School of Economics, Revised February 2026.

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 | 11.5 | 62.5 |
| Midrange | 297 | 53.5 | 30.0 |
| Premium | 722 | 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 | 1.2 | $B per year, average 2005–2020 |
| R&D-to-revenue ratio | 0.09 | pooled sample |
| Frontier benchmark growth | 98.7 | multiple, 2005–2020 |
| Premium-tier price growth | 1.9 | multiple, 2005–2020 |

*Notes:* Panel A reports means over the cleaned panel, weighted by units sold. Relative performance is normalized to the premium tier at 100. In Panel B, GPU-segment R&D is firm-reported spending allocated to the segment in proportion to segment revenue (Section 2).
*Sources:* Jon Peddie Research (JPR) for discrete-GPU units and desktop CPU shipments; firm filings and Compustat for financials; PassMark for benchmark performance; launch and retail-price records for price; and the Bureau of Labor Statistics for the semiconductor-wafer producer-price index that enters the price instruments (Section 4.1).

# A The Desktop GPU Market, 2005–2020

Table A.1 summarizes the product line and the market aggregates. 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 $722, with the two lower tiers delivering about 12 and 54 percent of premium benchmark performance.

Figure A.1 plots the premium-tier price and the observed quality frontier year by year. The frontier in panel (b), the best 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 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 business-facing, with data-center accelerators that sell for tens of thousands of dollars apiece and are bought in batches of tens of thousands by large AI developers.

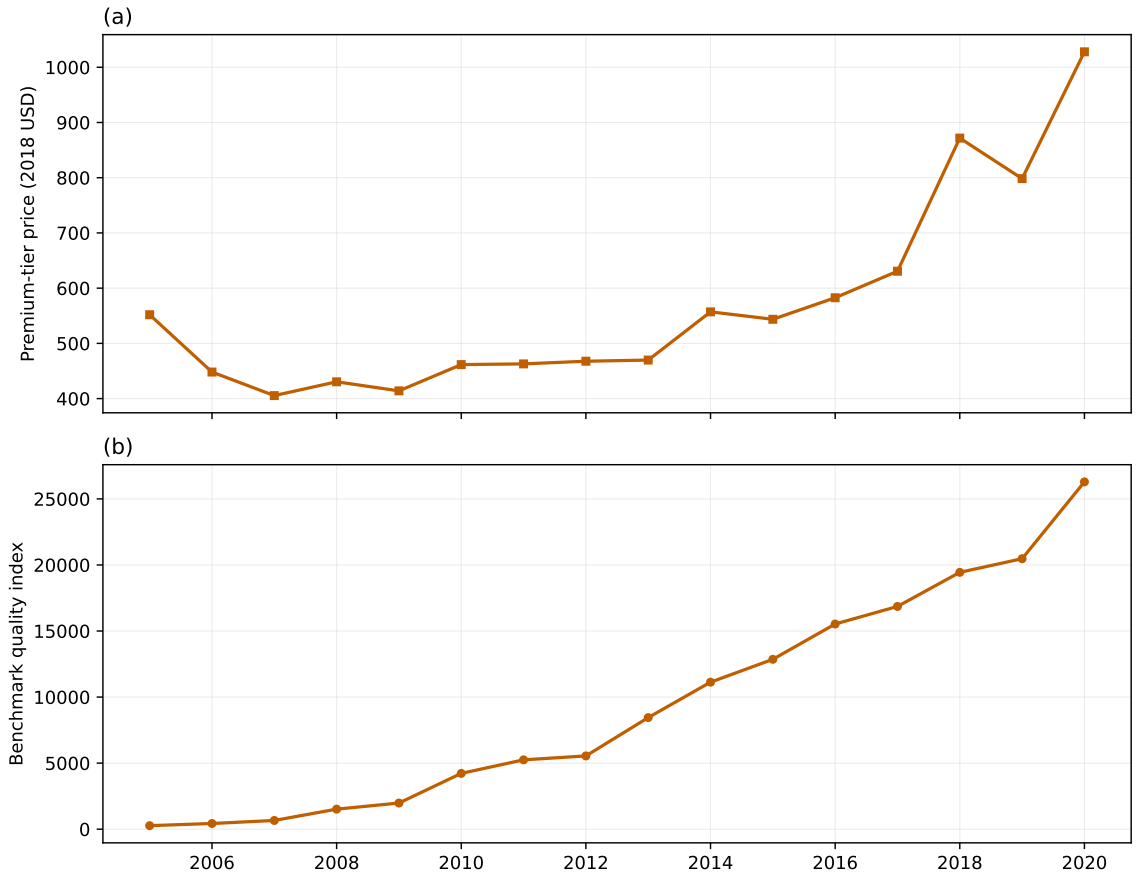



**Figure A.1.** Premium-tier prices and observed frontier quality

*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 benchmark index across the two firms' products in each year. The 2005–2020 endpoint ratios of the two series equal the growth multiples in Table A.1, Panel B.

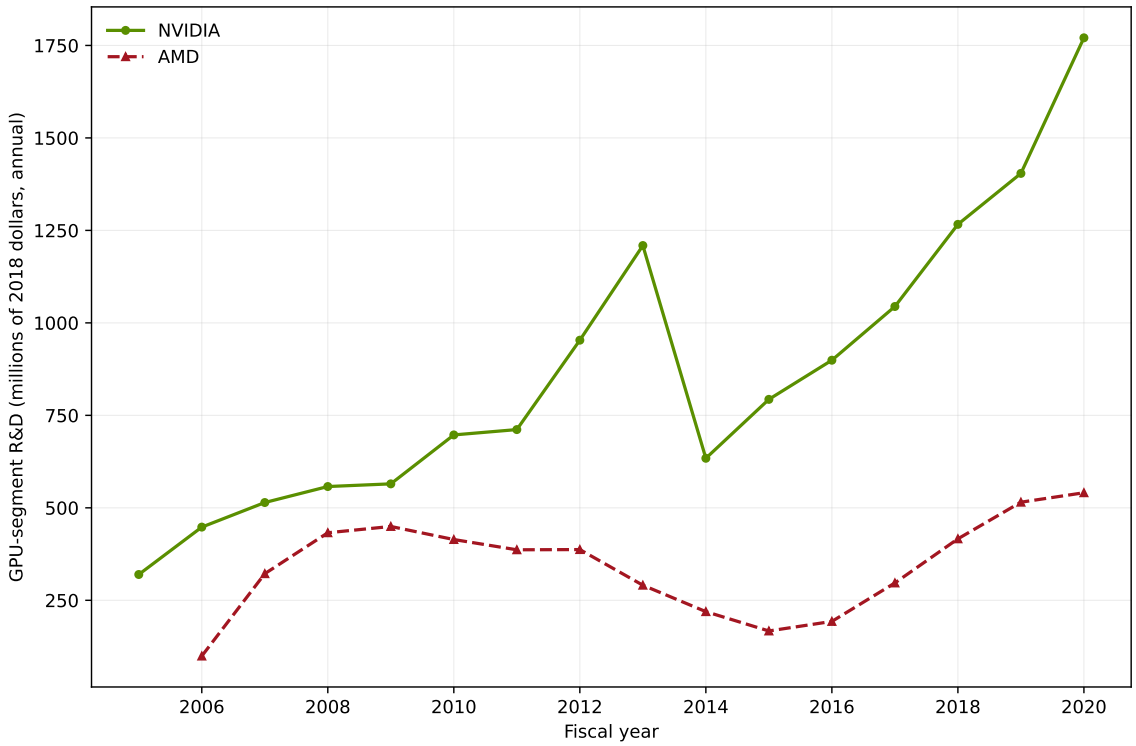



**Figure A.2.** Allocated GPU-segment R&D

*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 2018 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 R&D behind the race appears in Figure A.2, firm by firm. NVIDIA's GPU-segment R&D grows from roughly $0.32 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 $0.54 billion by fiscal 2020. 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 low point of its full-year R&D series in Figure A.2, and recovered to roughly 30 percent in the late 2010s. This 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.

Adoption and market size complete the picture. The attach rate, discrete-GPU units as a share of desktop CPU units in each quarter, fluctuates around its pooled 46.2 percent (Table A.1, Panel B). The denominator itself shrinks: 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 one sharp break in these series is the cryptocurrency-mining episode, the event behind the 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. Within the episode, prices peaked in the same quarter: in 2018Q1 NVIDIA's midrange and premium prices stood 68 and 51 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 for the price variable; Appendix B.2 reports the full elasticity matrix; 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.

## B.1 The Demand First Stage

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, consistent with the die-size and yield interpretation of the tiers' differential exposure to wafer costs (Section 4.1).

**Table B.1.** Demand first-stage estimates

| | Estimate (SE) |
|---|---|
| *Excluded instruments* | |
| Wafer cost shifter × premium tier | 1.631 (0.217) |
| Wafer cost shifter × midrange tier | 0.628 (0.122) |
| Differentiation instrument, own quality | 6.957 (2.434) |
| Differentiation instrument, rival quality | 6.732 (1.236) |
| Square root of own quality | 1.433 (0.455) |
| *Included effects* | |
| Quality | Yes |
| Firm fixed effects | Yes |
| Quarter fixed effects | Yes |
| *Diagnostics* | |
| Joint excluded-instrument $F$ statistic | 48.7 |
| Partial $R^2$ of excluded instruments | 0.44 |
| 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.

## 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 0.72 percent when that price rises by one percent; in the AMD premium column, no product's demand rises by more than 0.20 percent.

## 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 $\mathrm{MC}_t$ stack the six market shares, prices, and marginal costs, $\Delta_t$ 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,

**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 | −1.87 | 0.20 | 0.06 | 0.38 | 0.51 | 0.12 |
| AMD midrange | 0.20 | −3.24 | 0.11 | 0.36 | 0.66 | 0.24 |
| AMD premium | 0.14 | 0.20 | −4.65 | 0.26 | 0.72 | 0.62 |
| NVIDIA low-end | 0.22 | 0.21 | 0.07 | −2.01 | 0.56 | 0.15 |
| NVIDIA midrange | 0.19 | 0.23 | 0.13 | 0.34 | −3.34 | 0.39 |
| NVIDIA premium | 0.11 | 0.17 | 0.20 | 0.19 | 0.60 | −4.18 |

*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.

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

| Product | Price ($) | Marginal cost ($) | Markup ($) | Lerner index |
|---|---|---|---|---|
| AMD low-end | 121 | 50 | 71 | 0.585 |
| AMD midrange | 234 | 152 | 82 | 0.352 |
| AMD premium | 482 | 361 | 121 | 0.251 |
| NVIDIA low-end | 143 | 53 | 90 | 0.627 |
| NVIDIA midrange | 299 | 182 | 117 | 0.391 |
| NVIDIA premium | 658 | 449 | 209 | 0.317 |

*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.

$\Omega_t$ 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 to obtain the marginal-cost estimates:

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

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 $71 at the low end, $117 against $82 at the midrange, and $209 against $121 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 brand effect and its higher average quality within each tier.

### B.4 The Implied Marginal Fixed Cost of Quality

The series plotted in Figure B.1 is the implied marginal fixed cost of quality introduced in Section 4.2: the left-hand side of the quality first-order condition, equation (3.6), evaluated at the observed menu for the four interior products, each firm's low-end and midrange tiers. Under optimal quality choice this marginal variable profit would equal any marginal fixed cost of quality, so the evaluated series is the marginal fixed cost the observed menu implies. The evaluation includes the induced response of all six equilibrium prices to each quality change. Expectations over the demand and marginal-cost shocks, $\xi_{jt}$ and $\varepsilon_{jt}^{\mathrm{MC}}$, are computed by sparse Kronrod–Patterson quadrature, with both shocks taken as mean-zero normal, independent across products, and scaled by the standard deviations of their 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 75-fold, and almost 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 states to solve the game at each one (Section 4.3.2). Instead, I solve the stage game at a large sample of states, train constrained neural networks on those solutions, and validate the approximation in two layers: held-out accuracy (Appendix C.6) and coverage of the support the dynamic model actually visits (Appendix C.7). Appendices C.1–C.5 lay out the construction in pipeline order.

### C.1 The 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

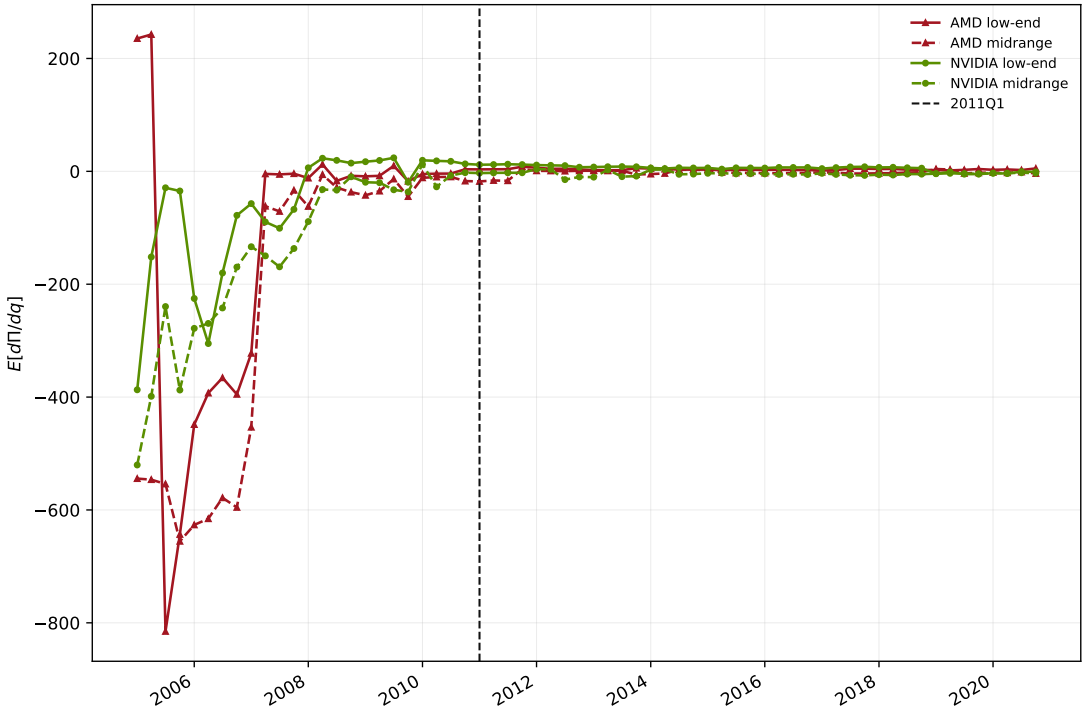



**Figure B.1.** Implied marginal fixed costs of quality

*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.

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 needs expected profits over the shock distribution.

### C.2 Solving the Stage Game at One State

The solver first 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. Shock nodes at which the solver finds no admissible pricing equilibrium are dropped, with the remaining quadrature weights renormalized.

Sparse quadrature is the right tool for this integral for two reasons. First, cost: the demand and marginal-cost shocks jointly span twelve dimensions, and every integration node requires solving the pricing game, so simulation-based integration, with the many draws it would need, is expensive; the sparse rule delivers the expectation from a small, deterministic node set. Second, the normal weight approximates the shock distributions well: the estimated demand and cost equations absorb most of the variation in their objects, so the shocks are modest in scale, and outside a handful of quarters tied to the cryptocurrency episode and late-sample product launches, their empirical distributions are near-symmetric with modest excess kurtosis, though formal normality tests reject. I nonetheless maintain the normal approximation, scaling each shock to the standard deviation of its estimated residuals.

The solver then selects an equilibrium. Both the pricing and the quality games are solved from multiple starts. A candidate is admissible only if prices are positive, product shares are feasible, qualities are ordered and inside their frontiers, and the equilibrium conditions hold up to the solver's tolerance; positive prices and feasible shares hold by construction. 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. Multiplicity is empirically rare: at most states the admissible candidates coincide or differ negligibly, and the selection rule decides a small residual of cases. The common rule makes equilibrium selection an explicit part of the model rather than an unrecorded numerical accident. Solved games are what the networks train on, and the design question is where to run them; Appendix C.3 records which training targets are direct solutions and which embed the fitted pricing map.

### C.3 Designing the Training Sample

The base design covers two rectangular regions of states and quality menus, one bracketing the observed path and one extending beyond it, 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, states bridging the two sampling regions, and the boundary regions of the state box the dynamic solvers sample, 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 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 the costly game solving 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, I augment the design: the states visited by the estimated equilibrium's simulated paths are appended, so each re-fit trains densely where the policy paths travel. Throughout, I set the training bounds 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-network ensembles are fitted to these training sets, one for quality and one for pricing. Each member is a small feed-forward network trained to minimize squared prediction error, with architecture chosen by cross-validated search on the training data, and each ensemble averages 20 independently initialized members with equal weight. The quality ensemble $Q$ 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 ensemble $\mathcal{P}$ 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):

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

Economic constraints are built into the output layers. The quality ensemble 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 ensemble returns the six inside shares and the outside share through a softmax over seven logits; shares are nonnegative and sum to one. 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 re-solve the 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 validation, in two layers.

**Table C.1.** Held-out surrogate accuracy

| Output block | Neural-network $R^2$ | Best cubic $R^2$ | Best LASSO $R^2$ |
|---|---|---|---|
| AMD low-end quality | 0.980 | 0.965 | 0.890 |
| AMD midrange quality | 0.967 | 0.947 | 0.880 |
| NVIDIA low-end quality | 0.985 | 0.971 | 0.894 |
| NVIDIA midrange quality | 0.956 | 0.941 | 0.866 |
| Interior quality, average | 0.972 | 0.956 | 0.882 |
| Prices, average | 0.997 | 0.997 | 0.994 |
| Shares, average | 0.993 | 0.716 | 0.615 |
| Variable profits, average | 0.981 | 0.695 | 0.587 |
| Consumer surplus | 0.998 | 0.970 | 0.953 |
| All pricing outputs, average | 0.991 | 0.811 | 0.743 |

*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 training sample and scored on the same held-out split; the LASSO benchmark is the $\ell_1$-penalized polynomial whose degree and penalty give the highest mean $R^2$ within each displayed output block.

## C.6 Held-Out Accuracy

I check accuracy 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 training sample and scored on the same held-out split, the best cubic polynomial and the best LASSO. The four interior qualities are predicted with $R^2$ between 0.956 and 0.985; on the pricing side, $R^2$ is about 1.00 for prices, 0.99 for shares, 0.98 for variable profits, and about 1.00 for consumer surplus. The block averages put the networks ahead of both benchmarks, 0.97 against 0.96 and 0.88 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.

## C.7 Coverage of the Visited Support

Held-out accuracy certifies the fit on states like those the networks were trained on. Does the dynamic model ever evaluate the mappings anywhere else? For every simulated state, the code 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

1.7 percent of simulated firm-quarters come within one ten-thousandth of a bound, split roughly evenly between the zero-investment and full-reinvestment ends, 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 to 0.02 percent, in which a chosen quality happens to coincide with the rival's frontier, a point the menu crosses freely rather than a constraint. These records complete this appendix's approximation audit.

# D Details for Dynamic Estimation

## 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 \$24.08 per year, with $R^2 = 0.95$ over 60 quarters; converted to utility units, it gives a quarterly drift of 0.11, the rounded value at which the model is solved (Panel D of Table D.1). 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 0.04 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 model is solved at the rounded long-run level of 28.0 million. The

**Table D.1.** Dynamic parameters and calibrated processes

| *Panel A. Estimated in the solved dynamic game* | | |
|---|---|---|
| Object | Parameter | Estimate (SE) |
| R&D-cost scale | $\chi$ | 5.13 (0.66) |
| *Panel B. Innovation arrival and jump distribution, estimated offline* | | |
| Object | Parameter | Estimate (SE) |
| AMD success-rate parameter | $a_{\text{AMD}}$ | −1.20 (0.43) |
| NVIDIA success-rate parameter | $a_{\text{NVIDIA}}$ | −1.61 (0.46) |
| Fishing-out coefficient | $\varphi$ | 0.36 (3.29) |
| Mean log jump | $m_J$ | −4.20 (0.17) |
| Variance of log jump | $v_J$ | 0.66 (0.15) |
| *Panel C. Calibrated discount factor* | | |
| Object | Parameter | Value |
| Quarterly discount factor | $\beta$ | 0.99 |
| *Panel D. Fitted exogenous-state processes* | | |
| Object | Parameter | Value |
| Outside-option quarterly drift | $\Delta_u$ | 0.11 |
| Outside-option ceiling | $u_{0,\max}$ | 7.0 |
| Process-innovation quarterly drift | $c_2$ | −0.04 |
| Process-innovation floor | $z_{\min}$ | −2.5 |
| Market-size persistence | $\rho_M$ | 0.85 |
| Long-run market size | $M_\infty$ | 28.0 million |

*Notes:* The R&D-cost scale $\chi$ is estimated from a single moment, the mean R&D intensity of the post-2010 sample (Appendix D.3); its standard error is asymptotic, computed by the delta method, and a parametric bootstrap that redraws innovation arrivals and jumps over the sample window yields a larger standard error of 3.33. Both constructions hold 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. Panel B standard errors come from the offline arrival-rate likelihood and sample-moment asymptotics. The fishing-out coefficient is imprecisely estimated, a consequence of the fact that frontier advances are rare events; Appendix F.2 reports sensitivity to the hazard specification. In Panel D 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).

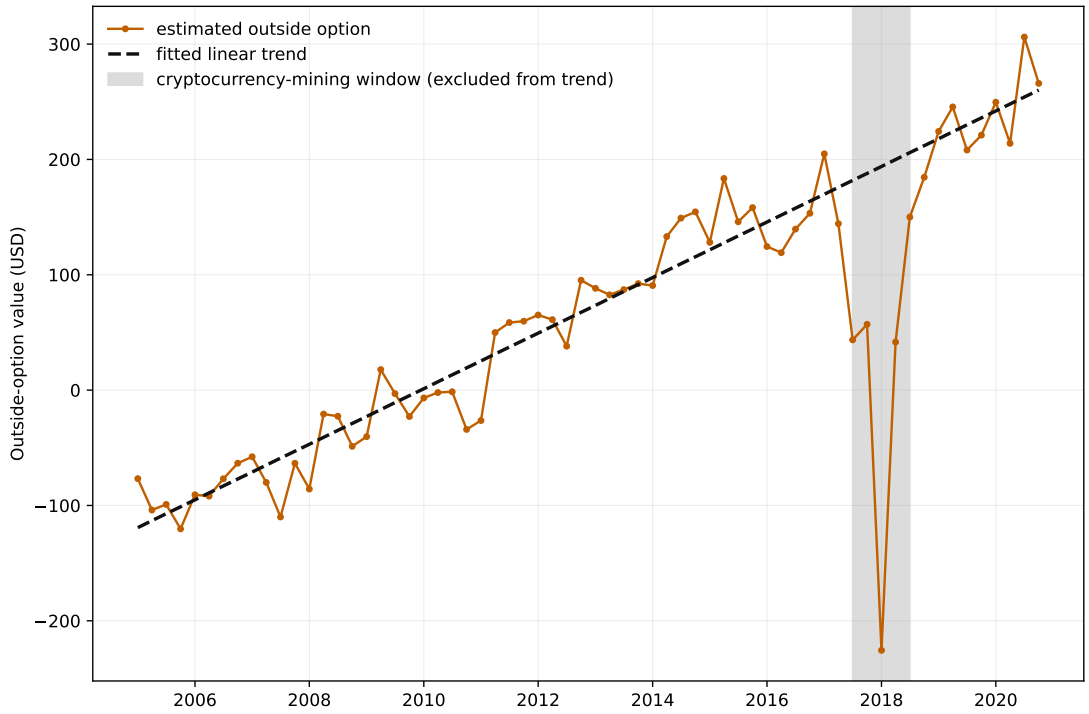



**Figure D.1.** Recovered outside option and fitted transition

*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.

initial market size is 36 million units, the rounded 2011–2013 mean.

## D.2 Solution Method and Simulation

This subsection records the solution method for the dynamic game and the simulation conventions the estimation and the counterfactuals share. 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 $\mathbb{E}[\pi_f^*(x_t)]$, and accounting-equivalent R&D spending is

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

the value that enters both the R&D-cost term of equation (3.8) and the innovation hazard. The same conversion is used when comparing simulated R&D with segment-allocated accounting R&D.

I solve the game 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.

The continuation value in the Bellman equation is computed as an exact sum rather than by simulation. In a quarter the frontiers can move in only four ways: neither firm advances, only AMD does, only NVIDIA does, or both do, so the solver sums over these four outcomes, weighting each by its probability, and integrates over the size of an advance by quadrature. 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 by matching the model-implied average R&D intensity to its observed value over 2011–2020. The objective function is the squared gap between the two:

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

with $T = 40$ the number of quarters in the 2011–2020 window. The targeted mean intensity is 0.21, 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 $\widehat{\chi} = 5.13$, the shape implied by a simulated mean intensity that declines in the scale (the 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. Model fit is thus demonstrated by how closely the model reproduces patterns in the data it was not asked to match. I provide that check for market share, frontier quality, and R&D intensity, for both firms over time.

Figure D.3 presents the check, with firms in columns and objects in rows. In the share panels (a), the dashed line is the model-implied expected share at the observed states, with market size

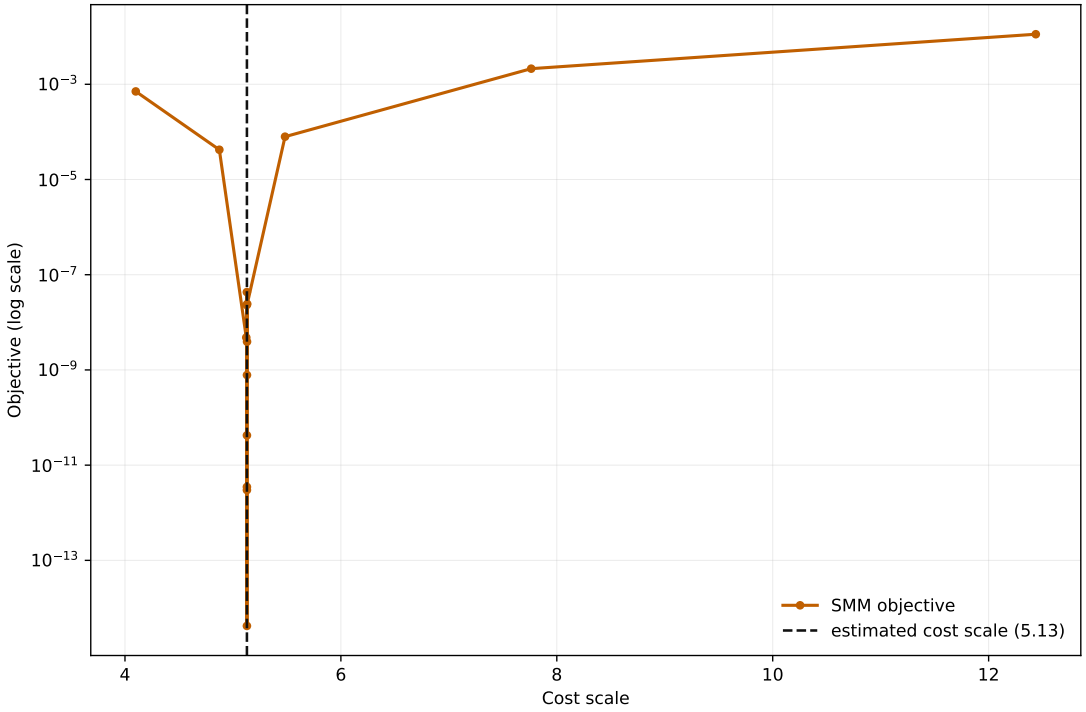



**Figure D.2.** The estimation objective for the dynamic R&D-cost scale

*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, $\widehat{\chi} = 5.13$.

on the fitted AR(1) law; in the frontier panels (b), it is the observed frontier quality; and in the R&D-intensity panels (c), it is the series that constructs the data moment, observed R&D spending over model-implied expected profit. The solid line and the shaded band are the mean and the 10th–90th percentile range across 256 simulations at the point estimate. Because the Markov-perfect equilibrium is in pure strategies, the band reflects randomness in innovation successes and jump sizes rather than parameter uncertainty. The mean series is smoother than the dashed series by construction: successes and jumps are discrete, so any single path, observed or simulated, moves in steps, and averaging across simulations smooths those steps out. Across the six panels, the dashed series lie inside the band in 92 percent of firm-quarter observations.

# 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. Counterfactual experiments change one economic primitive at a time, re-solve the relevant equilibrium, and compare 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. Appendix E.1 records the common protocol; Appendix E.2 reports the by-firm returns under a re-optimizing rival; Appendix E.3 details the policy implementation; Appendix E.4 constructs the static and dynamic welfare components; and Appendices E.5–E.7 carry the product-, firm-, and role-level results.

## 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

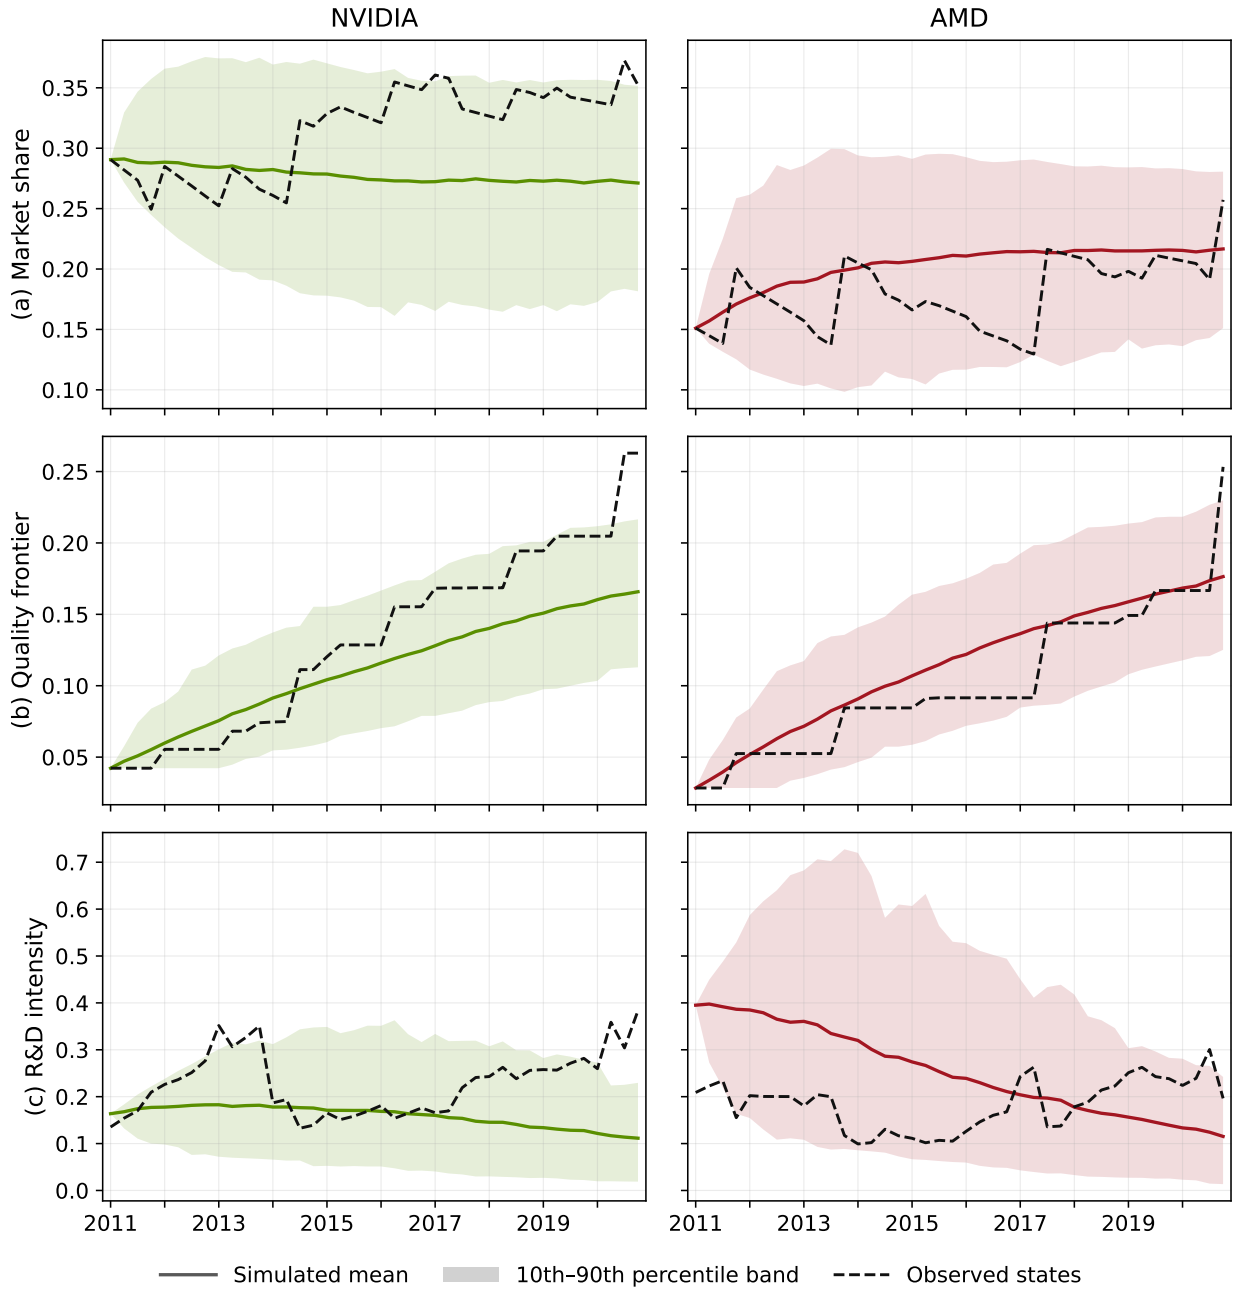



**Figure D.3.** Model-implied paths at the observed state path and simulated point-estimate bands

*Notes:* Columns are firms; rows show market share (a), the quality frontier (b), and R&D intensity (c), with the vertical scale shared within each row. Dashed paths are constructed at the observed state path: the model-implied expected share, with market size on the fitted AR(1) law, in the share panels (a); the observed frontier quality in the frontier panels (b); and observed R&D spending over model-implied expected profit in the R&D-intensity panels (c). Solid paths show the simulated mean, with the shaded 10th–90th percentile band across 256 simulations at the point estimate; the bands reflect innovation and jump randomness rather than parameter uncertainty. Across the six panels, 92 percent of the dashed firm-quarter values lie inside the bands.

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

| Firm | Return object | Fixed-rival policy | Rival re-optimized | Strategic gap |
|---|---|---|---|---|
| AMD | Private | 3.33 | 3.31 | +0.0278 |
| AMD | Social | 6.00 | 6.69 | −0.6849 |
| NVIDIA | Private | 3.81 | 3.82 | −0.0149 |
| NVIDIA | Social | 7.57 | 7.40 | +0.1749 |

*Notes:* The rival re-optimized 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 re-optimized return, computed from the unrounded returns.

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 provides sensitivity analysis for this closure assumption.

## E.2 Strategic Adaptation of the Rival

The by-firm returns of Section 5 shut one firm's R&D down while the rival keeps the R&D policy rule it follows in the estimated equilibrium. This fixed-rival convention makes the experiment a clean unilateral deviation, but the return it measures need not equal the long-run value of an active innovator when the rival understands that the shutdown is permanent. This subsection reports the alternative, a constrained re-equilibrium: the shutdown firm's R&D stays at zero, and the rival re-optimizes its dynamic policy against a permanently non-innovating competitor. Writing $r_f$ for a fixed-rival return of Section 5 and $\widetilde{r}_f$ for its rival re-optimized counterpart, the strategic-adaptation component is their difference, $r_f - \widetilde{r}_f$.

Table E.1 reports the comparison by firm and return object. The private returns barely move: AMD's changes from \$3.33 to \$3.31 per R&D dollar and NVIDIA's from \$3.81 to \$3.82, gaps below \$0.03 in absolute value. The social returns move more, and in opposite directions: AMD's rises from \$6.00 to \$6.69 when NVIDIA re-optimizes against its shutdown, while NVIDIA's falls from \$7.57 to \$7.40 when AMD does. The sign of the gap alone does not reveal whether the rival raises or lowers its R&D, because re-optimization also changes the state and product-market paths. Overall, the private returns are insensitive to the rival convention; the social returns shift by up to about \$0.7 per R&D dollar, but the conclusion that the social return sits well above the private return for both firms holds under either convention.

### E.3 Policy Implementation

The policy experiments add a new layer of support on top of the estimated equilibrium. 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 $p_{kt}$. In the pricing game, the rebate enters as the equivalent per-unit reduction of $b_d$ in the covered products' marginal cost, and the equilibrium price the solver returns is the consumer price $p_{kt}^{\text{net}}$. All prices reported under the two rebates are thus consumer prices with the rebate already netted out: a covered product's reported price change is the pass-through object itself, and reported markups are the consumer price minus marginal cost, the margin net of the rebate transfer. For uncovered products, consumer and posted prices coincide.

### 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, 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 is divided by the replay's own realized outlay. The two components sum to the total by construction, with the dynamic state response defined as the remainder.

### 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

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

| Product | Premium rebate (%) | Non-premium rebate (%) |
|---|---|---|
| NVIDIA low-end | −4.8 | −26.7 |
| NVIDIA midrange | −19.2 | −41.0 |
| NVIDIA premium | −63.7 | −14.5 |
| AMD low-end | +8.3 | −14.4 |
| AMD midrange | −12.3 | −36.4 |
| AMD premium | −67.6 | +4.3 |

*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.

over the 40 periods.

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

### E.6 Profit Incidence Across Firms and Products

This subsection reports the profit incidence behind two claims of Section 6.2: both firms are better off under either rebate once transfers are counted, and the product-level gains and losses do not follow the coverage sets. Throughout, 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 entries measure market-generated producer surplus; transfer-inclusive payoffs are reported separately below. Table E.3 reports the firm-level present values. The R&D credit raises operating profit at both firms, +$0.58 billion at AMD and +$0.52 billion at NVIDIA. The premium rebate lowers the ex-transfer component at both firms, by $1.18 billion at AMD and $3.19 billion at NVIDIA; the non-premium rebate lowers AMD's by $0.18 billion while NVIDIA's rises by $0.47 billion. Adding the subsidy receipts back recovers transfer-inclusive operating profit, which rises at both firms under both rebates: by $2.30 billion at AMD and $2.09 billion at NVIDIA under the premium rebate, and by $1.62 billion and $3.25 billion under the non-premium rebate. The negative entries are the discounted counterpart of the markup compression in Table E.2.

The firm totals conceal within-firm reallocation. Table E.4 carries the product-level breakdown. Under the premium rebate five of the six products lose ex-transfer operating profit, and the one gain, at AMD's uncovered low-end, rounds to zero. The coverage set does not order those losses: NVIDIA's uncovered midrange loses the most, $1.3 billion, more than either covered premium

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

| Instrument | AMD (USDb) | NVIDIA (USDb) | Total (USDb) |
|---|---|---|---|
| R&D credit | +0.58 | +0.52 | +1.10 |
| Premium rebate | −1.18 | −3.19 | −4.37 |
| Non-premium rebate | −0.18 | +0.47 | +0.28 |

*Notes:* Entries are 40-quarter discounted operating profit net of demand-rebate transfers, relative to the estimated equilibrium, in 2018 USD billions; totals may differ from the sum of rounded firm entries. 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.

**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 | −0.7 | +0.9 |
| NVIDIA midrange | −1.3 | 0.0 |
| NVIDIA premium | −1.2 | −0.5 |
| AMD low-end | 0.0 | +0.4 |
| AMD midrange | −0.1 | −0.2 |
| AMD premium | −1.1 | −0.4 |

*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.

product. Under the non-premium rebate the gains arrive at the two covered low-ends, +$0.9 billion at NVIDIA's and +$0.4 billion at AMD's, yet the equally covered midranges show no such gains, NVIDIA's approximately flat and AMD's losing $0.2 billion. The non-premium rebate's deepest loss, $0.5 billion, lands on NVIDIA's uncovered premium product. The table thus supplies the magnitudes behind the takeaway of Section 6.2: equilibrium repricing of the whole menu, not the coverage set, decides each product's gain or loss.

## E.7 Induced R&D and Frontier Allocation

Table 7 in Section 6.3 reports each instrument's R&D response as a percent change by frontier role. This subsection adds the firm-identity partition of induced R&D and the terminal frontier, the formal definition of the role attribution, and the signed dollar levels behind the role percents. Table E.5 reports the firm-identity partition: induced R&D and the terminal frontier by fixed firm name, AMD and NVIDIA, not by frontier role; the role partition follows below. The joint-frontier

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

| Instrument | $\Delta I_{\text{AMD}}$ (USDb PV) | $\Delta I_{\text{NVIDIA}}$ (USDb PV) | $\Delta I$ (USDb PV) | $\Delta q^{\max}_{\text{AMD}}$ (%) | $\Delta q^{\max}_{\text{NVIDIA}}$ (%) | $\Delta q^{\max}_{\text{joint}}$ (%) |
|---|---|---|---|---|---|---|
| R&D credit | +0.23 | +0.37 | +0.60 | +4.0 | +3.8 | +3.6 |
| Premium rebate | +0.02 | −0.48 | −0.46 | +0.4 | −5.3 | −2.8 |
| Non-premium rebate | +0.09 | +0.03 | +0.11 | +1.4 | +0.4 | +0.9 |

*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.

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: a firm's advances move the maximum only where it holds or overtakes the frontier, which is why the joint column can sit below both firm columns.

The credit raises R&D at both firms, +$0.23 billion at AMD and +$0.37 billion at NVIDIA, and lifts both terminal frontiers roughly in step, +4.0 percent at AMD and +3.8 at NVIDIA. Under the premium rebate AMD's R&D edges up +$0.02 billion while NVIDIA's falls $0.48 billion, more than the whole net change. The terminal frontiers split the same way, AMD's up +0.4 percent and NVIDIA's down 5.3 percent. Total R&D and the expected joint frontier both fall. Under the non-premium rebate both firms raise R&D, +$0.09 billion at AMD and +$0.03 billion at NVIDIA, and both terminal frontiers edge up, +1.4 percent at AMD and +0.4 at NVIDIA; the total R&D change, +$0.1 billion, is small, and the expected joint frontier still edges up. The firm split already suggests that the premium rebate's frontier retreat is concentrated where the frontier is held; the role allocation below makes that reading exact.

The role allocation is a different accounting partition: when a policy changes R&D at a given path-period state, is that change attributed to the firm classified as the leader or the follower on the estimated-equilibrium path? For role $r \in \{L, F\}$,

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

where $\Delta I_{ft}$ 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; the levels are the numerators behind

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

| Instrument | Leader $\Delta I$ (USDb PV) | Follower $\Delta I$ (USDb PV) | Total (USDb PV) |
|---|---|---|---|
| R&D credit | +0.302 | +0.299 | +0.601 |
| Premium rebate | −0.314 | −0.141 | −0.455 |
| Non-premium rebate | +0.086 | +0.029 | +0.114 |

*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, and totals may differ from the sum of rounded role entries; Table 7 reports the same role allocation as a percent change in role-attributed R&D.

the role columns of Table 7. The credit raises R&D almost equally across the two roles, about \$0.30 billion at each. The premium rebate's cut lands more heavily on the leader role, \$0.31 billion of the \$0.46 billion withdrawn against \$0.14 billion on the follower, the tilt Section 6.3 traces to frontier position. Under the non-premium rebate both roles rise modestly, +\$0.09 billion at the leader and +\$0.03 billion at the follower, and the firm-identity partition of Table E.5 carries the same structure, a rise at both AMD and NVIDIA. 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

This appendix collects the robustness checks behind the demand estimates and the R&D-cost scale of Section 4 and the returns to R&D of Section 5. Each check changes one primitive of the maintained specification and re-estimates the block that owns it: Appendix F.1 varies the demand specification and carries each variant through the recovered marginal costs to the tier markups, and Appendix F.2 varies the dynamic specification, re-estimates the R&D-cost scale, and re-computes the returns to firm R&D under each variant. The estimates are robust overall: in every variant the mean own-price elasticity stays near 3, the R&D-cost scale keeps the economic cost of R&D well above the accounting cost, and, with the demand system held fixed across the dynamic variants, the social return to R&D remains roughly three times the net private return.

## F.1 Demand-Specification Robustness

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, the recovered marginal costs are projected on the rich cost specification of Section 4.2.

Table F.1 varies these choices one at a time. 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. I judge the variants on the downstream objects they feed rather than on coefficient stability alone, so the table carries each through the recovered marginal costs and reports the tier Lerner indices.

The maintained specification estimates a price coefficient of $-17.83$ (standard error 5.30) and a quality coefficient of 60.51 (12.01), a mean own-price elasticity of 3.21, and tier Lerner indices of 0.608, 0.374, and 0.289 from low-end to premium. Dropping the own-quality moment moves the coefficient pair most, to $-25.07$ and 76.86, yet the mean elasticity is essentially unchanged at 3.26. The estimation-sample and market-size variants stay within one standard error of the maintained coefficients and give elasticities of 2.95 and 2.99. The Lerner indices move more: across the four specifications the low-end rate runs from 0.48 to 0.64, the midrange rate from 0.34 to 0.40, and the premium rate from 0.29 to 0.42. The pattern is that the variants move the scale of the coefficient pair more than anything the rest of the paper consumes: the elasticities that discipline the recovered marginal costs stay between 2.95 and 3.26, and the tier markups move within moderate ranges around the maintained values. The demand-stage inputs to the rest of the paper do not hinge on the moment set, the estimation sample, or the market-size path.

## F.2 Dynamic Specification and Estimation

This subsection reports the sensitivity behind the R&D-cost scale of Section 4.3.4 and the returns to R&D of Section 5: each variant changes one primitive of the dynamic specification, re-estimates the scale, and repeats the full-industry return experiment. The sensitivity analysis covers the cost scale and returns; it does not re-solve the policy exercises of Section 6 under each variant, as the lessons drawn there concern the underlying mechanism rather than the specific magnitudes and ranking. The maintained specification pairs the arrival-rate hazard with log-concave R&D exposure over the

**Table F.1.** Demand-specification robustness

| Specification | Price coef. (SE) | Quality coef. (SE) | Mean own elasticity $\lvert\varepsilon_{jj}\rvert$ | Lerner index low-end / midrange / premium |
|---|---|---|---|---|
| Maintained specification | −17.83 (5.30) | 60.51 (12.01) | 3.21 | 0.608 / 0.374 / 0.289 |
| Alternative 1 | −25.07 (7.08) | 76.86 (16.67) | 3.26 | 0.479 / 0.344 / 0.424 |
| Alternative 2 | −17.38 (4.68) | 52.17 (8.95) | 2.95 | 0.635 / 0.404 / 0.319 |
| Alternative 3 | −18.85 (4.79) | 55.97 (8.50) | 2.99 | 0.595 / 0.384 / 0.339 |

*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; and Alternative 3 replaces the time-varying market size with a constant. 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.

2011Q1–2020Q4 estimation window and closes the exogenous states five years after the sample ends. Hazard and jump parameters are estimated offline on the full 2006–2020 innovation history.

Table F.2 varies one primitive per row. The curvature row lowers the exposure curvature to $\kappa = 0.25$ within the maintained hazard family. The Poisson row holds the curvature at its maintained $\kappa = 0.5$ and replaces the saturating map of equation 3.9, $h_f = r_{ft}/(1 + r_{ft})$, with the Poisson map $h_f = 1 - \exp(-r_{ft})$, the arrival rate $r_{ft}$ entering as before. Both change the mapping from R&D and frontier position to innovation arrival. The hazard-window row instead fits the hazard and the conditional jump distribution to the post-2010 history alone and rematches $\chi$ to the unchanged R&D-intensity moment. The plateau rows defer the exogenous-state plateau to seven and ten years after the sample end, which changes continuation values through the levels at which the exogenous states settle. The return columns repeat the full-industry experiment of Section 5, in which both firms' R&D shuts down at once, at each row's re-estimated scale and re-solved equilibrium.

The plateau rows vary an assumption rather than an estimate: the in-sample drifts of the outside option and the process-innovation state are data inputs, while 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.

Within the hazard family fit to the full history, the estimate barely moves. The maintained specification gives a scale of 5.13 with an asymptotic standard error of 0.66; lowering the exposure curvature gives 5.23 and the Poisson family 5.21, and the returns under both stay close to the

**Table F.2.** Sensitivity of the R&D-cost scale and the returns to R&D

| | | Full-industry return per R&D dollar | | |
|---|---|---|---|---|
| Specification | $\widehat{\chi}$ (SE) | Net private | Social | Wedge |
| Maintained specification | 5.13 (0.66) | 3.00 | 8.95 | 5.95 |
| Exposure curvature 0.25 | 5.23 (0.63) | 2.93 | 8.83 | 5.90 |
| Poisson arrival hazard | 5.21 (0.51) | 2.95 | 8.91 | 5.96 |
| Post-2010 hazard window | 3.15 (0.44) | 3.53 | 10.69 | 7.16 |
| Plateau seven years after the sample | 5.46 (1.00) | 2.90 | 8.67 | 5.77 |
| Plateau ten years after the sample | 5.40 (0.85) | 2.81 | 8.50 | 5.69 |

*Notes:* Each row re-estimates the R&D-cost scale $\chi$ from the average R&D-intensity moment after one change to the maintained specification, described in the text, and repeats the full-industry experiment of Table 3, in which both firms' R&D shuts down at once, at the row's re-estimated scale; the return columns are defined as in that table's note. Parentheses report asymptotic standard errors, computed by the delta method.

maintained values of \$3.00 net private, \$8.95 social, and a \$5.95 wedge per R&D dollar. Across the eight hazard specifications estimated on the full history, the three printed here and five further variants (exposure curvature 0.60, the Poisson hazard at exposure curvatures 0.25 and 1.00, and the identity and power exposure forms), the scale ranges from 4.9 to over 5.2, the net private return from about \$2.7 to \$3.0, the social return from about \$8.5 to \$9.0, and the wedge from about \$5.8 to \$6.0. The hazard window is the one primitive that moves the estimate visibly: fitting the hazard to the post-2010 history alone lowers the scale to 3.15, with standard error 0.44, and raises the returns to \$3.53 net private and \$10.69 social, a wedge of \$7.16, because a lower cost scale makes the observed R&D cheaper and each dollar of it buys more. Deferring the plateau to seven or ten years after the sample raises the scale slightly, to 5.46 and 5.40. I also estimated a version of the model with firm-specific R&D-cost scales; the data do not identify the firm-specific scales precisely, and I therefore keep the more parsimonious single-scale specification.

# 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 |
| $\gamma$; $\bar{\alpha}$, $\sigma_{\text{price}}$; $\psi_{f(j)}$ | Demand parameters: quality coefficient; mean price coefficient and price-sensitivity dispersion; firm fixed effect |
| $\xi_{jt}$, $\varepsilon_{jt}^{\text{MC}}$ | Unobserved demand and marginal-cost shocks ($\lvert\varepsilon_{jj}\rvert$ instead denotes the mean absolute own-price elasticity) |
| $\Delta_t$ | Share Jacobian, $(j,k)$ entry $\partial s_{kt}/\partial p_{jt}$ ($\Delta_u$ is the outside-option drift; a bare $\Delta$ marks a change from the estimated equilibrium) |
| $\Omega_t$ | Ownership matrix, $(j,k)$ entry one when $j$ and $k$ share a firm |
| $\text{MC}_{jt}$, $\widehat{\text{MC}}_t$ | 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 $z_t$ |
| $\pi_{jt}$ | 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, $x_t = (q_{1t}^{\max}, q_{2t}^{\max}, u_{0t}, z_t, M_t)$ |
| $q_{ft}^{\max}$; $q_t^{\max}$ | Firm $f$'s quality frontier; the industry frontier, the higher of the two firms' frontier qualities |
| $z_t$, $z_{\min}$ | 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 |
| $M_t$; $\rho_M$, $M_\infty$ | Addressable desktop market size, in millions of units; its persistence and long-run level |
| $\chi$ | R&D-cost scale |
| $\iota_{ft}$ | 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 dollars; it enters the innovation hazard measured in hundreds of millions, the scale on which the exposure curvature $\kappa$ is defined |
| $h_f$ | 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 |
| $a_f$, $\varphi$ | Firm-specific success-rate intercept and common fishing-out parameter |
| $g(\cdot)$; $\kappa$ | Log-concave R&D exposure in the innovation hazard, $g(x) = \kappa^{-1}\log(1 + \kappa x)$, and its curvature |
| $m_J$, $v_J$ | Log-mean and log-variance of the frontier jump-size distribution |
| $V_f$ | Firm $f$'s value function in the dynamic game (equation (3.8)) |
| $\beta$ | 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 |
| $\tau$; $b_d$ | 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.