Public Procurement, the Nature of Innovation, and Growth: Firm-Level Evidence from Germany[*]

Johar Arrieta Vidal[†] Jorge Zavala[‡]

June 15, 2026

Abstract

German firms invest in R&D, yet fewer introduce products new to the market. Public procurement may contribute to the pattern: contracts that reward specification-compliant delivery over frontier novelty may reduce the payoff to frontier innovation and redirect innovative effort toward incremental product lines. In the data, a one-standard-deviation increase in procurement exposure is associated with a 0.20 percentage-point decline in market-novel product introduction and a 0.38 percentage-point increase in existing-product revenue. Whether such composition shifts slow growth depends on general-equilibrium forces. We build and estimate a quality-ladder endogenous-growth model incorporating this mechanism. Redirecting procurement funds to an R&D tax credit raises growth by 12 basis points per year and delivers a welfare gain of 3.2 percent of permanent consumption.

JEL codes: D22, H57, O31, O38, O41.

Keywords: public procurement, innovation, quality ladders, novel product.

---

[*] We thank Anton Bobrov, Kathryn Dominguez, Bastian Krieger, Imryoung Jeong, John Leahy, Matthew Schwartzman, Stephen Terry, Toni Whited, Wei Xiang and seminar participants at the University of Michigan. We thank Jakob Ehlich, Heidrun Förster and Sandra Gottschalk for providing us with access to the Mannheim Innovation Panel data via ZEW Mannheim and to Stefanie Wolter for granting us access to the BHP via the IAB Research Data Centre. All errors are our own.

[†] University of Michigan, Department of Economics.

[‡] WZB Berlin / Humboldt University.

1 Introduction

German firms continue to invest in innovation, yet fewer introduce products new to the market (BMBF/ZEW, 2024). Figure 1 documents the pattern. The market-novelty rate falls from about 26 percent in the early 2000s to 8 percent by 2023, while innovation intensity remains between 3 and 5 percent of turnover.

Public procurement may help explain the pattern. At roughly 15 percent of gross domestic product (GDP), German public procurement gives the state a large role in private demand. Procurement contracts specify compatibility standards, certification requirements, delivery schedules, and technical thresholds. Contract requirements reward delivery of known products over frontier novelty (products not yet offered in the market). At that scale, procurement may shape not only the level of innovative activity but also its composition.

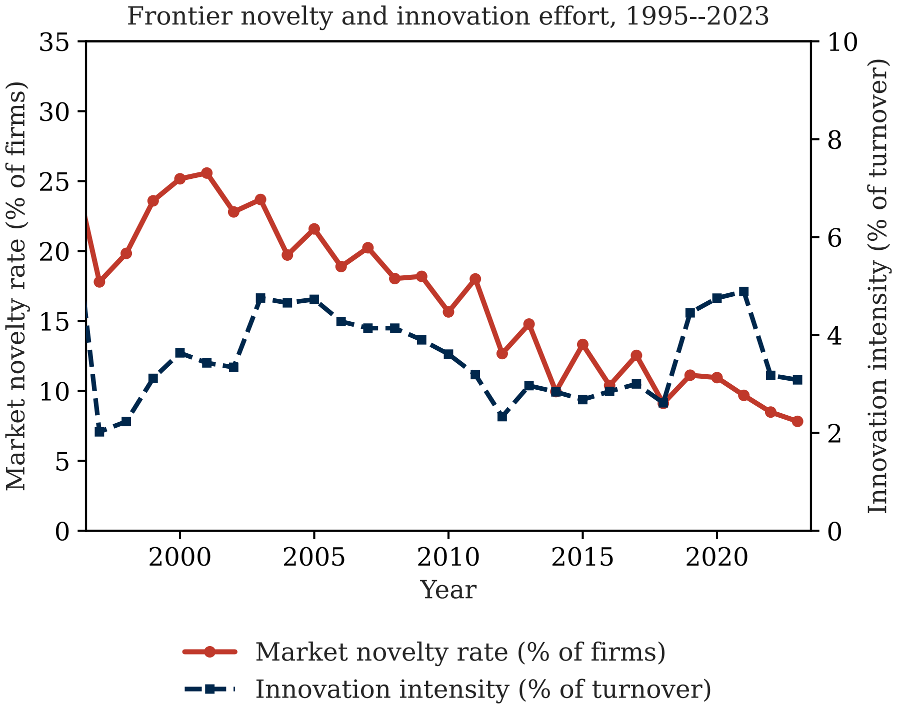

Figure 1: The motivating puzzle in German innovation data

Notes: The red line plots the market-novelty rate, defined as the share of firms introducing products new to the market, on the left axis. The blue dashed line plots innovation intensity, defined as innovation expenditure divided by turnover, on the right axis.

Source: Mannheim Innovation Panel, 1995–2023.

We link public procurement records to firm-level innovation surveys for German manufacturing and services. The Mannheim Innovation Panel (MIP), administered by ZEW Mannheim on behalf of the German Federal Ministry of Education and Research, provides firm-year measures of market novelty, R&D expenditure, process innovation, and new-product revenue for a representative sample of firms over 2009 to 2023. We combine the survey with Tenders Electronic Daily (TED) award records, which report publicly disclosed German procurement awards at the contract level: product category, award value, buyer identity, buyer category, and technical-requirement fields. The sample covers a large population of contract awards, one per contract and winning firm, over 2009–2023. We also merge to ORBIS, a firm-level financial database covering European companies, for roughly 15,000 matched contract winners, providing employment, turnover, productivity, and profitability over the same period. Each data source addresses a distinct margin. The MIP distinguishes market-novel product innovation from revenue from existing products, process innovation, and R&D inputs. The contract records identify which sectors receive public demand and which buyer types and contract categories generate that demand. ORBIS lets us examine whether contract winners expand in scale, productivity, or profitability.

Empirically, we estimate panel regressions with firm and year fixed effects relating sector-level procurement exposure to firm-year innovation outcomes. Variation in sector-level exposure over time, within firms, drives the estimates. Three facts emerge from the analysis.

First, procurement exposure is associated with lower frontier product innovation and higher revenue from existing products. Firms in more exposed sectors introduce market-novel products less often, and the same exposure is associated with a higher share of turnover from existing products. The pattern extends to cost-reducing innovations: the revenue share from cost-reducing products is also lower in more exposed sectors, while the rate of cost-reducing activity is unchanged. The results document a composition shift in commercial orientation, not a contraction of innovative effort.

Second, the composition pattern concentrates in supplies and services contracts and among central-government buyers; works contracts and regional or local buyers show no significant effect on either margin. The absence of the pattern in public works contracts, where the deliverable is a physical structure rather than a reproducible product, is consistent with product specification, rather than public demand in general, driving the composition shift.

Third, contract winners expand in scale without productivity or profitability gains: employment and turnover rise, while revenue per worker and return on capital show no statistically significant change. Contract winners are also less frontier-oriented than the average firm at entry, and additional contracts are associated with scale expansion without efficiency gains.

Contract design offers an interpretation. Procurement specifications that require certification, compatibility standards, and delivery schedules favor existing products over frontier novelty. In the most specification-intensive categories in the data (medical equipment, software and IT services, transport vehicles), contract specifications favor certifiability over diagnostic performance, backward compatibility over new architectures, and proven capacity over novel technology. The effect runs through industries where the same firms serve both public and private buyers.

The within-firm estimates document the composition shift but do not reveal its aggregate growth implications or the welfare consequences of alternative procurement designs. To formalize the mechanism, we build on the Klette and Kortum (2004) endogenous-growth framework, in which firms hold portfolios of product lines and compete to improve product quality step by step (quality-ladder innovation).[1] We add one margin: each period, firms decide not only how much to innovate but also what kind of innovation to pursue. Procurement enters as a shift in the relative payoff: frontier product innovation becomes less profitable than specification-compliant delivery, moving effort toward incremental improvements. The model therefore rationalizes the patterns in the German innovation and procurement microdata: procurement raises equilibrium innovation effort while narrowing the frontier share, matching the observed pattern: lower market novelty and higher existing-product revenue in more exposed sectors.

After building the model, we turn to quantification using the same microdata. We estimate four structural parameters by overidentified generalized method of moments (GMM), targeting five empirical statistics from the microdata: the frontier share of innovation, the frontier revenue share, firm size, creative destruction (the rate at which new innovations displace existing product lines), and the R&D personnel share. Estimated quality steps and direction costs fall within the ranges reported in García-Macià et al. (2019) and Akcigit and Kerr (2018). At the estimated parameters, the procurement wedge accounts for the observed composition redirection toward incremental lines.

We conduct four counterfactual exercises to quantify the growth and welfare consequences of procurement design. The primary exercise redirects procurement funds to a broad-based R&D tax credit modeled on Germany's *Forschungszulage* (a rule-based credit on eligible R&D labor costs introduced in January 2020). Growth rises by 12 basis points per year, and welfare rises by 3.21 percent of permanent consumption; the gain flows through long-run growth because removing procurement also eliminates the fiscal cost. Second, a budget-neutral specification reform replaces certification-based contract requirements with open performance criteria, allowing frontier products to compete for public contracts on equal terms; growth rises by 2.81 basis points. Third, a monitoring reform that reduces buyer-side uncertainty about frontier-product performance, making novel proposals less risky to contract for, raises growth by 1.71 basis points. Finally, reducing procurement demand to its pre-2014 level restores part of the frontier premium, with a growth gain of 1.23 basis points. These gains are moderately large. Against a German TFP growth baseline

---

[1] Akcigit and Kerr (2018), Akcigit and Ates (2023), and Acemoglu et al. (2018) extend the framework.

of roughly 50 basis points per year, the 12-basis-point growth gain represents about one-quarter of trend growth; the 3.21 percent welfare gain exceeds the estimated welfare cost of business cycles (Krusell et al., 2009) and falls within the range of trade-liberalization gains for small open economies. Across exercises, how procurement is designed shapes both the scale and the direction of innovative activity.

We contribute to four strands of the literature. First, we add a demand-composition mechanism to work on innovation composition and growth (Klette and Kortum, 2004; Aghion and Howitt, 1992; García-Macià et al., 2019; Akcigit and Kerr, 2018; Akcigit and Ates, 2023; Acemoglu et al., 2022; Aghion et al., 2023; Juhász et al., 2024). García-Macià et al. (2019) document that much aggregate growth comes from incumbents improving existing products rather than from creative destruction; Acemoglu et al. (2022) establish that the frontier-versus-incremental margin is central for long-run growth when incremental innovations face diminishing returns within a technology cluster. Public procurement can shift the innovation mix toward incremental activity through a demand channel, complementing the supply-side mechanisms emphasized in Bloom et al. (2020).

Second, we contribute to the literature on demand-side innovation policy (Griffith et al., 2010; Slavtchev and Wiederhold, 2016; Czarnitzki et al., 2020; Krieger and Zipperer, 2022; Krieger et al., 2024; Moretti et al., 2025; Chiappinelli et al., 2025; Lerche, 2025; Dechezleprêtre et al., 2023; Takalo et al., 2026; Chen et al., 2021; Chen and Xu, 2023). On the procurement side, Krieger et al. (2024) document that price-criterion tenders reduce product innovation among German firms, and Czarnitzki et al. (2020) show that even innovation-directed contracts tilt activity toward incremental rather than frontier novelty. On the subsidies side, Takalo et al. (2026) establish that rule-based tax credits reach a broader firm base than discretionary grants, motivating our counterfactual of redirecting procurement funds to Germany's *Forschungszulage*; in a related German setting, Lerche (2025) reports that investment tax incentives generate employment multipliers as large as the direct effect through local spillovers, indicating that tax instruments carry broader aggregate effects than the direct innovation margin. For German firms, procurement and subsidies operate on distinct margins: the composition gap is the relevant welfare dimension.

Third, we build on work on procurement contract design and innovation (Howell et al., 2025; Clemens and Rogers, 2026). Howell et al. (2025) document that open-format defense awards raise patenting and commercial adoption while specification-based awards generate lock-in; Clemens and Rogers (2026) demonstrate historically that fixed-price contracts redirect invention toward cost reduction while quality-focused contracts shift patenting toward buyer-preferred attributes. We extend both mechanisms from individual contract variation to aggregate sector-level exposure, measuring the economy-wide composition effect.

Finally, we contribute a demand-side channel to the misallocation literature (Acemoglu et al., 2018; Terry, 2023; Bachmann et al., 2026; Lehr, 2024, 2025). Government procurement changes not only the level of innovative effort but also its direction, adding a public-demand mechanism to accounts of how frictions in factor markets, R&D allocation, and market power slow long-run growth.

The paper proceeds as follows. Section 2 presents the data and empirical evidence. Section 3 develops the model. Section 4 reports the quantitative analysis. Section 5 presents the counterfactuals and robustness exercises. Section 6 concludes.

2 Data and Empirical Evidence

The empirical strategy links public procurement records to firm-level innovation surveys and financial accounts. First, the Mannheim Innovation Panel (MIP) tracks how innovation composition co-varies with sector-level procurement exposure over time. Second, Tenders Electronic Daily (TED) supplies the contract records from which we construct the exposure measure. Third, ORBIS matches procurement award records to firm-level financial accounts, documenting how winners differ in scale, productivity, and profitability. Fourth, the IAB Establishment History Panel provides entry and exit rates for the structural model's creative-destruction calibration. The MIP and ORBIS evidence serve distinct roles: the MIP captures average patterns across all firms in a sector (winners, unsuccessful bidders, and non-participants alike), while ORBIS describes winner outcomes directly.

2.1 Data Sources

2.1.1 Mannheim Innovation Panel

The Mannheim Innovation Panel is an annual firm-level survey administered by ZEW Mannheim on behalf of the German Federal Ministry of Education and Research. We use the 2009–2023 releases, which yield 80,028 firm-year observations across 23,547 firms. Main outcomes cover market-novel product innovation, unchanged-product revenue share, the market-novel share of new-product revenue, process innovation, and cost-reduction activity and revenue share. Process-innovation regressions use the full 2009–2023 sample.[2] Appendix Table A2 reports summary statistics for all variables, including world-first product innovation (2009–2018), R&D intensity, R&D personnel share, new-product revenue share, and public R&D subsidies used in robustness analyses.

---

[2] In 2019, the MIP adopted the Oslo-IV definition of process innovation, which expanded the concept and raised the reported rate from roughly 28 to 53 percent. Year fixed effects absorb this level shift in our analysis below.

2.1.2 Public Procurement Contracts

Tenders Electronic Daily (TED), the European Commission's official procurement journal, publishes award notices from German public buyers (federal, state, and municipal) that EU transparency rules require to be disclosed above specified value thresholds. The main sample contains 519,990 contract-award observations, one per contract and winning firm, covering 508,695 unique contracts from 2009 to 2023; 271,713 map to MIP-covered industries. Each observation reports the product code, award year, buyer region, buyer category, and number of submitted bids. The records also retain contract-design fields such as procedure type and technical-requirement text length. For the pre-period, we use TED award notices from 2006–2008 to construct pre-period exposure.

2.1.3 ORBIS

ORBIS is a firm-level financial database covering European companies, maintained by Bureau van Dijk. It contains annual balance sheet and profit-and-loss accounts, employment, and ownership structure for several million firms, with broad coverage of German companies across size classes and industries. We use it in two ways. First, we construct a German firm panel over 2007–2023 with employment, turnover, labor productivity, profitability, age, foreign ownership, and multinational status. Second, we match procurement winners to firm identifiers using a fuzzy name-and-postcode crosswalk, yielding 41,605 matched bidder–firm pairs. The matched panel contains 126,700 firm-year observations across 15,493 firms and supports the scale and selection analyses. Appendix Tables B1–B2 report the evidence by firm size, procurement dependency, and ownership structure.

2.1.4 IAB Establishment History Panel

The IAB Establishment History Panel (BHP) is an annual administrative panel covering the near-universe of German establishments with at least one employee subject to social security contributions.[3] We use the entry and exit event files for 2009–2023; across 22.5 million establishment-years, the annual exit rate is 5.1 percent. We use the exit rate to set the creative destruction rate in the quality-ladder model (Section 3).

2.2 Empirical Strategy

Procurement markets differ in the degree to which they draw firms that also serve private buyers. Contracts for defense platforms or classified services draw on a different firm base than contracts for engineering services, software, or industrial equipment, industries where the same firms serve

---

[3] We obtain the data through the IAB Research Data Centre under a restricted data use agreement. The BHP uses establishment-type classification codes to separate genuine new foundings and full closures from administrative re-identifications, spin-offs, and corporate restructurings.

both government and private buyers. In such sectors, procurement exposure is associated with lower frontier innovation (Krieger et al., 2024; Chiappinelli et al., 2025). The analysis uses variation in sector-level procurement exposure across these contestable industries and over time within sectors.

We assign each contract to one of the twenty-one MIP industries by mapping its EU product code to the corresponding two-digit German industry classification. Procurement exposure is $\operatorname{arcsinh}(\text{contracts}_{s,t})$, where $\text{contracts}_{s,t}$ is the count of awarded contracts in sector s and year t; the distribution is right-skewed, and the arcsinh transformation behaves similarly to a logarithm for large values. Current MIP releases do not link survey firms to individual contract awards; hence we aggregate contracts to sector-year counts and assign exposure at the industry level. Figure 2 documents concentration in manufacturing and professional services and substantial within-sector time variation; across 315 sector-years the measure has a mean of 3.1 and a standard deviation of 3.5.

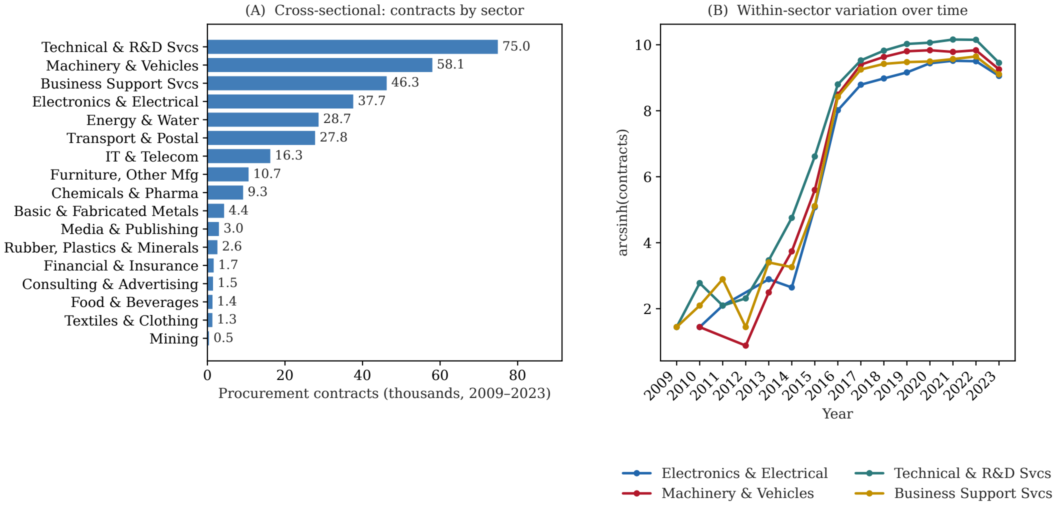

Figure 2: Procurement exposure across MIP-covered industries, 2009–2023

Notes: Panel A reports total awarded contracts, in thousands, across MIP-covered industries, aggregated over 2009–2023. Panel B plots arcsinh(contracts), the exposure measure, by year for the four largest sectors by total contract volume.

Source: Authors' calculations using TED, 2009–2023.

The baseline specification is

$$Y_{ist} = \alpha_i + \delta_t + \beta \operatorname{arcsinh}(\text{contracts}_{s,t}) + \varepsilon_{ist}, \tag{1}$$

where Y_{ist} denotes a firm-level innovation outcome for firm i in industry s and year t: binary indicators for product or process innovation, or continuous measures such as revenue shares; α_i is a firm fixed effect and δ_t a year fixed effect.[4] The coefficient β has percentage-point units for binary outcomes and natural units for continuous outcomes.

Firm fixed effects absorb time-invariant firm characteristics, including the permanent propensity to serve public buyers; year effects absorb aggregate shocks common to all firms. We cluster standard errors at the industry level, yielding twenty-one clusters. Since exposure varies at the sector-year level, β measures the average within-firm association between industry procurement exposure and innovation outcomes across all firms in a sector (contract winners, unsuccessful bidders, and non-participants alike).

2.3 Main Results

Table 1 reports the baseline specification across three innovation outcomes.

Table 1: Procurement Exposure and Innovation Outcomes

	(1) Market novelty (pp)	(2) Unchanged-product revenue (%)	(3) Cost-reduction revenue (%)
arcsinh(contracts$_{st}$)	-0.199^{*}	0.379^{**}	-0.047^{*}
	(0.101)	(0.137)	(0.024)
Observations	66,980	25,962	40,635
Industries	21	21	21
Firm FE	✓	✓	✓
Year FE	✓	✓	✓

Notes: Each cell reports a coefficient from a within-firm OLS regression with firm and year fixed effects. $\text{arcsinh}(\text{contracts}_{st})$ is the inverse hyperbolic sine of awarded contract counts in the firm's two-digit MIP sector and year (2009–2023). Standard errors, in parentheses, are clustered by MIP sector (21 industries). Binary outcomes are multiplied by 100 (pp). Col. (1): binary for a product new to the market. Col. (2): share of total turnover from products sold without innovation in the reference year. Col. (3): share of turnover from innovations that reduce unit production costs. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.
Source: Authors' calculations using the Mannheim Innovation Panel and TED.

Column (1) shows that a one-unit increase in $\text{arcsinh}(\text{contracts})$ corresponds to a 0.199 percentage-point lower probability of introducing a market-novel product. Against a sample mean of 16.2 percent, a one-standard-deviation increase in exposure (3.5 units) corresponds to a 0.70 percentage-point decline, about 4 percent of the mean. Column (2) shows the complementary revenue pattern: the share of turnover from existing products is 0.379 percentage points higher per unit of exposure; against a sample mean of 84.5 percent, a one-standard-deviation increase corresponds to a 1.3

---

[4] Since exposure varies at the sector-year level, the industry subscript s denotes the sector of firm i.

percentage-point difference, about 2 percent of the mean. Together, columns (1) and (2) document a shift in the product mix toward existing products and away from frontier novelty.

The shift operates through the extensive margin (whether a firm introduces any market-novel product) and the revenue mix; the market-novel share within new-product revenue and process innovation are not significantly associated with procurement exposure, ruling out a reallocation toward process activity.

Column (3) traces the cost-reduction margin: the share of turnover from cost-reducing innovations is 0.047 percentage points lower per unit of exposure; against a sample mean of 2.5 percent, a one-standard-deviation increase in exposure corresponds to a 0.16 percentage-point decline, about 7 percent of the mean. Higher procurement exposure is associated with lower revenue from both frontier and cost-reducing innovative lines. Appendix Table A1 shows that the broader product-innovation indicator of Krieger et al. (2024) (any new or improved product) is unrelated to procurement exposure, while market novelty falls. The pattern concerns what firms innovate, not whether they innovate.

2.4 Heterogeneity Across Procurement Environments

Table 2 reports the two main outcomes from Table 1 separately for different contract types and buyer categories. The splits are descriptive: procurement environments co-vary along multiple dimensions, and the patterns do not isolate a single mechanism.

Table 2: Heterogeneity Across Procurement Environments

	(1) Market novelty (pp)	(2) Unchanged-product revenue (%)
Panel A. Contract type		
Supplies & services	-0.211^{**}	0.393^{***}
	(0.099)	(0.137)
Works	-0.010	0.180
	(0.142)	(0.120)
Panel B. Buyer type		
Central government	-0.291^{**}	0.518^{**}
	(0.131)	(0.194)
Regional / local	-0.043	0.232
	(0.137)	(0.148)
Utilities	-0.086	0.313
	(0.146)	(0.195)
Firm and year FE	Yes	Yes
Observations	66,980	

Notes: Each cell reports the coefficient from a separate within-firm OLS regression with firm and year fixed effects. In Panel A, the exposure measure is arcsinh(sector-year contract count) for the indicated contract type. In Panel B, the exposure measure is arcsinh(sector-year contract count) for the indicated buyer category. All estimates use the final 21-sector, 2009–2023 panel; observations report the count for the market-novelty outcome (other outcomes have smaller samples due to wave availability). Standard errors in parentheses are clustered by industry (21 industries). Supplies and services account for 97.8 % of contracts in the MIP-covered sample, so the supply-type contrast has limited discriminating power. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.
Source: Authors' calculations using the Mannheim Innovation Panel and TED.

Panel A splits by contract type. For supplies-and-services contracts, market novelty is lower (-0.21) and existing-product revenue is higher (0.39), matching the aggregate pattern in Table 1. For works contracts, market novelty is close to zero (-0.01), and existing-product revenue is not statistically significant. Supplies and services account for 97.8 percent of contracts in the MIP-covered sample; the null result for works contracts shows the aggregate pattern does not reflect construction or infrastructure procurement.

Panel B splits by buyer category. The composition pattern is stronger for central-government procurement: market novelty is 0.29 percentage points lower and existing-product revenue is 0.52 percentage points higher. Regional and local government procurement and utilities show no significant effect on either margin. Central-government procurement is likely more administratively standardized than regional or local procurement, but it also differs in contract size, buyer composition, and product mix.

2.5 Winner Outcomes: Scale without Productivity

The sector-level regressions average over winners, unsuccessful bidders, and non-participants. ORBIS lets us examine winners directly. Table 3 reports within-firm estimates for the matched ORBIS panel of procurement winners over 2009–2023.

Table 3: Winner Outcomes: Scale without Productivity

	(1) Employment (log)	(2) Turnover (log)	(3) Labor productivity (log)	(4) Return on capital
arcsinh(contracts won)	0.025***	0.028***	0.007	0.001
	(0.002)	(0.007)	(0.007)	(0.001)
Obs.	88,363	38,525	33,846	28,470

Notes: Sample is the ORBIS winner panel, 2009–2023, comprising 15,493 firms matched to procurement award records. Each column reports a within-firm OLS regression of the indicated outcome on $\operatorname{arcsinh}(\text{contracts won})$, with firm and year fixed effects. Standard errors, in parentheses, are clustered by NACE section. Employment, turnover, and labor productivity are in natural logarithms; labor productivity is turnover divided by the number of employees. Return on capital is pretax profit divided by total assets. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.
Source: Authors' calculations using TED matched to ORBIS.

Columns (1) and (2) show that a one-unit increase in $\operatorname{arcsinh}(\text{contracts won})$ corresponds to 2.5 percent higher employment and 2.8 percent higher turnover. Columns (3) and (4) reveal no statistically significant change in labor productivity or return on capital. Winners have higher scale without higher productivity or profitability. The within-firm estimates capture the response to additional contracts among firms already participating in procurement; the effect of a firm's first contract award requires a separate accounting for selection into the winner pool.

The selection pattern is consistent with the scale-without-novelty interpretation: both young and established firms enter procurement with market novelty below the sample mean, and the pattern holds across firm age, size, and ownership structure (Appendix Figure B1; Appendix Tables B1–B2).

The evidence establishes three facts: (i) firms in sectors with higher procurement exposure are less likely to introduce market-novel products; (ii) they have higher revenue from existing products; and (iii) those that win contracts have higher employment and turnover without productivity gains. Reduced-form sector-level associations do not provide the general-equilibrium infrastructure needed to evaluate procurement design. Section 3 builds a quantitative model that formalizes the mechanism and evaluates procurement design counterfactually.

3 A Quality-Ladder Model with Procurement

We develop a quality-ladder model to organize the empirical pattern: procurement moves innovation away from frontier advances and toward incremental improvements. The framework builds on Klette and Kortum (2004), in which firms operate portfolios of product lines subject to creative destruction, and on the quality-ladder tradition of Aghion and Howitt (1992); Aghion et al. (2023). The choice between radical (frontier) and incremental innovation is related to the framework of Acemoglu et al. (2022), although applied here to a demand-composition setting. Government procurement enters through two channels. It increases the return to innovation effort, and it narrows the frontier premium because incremental lines serve both private and government demand while frontier lines serve private demand only.

3.1 Environment

Time is continuous, $t \in [0, \infty)$. A unit mass of product lines exists, and a single incumbent firm operates each line.

Households. A representative household supplies labor inelastically, owns all firms, and has preferences

$$\int_0^\infty e^{-\rho t} \log C_t \, dt.$$

We choose aggregate private expenditure as the numeraire; labor market clearing determines the wage w. On the balanced-growth path the household budget constraint is

$$C_0 = w\bar{L} + \Pi_{\text{net}} - T, \tag{2}$$

where $w\bar{L}$ is labor income, Π_{net} is aggregate profit income net of innovation costs, and T is the procurement tax defined below.

Production and firm state. A competitive final-good sector aggregates active product lines:

$$Y_t = \left(\int_{\mathcal{N}_t} y_j(t)^{\frac{\varepsilon-1}{\varepsilon}} \, dj \right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{3}$$

where $\mathcal{N}_t$ is the set of active product lines and $\varepsilon > 1$. The CES elasticity does not affect equilibrium markups; Bertrand competition at the quality lead determines the effective markup. Firms produce line-specific output with linear technology:

$$y_j = q_j \ell_j^P.$$

The economy has a fixed labor endowment $\bar{L}$. A firm is described by its portfolio (n_I, n_F), where $n_I \geq 0$ is the number of incremental lines and $n_F \geq 0$ is the number of frontier lines. The two types differ in their quality lead over the next-best competitor:

$$\text{incremental line: } \lambda_I, \qquad \text{frontier line: } \lambda_F > \lambda_I > 1.$$

Government demand and fiscal balance. Government demand $G \geq 0$ is the ratio of government spending per incremental line to private spending, measured in units of private-line output; it is exogenous. Government purchases flow to incremental lines only: public buyers reward compliance with a fixed specification, not quality above the threshold; Appendix D.7 microfounds the asymmetry.

The government produces output with the same technology as private output: one unit of labor at wage w yields one unit of quality-adjusted output. Firms fill government and private orders from the same workforce; the only resource cost of procurement is the production labor required to satisfy government demand.

The government finances procurement through a lump-sum tax on households. The incremental segment has measure $(1 - \mu^*)$, so total government expenditure is

$$T \equiv (1 - \mu^*)G,$$

measured as a fraction of aggregate private output. T is the procurement tax in the household budget constraint (2).

Innovation. The innovation effort rate $x \geq 0$ is the Poisson arrival rate of innovation attempts per product line. The frontier share $\mu \in [0, 1]$ is the probability that a successful innovation creates a frontier line rather than an incremental line. A firm with $N = n_I + n_F$ lines generates new lines at rate xN. Incremental and frontier lines arrive at rates

$$x(1 - \mu)N \qquad \text{and} \qquad x\mu N,$$

respectively. Creative destruction displaces each existing line at rate $\tau > 0$, determined in equilibrium by the aggregate innovation activity of all firms.

Innovation costs. Per-line innovation costs are

$$\Phi(x) = c_x x^{\zeta} w, \qquad c_x > 0,\ \zeta > 1, \tag{4}$$

$$\Psi(\mu) = \frac{\kappa}{\eta}\mu^{\eta}, \qquad \kappa > 0,\ \eta > 1. \tag{5}$$

$\Phi(x)$ is the cost of innovation effort and $\Psi(\mu)$ is the cost of frontier orientation.

Labor allocation. Firms allocate labor across three activities: production, incumbent R&D, and entry. Production labor per line satisfies $\ell_j^P = y_j/q_j$. Incumbent R&D labor per line is

$$\ell_R(x, \mu, w) \equiv \frac{\Phi(x) + \Psi(\mu)}{w}.$$

Outside entrants can displace any incumbent line by paying an entry cost of h_E units of labor; free entry pins $h_E w$ to the expected value of an entrant line. Entry governs the scale of creative destruction; it is separate from the incumbent innovation margin μ, which governs the composition of quality improvements across radical and incremental lines.

3.2 Agent Problems

Pricing. Quality leaders set prices through Bertrand competition against the next-best rival. We assume the Bertrand limit price binds; the quality lead, rather than the unconstrained CES markup, therefore determines the effective markup. A line with quality lead λ_k earns the per-line profit share

$$\bar{\pi}_k = \frac{\lambda_k - 1}{\lambda_k}, \qquad k \in \{I, F\}, \tag{6}$$

with $\bar{\pi}_F > \bar{\pi}_I > 0$.

Private buyers reward quality; both line types serve private demand. Government buyers reward specification compliance: incremental lines serve both private and government demand, while frontier lines serve private demand only. Per-line flow profits are

$$\pi_I(G) = \bar{\pi}_I(1 + G), \tag{7}$$

$$\pi_F(G) = \bar{\pi}_F. \tag{8}$$

The frontier premium equals

$$\pi_F(G) - \pi_I(G) = \bar{\pi}_F - \bar{\pi}_I(1 + G). \tag{9}$$

Higher government demand increases incremental profits and leaves frontier profits unchanged. The frontier premium remains positive for $G < G^* \equiv \bar{\pi}_F/\bar{\pi}_I - 1$.

Firm's Problem. A line's continuation value depends on its type and on G. Appendix D.2 establishes that the firm value function is additive:

$$V(n_I, n_F; G) = n_I v_I(G) + n_F v_F(G). \tag{10}$$

Separability follows from the absence of cross-line interactions: profits, innovation arrivals, and destruction risks operate independently across lines, and the model has no firm-level resource constraint. Portfolio composition affects firm value through (n_I, n_F) because $v_F > v_I$.

Let

$$\bar{v}(\mu; G) \equiv (1-\mu) v_I(G) + \mu v_F(G) \tag{11}$$

denote the expected value of a new line.

A firm with state (n_I, n_F) earns flow profits

$$n_I \pi_I(G) + n_F \pi_F(G) - (n_I + n_F)\Big[\Phi(x) + \Psi(\mu)\Big].$$

Each existing line generates innovation opportunities at rate x. Conditional on arrival, the firm creates an incremental line with probability $1-\mu$ and a frontier line with probability μ. Rivals displace each existing line at rate τ.

The portfolio HJB is

$$\begin{aligned} rV(n_I, n_F; G) = \max_{x,\mu} \Big\{ & \underbrace{n_I \pi_I(G) + n_F \pi_F(G)}_{\text{flow operating profits}} - \underbrace{(n_I + n_F)\Big[\Phi(x) + \Psi(\mu)\Big]}_{\text{innovation and frontier-orientation costs}} \\ & + (n_I + n_F) x \Big[(1-\mu) \underbrace{\Big(V(n_I + 1, n_F; G) - V(n_I, n_F; G)\Big)}_{\text{arrival of a new incremental line}} \\ & + \mu \underbrace{\Big(V(n_I, n_F + 1; G) - V(n_I, n_F; G)\Big)}_{\text{arrival of a new frontier line}} \Big] \\ & + \tau n_I \underbrace{\Big(V(n_I - 1, n_F; G) - V(n_I, n_F; G)\Big)}_{\text{destruction of an incremental line}} \\ & + \tau n_F \underbrace{\Big(V(n_I, n_F - 1; G) - V(n_I, n_F; G)\Big)}_{\text{destruction of a frontier line}} \Big\}. \end{aligned} \tag{12}$$

Substituting (10) into the portfolio HJB reduces the problem to two value equations:

$$(r+\tau)v_I(G) = \pi_I(G) - \Phi(x) - \Psi(\mu) + x\bar{v}(\mu; G), \tag{13}$$

$$(r+\tau)v_F(G) = \pi_F(G) - \Phi(x) - \Psi(\mu) + x\bar{v}(\mu; G). \tag{14}$$

The equations differ only in flow profits. The firm chooses (x, μ) given (v_I, v_F). Appendix D.2 derives the system.

Optimality. The effort first-order condition is

$$c_x \zeta x^{\zeta-1} w = \bar{v}(\mu; G). \tag{15}$$

The unconstrained frontier-choice condition is

$$\kappa \mu^{\eta-1} = x(v_F - v_I). \tag{16}$$

Equation (15) says that firms raise innovation effort until the marginal cost of another arrival equals the expected value of the new line. Equation (16) requires the marginal cost of frontier orientation to equal the value premium from holding a frontier rather than an incremental line. The condition (16) holds on the interior region $\mu^* \in (0,1)$; otherwise $\mu^* \in \{0,1\}$.

3.3 General Equilibrium

The balanced-growth-path equilibrium requires household utility maximization, firm-level effort and innovation-type optimization, free entry, and labor-market clearing.

Log utility gives the balanced-growth Euler equation

$$r = \rho + g. \tag{17}$$

The unit mass of product lines is fixed. Entrants do not create new varieties; they leapfrog incumbents and take over existing product lines. Let τ denote the economy-wide rate at which any given line changes hands through creative destruction, and let e denote the rate at which outside entrants displace incumbents. Since incumbents also displace rivals at rate x per line they hold, the total displacement rate satisfies

$$x + e = \tau. \tag{18}$$

Hence $e = \tau - x$, aggregate entry labor is $h_E(\tau - x)$, and equilibrium requires $x < \tau$.

Entrants inherit the equilibrium composition: the displaced line is upgraded to frontier quality with probability μ^*. Free entry implies

$$\bar{v}(\mu^*; G) = wh_E. \tag{19}$$

Production labor accounts for private and government demand. Incremental lines serve private and government demand; frontier lines serve private demand only. Expected production labor per line is

$$L_P(\mu, G, w) = \frac{(1-\mu)(1-\bar{\pi}_I)(1+G) + \mu(1-\bar{\pi}_F)}{w}. \tag{20}$$

R&D labor per line is $\ell_R(x, \mu, w) = [\Phi(x) + \Psi(\mu)]/w$. Labor-market clearing requires

$$\bar{L} = L_P + \ell_R + h_E(\tau - x). \tag{21}$$

Balanced-growth output growth g is

$$g = x\Big[(1-\mu)\log\lambda_I + \mu\log\lambda_F\Big]. \tag{22}$$

Definition 3.1 (General equilibrium)**.** Given government spending G, labor endowment $\bar{L}$, household discount rate ρ, and innovation cost scale c_x, with the interest rate $r = \rho + g$ from (17), a balanced-growth-path general equilibrium is a collection of value functions (v_I, v_F) for incremental and frontier product lines, a firm innovation effort rate x and frontier-innovation probability μ, an economy-wide wage w, and a creative destruction rate τ, with $x < \tau$, such that:

(i) *Firm value functions.* Incremental and frontier line values satisfy the Bellman equations

$$(r+\tau)v_I = \pi_I(G) - \Phi(x) - \Psi(\mu) + x\bar{v}(\mu; G), \tag{23}$$
$$(r+\tau)v_F = \pi_F(G) - \Phi(x) - \Psi(\mu) + x\bar{v}(\mu; G). \tag{24}$$

(ii) *Optimal innovation effort.* The marginal cost of an additional unit of effort equals the expected value of a new line:

$$c_x\zeta x^{\zeta-1}w = \bar{v}(\mu; G). \tag{25}$$

(iii) *Optimal innovation composition.* The marginal cost of increasing the frontier-innovation probability equals the value premium from holding a frontier rather than an incremental line:

$$\kappa\mu^{\eta-1} = x(v_F - v_I), \quad \mu \in [0, 1]. \tag{26}$$

(iv) *Free entry.* The expected value of a newly created line equals its creation cost:

$$\bar{v}(\mu; G) = wh_E. \tag{27}$$

(v) *Labor-market clearing.* Production labor, R&D labor, and entry labor exhaust the endowment:

$$\bar{L} = L_P(\mu, G, w) + \ell_R(x, \mu, w) + h_E(\tau - x). \tag{28}$$

Computationally, the equilibrium is a system of six equations in six unknowns $(v_I, v_F, x, \mu, w, \tau)$. Given a trial wage w, the Bellman equations (23)–(24), the optimality conditions (25)–(26), and the free-entry condition (27) jointly determine $(v_I, v_F, x^*, \mu^*, \tau)$ in closed form. Bisection over w then satisfies labor-market clearing (28). Appendix D.3 describes the solution algorithm. Counterfactuals vary G and other objects, and re-solve the full system, allowing μ^*, x^*, and g to respond endogenously.

3.4 Characterization

The Compression Mechanism. The central result is that procurement compresses the frontier premium: government demand raises incremental-line profits $\pi_I(G) = \bar{\pi}_I(1 + G)$ without changing frontier-line profits $\pi_F(G) = \bar{\pi}_F$, narrowing the value gap $v_F - v_I$ that incentivizes frontier investment. Subtracting (13) from (14) makes this precise:

$$v_F(G) - v_I(G) = \frac{\pi_F(G) - \pi_I(G)}{r + \tau} = \frac{\bar{\pi}_F - \bar{\pi}_I(1 + G)}{r + \tau}. \tag{29}$$

As G rises, the numerator shrinks, so the frontier premium falls. On the interior region, substituting (29) into (16) gives

$$\mu(G) = \left[\frac{x^*(G)}{\kappa} \cdot \frac{\bar{\pi}_F - \bar{\pi}_I(1 + G)}{r + \tau}\right]^{1/(\eta-1)}. \tag{30}$$

The constrained policy is the projection of $\mu(G)$ onto $[0, 1]$.[5]

When $\mu^*(G)$ is interior, differentiating (30) yields

$$\frac{\partial \mu^*}{\partial G} = \frac{1}{\kappa(\eta - 1)(\mu^*)^{\eta-2}}\left[\frac{\bar{\pi}_F - \bar{\pi}_I(1 + G)}{r + \tau}\frac{\partial x^*}{\partial G} - \frac{x^*(G)\bar{\pi}_I}{r + \tau}\right]. \tag{31}$$

---

[5] Both line types generate positive flow profits for all $G \geq 0$: $\pi_I(G) = \bar{\pi}_I(1 + G) > 0$ and $\pi_F(G) = \bar{\pi}_F > 0$. Firms accept procurement demand without exit. The interior condition $G < G^*$ guarantees a positive frontier premium but does not guarantee $\mu^* < 1$; Appendix D.2 states the optimality condition used when counterfactuals reach a corner at $\mu^* = 1$.

The derivative separates two forces. The first term is the effort force: higher G can increase effort x^*. The second term is the composition force: higher G narrows the frontier premium.

Proposition 3.2 (Procurement compresses the frontier share)**.** *Suppose* $\bar{\pi}_F > \bar{\pi}_I$, $r + \tau > 0$, $\kappa > 0$, $\eta > 1$, $G < G^* \equiv \bar{\pi}_F/\bar{\pi}_I - 1$, *and* $\mu^*(G) \in (0, 1)$. *If*

$$\frac{\partial x^*(G)}{\partial G} < \frac{x^*(G)\bar{\pi}_I}{\bar{\pi}_F - \bar{\pi}_I(1+G)},$$

then $\partial \mu^*(G)/\partial G < 0$.

Proof. On the interior region, substituting (29) into (16) yields (30). Differentiating gives (31). Therefore,

$$\frac{\partial \mu^*}{\partial G} < 0 \quad \Longleftrightarrow \quad \frac{\partial x^*(G)}{\partial G} < \frac{x^*(G)\bar{\pi}_I}{\bar{\pi}_F - \bar{\pi}_I(1+G)}.$$

The stated condition ensures that frontier-premium erosion dominates the effort response. ☐

Proposition 3.2 rationalizes the pattern in Table 1.[6] Procurement environments that place more weight on specified compliance and implementation reliability correspond to a larger effective demand wedge. As the wedge rises, the frontier premium narrows and the equilibrium frontier share falls.

3.5 Stationary Distribution

The stationary distribution maps equilibrium objects into a model-implied cross section of firms. Let $N = n_I + n_F$ denote total firm scope. A firm gains lines (displacing rivals) at rate $x^* N$ and loses lines (displaced by rivals) at rate τN. Entrants enter with a single displaced line ($N = 1$). Since the unit mass of product lines is fixed, the entrant displacement rate satisfies $e = \tau - x^*$ and $x^* < \tau$. Define

$$q \equiv \frac{x^*}{\tau} \in (0, 1).$$

The stationary distribution of firm scope is

$$\Pr(N = n) = \frac{q^n}{n[-\log(1-q)]}, \qquad n \geq 1. \tag{32}$$

---

[6] Procurement affects growth through two margins. Higher G can increase innovation effort x^* (the effort channel) and lower frontier orientation μ^* (the composition channel). The marginal growth effect decomposes as $\partial g/\partial G = (\partial g/\partial x^*)(\partial x^*/\partial G) + (\partial g/\partial \mu^*)(\partial \mu^*/\partial G)$. Proposition 3.2 gives a sufficient condition under which the composition channel dominates.

The logarithmic form is the steady-state solution to the birth–death process in which a firm gains lines at rate x^* and loses them at rate τ; the ratio $q = x^*/\tau$ is the single sufficient statistic for the scope distribution.

Conditional on $N = n$, each line is frontier with probability μ^*:

$$\Pr(n_F = b \mid N = n) = \binom{n}{b}(\mu^*)^b(1-\mu^*)^{n-b}. \tag{33}$$

The joint distribution Γ combines scope heterogeneity with composition heterogeneity; it is the object used to compute aggregate model moments such as the frontier share and the large-firm share.

$$\Gamma(n_I, n_F; G, \boldsymbol{\psi}) = \Pr(N = n_I + n_F)\binom{n_I + n_F}{n_F}(\mu^*)^{n_F}(1-\mu^*)^{n_I}. \tag{34}$$

Proposition 3.3 (Cross-sectional implications of procurement)**.** *Let* $q = x^*/\tau$ *and* $s_F \equiv n_F/N$. *Under* Γ*:*

(i) $\mathbb{E}[s_F] = \mu^*$ *and*

$$\operatorname{Var}(s_F) = \mu^*(1-\mu^*)\frac{\sum_{n=1}^{\infty} q^n/n^2}{-\log(1-q)}.$$

(ii)

$$\Pr(n_F = 0) = \frac{-\log\left(1 - q(1-\mu^*)\right)}{-\log(1-q)}.$$

(iii) $\mathbb{E}[\pi(G)] = \mu^*\bar{\pi}_F + (1-\mu^*)\bar{\pi}_I(1+G)$.

A decline in the frontier share shifts the cross section toward incremental portfolios and lowers average profits per line.

Appendix Figure D2 plots the marginal distributions of firm scope and frontier-line share at the benchmark calibration. Two equilibrium scalars determine the stationary distribution: $q = x^*/\tau$ governs portfolio size through $\Pr(N = n)$, and μ^* governs the within-firm frontier share through the binomial. The structural parameters $(\kappa, \lambda_I, \lambda_F, h_E)$ are free; the next section pins them down by matching five moments from the MIP and BHP microdata.

4 Quantitative Analysis

Our quantitative exercise asks whether the model matches the observed composition of innovation and scale of firm activity, then uses the fitted model to compare reforms along the composition and effort margins. The parameterization has two blocks: externally calibrated parameters and the jointly estimated structural parameters.

4.1 Parameterization

Empirical sample. Model moments come from the 2009–2018 MIP sample (52,518 observations, 18,824 firms). The sample starts in 2009, when frontier-share, firm-size, and new-product variables first become jointly available, and ends in 2018.

Externally calibrated objects. We set the household discount rate to $\rho = 0.04$ and use annual TFP growth $\bar{g} = 0.5\%$ as the benchmark growth normalization, consistent with Penn World Tables 10.01 and Bergeaud et al. (2016). The balanced-growth Euler equation implies $r = \rho + \bar{g} = 0.045$. We set the innovation-effort and frontier-orientation-cost curvatures to $\zeta = \eta = 2$.

Benchmark government demand is $\bar{G} = 0.117$, implying economy-wide procurement of $(1 - \mu^*)\bar{G} \approx 6$ percent of output, consistent with Eurostat Government Procurement Statistics for above-threshold contracts in German manufacturing and services over 2009–2018. The regressions measure within-sector co-movement between contract volume and frontier innovation; the model translates that directional evidence into aggregate growth and welfare implications of a permanent change in G.

4.2 Estimation

We estimate the structural parameter vector $\boldsymbol{\psi} = (\lambda_I, \lambda_F, \kappa, h_E)$ by overidentified GMM, targeting five moments with direct model counterparts: (i) the frontier innovation share, (ii) the frontier revenue share among new products, (iii) the large-firm share, (iv) the creative-destruction rate, and (v) the R&D personnel share. The GMM objective is

$$\hat{\boldsymbol{\psi}} = \arg\min_{\boldsymbol{\psi}} \Big[m(\boldsymbol{\psi}) - \hat{m}\Big]^\top W \Big[m(\boldsymbol{\psi}) - \hat{m}\Big],$$

where $m(\boldsymbol{\psi})$ is the vector of model-implied moments, $\hat{m}$ is the empirical moment vector, and W is set to the inverse of the estimated variance-covariance matrix of the empirical moments, constructed using influence functions. For each candidate $\boldsymbol{\psi}$, the inner normalization recovers $c_x(\boldsymbol{\psi})$ from $g(\boldsymbol{\psi}, c_x; \bar{G}) = \bar{g}$; hence macro growth enters as a normalization rather than a targeted moment. Given a trial set of structural parameters, the innovation cost scale is adjusted so that the model exactly reproduces the observed German long-run growth rate; the structural parameters discipline the composition and firm-dynamics margins rather than the overall level of innovative activity.

4.2.1 Identification

The system has four estimated parameters $(\kappa, \lambda_I, \lambda_F, h_E)$ and five targeted moments, yielding one overidentifying restriction. Each parameter has a primary identifying block.

The two frontier-composition moments (the frontier share among product innovators, model counterpart μ^*, and the frontier revenue share among new products, model counterpart the frontier-weighted share of newly created lines) are most informative about the frontier-choice block $(\kappa, \lambda_I, \lambda_F)$. A higher κ increases the cost of frontier-oriented innovation, lowers novelty innovation lines, and reduces both moments. The quality steps (λ_I, λ_F) shape the profit gap $\bar{\pi}_F - \bar{\pi}_I$ and therefore the revenue-share differential.

The three scope-and-entry moments (the large-firm share, the creative-destruction rate, and the R&D personnel share) are most informative about h_E. A higher h_E increases the labor cost of new-line creation, lowers equilibrium entry, shifts the scope distribution toward smaller firms, and increases innovation labor relative to production labor. The R&D personnel share provides an additional handle on h_E beyond the size and turnover margins. Table 4 presents the estimates.

Table 4: Model parameterization, estimates, and recovered objects

Parameter	**Description**	**Estimate**	**Source / target**
Normalizations and recovered objects			
w	Wage	0.994	Endogenous benchmark wage
$\bar{\pi}_F$	Frontier per-line profit	0.174	Implied by Bertrand pricing
$\bar{\pi}_I$	Incremental per-line profit	0.085	Implied by Bertrand pricing
$c_x(\boldsymbol{\psi})$	Innovation cost scale	21.360	Inner normalization: $g(\boldsymbol{\psi}, c_x; \bar{G}) = \bar{g}$
Externally calibrated			
ρ	Household discount rate	0.040	Literature value
r	Interest rate	0.045	Euler equation: $r = \rho + g$
$\bar{g}$	TFP growth normalization	0.005	Feenstra et al. (2015); Bergeaud et al. (2016)
$\bar{G}$	Benchmark government demand	0.117	Eurostat (2023)
ζ	Innovation-effort cost curvature	2.000	Literature value
η	Direction-cost curvature	2.000	Calibration
Jointly estimated (GMM, 5 moments, 4 parameters)			
λ_F	Frontier quality step	1.210 (0.003)	GMM estimate
λ_I	Incremental quality step	1.093 (0.002)	GMM estimate
κ	Direction cost	0.070 (0.004)	GMM estimate
h_E	Entry labor cost	1.600 (0.010)	GMM estimate

Notes: The model is overidentified, with five empirical moments and four free parameters. The targeted moments are the frontier share among product innovators, the frontier revenue share among new products, the large-firm share, the creative-destruction rate, and the R&D personnel share. Reported standard errors are GMM standard errors; the variance-covariance matrix of empirical moments is constructed using influence functions, with MIP moments clustered at the firm level and the creative-destruction moment using the cross-year standard deviation from BHP. Per-line profits follow from (λ_I, λ_F) under Bertrand pricing, with $\bar{\pi}_k = (\lambda_k - 1)/\lambda_k$. Benchmark government demand $\bar{G} = 0.117$ implies economy-wide procurement of $(1 - \mu^*)\bar{G} \approx 6\%$ of output, consistent with Eurostat GPS above-threshold procurement in German manufacturing and KIBS over 2009–2018. The cost scale $c_x(\boldsymbol{\psi})$ enters as an inner normalization rather than as an estimated parameter or target moment. *Source:* Authors' calculations using the Mannheim Innovation Panel, BHP, Eurostat, and Destatis.

The estimates carry intuitive magnitudes and align with the innovation literature. The frontier quality step $\hat{\lambda}_F = 1.210$ implies a 21% productivity improvement per creative-destruction event; the incremental step $\hat{\lambda}_I = 1.093$ implies a 9% improvement per own-product advance. The ordering $\lambda_F > \lambda_I$ matches the finding of García-Macià et al. (2019) that creative-destruction events carry a larger quality improvement than own-product improvements, and both steps fall within the range typically calibrated in Klette-Kortum models (Klette and Kortum, 2004; Akcigit and Kerr, 2018).

Under Bertrand competition, per-line profits follow from the quality steps: $\bar{\pi}_F = (\lambda_F - 1)/\lambda_F = 17.4\%$ and $\bar{\pi}_I = (\lambda_I - 1)/\lambda_I = 8.5\%$, a frontier premium of 8.9 percentage points. Procurement erodes the premium. Incremental lines serve both private and government demand, raising their profits by a factor of $(1 + \bar{G})$; frontier lines face private demand only. As $\bar{G}$ rises, the gap narrows and μ^* falls. The frontier-orientation cost $\hat{\kappa} = 0.070$ places adjustment frictions well below the profit premium; composition responds substantially to the procurement-induced gap narrowing, and the joint restriction from both composition moments identifies κ. The entry cost $\hat{h}_E = 1.600$ wage units clears the free-entry condition at the observed 5.1% creative-destruction rate; comparable models calibrated to US firm dynamics imply higher entry rates, reflecting Germany's lower business dynamism (Akcigit and Ates, 2023; Acemoglu et al., 2018).

Table 5: Model fit: empirical and model moments

Moment	**Data**	**SE**	**Model**	**Diff.**
Targeted moments				
Frontier share among product innovators	0.427	(0.005)	0.441	−0.014
Frontier revenue share among new products	58.249	(0.548)	59.141	−0.892
Large-firm share	0.104	(0.003)	0.090	0.014
Creative destruction	0.051	(<0.001)	0.051	0.000
R&D personnel share	0.029	(0.001)	0.037	−0.008

Notes: The table reports model fit on the 2009–2018 common MIP sample. Standard errors for MIP moments are cluster-robust at the firm level. The creative-destruction standard error is the cross-year standard deviation from BHP. The specification is overidentified, with five moments targeting four parameters. The overidentifying restriction is rejected (J = 23.5, one degree of freedom).

Source: Authors' calculations using the calibrated model, the Mannheim Innovation Panel, and BHP.

Table 5 reports the model's fit. The model fits the targeted moments closely. At the estimated parameters, 44% of product-innovating firms pursue market-novel advances and 56% pursue incremental improvements in the baseline. Creative destruction at 5.1% per year falls below comparable US estimates (Akcigit and Ates, 2023), consistent with Germany's lower business dynamism. The model overpredicts the R&D personnel share (0.037 vs. 0.029); part-time and multitasking R&D roles, excluded from the model, may be related to the gap.

4.2.2 Comparative Statics

Figure 3 traces the model's endogenous variables as government demand rises. The pattern is a composition shift: the frontier share and frontier revenue share fall, firm size rises, and creative destruction slows.

Higher government demand increases per-line profits for incremental lines, which serve both private and government buyers, while leaving frontier-line profits unchanged. The frontier premium narrows, reducing the return to novel innovation. Incumbents lower μ^*, the share of new lines targeted at the frontier; the fall in the frontier revenue share mirrors the reallocation within the product portfolio.

Firm dynamics reinforce the composition shift. As incumbents move toward incremental lines, the value of an existing portfolio rises relative to creative destruction; average firm scope expands and the creative-destruction rate slows, consistent with the ORBIS evidence on employment and turnover among winners. Free entry ties aggregate effort x^* to the entry-cost condition; procurement changes *what* firms innovate rather than how much. Because frontier innovations carry a larger quality step $\lambda_F > \lambda_I$, the shift toward incremental lines lowers the average quality improvement per innovation event and long-run growth g.

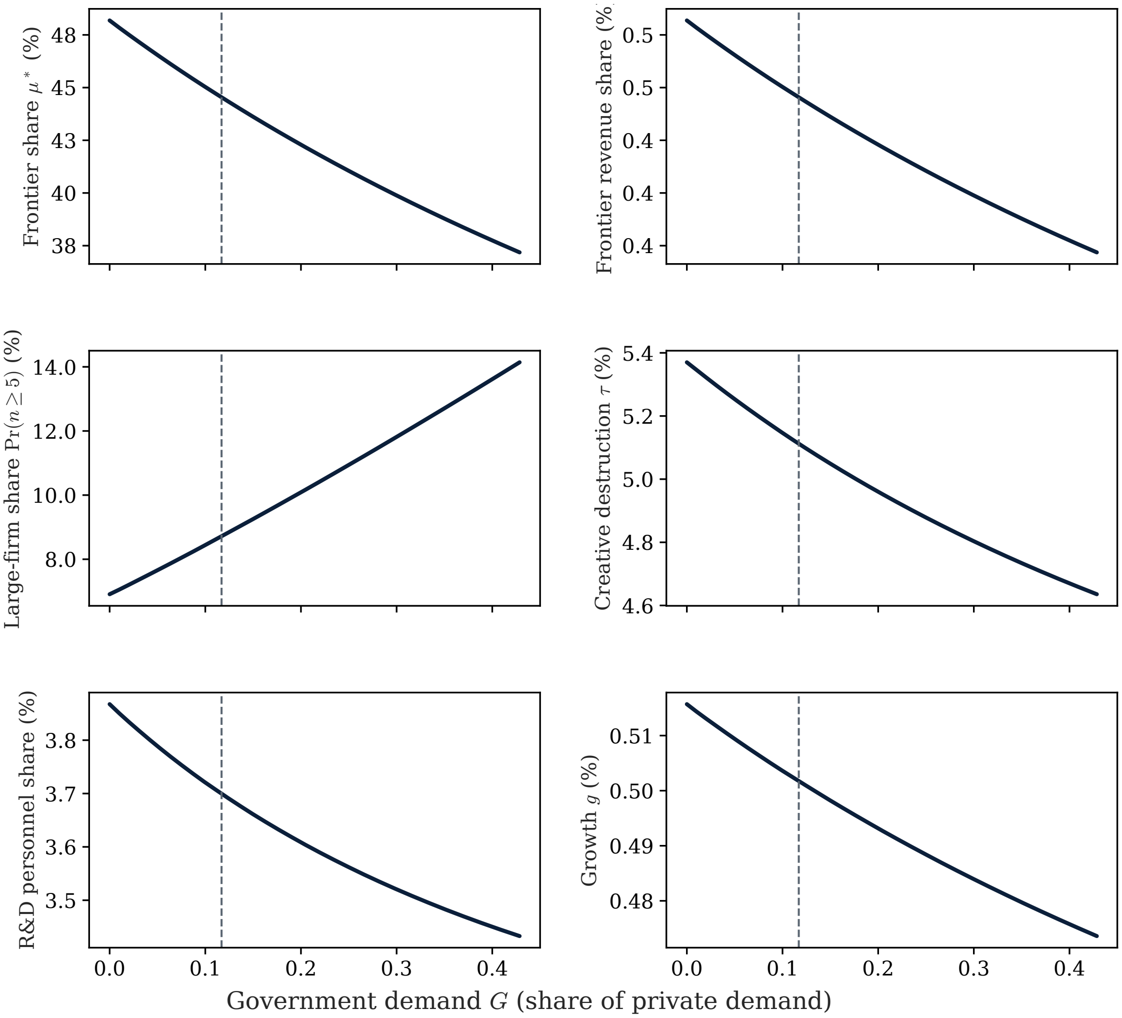

Figure 3: Comparative statics over government demand

Notes: The figure reports the five targeted GMM moments and long-run growth as functions of government demand G. The dashed vertical line marks the benchmark $\bar{G} = 0.117$. Panels report the frontier share μ^*, frontier revenue share, large-firm share $\Pr(n \geq 5)$, creative destruction τ, R&D personnel share, and growth g. The horizontal axis reports government demand as a share of private demand.

Source: Authors' calculations from the calibrated model.

5 Policy Counterfactuals

Using four counterfactual experiments, we examine the implications of the estimated model. Each experiment changes one policy object and holds the calibrated parameters at their benchmark values.

Each experiment uses the same fiscal accounting. The government finances procurement through a lump-sum tax on households. At the benchmark, the household burden equals $T_{bm} = (1-\mu^*_{bm})\bar{G} \approx 0.065$. Because a lump-sum transfer is neutral in isolation, the welfare effect depends on how the resources are used. In the benchmark, procurement directs resources to incremental demand, which increases $\pi_I(\bar{G})$ and narrows the frontier premium. In all counterfactuals, the household tax adjusts with procurement demand: $T = (1-\mu^*)G$. Removing procurement eliminates both the composition wedge and the fiscal burden.

Table 6 summarizes four exercises. The primary counterfactual removes procurement and redirects the freed funds to an R&D tax credit, motivated by Germany's *Forschungszulage* (Dechezleprêtre et al., 2023). The exercise asks how outcomes change when the same fiscal resources support innovation effort directly rather than incremental demand. Two design reforms keep total procurement spending fixed but change how awards are allocated: one realigns contract specifications toward open awards, and one improves buyer screening capacity. A fourth exercise rolls G back to its pre-2014 level to summarize how much the model attributes to the post-2014 increase in procurement demand.

Welfare metric. We report a permanent-consumption equivalent (CEV). On a balanced-growth path,

$$W = \frac{\log C_0}{\rho} + \frac{g}{\rho^2},$$

and the welfare difference between policy B and benchmark A is

$$\text{CEV}_{B,A} = \Delta \log C_0 + \frac{g_B - g_A}{\rho}, \tag{35}$$

with C_0 from the household budget (2) and $\rho = 0.04$. In all scenarios, the household tax adjusts with G: $T = (1-\mu^*)G$. The first term captures the consumption gain or loss as G and μ^* adjust; the second captures the growth gain. Both components are expressed in percent of benchmark consumption.

5.1 Main Counterfactual: Proportional R&D Tax Credit

The main counterfactual redirects procurement resources to a proportional R&D tax credit. R&D tax credits are a standard policy instrument for supporting private innovation effort (Hall and Van Reenen, 2000; Bloom et al., 2002). The exercise compares the benchmark allocation with one in which the full procurement-tax envelope lowers innovation costs. The government removes procurement ($G = 0$) and introduces a tax credit at rate s, reducing the effective innovation cost to

$(1-s)c_x$. The incumbent Bellman equations become

$$(r+\tau)v_k = \pi_k(0) - (1-s)\Phi(x) - \Psi(\mu) + x\bar{v}(\mu;0), \qquad k \in \{I, F\}, \tag{36}$$

where the cost reduction increases equilibrium innovation effort and, through the frontier-choice condition (16), moves the innovation mix toward frontier lines. The tax-credit fiscal bill is $T_s = s \cdot wL_R$.

Germany's *Forschungszulage*, introduced in January 2020, provides a 25 percent credit on eligible R&D labor costs (wages, employer social security contributions, and 60 percent of contracted R&D) and motivates the instrument modeled here. The counterfactual redirects the procurement tax $T_{bm} = 0.065$ to a proportional R&D credit at rate $s = 0.141$. The baseline extends eligibility beyond R&D wages to entry costs and a production-labor component equal to 46.1% of production labor, the model-implied share covering prototyping, process redesign, and adoption costs that credit designs routinely include and that firms reclassify under wide-eligibility rules.[7] As a lower bound, removing procurement without an R&D credit yields 1.57 basis points (0.69% CEV); a narrow-base credit restricted to R&D wages absorbs only 9.1% of the benchmark tax and yields the same amount, since the credit base is too small to shift the innovation margin further. Table 6 reports the main exercise.

Under the full-budget implementation, the frontier share increases from 0.441 to 0.534 ($\Delta\mu^* = 0.093$), and innovation effort is 16% higher. Growth rises by 12.33 basis points per year. García-Macià et al. (2019) estimate that the US composition drift away from creative destruction accounts for 16 of the 34 basis points of TFP growth lost between the 1980s and the 2000s; the growth change here is of the same order and operates through the same composition margin. Removing procurement eliminates the lump-sum tax; those freed funds finance the R&D credit, so the household tax falls to zero. The composition gains keep equilibrium profits roughly constant, and the direct consumption effect is minor ($\Delta \log C_0 \approx 0.13\%$). Most of the welfare change comes from higher growth: $\Delta g/\rho = 3.08\%$ at $\rho = 0.04$, which brings the full CEV to 3.21%.

5.2 Alternative Instruments and Design Reforms

The remaining exercises separate two margins. The first two ask how much changes when procurement design is altered but total procurement spending is held fixed. The final exercise changes the level of procurement demand by returning G to its pre-2014 value.

---

[7] Guceri and Liu (2019) show the UK Enhanced R&D Credit explicitly covers production-adjacent labor; Chen et al. (2021) estimate that 24.2% of reported R&D growth under China's super-deduction reflects relabeling. The 46.1% component is the model-implied residual, not a separate data target.

Specification realignment. Specification-intensive procurement concentrates government demand on incremental lines, narrowing the frontier premium. Realigned specifications allocate awards toward each innovation type in proportion to the market mix at zero additional fiscal cost. The exercise corresponds to a shift from product requirements toward open, functional awards that ask for performance objectives rather than predetermined specifications (Howell et al., 2025). Open awards raise frontier innovation rates, and functional specifications, rather than product-specific requirements, drive the effect (Edquist and Zabala-Iturriagagoitia, 2020; Uyarra et al., 2014). In the model, the counterfactual distributes government demand across both line types in proportion to their equilibrium shares (μ^* to frontier lines and $(1 - \mu^*)$ to incremental lines) rather than concentrating all procurement on incremental lines. The total fiscal bill is unchanged, but the asymmetric demand wedge is removed. The frontier share rises by $\Delta\mu^* = 0.072$ and growth increases by 2.81 basis points per year, the largest budget-neutral gain across all exercises. Appendix D.8 formally derives the equilibrium response.

Monitoring-capacity reform. When implementation risk weighs heavily in a buyer's assessment, specification-compliant incremental suppliers become more attractive than frontier innovators with uncertain delivery. Better screening capacity lowers the uncertainty attached to frontier awards and moves the buyer's preference margin toward frontier procurement at zero additional fiscal cost. Evidence links contracting officer capability and buyer capacity to innovation outcomes in procurement settings (Decarolis et al., 2021; Uyarra et al., 2014; OECD, 2017). In the model, we parameterize improved screening capacity as a reduction in the buyer's effective preference for specification compliance, implemented as a decrease in the demand wedge G on incremental lines while holding total procurement volume constant. Appendix D.7 derives the preference threshold and the mapping from screening improvements to the equivalent demand wedge. The frontier share rises by $\Delta\mu^* = 0.053$ and growth by 1.71 basis points, almost entirely through the composition margin.

Historical rollback. The post-2014 increase in procurement moves G from 0.026 to 0.117 in the model, narrowing the frontier premium and shifting innovation toward incremental lines. At the pre-2014 value of G (Eurostat GPS 2006–2008 average), incremental-line profits are lower, part of the frontier premium is restored, and μ^* is higher. Unlike the design reforms, the rollback also reduces the household tax $T = (1 - \mu^*)G$, so households benefit from both improved composition and a lower fiscal burden. Growth rises by 1.23 basis points per year with a welfare gain of 0.55%. The gain is smaller than in the R&D credit exercise because procurement demand is roughly 6% of output at the benchmark and the composition channel operates on that share. All exercises operate through the same margin: reforms that shift demand away from incremental lines raise the frontier share and model-implied growth.

Table 6: Policy counterfactuals from the GMM benchmark

Counterfactual	Δx^* (%)	$\Delta\mu^*$	Δg (bp/yr)	CEV (%)
Broad-based proportional R&D tax credit†	16	0.093	12.33	3.21
Specification realignment	0.00	0.072	2.81	0.70
Monitoring-capacity reform	0.01	0.053	1.71	0.43
Historical rollback ($G : 0.117 \to 0.026$)	0.01	0.032	1.23	0.55

Notes: Δx^* is the percent change in equilibrium innovation effort. CEV is the permanent-consumption equivalent $\Delta \log C_0 + \Delta g/\rho$ with $\rho = 0.04$; C_0 follows the household budget constraint (2). The household tax adjusts endogenously: $T = (1 - \mu^*)G$. Row 1 removes procurement ($G = 0$, $T = 0$) and redeploys the funds to a broad-based R&D credit at rate $s = 14.1\%$. †The credit base covers R&D wages, entry costs, and the model-implied implementation component (46.1% of production labor); removing procurement without any credit yields $\Delta g = +1.57$ bp and CEV $= +0.69\%$. Rows 2–3 hold G fixed and realign contract specifications or buyer screening capacity. Row 4 reduces G to its pre-2014 level and returns the procurement funds to households. *Source:* Authors' calculations from the calibrated model.

The welfare effects are comparable in scale to other quantitative exercises in endogenous growth and related literatures. The main counterfactual (3.21% CEV) is larger than estimated welfare effects from short-termism distortions (Terry, 2023) and smaller than estimates for R&D misallocation and monopsony costs (Lehr, 2024, 2025). It also exceeds business-cycle costs of around 2% (Krusell et al., 2009) and falls within the 1–5% range of trade-liberalization gains for small open economies (Arkolakis et al., 2012; Caliendo and Parro, 2015; Melitz and Redding, 2015; Costinot and Rodríguez-Clare, 2015). Budget-neutral design reforms (0.43–0.70%) are smaller, close to Lucas (2000)'s inflation-cost estimate of 1%. The difference between the budget-neutral and full-reform gains reflects the fiscal channel: removing procurement eliminates both the composition wedge and the household tax, with procurement funds financing the R&D credit, while design reforms operate on composition alone.

These basis-point gains map directly into long-run TFP and GDP per capita growth: in quality-ladder models, output expansion comes entirely from quality improvements, which appear as total factor productivity in the data. Against a German baseline of $\bar{g} \approx 50$ basis points per year (Bergeaud et al., 2016), the main counterfactual adds 12.33 basis points and the zero-fiscal-cost specification reform adds 2.81 basis points. Both gains run through composition: procurement reform moves the innovation mix toward creative destruction, the highest-step margin. By contrast, Atkeson and Burstein (2019) find that scale-neutral innovation subsidies yield limited productivity gains, reinforcing the model's emphasis on composition rather than scale.

5.3 Robustness Checks

We conduct two robustness checks. First, we perturb each structural parameter by one standard error around the calibrated baseline, holding the remaining parameters constant, and recompute the main outcomes. The exercise changes the quantitative magnitudes but not the overall conclusions. Holding the benchmark R&D cost scale fixed, annual growth moves by at most 0.70 basis points in either direction, and the welfare change ranges from -0.286 to 0.318 percent relative to the baseline. The frontier-orientation-cost parameter κ generates the largest welfare movement, while λ_I generates the largest movement in the frontier-innovation share. Table D4 in Appendix D.11 reports the full set of results.

Second, we re-estimate the model using sectoral procurement frontier shares G_s instead of the economy-wide measure $\bar{G} = 0.117$. The exercise allows procurement exposure to vary with each sector's own composition of frontier and non-frontier lines. Sectoral procurement exposure ranges from $G_s = 0.045$ to $G_s = 0.135$. The frontier share in each sector is inversely related to that sector's procurement exposure, matching the observed cross-sectoral pattern: the largest declines in μ_s^* occur in the most exposed sectors, while the smallest occur in the least exposed. The pattern holds across the full range of procurement exposure, from manufacturing to mining. Allowing procurement exposure to vary by sector changes the sector-level magnitudes but is consistent with the main interpretation.

6 Conclusion

German R&D spending holds steady while the share of firms introducing market-novel products falls from 26 percent in the early 2000s to 8 percent by 2023. We study public procurement as one potential channel explaining these patterns. Government contracts write specifications around fixed technical standards, certification requirements, and delivery schedules; meeting those requirements is more profitable for existing products than for frontier novelty. Sectors with higher procurement exposure introduce market-novel products less often and generate more revenue from existing products.

Three facts characterize the evidence. Sectors with higher procurement exposure have a 0.20 percentage-point lower probability of market-novel introductions (about four percent of the sample mean) and a 0.38 percentage-point higher share of revenue from existing products. Procurement winners expand in employment and turnover without corresponding gains in productivity or profitability. The heterogeneity results point in the same direction: the market-novelty decline concentrates in supplies and services contracts (-0.21 pp), not in public works contracts (-0.01 pp), and the pattern is stronger for central-government buyers than for regional, local, or utility buyers. The evidence documents a shift away from frontier novelty and toward scale-oriented delivery, not a reduction in innovation effort overall.

We build a model that extends Klette and Kortum (2004) by adding a direction-choice margin between frontier and incremental innovation, with separate quality steps and a reallocation cost. We estimate structural parameters by GMM, targeting empirical statistics from the microdata. Procurement enters as a demand wedge that narrows the frontier premium. The fitted model attributes the growth cost to composition rather than the level of public spending.

The primary counterfactual redirects procurement funds to an R&D tax credit, modeled on Germany's *Forschungszulage* (a rule-based credit on eligible R&D labor costs introduced in January 2020). Growth rises by 12 basis points per year and welfare rises by 3.21 percent of permanent consumption. Budget-neutral design reforms raise growth by 1.71–2.81 basis points at zero fiscal cost. Reducing procurement back to its pre-2014 level adds 1.23 basis points and a welfare gain of 0.55 percent. These effects are moderately large and in line with other quantitative exercises in the endogenous growth and related literatures: the main counterfactual exceeds the welfare cost of business cycles (Krusell et al., 2009), falls within the range of trade-liberalization gains for small open economies (Arkolakis et al., 2012; Caliendo and Parro, 2015), and is larger than the short-termism cost in Terry (2023) but smaller than the misallocation costs in Lehr (2024, 2025). Procurement creates a wedge between private and aggregate returns to frontier innovation: contract winners expand in scale, while the aggregate growth cost flows through composition loss rather than a contraction of effort. In general equilibrium, the reduction in frontier innovation by procurement-exposed firms narrows competition at the frontier, which raises the payoff to frontier innovation for other firms. This attenuation force is present in the model; the counterfactual estimates therefore represent a lower bound on the aggregate frontier-innovation loss if non-procurement firms do not fully fill the frontier gap created by procurement recomposition.

The results highlight a distinction between the private and aggregate consequences of procurement. For individual firms, contracts improve scale and revenue stability, consistent with the employment and turnover patterns in the data. At the aggregate level, the composition shift toward specification-compliant delivery lowers growth because the social return to frontier innovation exceeds its private return in the presence of a procurement wedge. The model therefore gives quantitative grounding to long-standing practitioner concerns that specification-heavy procurement creates barriers to frontier innovation (Edquist and Zabala-Iturriagagoitia, 2020; Uyarra et al., 2014; OECD, 2017). Our mechanism may not be specific to Germany: similar channels appear active in US defense contracting (Howell et al., 2025), historical procurement (Clemens and Rogers, 2026), and government R&D programs (Moretti et al., 2025). Future work may exploit these diverse contexts to test when specification-based systems erode the frontier premium and when functional criteria preserve it. A complementary direction, following García-Macià et al. (2019), would decompose the composition margin into vertical quality upgrading and horizontal product-line expansion using establishment-level entry and exit records, and quantify how procurement recomposition propagates through the firm size distribution.

References

Acemoglu, Daron, Ufuk Akcigit, and Murat Alp Celik, "Radical and Incremental Innovation: The Roles of Firms, Managers, and Innovators," *American Economic Journal: Macroeconomics*, 2022, *14* (3), 199–249.

— , — , Harun Alp, Nicholas Bloom, and William Kerr, "Innovation, Reallocation, and Growth," *American Economic Review*, November 2018, *108* (11), 3450–3491.

Aghion, Philippe and Peter Howitt, "A Model of Growth Through Creative Destruction," *Econometrica*, 1992, *60* (2), 323–351. Publisher: [Wiley, Econometric Society].

— , Antonin Bergeaud, Timo Boppart, Peter J Klenow, and Huiyu Li, "A Theory of Falling Growth and Rising Rents," *Review of Economic Studies*, November 2023, *90* (6), 2675–2702.

Akcigit, Ufuk and Sina T. Ates, "What Happened to US Business Dynamism?," *American Economic Journal: Macroeconomics*, 2023, *15* (2), 24–71.

— and William R. Kerr, "Growth through Heterogeneous Innovations," *Journal of Political Economy*, 2018, *126* (4), 1374–1443.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, 2012, *102* (1), 94–130.

Atkeson, Andrew and Ariel Burstein, "Aggregate Implications of Innovation Policy.," *Journal of Political Economy*, December 2019, *127* (6), 2625–2683. Publisher: University of Chicago Press.

Bachmann, Rüdiger, Christian Bayer, Heiko Stüber, and Felix Wellschmied, "Monopsony Makes Firms not Only Small but Also Unproductive: Why East Germany Has Not Converged," *Review of Economic Studies*, 2026, *0*, 1–0.

Bergeaud, Antonin, Gilbert Cette, and Rémy Lecat, "Productivity Trends in Advanced Countries between 1890 and 2012," *Review of Income and Wealth*, 2016, *62* (3), 420–444.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," *American Economic Review*, 2020, *110* (4), 1104–1144.

— , Rachel Griffith, and John Van Reenen, "Do R&D Tax Credits Work? Evidence from a Panel of Countries 1979–1997," *Journal of Public Economics*, 2002, *85* (1), 1–31.

BMBF/ZEW, "Innovationsindikator Deutschland 2024," Technical Report, ZEW – Leibniz Centre for European Economic Research, Mannheim 2024.

Caliendo, Lorenzo and Fernando Parro, “Estimates of the Trade and Welfare Effects of NAFTA,” *Review of Economic Studies*, 2015, *82* (1), 1–44.

Chen, Yanyou and Daniel Yi Xu, “A Structural Empirical Model of R&D Investment, Firm Heterogeneity, and Industry Evolution,” *Journal of Industrial Economics*, 2023, *71* (2), 285–333.

Chen, Zhao, Zhikuo Liu, Juan Carlos Suárez Serrato, and Daniel Yi Xu, “Notching R&D Investment with Corporate Income Tax Cuts in China,” *American Economic Review*, 2021, *111* (7), 2065–2100.

Chiappinelli, Olga, Leonardo Maria Giuffrida, and Giancarlo Spagnolo, “Public Procurement as an Innovation Policy: Where Do We Stand?,” *International Journal of Industrial Organization*, 2025, *100*, 103157.

Clemens, Jeffrey P. and Parker Rogers, “Demand Shocks, Procurement Policies, and the Nature of Medical Innovation: Evidence from Wartime Prosthetic Device Patents,” *Review of Economics and Statistics*, 2026, *108* (1), 75–89.

Costinot, Arnaud and Andrés Rodríguez-Clare, “Trade Theory with Numbers: Quantifying the Consequences of Globalization,” in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics*, Vol. 4, Elsevier, 2015, pp. 197–261.

Czarnitzki, Dirk, Paul Hünermund, and Nima Moshgbar, “Public Procurement of Innovation: Evidence from a German Legislative Reform,” *International Journal of Industrial Organization*, 2020, *71*, 102620.

Decarolis, Francesco, Gaétan de Rassenfosse, Leonardo M. Giuffrida, Elisabetta Iossa, Vincenzo Mollisi, Emilio Raiteri, and Giancarlo Spagnolo, “Buyers’ Role in Innovation Procurement: Evidence from US Military R&D Contracts,” *Journal of Economics & Management Strategy*, 2021, *30* (4), 697–720.

Dechezleprêtre, Antoine, Elias Einiö, Ralf Martin, Kieu-Trang Nguyen, and John Van Reenen, “Do Tax Incentives Increase Firm Innovation? An RD Design for R&D, Patents, and Spillovers,” *American Economic Journal: Economic Policy*, November 2023, *15* (4), 486–521.

Edquist, Charles and Jon Mikel Zabala-Iturriagagoitia, “Functional Procurement for Innovation, Welfare, and the Environment,” *Science and Public Policy*, 2020, *47* (5), 595–603.

Eurostat, “Government Procurement Statistics: Public Procurement Indicators 2022,” Technical Report, Publications Office of the European Union, Luxembourg 2023.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer, “The Next Generation of the Penn World Table,” *American Economic Review*, 2015, *105* (10), 3150–3182.

García-Macià, Daniel, Chang-Tai Hsieh, and Peter J. Klenow, “How Destructive Is Innovation?,” *Econometrica*, 2019, *87* (5), 1507–1541.

Griffith, Rachel, Rupert Harrison, and Helen Simpson, “Product Market Reforms and Innovation in the EU,” *Scandinavian Journal of Economics*, 2010, *112* (2), 389–415.

Guceri, Irem and Li Liu, “Effectiveness of Fiscal Incentives for R&D: Quasi-Experimental Evidence,” *American Economic Journal: Economic Policy*, 2019, *11* (1), 266–291.

Hall, Bronwyn H. and John Van Reenen, “How Effective Are Fiscal Incentives for R&D? A Review of the Evidence,” *Research Policy*, 2000, *29* (4–5), 449–469.

Howell, Sabrina T., Jason Rathje, John Van Reenen, and Jun Wong, “Opening Up Military Innovation: Causal Effects of Reforms to US Defense Research,” *Journal of Political Economy*, 2025, *133* (11), 3605–3658.

Juhász, Réka, Mara P. Squicciarini, and Nico Voigtländer, “Technology Adoption and Productivity Growth: Evidence from Industrialization in France.,” *Journal of Political Economy*, October 2024, *132* (10), 3215–3259. Publisher: University of Chicago Press.

Klette, Tor Jakob and Samuel Kortum, “Innovating Firms and Aggregate Innovation,” *Journal of Political Economy*, 2004, *112* (5), 986–1018.

Krieger, Bastian and Vera Zipperer, “Does Green Public Procurement Trigger Environmental Innovations?,” *Research Policy*, 2022, *51* (6), 104516.

— , Malte Prüfer, and Linus Strecke, “Public Procurement Can Hinder Innovation,” MSI Discussion Paper 2403, KU Leuven, Faculty of Economics and Business, Management, Strategy and Innovation 2024.

Krusell, Per, Toshihiko Mukoyama, Ayşegül Şahin, and Anthony A. Smith, “Revisiting the Welfare Effects of Eliminating Business Cycles,” *Review of Economic Dynamics*, 2009, *12* (3), 393–402.

Lehr, Nils H., “Did R&D Misallocation Contribute to Slower Growth?,” 2024. IMF Working Paper No. 2025/183; revise and resubmit, *Review of Economic Studies*.

— , “Does Monopsony Matter for Innovation?,” 2025. Revise and resubmit, *Econometrica*.

Lerche, Adrian, "Direct and Indirect Effects of Investment Tax Incentives," *American Economic Review*, August 2025, *115* (8), 2781–2818.

Lucas, Robert E., "Inflation and Welfare," *Econometrica*, 2000, *68* (2), 247–274.

Melitz, Marc J. and Stephen J. Redding, "New Trade Models, New Welfare Implications," *American Economic Review*, 2015, *105* (3), 1105–1146.

Moretti, Enrico, Claudia Steinwender, and John Van Reenen, "The Intellectual Spoils of War? Defense R&D, Productivity, and International Spillovers," *The Review of Economics and Statistics*, 2025, *107* (1), 14–27.

OECD, *Public Procurement for Innovation: Good Practices and Strategies*, Paris: OECD Publishing, 2017.

Roodman, David, Morten Ørregaard Nielsen, James G. MacKinnon, and Matthew D. Webb, "Fast and Wild: Bootstrap Inference in Stata Using `boottest`," *Stata Journal*, 2019, *19* (1), 4–60.

Slavtchev, Viktor and Simon Wiederhold, "Does the Technological Content of Government Demand Matter for Private R&D? Evidence from US States," *American Economic Journal: Macroeconomics*, 2016, *8* (2), 45–84.

Takalo, Tuomas, Otto Toivanen, and Toni Wahlberg, "Welfare Effects of R&D Support Policies," *RAND Journal of Economics*, 2026. Forthcoming.

Terry, Stephen J., "The Macro Impact of Short-Termism," *Econometrica*, 2023, *91* (5), 1881–1912.

Uyarra, Elvira, Jakob Edler, Javier Garcia-Estevez, Luke Georghiou, and Jillian Yeow, "Barriers to Innovation through Public Procurement: A Supplier Perspective," *Technovation*, 2014, *34* (10), 631–645.

Appendix

The appendix has four parts. Appendix A documents the merged MIP–procurement panel and exposure measures. Appendix B reports ORBIS validation exercises on winner scale, profitability, and selection. Appendix C presents additional empirical robustness checks, including public R&D subsidies, alternative exposure definitions, and common-sample specifications. Appendix D contains model derivations, numerical solution details, counterfactual diagnostics, and sector-level direction-cost exercises.

A Additional Procurement Descriptives

A.1 Variable Construction

Innovation outcomes (MIP). The primary outcome is market novelty: a binary indicator equal to one if a firm introduced at least one product new to the market in the survey year. The market-novelty measure is distinct from the broader product-innovation indicator used by Krieger et al. (2024), which equals one if a firm introduced any new or improved product—including improvements to existing product lines that are not new to the market. Appendix Table A1 reports both measures side by side.

The complementary revenue outcome is the share of total turnover from existing (unchanged) products, defined as one hundred minus the MIP share of turnover from products new or improved in the last three years. We use the complement because our hypothesis concerns the reallocation of commercial activity toward existing lines. The process-innovation outcome is a binary for any new process introduced in the survey year; year fixed effects absorb the 2019 definitional expansion that raised the reported rate from 28 to 53 percent.

Additional outcomes—the market-novel revenue share, cost-reducing process innovation, geographic novelty, and R&D inputs—follow their standard MIP questionnaire definitions and are described where first used in the main text.

Procurement exposure (TED). The treatment variable is $\operatorname{arcsinh}(\text{contracts}_{s,t})$, where $\text{contracts}_{s,t}$ is the count of contract award notices in TED for sector s in year t. Contracts are assigned to MIP sectors by mapping each contract's EU product code (CPV division) to the corresponding two-digit German industry classification, aggregated to the 21-sector industry scheme of Czarnitzki et al. (2020)—the same classification used by Krieger et al. (2024).

A.2 Comparison with Krieger, Pruefer, and Strecke (2024)

Krieger et al. study price-criterion tenders using a binary for *any* new or improved product—a measure that pools frontier novelty and incremental improvements to existing lines. We focus on *market novelty* alone, precisely because the composition mechanism predicts that procurement reallocates activity between these two types rather than suppressing innovation overall.

The composition mechanism predicts exactly this difference in results. If procurement shifts firms from frontier novelty toward market-known improvements—leaving total product-innovation activity roughly unchanged—then a measure that aggregates both types should be near zero. Table A1 confirms it: market novelty falls with procurement exposure (-0.199^{*}, col. 1), while the Krieger et al. aggregate indicator is near zero (-0.009, n.s., col. 3). The near-zero aggregate is a consequence of the composition shift, not evidence against it.

Table A1: Variable Definition Comparison Krieger, Pruefer, and Strecke (2024)

	Baseline		KPS (2024) comparable	
	(1) Mkt. novelty	(2) Existing-prod. rev. (%)	(3) Prod. innov.	(4) New-prod. rev. (%)
arcsinh(contracts$_{st}$)	-0.199*	0.379**	-0.009	-0.122
	(0.101)	(0.137)	(0.279)	(0.102)
Observations	66,980	25,962	70,462	66,367
Industries	21	21	21	21
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓

Notes: Each cell reports a coefficient from a within-firm OLS regression with firm and year fixed effects. Treatment: $\mathrm{arcsinh}(\text{contracts}_{st})$, the count of contract award notices in TED for sector s in year t (2009–2023). Standard errors, in parentheses, are clustered by industry (21 industries). Cols. (1)–(2): this paper's outcomes. Market novelty is a binary for products new to the market; existing-product revenue is the share of turnover from unchanged products (multiplied by 100, pp). Cols. (3)–(4): outcomes following Krieger et al. (2024). Product innovation is a binary for any new or improved product (multiplied by 100, pp); new-product revenue is the share of turnover from new or improved products. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Source: Authors' calculations using the Mannheim Innovation Panel and TED.

A.3 Merged Panel and Exposure Measures

The appendix reports the merged MIP–procurement panel and the cross-sectional exposure contrast used in the main text.

Table A2: Summary Statistics: MIP–Procurement Merged Panel, 2000–2023

Variable	**N**	**Mean**	**SD**
Panel A. Innovation outcomes			
Market-novel product (=1)	132,311	16.19	36.83
World-first product (=1)	78,652	19.97	39.98
Product innovator (=1)	144,286	41.84	49.33
Process innovator (=1)	144,286	38.71	48.71
R&D active (=1)	140,780	30.64	46.10
New-product revenue (% of sales)	108,632	9.839	20.77
Market-novel share within new products (%)	33,925	25.77	37.05
Panel B. Procurement and firm controls			
$\operatorname{arcsinh}$(Contracts, sector × year)	144,287	2.765	3.637
Exporter (=1)	144,287	38.09	48.56
Export intensity (% of turnover)	112,858	14.48	24.16

Notes: Summary statistics use the MIP firm panel for 2000–2023 merged to sector-year procurement exposure from TED, with procurement set to zero before 2009. Binary variables are reported as percentages. Revenue-share variables are measured as shares of firm turnover. Market-novel share within new products is conditional on positive new-product revenue and is capped at 100 when survey coding implies market-novel revenue above total new-product revenue. *Source:* Mannheim Innovation Panel and TED.

Table A3: Innovation Outcomes by Procurement Exposure: High vs. Low Exposure Sector-Years

	Low exposure	**High exposure**	**Difference**
Innovation outcomes			
Market novelty (binary, %)	14.17	10.70	-3.47^{***}
World-novel product (binary, %)	18.54	15.63	-2.90^{***}
Any product innovator (%)	35.02	34.41	-0.61^{*}
Any process innovator (%)	29.50	42.64	13.14^{***}
New-product revenue (% turnover)	8.93	8.07	-0.86^{***}
Market-novel / new-prod revenue (%)	25.10	21.20	-3.90^{***}
R&D active (%)	28.56	25.09	-3.47^{***}
R&D intensity (% turnover)	1.05	1.99	0.94^{***}
Observations	40,029	39,999	

Notes: Unit of observation is a firm-year. Low-exposure sector-years have $\operatorname{arcsinh}(\text{contracts}_{s,t})$ at or below the sample median; high-exposure sector-years are above the median. Difference is the high-minus-low mean, with significance from a Welch two-sample t-test. Binary variables multiplied by 100 (%). R&D intensity is in percent of turnover, conditional on positive R&D activity. $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

A.4 Aggregate Procurement Trend

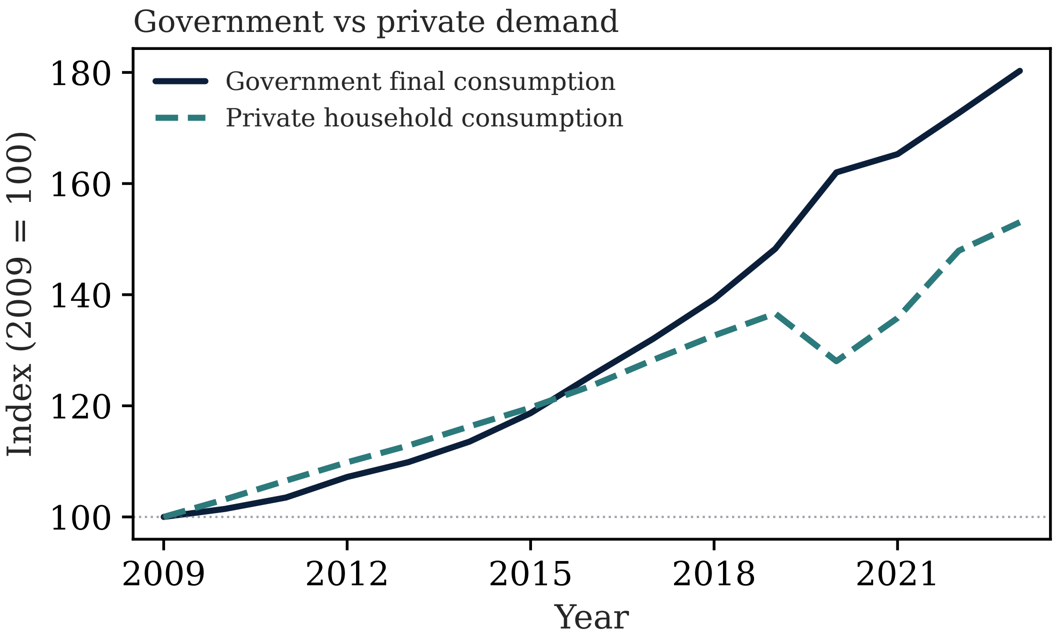

Figure A1: Government final consumption outpaced private household consumption, 2009–2023

Notes: German government final consumption expenditure and private household consumption, both indexed to 2009 = 100 (nominal EUR). Government consumption grew by 80 percent over the period, compared with 52 percent for private households. The divergence accelerates from 2015 onward.

Source: Destatis, Volkswirtschaftliche Gesamtrechnungen (VGR), annual national accounts.

B ORBIS Validation

The MIP results describe the innovation-composition margin. ORBIS provides evidence on the scale margin among matched procurement winners. In the matched winner panel, procurement exposure is associated with higher employment and turnover, while profitability effects remain small. The ORBIS evidence complements the MIP results: procurement winners have higher scale even as the survey evidence shows lower market novelty. Table B1 column (5) reports a one-year lagged treatment specification for the full sample; estimates are similar to the contemporaneous results in column (1), which addresses reverse-causality concerns for level outcomes.

The appendix focuses on the matched ORBIS winner panel. Table B1 reports the full sample and firm-size heterogeneity. In the full matched sample, a one-unit increase in $\operatorname{arcsinh}$(contracts won) is associated with 0.025 higher log employment, 0.008 higher employment growth, and 0.028 higher log turnover. Turnover growth, revenue labor productivity (RLP, revenue divided by employees) growth, and profit margins show no statistically significant association. The pattern is consistent with scale expansion rather than a profitability gain.

Table 3 in the main text reports the full-sample results. Table B2 below reports ownership heterogeneity. Employment responses are stronger for multinationals and foreign-owned firms; return on capital is statistically unchanged across ownership groups.

B.1 Size Heterogeneity

Table B1 reports the matched-winner estimates for the full sample and for small, medium, and large firms.

Table B1: Procurement and Firm Outcomes: Full Sample and by Firm Size

	Full (1)	Small ($<$ 50) (2)	Medium (50–249) (3)	Large ($\geq$ 250) (4)	Full (lag 1) (5)
Dep. var.: as labeled					
Log employment	0.026^{***}	0.025^{***}	0.031^{***}	0.031^{**}	0.025^{***}
	(0.002)	(0.004)	(0.004)	(0.015)	(0.003)
Employment growth (% y-o-y)	0.008^{***}	0.006^{***}	0.010^{***}	0.006	0.000
	(0.002)	(0.002)	(0.003)	(0.009)	(0.002)
Log turnover	0.050^{***}	0.049^{***}	0.041^{***}	0.048^{*}	0.038^{***}
	(0.009)	(0.007)	(0.011)	(0.025)	(0.012)
Turnover growth (% y-o-y)	-0.003	-0.007	-0.003	0.005	-0.004
	(0.006)	(0.009)	(0.006)	(0.011)	(0.006)
RLP growth (% y-o-y)	-0.005	-0.007	-0.002	-0.005	-0.002
	(0.007)	(0.011)	(0.005)	(0.010)	(0.003)
Profit margin	0.001	0.010	-0.000	-0.000^{*}	-0.001
	(0.001)	(0.006)	(0.000)	(0.000)	(0.001)
N (emp growth)	73,190	42,915	21,380	8,895	73,190
Firms	15,493	10,298	3,547	1,349	15,493
Firm FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓

Notes: Within-firm OLS. Sample: ORBIS-matched procurement winners, 2010–2024. Treatment: $\operatorname{arcsinh}$(contracts won) at firm × year. Column (5) uses one-year lagged contract counts (full sample, same firms as column (1)) to address reverse-causality concerns for level outcomes; results are similar. Size groups based on first observed employment. Growth rates: year-on-year log-differences (consecutive years only). RLP = revenue / employees. Profit margin = pretax profit / turnover, winsorised $[-5, 5]$. Standard errors clustered by NACE section. ${}^{*}p<0.1$, ${}^{**}p<0.05$, ${}^{***}p<0.01$.

B.2 Profitability and Ownership Heterogeneity

Table 3 in the main text reports the full-sample results. Employment and turnover are higher among heavier contract winners, while revenue labor productivity (turnover per employee) and return on capital employed are statistically unchanged. The evidence is consistent with capacity expansion rather than efficiency gains or rents.

Table B2 examines ownership heterogeneity. Multinational enterprises, defined as firms with a foreign ultimate owner, have larger employment responses: log employment is 0.049 higher for multinational enterprises (MNEs) and 0.025 higher for domestic firms. Foreign-owned firms have a similar response, 0.057. Return on capital is not statistically significant in any ownership subgroup.

The ownership split reinforces the scale interpretation. International firms have larger employment responses to contract volume, but profitability effects are not statistically significant, consistent with capacity growth rather than a pure rent channel.

Table B2: Winner Outcomes by Ownership

	MNE	Domestic	Foreign-owned
	(1) Log emp.	(2) Log emp.	(3) Log emp.
arcsinh(contracts won)	0.049***	0.025***	0.057***
	(0.011)	(0.003)	(0.013)
Obs.	9,111	60,952	6,475
Industries	18	20	19

Notes: The sample is the ORBIS winner panel over 2009–2023. Each column reports a separate within-firm regression of the indicated outcome on arcsinh(contracts won), with firm and year fixed effects. Standard errors in parentheses are clustered by NACE section. MNE denotes multinational enterprise status. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.
Source: Authors' calculations using TED matched to ORBIS.

B.3 Repeat Contractors

Table B3 splits the winner panel by repeat-contractor status. Repeat contractors are firms winning contracts in three or more sample years; one-time winners win in exactly one sample year. Employment is higher in both groups. Turnover is significantly higher among one-time winners and is positive but imprecise among repeat contractors. Return on capital remains statistically unchanged in both groups, and revenue labor productivity is not higher among repeat contractors.

Table B3: Winner Outcomes by Repeat-Contractor Status

	Repeat contractors				One-time winners			
	(1) Log emp.	(2) Log turnover	(3) Log RLP	(4) Return on capital	(5) Log emp.	(6) Log turnover	(7) Log RLP	(8) Return on capital
arcsinh(contracts won)	0.027***	0.020	-0.007	0.001	0.018***	0.036***	0.021*	-0.001
	(0.003)	(0.013)	(0.012)	(0.001)	(0.003)	(0.012)	(0.011)	(0.004)
Obs.	30,012	14,232	12,749	12,404	33,266	13,524	11,929	9,026
Industries	19	19	19	19	19	19	19	19

Notes: The sample is the ORBIS winner panel over 2009–2023. Repeat contractors are firms winning contracts in three or more sample years; one-time winners win in exactly one sample year. Each column reports a separate within-firm regression of the indicated outcome on $\mathrm{arcsinh}(\text{contracts won})$, with firm and year fixed effects. Standard errors in parentheses are clustered by NACE section. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Source: Authors' calculations using TED matched to ORBIS.

B.4 Firm Dynamics Around Procurement Entry

Figure B1 traces firm outcomes relative to the first observed procurement win in the ORBIS-matched winner panel. Employment is higher after procurement entry, while profitability shows no significant change. Both young entrants (firms under five years old at the first win, 59 percent of matched winners) and established firms enter with market novelty below the sample mean—winners are less frontier-oriented before contract entry.

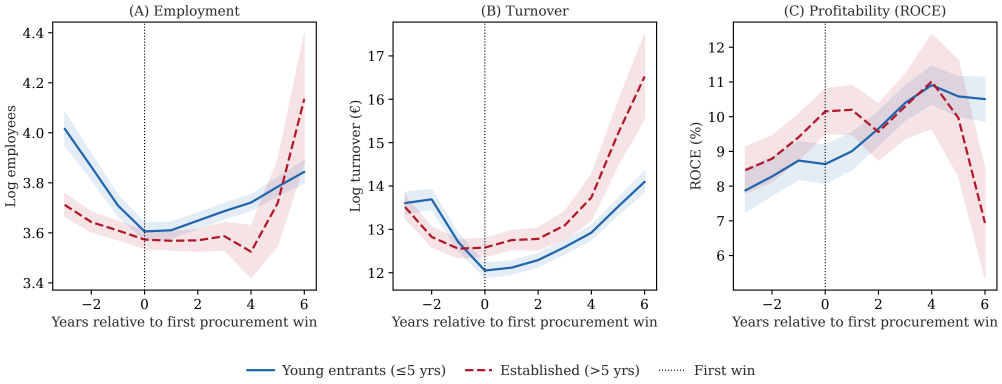

Figure B1: Firm outcomes around the first procurement win

Notes: Event time is measured relative to the year of first contract award in the ORBIS-matched winner panel over 2009–2023. Employment is higher after procurement entry; profitability shows no significant change. Young entrants, defined as firms under five years old at the first win, account for 59 percent of matched winners. Both young and established winners enter procurement with market novelty below the sample mean.

Source: Authors' calculations using TED matched to ORBIS.

C Mechanism Robustness

The appendix reports supplementary robustness results for the empirical evidence in Section 2.

C.1 Heterogeneity Table Construction

Table 2 in the main text reports descriptive splits by contract type and buyer category using the final 21-sector panel. We do not interpret these splits as isolating a single mechanism: contract type, buyer organization, contract size, and administrative procedures co-vary in the procurement data.

C.2 Channels and Process Sub-types

Table C1 estimates each channel on three composition outcomes: market novelty (binary, pp), market-novel share (fraction of new-product revenue that is new to the market), and existing-product revenue share. Procurement raises existing-product revenue ($+0.38$ pp, $p < 0.05$) and reduces market novelty (-0.20 pp, $p < 0.10$), consistent with a within-firm composition shift. Public R&D subsidies have the opposite pattern: they raise market novelty by over 5 pp and reduce existing-product revenue by nearly 4 pp ($p < 0.01$), consistent with an overall innovation stimulus. The contrast confirms that the two policies shift the *direction* of innovation in opposite directions.

Table C1: Procurement vs. Public Subsidies: Two Channels

	(1) Market novelty (pp)	(2) Market-novel share (% of new products)	(3) Existing-product revenue (pp)
Procurement exposure, 2009–2023			
$\text{arcsinh}(\text{contracts})$	-0.199^{*}	-0.099	0.379^{**}
	(0.101)	(0.198)	(0.137)
N	66,980	15,830	25,962
Public R&D subsidy (subsidy-wave years)			
Public R&D subsidy	5.133^{***}	1.152	-3.819^{***}
	(1.258)	(1.801)	(1.103)
N	26,936	5,396	14,112
Firm & year FE	Yes	Yes	Yes
Industries	21	21	21

Notes: Each channel estimated separately on its available sample; firm and year fixed effects throughout; standard errors clustered by industry (21 industries). The procurement rows use the full 2009–2023 panel (same specification as Table 1); the subsidy rows restrict to waves in which *oeff* is observed. Market novelty: binary ×100 (pp); market-novel share: share of new-product revenue that is new to the market (%); existing-product revenue: share of total turnover from unchanged products (pp). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

C.3 Manufacturing and Supplies-and-Services Subsamples

We also report earlier subsample checks for manufacturing and supplies-and-services procurement. These subsamples are not the primary mechanism evidence because supplies and services account for nearly all MIP-covered procurement contracts.

Table C2: Manufacturing and Supplies-and-Services Subsamples

	(1) Manufacturing	(2) Supplies & services
Market novelty (pp)	-0.951^{***}	-0.618^{***}
	(0.114)	(0.104)
Market-novel share (% of new products)	-0.578^{***}	-0.559^{***}
	(0.147)	(0.111)
World-first product (pp)	-0.766^{***}	-0.406^{***}
	(0.173)	(0.122)
Process innovator (pp)	1.666^{***}	1.550^{***}
	(0.236)	(0.176)
Observations by outcome		
Market novelty	33,866	76,468
Market-novel share	12,047	21,575
World-first product	23,343	52,452
Process innovator	35,615	80,028
Firm fixed effects	Yes	Yes
Year fixed effects	Yes	Yes

Notes: We use the finer 21-industry procurement treatment in both columns. Column (1) restricts the sample to the ten manufacturing MIP sectors; standard errors are clustered at the industry level. Column (2) restricts procurement exposure to supplies and services contracts only, excluding works entirely; standard errors are clustered by industry (21 industries). Binary outcomes are multiplied by 100. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.
Source: Authors' calculations using the Mannheim Innovation Panel and TED.

C.4 Denominator and Common-Sample Robustness

A potential concern is a denominator effect: procurement raises total turnover, mechanically raising the existing-product share without any change in absolute new-product revenue. Against this, total new-product revenue share falls by 0.12 pp per unit of exposure ($p = 0.24$) and the market-novelty revenue share falls by 0.048 pp ($p = 0.09$). A pure denominator story would leave new-product shares flat; instead all three innovation-revenue margins fall directionally, consistent with a genuine composition shift.

The three columns of Table 1 use different samples because outcome availability differs across waves. Table C3 re-estimates the core outcomes on a common sample where feasible. The market-novelty estimate (-0.199^{*}) is unchanged, while process innovation (0.092, n.s.) and world-first product (0.056, n.s.) remain statistically zero. The pattern confirms that the composition effect is concentrated in market novelty, not a residue of sample differences across outcome waves.

Table C3: Common-Sample Robustness

Panel A: Product novelty	(1) Mkt. novelty (pp)	(2) Mkt-novel shr. % new-prod. rev.
arcsinh(contracts$_{st}$)	-0.199^{*}	-0.099
	(0.101)	(0.198)
Observations	66,980	15,830
Industries	21	21
Panel B: Other margins	(3) World-first (pp)	(4) Proc. innov. (pp)
arcsinh(contracts$_{st}$)	0.056	0.092
	(0.142)	(0.161)
Observations	43,787	66,980
Industries	21	21

Common sample: both market novelty and process innovation observed.
Firm and year FEs. SE clustered by industry (21 industries).
* p<0.10, ** p<0.05, *** p<0.01.

C.5 Wild Bootstrap Inference

With only 21 MIP-sector clusters, conventional cluster-robust standard errors may over-reject. Table C4 reports wild cluster bootstrap p-values alongside the cluster-t p-values for three key outcomes, using 9,999 Rademacher Monte Carlo draws (Roodman et al., 2019) (full enumeration of 2^{21} patterns is infeasible at $G = 21$). The wild bootstrap p-values are close to the cluster-t values: for market novelty, $p_{\text{wb}} = 0.059$ versus $p_{\text{cluster}} = 0.063$; for existing-product revenue, $p_{\text{wb}} = 0.007$ versus $p_{\text{cluster}} = 0.012$; and for process innovation, $p_{\text{wb}} = 0.504$ versus $p_{\text{cluster}} = 0.555$. The primary results are not an artefact of few-cluster asymptotics.

Table C4: Wild Bootstrap Robustness

Outcome	β	SE (cluster)	p (cluster)	p (wild bootstrap)
Mkt. novelty (2009–2023)	-0.199*	(0.101)	0.063	0.059
Unchanged-prod. rev. (%) (2009–2023)	0.379**	(0.137)	0.012	0.007
Proc. innov. (2009–2023)	0.101	(0.169)	0.555	0.504

Notes: Firm and year FEs. SE clustered by industry (21 industries). Wild cluster bootstrap with 9,999 Monte Carlo Rademacher draws (Roodman et al., 2019); full enumeration (2^{21} patterns) is infeasible for $G = 21$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ (cluster SE).

C.6 Pre-Treatment Trends

We assess pre-trends using an event-study dose variable S_s, defined as the average arcsinh lot count in sector s over 2009–2011 (early OpenTender intensity), normalized to sum to one. Sectors with higher early-period exposure show no significant pre-trend difference in market novelty over 2003–2005 (all $p > 0.50$). The year 2006 alone is significant ($\hat{\beta}_{2006} = -20.6$, $p = 0.012$), coinciding with the initial TED period when above-threshold contracts were first systematically recorded and may have pre-selected high-innovation sectors. The joint F-test for 2003–2006 is $F(4, 22) = 3.08$ ($p = 0.037$), driven predominantly by the 2006 term.

Table C5 reports the full event-study coefficients. Firm and year fixed effects absorb firm-specific baselines and common macro shocks.

Table C5: Pre-trends Test: Event Study around Procurement Exposure Onset

Year	$\hat{\beta}$ (pp)	SE	p-value
Pre-treatment period (relative to 2007)			
2003	-4.603	(12.450)	0.715
2004	-6.929	(12.305)	0.579
2005	-6.315	(12.091)	0.607
2006	-20.571^{**}	(7.513)	0.012
2007	0.000	—	—
Joint F-test: H_0*:* $\beta_{2003} = \cdots = \beta_{2006} = 0$			
$F(4, 22)$	3.084^{**}	($p = 0.037$)	
Post-reference (selected years)			
2008 *(crisis)*	-17.719^{***}	(4.991)	0.002
2009	7.490	(6.676)	0.274
2012	7.039	(6.840)	0.315
2015	6.374	(8.047)	0.437
2018	14.787^{*}	(7.211)	0.052
2021	26.776^{**}	(10.169)	0.015
N	91,516 firm-year obs		
Sector clusters	23		

Notes: Dose variable S_s = sector s's average $\operatorname{arcsinh}(\text{contracts}_{s,t})$ over 2009–2011 (early-treatment intensity at MIP-sector level), normalized to sum to one. Outcome: market novelty $\times 100$ (pp). Firm and year fixed effects. Standard errors clustered by industry (23 industries). Reference year: 2007 (last pre-OpenTender year; avoids the 2008 financial-crisis trough, which disproportionately depressed innovation in high-procurement sectors). The pre-period F-test covers 2003–2006; 2006 overlaps with TED above-threshold procurement, so the truly pre-treatment window is 2003–2005. Post-reference years 2009–2023 correspond to the OpenTender sample period. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D Technical Derivations for the Klette–Kortum Model

The appendix collects the technical material behind the model: the Bertrand profit mapping, the incumbent fixed-point reduction, the numerical solution, the stationary distribution, and several extensions. It also clarifies the entry and creative-destruction accounting used in the quantitative block.

D.1 Bertrand Limit Pricing and Per-Line Profits

Each product line is a vertically differentiated market. An incumbent with quality step $\lambda_j > 1$, where $j \in \{I, F\}$ denotes incremental or frontier innovation, prices just below the nearest rival's effective quality. We assume the Bertrand limit price binds, so the quality lead, rather than the unconstrained CES markup, determines the effective markup.

Derivation. Normalize aggregate expenditure per line to one and the follower's marginal cost to one. The leader with quality step λ_j sets the limit price $p_j^* = \lambda_j$, leaving the consumer indifferent between leader and follower on a quality-adjusted basis ($q/p = 1$ for both). The leader captures the full unit of expenditure. Variable cost per unit produced is $1/\lambda_j$ (one quality unit costs $1/\lambda_j$ at the follower's marginal cost). Per-line operating profit is

$$\bar{\pi}_j = 1 - \frac{1}{\lambda_j} = \frac{\lambda_j - 1}{\lambda_j}, \qquad j \in \{I, F\}. \tag{37}$$

Since $\lambda_F > \lambda_I > 1$, frontier lines earn higher private-market profits than incremental lines. Aggregate expenditure is normalized, while the wage is determined in general equilibrium. In the benchmark, $\lambda_F = 1.210$ and $\lambda_I = 1.093$, implying $\bar{\pi}_F \simeq 0.174$ and $\bar{\pi}_I \simeq 0.085$.

D.2 Fixed-Point Reduction

The incumbent block is governed by

$$(r + \tau)v_I = \pi_I(G) - \Phi(x) - \Psi(\mu) + x\bar{v}(\mu; G), \tag{38}$$

$$(r + \tau)v_F = \pi_F(G) - \Phi(x) - \Psi(\mu) + x\bar{v}(\mu; G), \tag{39}$$

where

$$\bar{v}(\mu; G) \equiv (1 - \mu)v_I + \mu v_F.$$

We restrict attention to equilibria with

$$r + \tau - x > 0, \tag{40}$$

so that line values are finite.

Subtracting (38) from (39) gives the value gap:

$$v_F - v_I = \frac{\pi_F(G) - \pi_I(G)}{r+\tau} = \frac{\bar{\pi}_F - \bar{\pi}_I(1+G)}{r+\tau}. \tag{41}$$

Premultiplying (38) by $(1-\mu)$ and (39) by μ, then adding, gives the average line value:

$$\bar{v}(\mu;G) = \frac{\bar{\pi}(\mu;G) - \Phi(x) - \Psi(\mu)}{r+\tau-x}, \tag{42}$$

where

$$\bar{\pi}(\mu;G) \equiv (1-\mu)\pi_I(G) + \mu\pi_F(G).$$

The effort and direction first-order conditions are

$$c_x \zeta x^{\zeta-1} w = \bar{v}(\mu;G), \tag{43}$$

$$\kappa\mu^{\eta-1} = x(v_F - v_I). \tag{44}$$

Using (41), the interior direction condition becomes

$$\mu^{\text{int}} = \left[\frac{x}{\kappa}\frac{\bar{\pi}_F - \bar{\pi}_I(1+G)}{r+\tau}\right]^{1/(\eta-1)}. \tag{45}$$

Because $\mu \in [0,1]$, the constrained policy is

$$\mu^* = \min\left\{1, \max\left\{0, \mu^{\text{int}}\right\}\right\}. \tag{46}$$

The interior first-order condition applies only when $\mu^* \in (0,1)$. Counterfactuals that reach $\mu^* = 1$ use the corresponding KKT condition.

Equations (42), (43), and (46) reduce the incumbent problem to a two-dimensional fixed point in (x,μ). The first equation gives the value of an average new line, the second pins down innovation effort, and the third pins down the frontier share.

D.3 Numerical Solution

Let $\mathbf{z} \equiv (x, \mu)^\top$ denote the common per-line innovation choice. Using the Bellman equations and the closed-form value gap, the incumbent block can be written as

$$F(\mathbf{z}; G, \boldsymbol{\psi}) = 0, \tag{47}$$

where $\boldsymbol{\psi}$ collects the parameters held fixed in the incumbent problem. In the interior region, the two residuals are

$$F_1(\mathbf{z}; \cdot) \equiv c_x \zeta x^{\zeta-1} w - \bar{v}(\mathbf{z}; G) = 0, \tag{48}$$

$$F_2(\mathbf{z}; \cdot) \equiv \kappa \mu^{\eta-1} - \frac{x[\bar{\pi}_F - \bar{\pi}_I(1+G)]}{r+\tau} = 0. \tag{49}$$

At the boundaries $\mu = 0$ or $\mu = 1$, the algorithm applies the projection in (46).

We solve the system by damped fixed-point iteration. Given an initial guess $\mathbf{z}^{(0)}$, the algorithm updates μ from the projected direction condition and x from the effort condition until

$$\|\mathbf{z}^{(k+1)} - \mathbf{z}^{(k)}\| < 10^{-5}.$$

Once (x^*, μ^*) is obtained, (v_I, v_F, g) follow analytically.

The outer loop updates wages and entry until labor-market clearing and free entry hold jointly. Because the benchmark shuts down cross-line cost interactions, the policy functions are flat in the firm state (n_I, n_F): all lines use the same (x^*, μ^*) in a given equilibrium.

At the benchmark, the numerical routine nests the general-equilibrium solve inside a scalar search over the cost scale c_x. For each candidate structural vector $\boldsymbol{\psi}$, the algorithm recovers the unique $c_x(\boldsymbol{\psi})$ that delivers the benchmark growth target. Macro growth is a normalization, not a GMM target moment.

D.4 Value and Policy Functions in the Firm State

The value function varies with portfolio composition:

$$V(n_I, n_F; G) = n_I v_I(G) + n_F v_F(G).$$

Firms with the same total scope can have different values if their frontier shares differ. Innovation policies, by contrast, are flat in the firm state because the benchmark abstracts from cross-line interactions.

Figure D1 illustrates the result at $\bar{G} = 0.117$. The left panel reports $V(n_I, n_F)$. The middle and right panels report the effort and direction policies. The flat policy surfaces are the graphical counterpart of the model's linearity result.

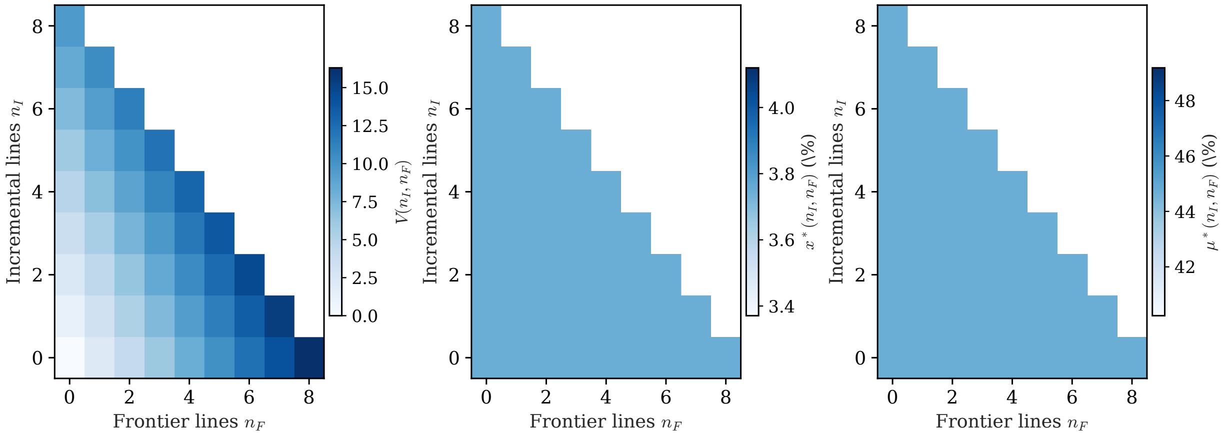

Figure D1: Benchmark value and policy functions in the firm state

Notes: The left panel reports firm value $V(n_I, n_F)$ at the benchmark $\bar{G} = 0.117$. The middle and right panels report effort and direction policies over the same state space. Policy surfaces are flat across firm states because the benchmark shuts down cross-line interactions.

Source: Authors' calculations from the calibrated model.

D.5 Entry, Creative Destruction, and Stationarity

The model has two sources of new product lines: incumbent innovation and entrant innovation. Incumbents create new lines at rate x^* per existing line. Entrants create new firms with one line. Existing lines are displaced at total hazard τ. With a unit mass of product lines, stationarity of the product-line mass requires

$$x^* + e = \tau, \tag{50}$$

where e is the entrant line-creation rate. Hence,

$$e = \tau - x^*. \tag{51}$$

The entry-labor requirement is

$$L_E = h_E e = h_E(\tau - x^*), \tag{52}$$

rather than $h_E \tau$, whenever τ denotes total line destruction.

Free entry determines the expected value of an entrant line:

$$\bar{v}(\mu^*; G) = w h_E. \tag{53}$$

Labor-market clearing uses production labor, incumbent R&D labor, and entrant labor:

$$\bar{L} = L_P(\mu, G, w) + \ell_R(x, \mu, w) + h_E(\tau - x). \tag{54}$$

D.6 Stationary Firm-Size Distribution

Let $N = n_I + n_F$ denote firm scope. Conditional on a firm having $N = n$ lines, births arrive at rate $x^* n$ and deaths arrive at rate τn. Entrants enter with one line at rate $e = \tau - x^*$ under the unit line-mass normalization. We require $x^* < \tau$ for a stationary firm-size distribution.

Let m_n denote the mass of firms with n product lines. For $n \geq 2$, stationarity implies

$$0 = x^*(n-1)m_{n-1} + \tau(n+1)m_{n+1} - (x^* + \tau)nm_n. \tag{55}$$

The solution is

$$m_n = \frac{e}{x^*}\frac{1}{n}\left(\frac{x^*}{\tau}\right)^n, \qquad n \geq 1. \tag{56}$$

Let

$$q \equiv \frac{x^*}{\tau} \in (0, 1).$$

Conditional on observing an active firm, the firm-scope distribution is

$$\Pr(N = n) = \frac{q^n}{n[-\log(1-q)]}, \qquad n \geq 1. \tag{57}$$

Equation (57) gives the Klette–Kortum/Yule distribution, which differs from a geometric distribution by the factor $1/n$.

Conditional on $N = n$, each line is frontier with probability μ^*:

$$\Pr(n_F = b \mid N = n) = \binom{n}{b}(\mu^*)^b(1-\mu^*)^{n-b}. \tag{58}$$

The joint distribution is

$$\Gamma(n_I, n_F; G) = \Pr(N = n_I + n_F)\binom{n_I + n_F}{n_F}(\mu^*)^{n_F}(1-\mu^*)^{n_I}. \tag{59}$$

The product-innovation probability over a horizon T is

$$p_{\text{innov}} = 1 - \sum_{n=1}^{\infty} \Pr(N = n)e^{-Tx^*n}. \tag{60}$$

Using (57),

$$p_{\text{innov}} = 1 - \frac{-\log\left(1 - qe^{-Tx^*}\right)}{-\log(1-q)}. \tag{61}$$

Proof of Proposition 3.3. (i) By iterated expectations $\mathbb{E}[s_F] = \mu^*$. For the variance, $\mathbb{E}[(n_F/N)^2 \mid N] = \mu^*(1-\mu^*)/N + \mu^{*2}$, so $\mathrm{Var}(s_F) = \mu^*(1-\mu^*)\mathbb{E}[1/N]$. Under (57),

$$\mathbb{E}[1/N] = \frac{\sum_{n=1}^{\infty} q^n/n^2}{-\log(1-q)},$$

giving the stated formula.

(ii) $\Pr(n_F = 0) = \sum_{n \geq 1} \Pr(N = n)(1-\mu^*)^n = -\log(1 - q(1-\mu^*))/(-\log(1-q))$, which follows by substituting (57) and recognising the Taylor series of $\log(1-z)$.

(iii) $\mathbb{E}[\pi(G)] = \mu^*\bar{\pi}_F + (1-\mu^*)\bar{\pi}_I(1+G)$ by linearity of expectation over the portfolio and the per-line profit expressions. □

Figure D2 plots both marginal distributions at the GMM benchmark.

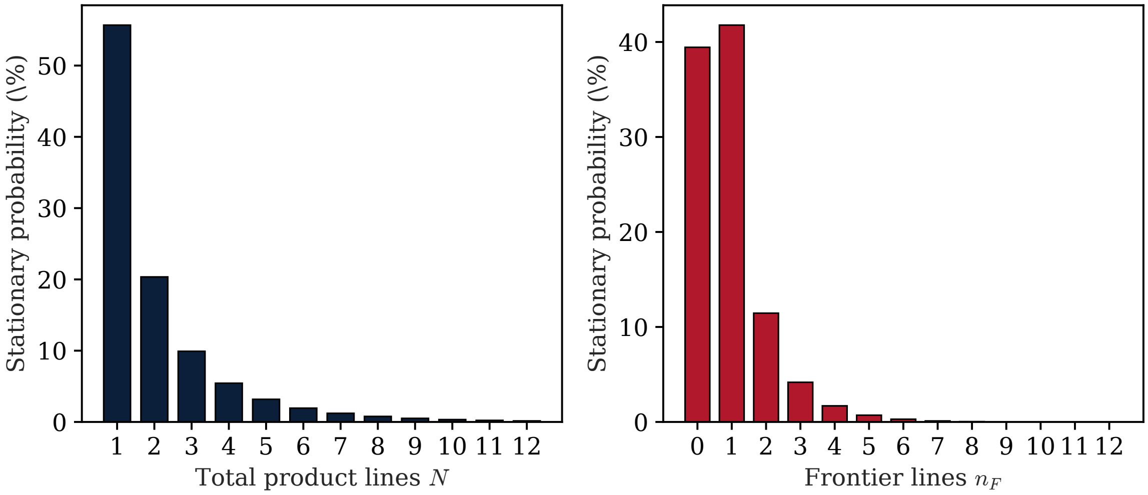

Figure D2: Model-implied stationary distributions at the GMM benchmark

Notes: Panel A reports the firm-scope distribution in (57). Panel B reports the marginal frontier-line distribution implied by binomial mixing under the benchmark equilibrium.

Source: Authors' calculations from the calibrated model.

D.7 A Government-Objective Extension

The active model treats procurement demand as reduced form. The extension microfounds the demand wedge through a government objective over expected delivered quality and implementation risk. It does not add a second quantitative block; instead, it states conditions under which the baseline demand asymmetry arises from a simple buyer-side screening problem.

Suppose the government procures one line at a time. An incremental line delivers

$$q_I \sim \mathcal{N}(\lambda_I, \sigma_I^2),$$

while a frontier line delivers

$$q_F \sim \mathcal{N}(\lambda_F, \sigma_F^2), \qquad \lambda_F > \lambda_I, \qquad \sigma_F^2 > \sigma_I^2.$$

The frontier line has higher expected quality but also higher delivery risk—frontier projects are harder to specify, verify, and monitor ex ante.

The government values expected quality and dislikes implementation risk:

$$U_G(k) = \mathbb{E}[q_k] - \frac{\gamma_G}{2}\mathrm{Var}(q_k), \qquad k \in \{I, F\},$$

where $\gamma_G \geq 0$ is the government's risk-aversion parameter. Under the Gaussian specification,

$$U_G(I) = \lambda_I - \frac{\gamma_G}{2}\sigma_I^2, \qquad U_G(F) = \lambda_F - \frac{\gamma_G}{2}\sigma_F^2.$$

Proposition D.1 (Government preference threshold)**.** *Suppose* $\lambda_F > \lambda_I$ *and* $\sigma_F^2 > \sigma_I^2$*. The government prefers incremental procurement to frontier procurement if and only if*

$$\gamma_G > \gamma_G^* \equiv \frac{2(\lambda_F - \lambda_I)}{\sigma_F^2 - \sigma_I^2}.$$

Proof. The government prefers incremental procurement when $U_G(I) \geq U_G(F)$. Using the expressions above,

$$\lambda_I - \frac{\gamma_G}{2}\sigma_I^2 \geq \lambda_F - \frac{\gamma_G}{2}\sigma_F^2$$

if and only if

$$\gamma_G(\sigma_F^2 - \sigma_I^2) \geq 2(\lambda_F - \lambda_I).$$

Since $\sigma_F^2 > \sigma_I^2$, division gives the threshold. □

Proposition D.1 gives the buyer-side logic behind the demand wedge. When the government is sufficiently risk averse, it forgoes the expected quality advantage of frontier lines and steers procurement toward specification-compliant incremental supply. In this extension, G is not primitive;

it is generated by preferences over quality and implementation risk. Higher γ_G corresponds to a larger effective incremental-demand wedge in the active model.

The extension maps to the contract-design interpretation in a limited way. The main text reports descriptive splits by contract type and buyer category. Those splits are consistent with the idea that some procurement environments place more weight on compliance and implementation reliability than on frontier upside, but they do not isolate the buyer-side preference parameters in the model.

If frontier projects also carry higher failure risk, the government objective becomes

$$U_G(k) = \lambda_k - \frac{\gamma_G}{2}\sigma_k^2 - \chi_G p_k^{\text{fail}},$$

where $p_k^{\text{fail}} \in [0,1]$ is the probability that a type-k project fails to deliver on its specification (cost overrun, incomplete delivery, or quality shortfall) and $\chi_G \geq 0$ is the government's penalty for procurement failure. Variance aversion and failure penalties both move procurement away from frontier delivery.

Remark D.2 (Monitoring capacity as a procurement-design reform)**.** The threshold γ_G^* rises when monitoring or verification technology reduces σ_F^2. For a fixed γ_G, better monitoring makes the government less likely to prefer incremental procurement over frontier procurement. In the active model, lower frontier-delivery uncertainty works like a reduction in the effective demand wedge G.

Table D1 evaluates the threshold using the GMM estimates $\lambda_F = 1.210$ and $\lambda_I = 1.093$. Setting $\sigma_I = 0$ and varying σ_F gives the following values:

Table D1: Government preference threshold at the GMM benchmark

σ_F	γ_G^*
0.15	10.40
0.20	5.85
0.25	3.74
0.30	2.60
0.35	1.91

Notes: Computed from $\gamma_G^* = 2(\lambda_F - \lambda_I)/\sigma_F^2$ with $\lambda_F - \lambda_I = 0.117$ and $\sigma_I = 0$. *Source:* Authors' calculations from the calibrated model.

The incremental-demand regime holds when frontier projects carry sufficient delivery uncertainty relative to the quality-step gap. Table D1 shows the frontier delivery standard deviation σ_F required for the regime to hold at a given risk-aversion level. We do not estimate σ_F or γ_G separately; the table provides a reference for how delivery uncertainty maps to the preference threshold given the calibrated quality steps $\lambda_F = 1.210$ and $\lambda_I = 1.093$.

The procurement data are qualitatively consistent with the mechanism, but do not identify the buyer-side parameters directly. Award prices vary around buyer estimates, but the variance does not identify delivered frontier quality. Award-to-estimate variation mixes buyer estimation error, procurement procedure, contract heterogeneity, and bidder composition, so it cannot calibrate γ_G or σ_F separately.

Identifying the two preference parameters would require contract-level ex post performance data—completion rates, cost overruns, or independent quality assessments—not available in the current administrative records. The extension serves as a microfoundation for the baseline demand asymmetry rather than as an independently estimated buyer-side model.

D.8 Specification Design Reform: Formal Derivations

The specification design reform modifies per-line profit expressions so that a fraction $\alpha \in [0, 1]$ of government demand flows to frontier lines rather than all demand flowing to incremental lines. The modified profit expressions are

$$\pi_I(G, \alpha) = \bar{\pi}_I\Big(1 + (1 - \alpha)G\Big), \tag{62}$$

$$\pi_F(G, \alpha) = \bar{\pi}_F\Big(1 + \alpha G\Big). \tag{63}$$

At $\alpha = 0$ all demand flows to incremental lines (the baseline). The value gap under the reform is

$$v_F(G, \alpha) - v_I(G, \alpha) = \frac{\bar{\pi}_F(1 + \alpha G) - \bar{\pi}_I(1 + (1 - \alpha)G)}{r + \tau}. \tag{64}$$

Since effort satisfies $x^* = h_E/(c_x\zeta)$ under $\zeta = 2$ and free entry, the effort margin remains unchanged ($\partial x^*/\partial \alpha = 0$) and the direction condition gives

$$\mu^*(G, \alpha) = \frac{x^*}{\kappa} \cdot \frac{\bar{\pi}_F(1 + \alpha G) - \bar{\pi}_I(1 + (1 - \alpha)G)}{r + \tau}. \tag{65}$$

Proposition D.3 (Specification design raises frontier orientation)**.** *Suppose $G > 0$, $\mu^*(G, \alpha) \in (0, 1)$, and all parameters are positive. The frontier share is strictly increasing in α:*

$$\frac{\partial \mu^*}{\partial \alpha} = \frac{x^* G(\bar{\pi}_F + \bar{\pi}_I)}{\kappa(r + \tau)} > 0.$$

The gain is proportional to G: specification reform has a larger effect when aggregate procurement demand is larger.

Proof. Raising α by a unit shifts government demand away from incremental lines and toward frontier lines. From (62)–(63), frontier profits rise by $\bar{\pi}_F G$ and incremental profits fall by $\bar{\pi}_I G$. Both changes widen the frontier value gap $(v_F - v_I)$ in (64), so the direction first-order condition moves μ^* toward frontier orientation. Formally, differentiating (65) with respect to α:

$$\frac{\partial \mu^*}{\partial \alpha} = \frac{x^*}{\kappa(r+\tau)} \cdot \frac{\partial}{\partial \alpha}\Big[\bar{\pi}_F(1+\alpha G) - \bar{\pi}_I(1+(1-\alpha)G)\Big] = \frac{x^* G(\bar{\pi}_F + \bar{\pi}_I)}{\kappa(r+\tau)}.$$

Since $G, \bar{\pi}_F, \bar{\pi}_I, \kappa, r+\tau, x^* > 0$, this is strictly positive. The factor G enters multiplicatively because specification reform can only redirect existing procurement demand: the larger the procurement footprint, the more demand is available to realign toward frontier lines. ☐

D.9 Heterogeneous Direction Costs

The benchmark estimates a single direction-cost parameter κ common to all firms and sectors. Allowing sector-specific direction costs $\kappa_s > 0$ leaves the *proportional* compression of the frontier share unchanged, separating the demand-composition mechanism from the supply-side forces documented by Bloom et al. (2020).

Let sector s have direction cost $\Psi_s(\mu) = (\kappa_s/2)\mu^2$, with $\kappa_s > 0$, and let all other parameters be common across sectors. The direction first-order condition in sector s is

$$\kappa_s \mu_s^* = x^*(v_F - v_I) = x^* \frac{\bar{\pi}_F - \bar{\pi}_I(1+G)}{r+\tau}, \tag{66}$$

so the equilibrium frontier share is

$$\mu_s^* = \frac{x^*}{\kappa_s} \cdot \frac{\bar{\pi}_F - \bar{\pi}_I(1+G)}{r+\tau}. \tag{67}$$

Sectors with larger κ_s (frontier innovations costlier to orient toward) have lower frontier shares at any given G, while their innovation behavior is otherwise governed by the same demand-payoff margin.

Proposition D.4 (Supply-side invariance of proportional compression)**.** *Let direction costs be sector-specific:* $\kappa_s > 0$*. The proportional compression of the frontier share,*

$$\frac{\partial \ln \mu_s^*}{\partial G} = -\frac{\bar{\pi}_I}{\bar{\pi}_F - \bar{\pi}_I(1+G)},$$

is identical across all sectors and equal to the homogeneous-κ *value. The demand-composition mechanism operates in proportion to the current frontier share, regardless of sector-level direction costs.*

Proof. Taking logs of (67):

$$\ln \mu_s^* = \ln x^* - \ln \kappa_s + \ln\Big[\bar{\pi}_F - \bar{\pi}_I(1+G)\Big] - \ln(r+\tau).$$

Differentiating with respect to G, using $\partial x^*/\partial G = 0$ (effort is fixed by free entry at $\zeta = 2$):

$$\frac{\partial \ln \mu_s^*}{\partial G} = \frac{-\bar{\pi}_I}{\bar{\pi}_F - \bar{\pi}_I(1+G)},$$

which does not depend on κ_s. □

Proposition D.4 has two implications. First, the absolute compression $\partial \mu_s^*/\partial G = \mu_s^* \times (-\bar{\pi}_I/[\bar{\pi}_F - \bar{\pi}_I(1+G)])$ is smaller in sectors where frontier orientation is already low due to high κ_s, but the proportional rate is identical. Second, rising research productivity requirements—as documented in the aggregate by Bloom et al. (2020)—shift the level of μ_s^* downward over time but do not alter the sensitivity of frontier orientation to procurement demand. The demand-composition and supply-cost channels operate along orthogonal margins: the former governs innovation composition conditional on costs; the latter governs the cost of achieving a given frontier share.

Empirical range of direction costs. Table D2 recovers $\hat{\kappa}_s$ for each WZ section and for the 21 MIP subsectors by inverting (67) at the GMM benchmark, using the sector-level mean frontier share among product innovators as the empirical counterpart of μ_s^*. The recovered values range from 0.055 (Chemicals and Pharmaceuticals; Glass and Ceramics) to 0.128 (Transport), roughly a factor of two around the GMM homogeneous benchmark $\hat{\kappa} = 0.070$. Manufacturing subsectors—particularly chemicals, electrical equipment, and machinery—sit below the benchmark, reflecting higher baseline frontier shares; service sectors and transport sit above it. Figure D3 plots the full subsector distribution.

The heterogeneity is informative about the mechanism but does not change the main result. Procurement narrows the frontier premium through the same payoff-composition channel in all sectors.

Year fixed effects in the within-firm regressions absorb any common secular trend in research productivity, leaving procurement-demand variation to identify the composition pattern.

Table D2: Sector-level direction costs

WZ	Sector	$\bar{\mu}_s$	$\hat{\kappa}_s$ (SE)
G	Trade/Retail	0.513	0.060 (0.001)
C	Manufacturing	0.497	0.062 (0.001)
M	Prof. Services	0.440	0.070 (0.001)
K	Finance	0.374	0.083 (0.003)
D	Energy/Water	0.367	0.084 (0.004)
J	ICT/Media	0.334	0.092 (0.003)
N	Other Services	0.286	0.108 (0.006)
B	Mining	0.274	0.113 (0.010)
H	Transport	0.244	0.126 (0.008)
Homogeneous GMM benchmark			0.070 (0.004)

Notes: $\bar{\mu}_s$ is the mean market-novelty share among product innovators in sector s, MIP 2003–2018. $\hat{\kappa}_s$ is recovered from the direction first-order condition $\kappa_s\bar{\mu}_s = x^*D_{bm}/(r+\tau)$ using GMM benchmark parameters; standard errors propagate sampling uncertainty in $\bar{\mu}_s$ via the delta method. Proposition D.4 establishes that the proportional compression of the frontier share by procurement demand is the same across all sectors regardless of sector-level direction costs. *Source:* Authors' calculations using the Mannheim Innovation Panel and OpenTender procurement records.

Table D3: Direction costs at the 21-subsector level

Subsector	$\bar{\mu}_s$	N	$\hat{\kappa}_s$ (SE)
Mining (WZ B)			
Mining	0.274	339	0.113 (0.010)
Manufacturing (WZ C)			
Chemicals & Pharma	0.572	1637	0.054 (0.001)
Glass & Ceramics	0.545	721	0.057 (0.002)
Machinery & Vehicles	0.539	2291	0.057 (0.001)
Electrical Equip.	0.531	3077	0.058 (0.001)
Textiles	0.498	1000	0.062 (0.002)
Plastics	0.494	1103	0.062 (0.002)
Metals	0.472	1862	0.065 (0.002)
Furniture & Med. Tech.	0.460	1401	0.067 (0.002)
Wood & Paper	0.399	1141	0.077 (0.003)
Food & Tobacco	0.381	1202	0.081 (0.003)
Energy/Water (WZ D)			
Energy/Water	0.367	632	0.084 (0.004)
Trade/Retail (WZ G)			
Retail	0.576	1772	0.054 (0.001)
Wholesale	0.323	588	0.096 (0.006)
Transport (WZ H)			
Transport	0.244	704	0.126 (0.008)
ICT/Media (WZ J)			
IT/Telecom	0.348	1570	0.089 (0.003)
Media Services	0.313	1068	0.099 (0.004)
Finance (WZ K)			
Finance	0.374	1413	0.083 (0.003)
Prof. Services (WZ M)			
Prof. Services I	0.480	2044	0.064 (0.001)
Prof. Services II	0.375	1230	0.082 (0.003)
Other Services (WZ N)			
Other Services	0.286	766	0.108 (0.006)
Homogeneous GMM benchmark			0.070 (0.004)

Notes: $\bar{\mu}_s$ is the mean market-novelty share among product innovators in the subsector, MIP 2003–2018. $\hat{\kappa}_s$ is recovered from the direction first-order condition; SE propagates sampling error in $\bar{\mu}_s$ via the delta method. Procurement exposure in the regression analysis varies at the broader WZ section level; $\hat{\kappa}_s$ values within a section share the same underlying data variation. *Source:* Authors' calculations using the Mannheim Innovation Panel.

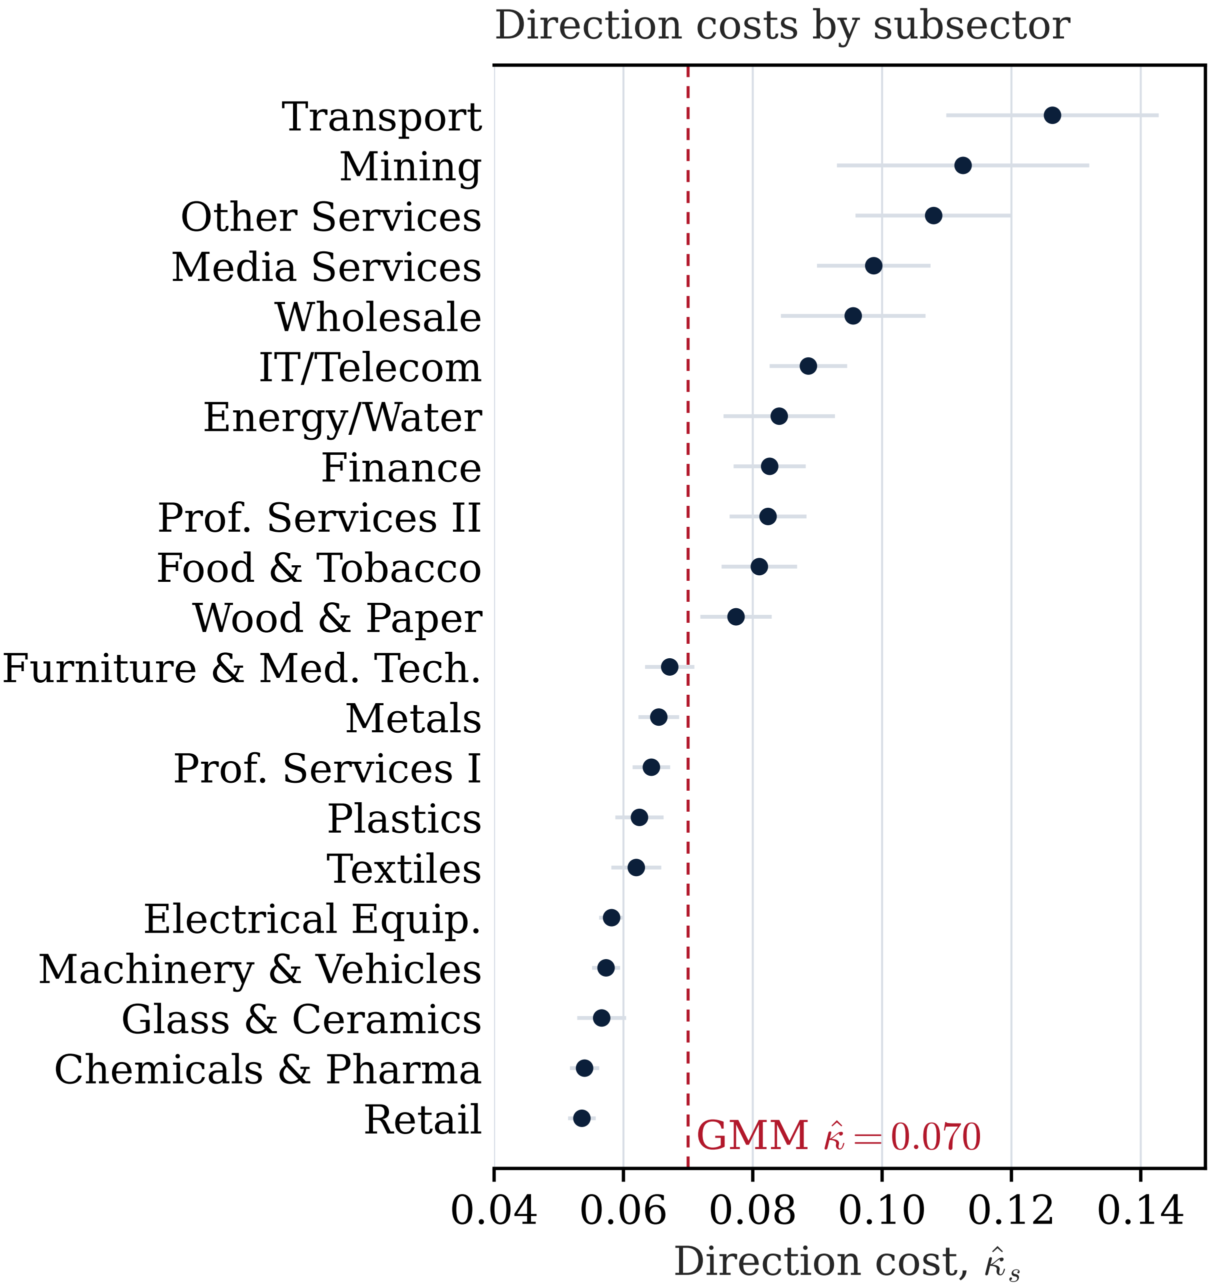

Figure D3: Sector-level direction costs

Notes: The figure plots $\hat{\kappa}_s$ for all 21 MIP subsectors, with 95% confidence intervals (delta-method SE). The dashed vertical line marks the homogeneous GMM benchmark $\hat{\kappa} = 0.070$. Proposition D.4 establishes that cost heterogeneity does not change the proportional compression of frontier orientation by procurement demand.

Source: Authors' calculations using the Mannheim Innovation Panel, 2003–2018.

D.10 Parameter Comparative Statics

Figure D4 shows how each targeted moment responds to individual parameter variation. Each parameter has a distinct signature: κ and λ_I strongly rotate the frontier-share moments (rows 1–2);

h_E primarily moves the large-firm share and creative-destruction rate (rows 4–5); λ_F shifts frontier revenue while leaving other moments relatively flat. Government demand G (rightmost column) reduces frontier composition and creative destruction monotonically, motivating the composition mechanism. The parameter effects are separable enough to identify all four structural parameters, though the quality steps share some influence over rows 1–2.

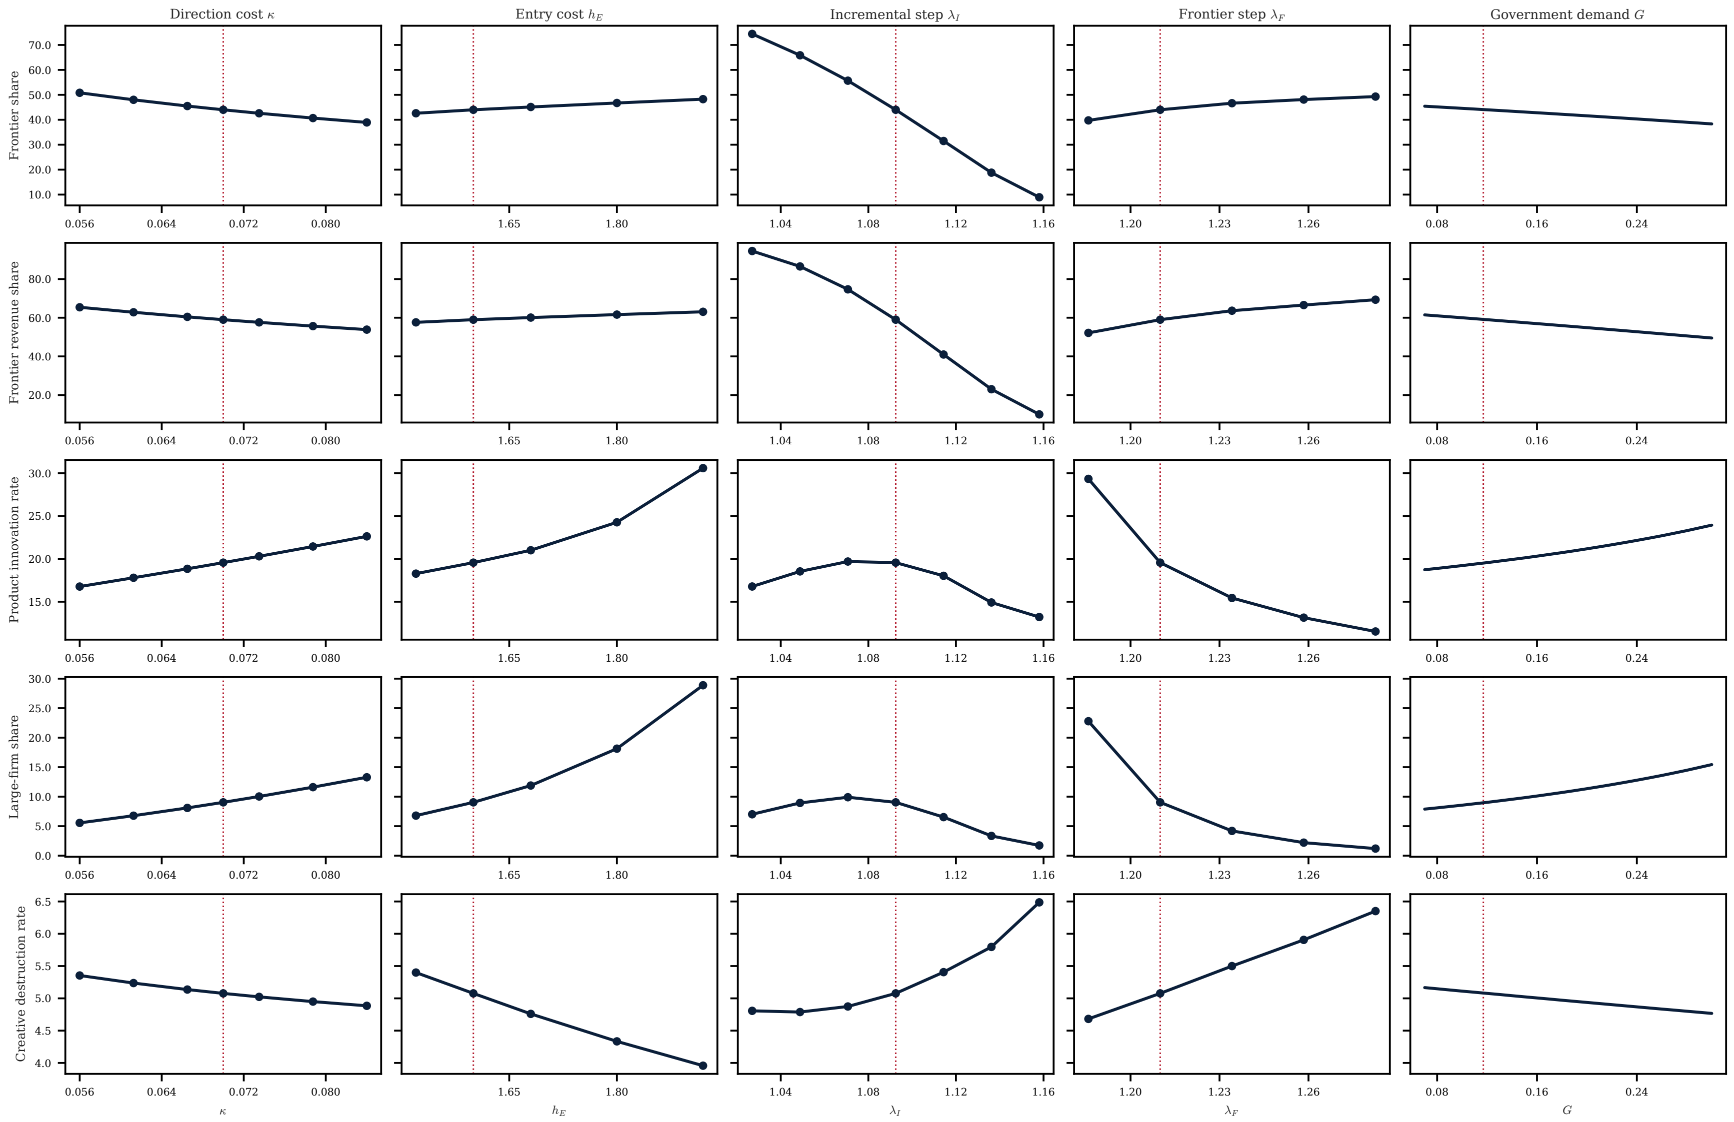

Figure D4: Moment comparative statics: identification

Notes: Rows are targeted moments; columns are parameters. For each structural parameter (κ, h_E, λ_I, λ_F), the curve evaluates the moment at seven values spanning $\pm 20\%$ of the calibrated estimate ($\pm 6\%$ for the quality steps), with all other parameters held fixed at their calibrated values and $G = \bar{G} = 0.117$; dots mark the computed equilibria. The G column (right) traces the full sweep $G \in [0.07, 0.30]$ at benchmark structural parameters. The dotted red vertical line marks the calibrated benchmark of each parameter.

Source: Authors' calculations from the calibrated model.

D.11 Robustness Exercises

Two exercises assess robustness of the quantitative results. The first perturbs each structural parameter by one posterior standard error to assess sensitivity of counterfactual outcomes to estimation uncertainty. The second varies exposure intensity across sectors to check whether the aggregate result masks offsetting heterogeneity.

D.11.1 Parameter Sensitivity

Table D4 perturbs each structural parameter by one posterior standard error, holding the remaining parameters at their GMM estimates, and re-solves the equilibrium at the benchmark procurement level. The exercise holds the unperturbed benchmark R&D cost scale fixed, so each entry is a local deviation from the calibrated baseline rather than a new counterfactual rollback. Across all eight perturbations, annual growth moves by at most 0.70 basis points in either direction and the full CEV ranges from -0.286 to 0.318 percent. The direction-cost parameter κ generates the largest welfare movement, while λ_I generates the largest movement in the frontier-innovation share.

Table D4: Baseline Equilibrium Sensitivity to Parameter Perturbations

	-1 SE shock			1 SE shock		
Parameter	$\Delta\mu^*$ (pp)	Δg (bp)	CEV (%)	$\Delta\mu^*$ (pp)	Δg (bp)	CEV (%)
λ_I	2.18	0.13	-0.030	-2.18	-0.07	0.056
λ_F	-0.50	-0.45	-0.200	0.49	0.46	0.208
κ	1.81	0.70	0.318	-1.67	-0.64	-0.286
h_E	-0.54	-0.67	-0.086	0.54	0.68	0.087

Notes: Entries are deviations from the calibrated baseline. Each row perturbs one parameter by one posterior SE in either direction, holds all other parameters at their GMM estimates, and re-solves the equilibrium at the benchmark procurement level with the unperturbed R&D cost scale fixed. $\Delta\mu^*$ is in percentage points; Δg is in basis points per year; CEV is the full permanent-consumption equivalent, $\Delta \log C_0 + \Delta g/\rho$. Standard errors: $\text{SE}(\lambda_I) = 0.0038$, $\text{SE}(\lambda_F) = 0.0020$, $\text{SE}(\kappa) = 0.0034$, $\text{SE}(h_E) = 0.0149$.

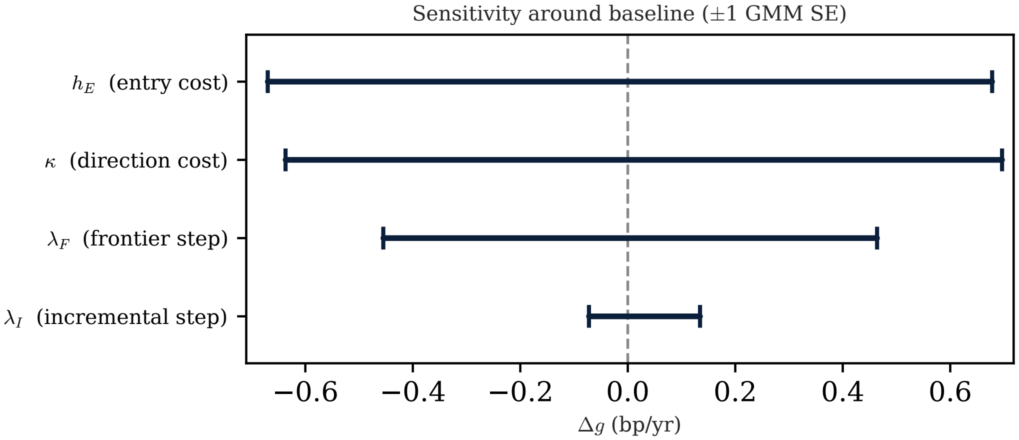

Figure D5: Baseline sensitivity (±1 GMM SE per parameter)

Notes: Each row perturbs one estimated parameter by one GMM standard error in each direction, holding the remaining parameters at their benchmark values. The horizontal segment spans $[\Delta g(-1 \text{ SE}), \Delta g(+1 \text{ SE})]$ relative to the unperturbed baseline; the dashed vertical line marks zero. The unperturbed benchmark R&D cost scale $c_x(\psi)$ is held fixed.
Source: Authors' calculations from the calibrated model.

D.11.2 Sector Exposure Heterogeneity

The composition wedge scales with the government demand parameter G. Mapping mean sector-year procurement exposure to model units proportionally (with the economy-wide mean corresponding to the benchmark $\bar{G} = 0.117$), manufacturing and professional services face the largest wedges: without procurement, the frontier-innovation share is 4.7 and 4.4 percentage points higher, respectively, in those sectors. Transport and IT services fall in the middle range (3.6–3.8 pp), while finance and mining face smaller wedges (1.6–2.2 pp). The pattern is the composition-channel counterpart to Proposition D.3: sectors with larger G face larger wedges and benefit more from specification reform. The aggregate 12-basis-point growth gain from the main counterfactual adds the R&D tax credit component—which benefits all firms regardless of sector—to the sector-level composition effects.