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

Johar Arrieta Vidal[†] Jorge Zavala[‡]

July 4, 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 shift the composition of innovation toward existing products. Linking firm-level innovation surveys to public contract award records, we find that a one-standard-deviation increase in sector procurement exposure is associated with a 0.70 percentage-point lower probability of market-novel product introduction and a 1.3 percentage-point higher share of revenue from existing products. 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 10 basis points per year and delivers a welfare gain of 1.18 percent of permanent consumption.

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

Keywords: public procurement, innovation composition, quality ladders, market novelty, demand-side policy.

---

[*] We thank Anton Bobrov, Kathryn Dominguez, Bastian Krieger, Imryoung Jeong, Nils Lehr, John Leahy, Matthew Schwartzman, Linus Strecke, 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

Public procurement channels government demand for goods and services through contracts that specify technical standards, certification requirements, and delivery schedules. Meeting those specifications is more profitable for existing products than for frontier novelty. In this paper, we study public procurement as a demand-side channel that may redirect the composition of private innovation and quantify the aggregate growth consequences of procurement design.

Germany provides our empirical laboratory. Between the early 2000s and 2023, the share of firms introducing products new to the market fell from 26 to 8 percent while innovation spending held steady at 3 to 5 percent of turnover (BMBF/ZEW, 2024) (Figure 1). The same firms that produce for private markets also compete for public contracts, suggesting a link between contract design and economy-wide innovation composition.

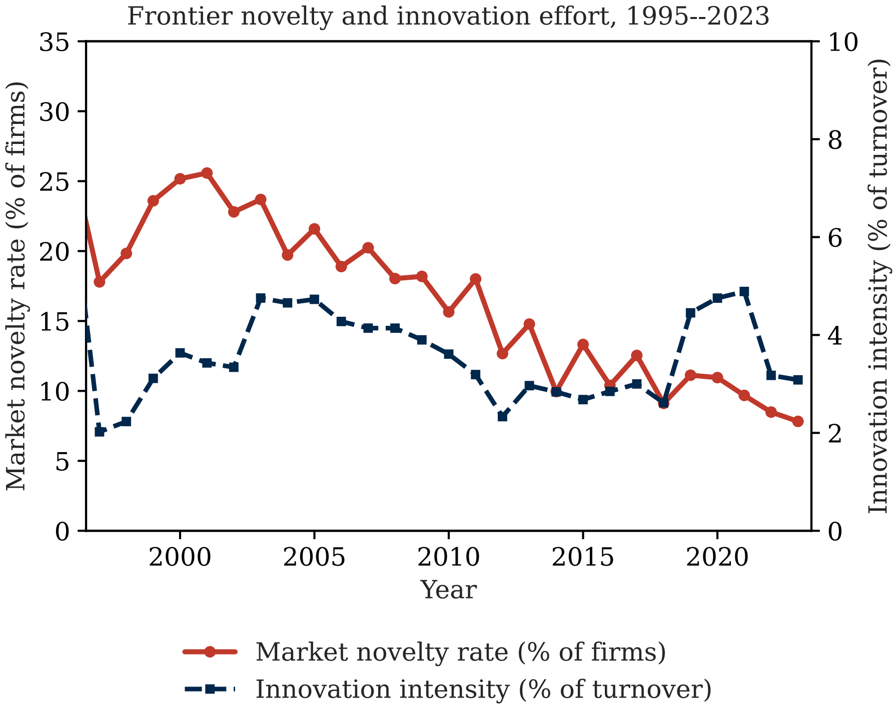

Figure 1: The motivating puzzle in German innovation data

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.

To test whether procurement design shapes the composition of innovation, we link German procurement contract records to firm-level innovation surveys and financial accounts. 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 from 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 with 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. The MIP captures innovation composition, the contract records identify sector-level exposure and contract design, and ORBIS describes how winning firms scale and perform.

Empirically, we estimate panel regressions relating sector-level procurement exposure to firm-year innovation outcomes, with firm and year fixed effects. Identifying variation comes from within-sector changes in procurement activity over time. Three facts emerge.

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; their share of turnover from existing products is correspondingly higher. The pattern extends to cost-reducing innovations: the revenue share from such 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, by contract type, the pattern concentrates in supplies and services; works contracts, where the deliverable is a physical structure rather than a reproducible product, show no significant association on either margin, consistent with product specification requirements rather than public demand as a driver. By category, the composition association concentrates in central-government procurement; regional and local government procurement and utilities show no significant association on either margin.

Third, contract winners expand in scale without productivity or profitability gains: employment and turnover rise with contract volume, while revenue per worker and return on capital show no statistically significant change. Contract winners are also less frontier-oriented than the average firm when first observed in procurement.

Contract design offers an interpretation. Procurement specifications that require certification, compatibility standards, and delivery schedules favor existing products over frontier novelty. In software procurement, integration and compatibility standards are associated with lower rates of new-architecture introduction; in public R&D service contracts, tightly specified deliverables are associated with incremental over exploratory innovation; in medical-device procurement, regulatory certification requirements may reward known-products over diagnostic innovators. The composition association concentrates in central-government procurement, where awards cluster in the specification-intensive supply-and-service categories that show the largest wedge, and is most pronounced in 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.[1] We add one novel 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; firms shift effort toward incremental improvements. The model rationalizes the patterns in the German innovation and procurement microdata: procurement raises equilibrium innovation effort while narrowing the frontier share, consistent with the observed decline in market novelty and the rise in 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 generalized method of moments (GMM), targeting five empirical moments: the frontier share of innovation, the frontier revenue share, firm size, creative destruction, 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 three counterfactual exercises to quantify the growth and welfare consequences of procurement design. The primary exercise 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 10.15 basis points per year, and welfare rises by 1.18 percent of permanent consumption. 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 0.38 basis points. Third, a historically scaled reduction in procurement demand to Germany's 2008 level restores part of the frontier premium, with a growth gain of 0.07 basis points.

Against a German TFP growth baseline of roughly 50 basis points per year, the R&D-credit growth gain represents about one-fifth of trend growth; the 1.18 percent welfare gain lies within the range

---

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

of trade-liberalization gains for small open economies and exceeds standard inflation-cost estimates. Across exercises, procurement design shapes both the scale and the direction of innovative activity.

The paper connects to four strands of the literature. First, we contribute to the endogenous-growth literature on innovation composition (Aghion and Howitt, 1992; Klette and Kortum, 2004; Akcigit and Kerr, 2018; García-Macià et al., 2019; Acemoglu et al., 2022; Aghion et al., 2023; Akcigit and Ates, 2023; Juhász et al., 2024). García-Macià et al. (2019) document that 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. We add a procurement-demand margin to this mechanism: government contracts shift the relative payoff between frontier and incremental activity, redirecting innovation composition without changing its level. The model complements supply-side accounts of composition in Bloom et al. (2020).

Second, we contribute to the empirical literature on demand-side innovation policy (Griffith et al., 2010; Slavtchev and Wiederhold, 2016; Czarnitzki et al., 2020; Chen et al., 2021; Krieger and Zipperer, 2022; Chen and Xu, 2023; Dechezleprêtre et al., 2023; Krieger et al., 2024; Moretti et al., 2025; Chiappinelli et al., 2025; Lerche, 2025; Takalo et al., 2026). Krieger et al. (2024) document that price-criterion tenders reduce product innovation among German firms; Czarnitzki et al. (2020) show that even innovation-directed contracts tilt activity toward incremental rather than frontier novelty. We extend both findings from individual contract variation to sector-level exposure and connect the evidence to quantified design margins. 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*; Lerche (2025) shows that investment tax incentives generate local spillovers as large as the direct effect. For German firms, procurement and subsidies operate on distinct margins: procurement redirects innovation composition while subsidies raise its level.

Third, we contribute to the literature 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 association.

Finally, we contribute a demand-side channel to the misallocation literature (Acemoglu et al., 2018; Terry, 2023; Lehr, 2024, 2025; Bachmann et al., 2026). 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.[2] We use the 2009–2023 releases, which yield 80,028 firm-year observations across 23,547 firms. The panel covers manufacturing and business-service industries; construction, agriculture, and the public sector are excluded. 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.[3] Appendix Table A1 reports summary statistics for all variables, including world-first product innovation, R&D intensity, R&D personnel share, new-product revenue share, and public R&D subsidies used in robustness analyses.

---

[2] We access the MIP through a research data use agreement with ZEW Mannheim. The panel uses a stratified random sampling frame, refreshed annually to maintain representativeness across firm size classes, industries, and regions.

[3] In 2019, the MIP adopted the Oslo-IV definition of process innovation, which expanded the concept and raised the reported rate from 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.[4] 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 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 participants to firm identifiers using direct identifiers, administrative identifiers, and fuzzy company names. The crosswalk matches 41,605 unique procurement winners to the identifiers; the resulting financial panel contains 126,700 firm-year observations across 15,493 firms and supports the winner scale and selection analyses. Appendix Tables B2–B3 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.[5] The annual exit hazard averages 5.07 percent over 2009–2023; we use this rate to calibrate creative destruction in the quality-ladder model (Section 3).

---

[4] Reporting compliance improved markedly after April 2016, when Germany transposed the 2014 EU Procurement Directives; the transposition modernized existing award-notice requirements and coincided with a step change in compliance rates; pre-2016 contract counts therefore understate true sector-level activity. Appendix A.5 details the coverage gap; Appendix Table A4 shows the baseline signs survive with sector-specific trends, and Appendix Table C2 restricts to investment-category CPV contracts.

[5] 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. We use the entry and exit event files for 2009–2023, covering 22.5 million establishment-years. The annual exit hazard is the number of establishments classified as genuine exits in year t divided by the stock active in $t - 1$.

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 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 $\mathrm{arcsinh}(\text{contracts}_{s,t})$, where $\text{contracts}_{s,t}$ is the count of awarded contracts reported 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. The post-2015 rise in contract counts visible in Panel B reflects the expansion of TED reporting coverage rather than proportional growth in actual procurement: German government procurement grew moderately as a share of GDP, while TED counts more than doubled, with the composition by contract type and buyer stable across years. Year fixed effects absorb common changes in reporting coverage; sector-specific trends absorb smooth differences in reporting growth across industries. Residual underreporting likely varies across sectors, buyers, and contract types, so its effect on the estimated coefficient is unknown a priori. We interpret the treatment as reported procurement exposure and assess sensitivity using post-2016, sector-trend, and stable-buyer specifications.

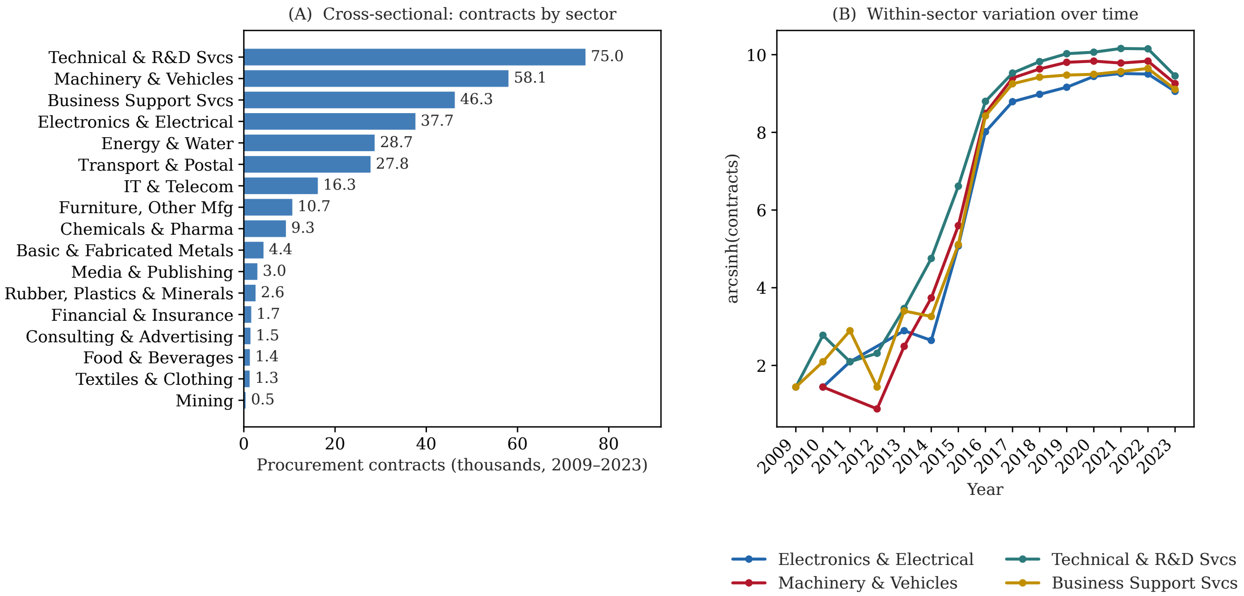

Figure 2: Procurement exposure across MIP-covered industries, 2009–2023

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

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

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.[6] 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 reported industry procurement exposure and innovation outcomes across all firms in a sector (contract winners, unsuccessful bidders, and non-participants alike). Identifying variation comes from changes within sectors over time: within a given year, sectors with above-average procurement exposure have above-average market novelty (cross-sectional selection), yet the within-firm estimate is negative. The specification

---

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

does not rule out time-varying sector shocks or differential reporting changes correlated with innovation.

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)
Effect of +1 SD (pp)	-0.709	1.352	-0.162
Observations	66,980	25,962	40,635
Industries	21	21	21
Firm FE	✓	✓	✓
Year FE	✓	✓	✓

Notes: Each column reports a within-firm OLS regression with firm and year fixed effects. Standard errors in parentheses are clustered by MIP sector (21 groups). The treatment is the inverse hyperbolic sine of awarded contract counts in the firm's MIP sector and year, 2009–2023. The standardized row multiplies the coefficient by the estimation-sample standard deviation of treatment (3.5 units). Binary outcomes are multiplied by 100. Column (2) is the share of turnover from unchanged products; column (3) is the share 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 $\operatorname{arcsinh}(\text{contracts})$ corresponds to a 0.199 percentage-point lower probability of introducing a market-novel product; a one-standard-deviation increase in exposure (3.5 units) corresponds to a 0.7 percentage-point decline. Column (2) shows the complementary revenue pattern: the share of turnover from existing products is 0.379 percentage points higher per unit of exposure; a one-standard-deviation increase corresponds to a 1.3 percentage-point difference. Together, columns (1) and (2) document a shift in the product mix toward existing products and away from frontier novelty.

The composition shift is concentrated at the extensive margin: firms in more-exposed sectors are less likely to introduce any market-novel product, while the market-novel share within new-product revenue and process innovation show no significant association with exposure.

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; a one-standard-deviation increase corresponds to a 0.16 percentage-point decline. Appendix Table A3 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 if they innovate.

2.4 Heterogeneity Across Procurement Environments

The baseline estimate pools all contract types and buyer categories. We now ask whether the composition shift concentrates in specific procurement environments. Figure 3 describes the procurement distribution in MIP-covered industries: services and supplies together account for 93 percent of contracts by count — services 51 percent, supplies 43 percent — with works contracts representing the remaining 7 percent. On the buyer side, regional and local government agencies issue 47 percent of contracts, other public bodies 45 percent, and central-government agencies 8 percent.

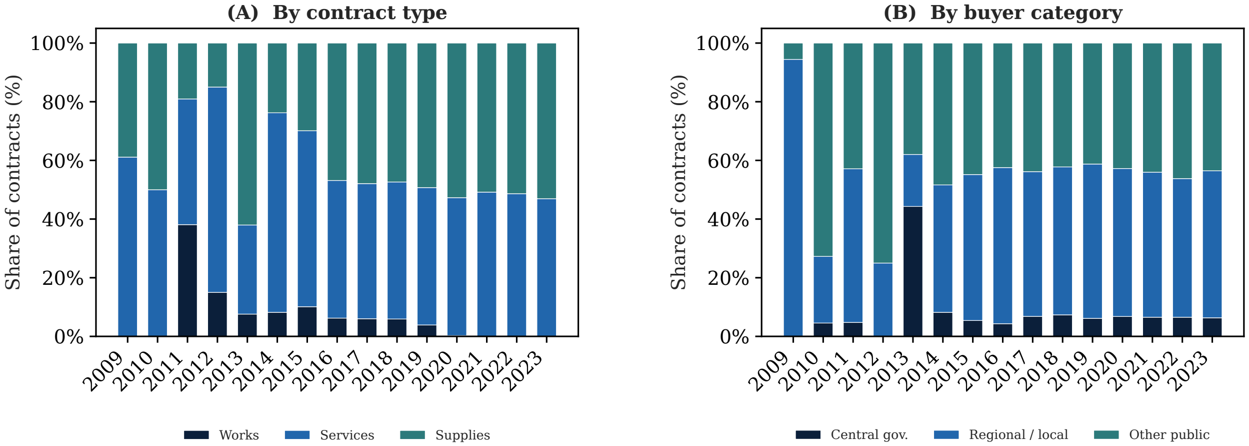

Figure 3: Procurement composition in MIP-covered industries, 2009–2023

Notes: Share of above-threshold contracts by count; Panel A by contract type (works, services, supplies), Panel B by buyer category (central government, regional and local government, other public bodies). Restricted to WZ sections covered by the Mannheim Innovation Panel (manufacturing and business services). Appendix Figure A1 reports the same decomposition by contract value.

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

Table 2 reports the two main innovation outcomes from Table 1 separately for contract types and buyer categories. The splits are descriptive: contract type, buyer category, and contract size co-vary in the procurement data, and the patterns do not isolate a single mechanism.

Table 2: Heterogeneity Across Procurement Environments

	(1) Market novelty (pp)	(2) Unchanged-product revenue (%)	(3) Cost-reduction revenue (%)
Panel A. Contract type			
Supplies & services	-0.211^{**}	$+0.393^{***}$	-0.046^{*}
	(0.099)	(0.137)	(0.023)
+1 SD effect (pp)	-0.752	1.403	-0.158
Works	-0.010	$+0.180$	$+0.024$
	(0.142)	(0.120)	(0.022)
+1 SD effect (pp)	-0.016	0.331	0.039
Panel B. Buyer type			
Central government	-0.291^{**}	$+0.518^{**}$	-0.045
	(0.131)	(0.194)	(0.037)
+1 SD effect (pp)	-0.739	1.329	-0.112
Regional / local	-0.043	$+0.232$	-0.028
	(0.137)	(0.148)	(0.034)
+1 SD effect (pp)	-0.154	0.810	-0.096
Utilities	-0.086	$+0.313$	-0.050
	(0.146)	(0.195)	(0.033)
+1 SD effect (pp)	-0.222	0.788	-0.125
Panel C. Contract type within long-specification sector-years			
Supplies & services	-0.728^{**}	-0.245	-0.158^{*}
	(0.269)	(0.311)	(0.090)
Works	-0.025	-0.258	$+0.092^{**}$
	(0.162)	(0.463)	(0.043)
Observations	$66,980$	$25,962$	$40,635$
Firm and year FE	Yes	Yes	Yes
MIP-sector clusters	21	21	21

Notes: Each cell reports the coefficient from a separate within-firm OLS regression with firm and year fixed effects. Standard errors are in parentheses. In Panels A and B, the treatment is arcsinh(sector-year lot count) for the indicated contract type or buyer category. In Panel C, the sample is restricted to sector-years with above-median technical-requirement text length (median 912 characters); the treatment is arcsinh(sector-year lot count) for the indicated contract type. All estimates use the final 21-sector, 2009–2023 panel. Standard errors are clustered by MIP sector. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

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

Panel A splits exposure by contract type. The supplies-and-services coefficient reproduces the aggregate pattern: market novelty is 0.21 percentage points lower and existing-product revenue is 0.39 percentage points higher. The works estimate is near zero and imprecise, as expected: construction and infrastructure contracts fall outside the MIP coverage frame and contribute little identifying variation.

Panel B splits exposure by buyer category. Despite representing only 8 percent of contracts, central-government procurement drives the aggregate pattern: market novelty falls by 0.29 percentage points and existing-product revenue rises by 0.52 percentage points, both significant at the 5 percent level. Regional and local government procurement (47 percent of contracts) and utilities show no significant effect on either margin. Central-government awards cluster in specification-intensive supply-and-service categories; the composition association reflects their procurement design standards rather than their volume.

Panel C re-estimates the contract-type split within sector-years with above-median technical-requirements text length. In this high-prescriptiveness subsample, services and supplies contracts are associated with a market-novelty decline of 0.728 percentage points, roughly three times the full-sample estimate of 0.211 percentage points in Panel A. Works contracts remain near zero and insignificant. The pattern is consistent with specification detail as a source of the composition association: the market-novelty decline concentrates in the contract type and procurement environment where technical requirements are most prescriptive.

2.5 Winner Outcomes: Scale and Productivity

The sector-level evidence averages over winners, unsuccessful bidders, and non-participants alike. ORBIS, matched to procurement award records for 41,605 winner firms, allows a separate look at winners directly: do firms that secure contracts grow, and does growth come at the cost of market-novel innovation? Table 3 reports within-firm estimates for the matched ORBIS panel over 2009–2023.

Table 3: Winner Outcomes

	(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)
Effect of +1 SD	0.023	0.027	0.006	0.001
Observations	88,363	38,525	33,846	28,470
NACE clusters	20	19	19	19
Firm and year FE	Yes	Yes	Yes	Yes

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. The standardized row multiplies each coefficient by the estimation-sample standard deviation of treatment. 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) show no statistically significant change in labor productivity or return on capital. Winners expand in scale without a corresponding efficiency gain. The estimates capture the intensive margin: the response to additional contracts among firms already in the winner pool.

Figure 4 traces employment, turnover, labor productivity, and return on capital around the year of the first observed contract win. Employment and turnover rise at $t = 0$ and reach 7.5 and 19.5 percentage points by $t = 2$. Labor productivity increases as well, but the gain is mechanical, so revenue per employee rises without a genuine efficiency improvement. Return on capital is statistically indistinguishable from zero throughout the window.

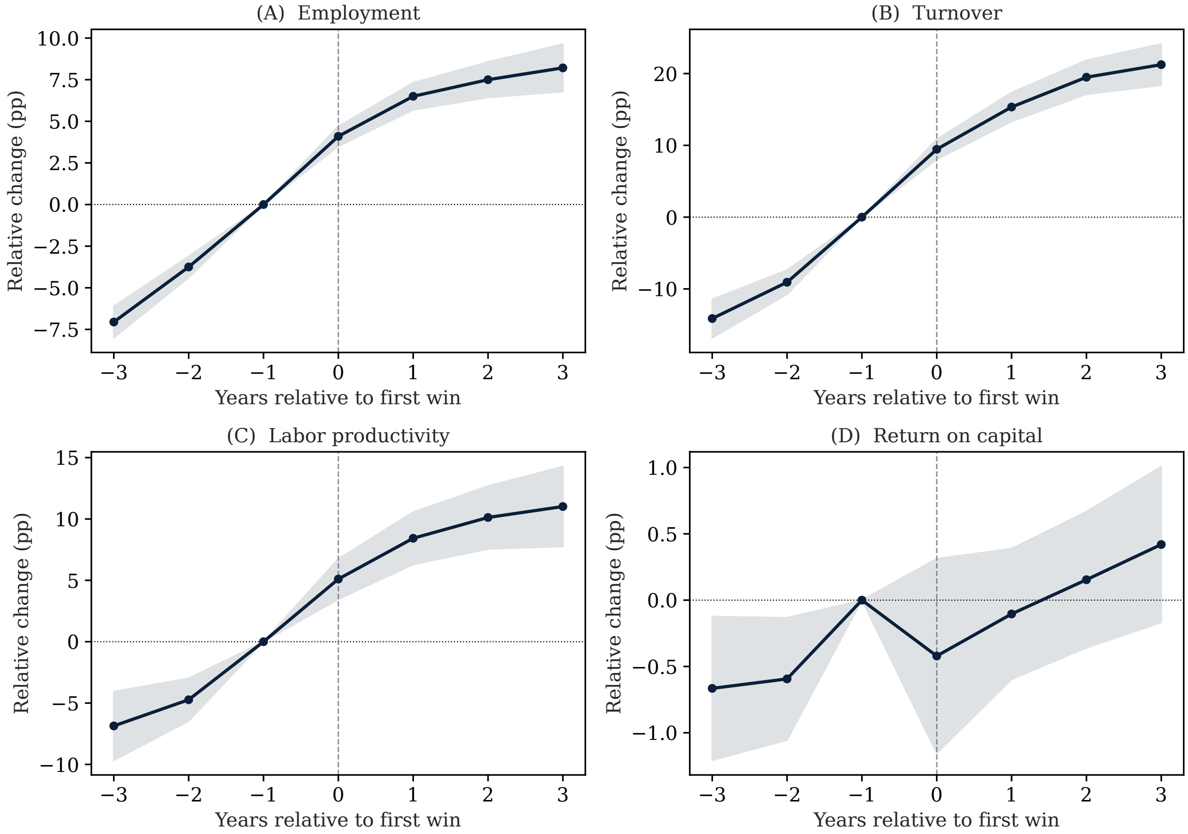

Figure 4: Scale and productivity around the first procurement win

Notes: Each panel shows the outcome demeaned relative to $t = -1$ (year before first win), expressed as percentage-point changes (log outcomes multiplied by 100). Labor productivity is revenue per employee. Return on capital is profit before tax divided by total assets. Shaded bands are 90% confidence intervals. The profile pools all matched winners and does not identify the causal effect of an award. Sample: ORBIS-matched winner panel, TED 2009–2023.
Source: Authors' calculations using TED matched to ORBIS.

The intensive-margin result is robust across firm age, size, and ownership structure (Appendix Tables B2–B3): scale rises and profitability is unchanged in every subgroup. Procurement also operates through a selection margin: the firms that enter the winner pool are mostly oriented toward scale and delivery rather than frontier innovation, compounding the composition shift measured at the sector level.

Together, the sector-level and winner-level evidence document a consistent pattern: procurement exposure shifts firm activity toward existing products, expands scale, and leaves productivity unchanged. This panel evidence identifies the sign and magnitude of the relationship but not the aggregate growth and welfare consequences of a sector facing a persistently different level of exposure. Section 3 builds a quantitative model that formalizes this mechanism and evaluates the aggregate implications of alternative procurement rules by comparing steady states at different, permanent levels of government demand rather than tracing the transition between them.

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 the quality-ladder tradition of Aghion and Howitt (1992); Aghion et al. (2023); the key extension is the composition choice μ, which lets incumbent firms select not only how much but also what type of innovation to pursue. The choice between radical (frontier) and incremental innovation follows Acemoglu et al. (2022), applied to a demand-composition setting. Government procurement enters through two channels: it raises 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 with a log aggregator over the unit-mass continuum:

$$Y = \exp\left(\int_0^1 \log y_j \, dj\right), \tag{3}$$

where the unit mass of lines is fixed. Bertrand competition at the quality lead determines each firm's 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 procurement raises demand for incremental product lines. The government spends a fraction $G \geq 0$ of private-sector output per incremental line, going exclusively to specification-compliant suppliers rather than frontier innovators; Appendix E.6 microfounds the asymmetry. The direct effect is to raise incremental-line profits by a factor $(1 + G)$, compressing the frontier premium.

The government produces output with the same technology as the private sector: one unit of labor at wage w yields one unit of quality-adjusted output, so procurement is real spending on real goods rather than waste. Firms fill government and private orders from the same workforce; procurement raises the demand for production labor without directly distorting R&D incentives.

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. The model has no knowledge spillovers, which rules out composition distortions from uninternalized knowledge benefits; procurement supplies the demand-side wedge studied here.

Entry. Outside entrants acquire an existing product line by using h_E units of setup labor, with monetary entry cost wh_E. Entry is untargeted: entrants arrive randomly and take ownership of an incumbent line with equal probability, independently of its type or quality level. Entry governs firm turnover and the scale of creative destruction; quality growth comes exclusively from incumbent innovation.

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

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

Entry labor is h_E per entrant.

3.2 Agent Problems

Pricing. Each quality leader sets the Bertrand limit price against the next-best rival; the quality lead λ_k, rather than demand elasticity, determines the 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. The firm's portfolio value combines operating profits from each line, net of innovation costs, with the option to expand the portfolio through incremental or frontier advances: both the total number of lines and their composition matter for value. Appendix D.1 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.1 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 $\mu^* \in \{0, 1\}$.

At corner solutions $\mu^* \in \{0, 1\}$, the KKT conditions replace the interior equality; Appendix D.1 states the full characterization.

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. In the baseline, entrants do not create varieties or quality improvements; they use h_E units of setup labor to take ownership of existing product lines. Let τ denote the economy-wide rate at which any given line changes hands, and let e denote the rate at which outside entrants acquire lines. Since incumbents also displace rivals through quality-improving innovation at rate x per line they hold, the total ownership-change rate satisfies

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

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

Potential entrants do not choose innovation effort or direction in the baseline. They draw the type of the line they acquire from the equilibrium line distribution, so the expected acquisition value is $\bar{v}(\mu^*; G)$. Free entry equates this value to the monetary setup cost:

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

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

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

Only incumbent innovation appears in (22); entrant acquisitions change ownership and firm demographics but add no quality steps. Appendix **??** considers the case where entrants also innovate.

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.* On the interior, the marginal cost of increasing the frontier-innovation probability equals the value premium from holding a frontier rather than an incremental line:

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

At $\mu \in \{0,1\}$, the KKT conditions in Appendix D.1 replace the equality.

(iv) *Free entry.* The expected value of an acquired line equals the entrant's setup 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.2 describes the solution algorithm. Counterfactuals vary G and related policy parameters and re-solve the full system.

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 $[0,1]$.[7]

Proposition 3.2 (Procurement compresses the frontier share)**.** *Let* $A(G) \equiv \bar{\pi}_F - \bar{\pi}_I(1+G)$. *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*

$$\frac{1}{x^*(G)}\frac{\partial x^*(G)}{\partial G} < \frac{\bar{\pi}_I}{A(G)} + \frac{1}{r+\tau}\left(\frac{\partial r}{\partial G} + \frac{\partial \tau}{\partial G}\right),$$

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

Proof. Substituting (29) into (16) gives the interior solution (30). Differentiating along the equilibrium path and rearranging,

$$\frac{\partial \mu^*}{\partial G} < 0 \iff \frac{1}{x^*(G)}\frac{\partial x^*(G)}{\partial G} < \frac{\bar{\pi}_I}{A(G)} + \frac{1}{r+\tau}\left(\frac{\partial r}{\partial G} + \frac{\partial \tau}{\partial G}\right).$$

---

[7] 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.1 states the optimality condition used when counterfactuals reach a corner at $\mu^* = 1$.

The stated condition ensures that direct frontier-premium compression and the equilibrium $r + \tau$ response dominate the effort response. $\square$

Proposition 3.2 rationalizes the pattern in Table 1.[8] 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.

Labor reallocation. When G rises, the labor market clearing condition (21) tightens against R&D: incremental lines serve both private and government demand, so production labor L_P rises with G, crowding out R&D labor ℓ_R. The compression mechanism in equation (29) captures the price channel (the frontier premium narrows) but not the quantity channel (fewer workers available for R&D). Both channels reduce frontier innovation. The policy counterfactuals reverse both: removing the demand wedge restores the frontier premium, and reducing government-driven production-labor demand frees resources for innovation. Table 6 reports ΔL_P and ΔL_R for each exercise, quantifying the size of the quantity channel relative to the price channel.

Markup distribution. The composition shift has a second implication. Under Bertrand pricing, frontier lines earn a higher operating-profit share than incremental lines ($\bar{\pi}_F > \bar{\pi}_I$); as G rises and μ^* falls, the portfolio shifts toward lower-markup lines. Markup heterogeneity of this kind generates endogenous misallocation (Peters, 2020): procurement design concentrated on incremental lines compresses average markups and amplifies productivity dispersion across product lines.

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. Each entrant begins 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

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

$q = x^*/\tau$ is the probability that when a product line changes ownership, it is acquired by an expanding incumbent rather than a new entrant; higher q implies larger average firm scope. Together with μ^*, it is a sufficient statistic for the full cross-sectional distribution of firm portfolios. The

---

[8] 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.

stationary distribution of firm scope is

$$\Pr(N = n) = \frac{q^n}{n[-\log(1-q)]}, \qquad n \geq 1. \tag{31}$$

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.[9]

The joint distribution Γ combines the scope distribution (31) and, conditional on N, a binomial over line types: each line independently becomes frontier with probability μ^*, regardless of portfolio size.

$$\Gamma(n_I, n_F; G, \boldsymbol{\psi}) = \Pr(N = n_I + n_F) \binom{n_I + n_F}{n_F} (\mu^*)^{n_F} (1-\mu^*)^{n_I}. \tag{32}$$

Conditional on $N = n$, the number of frontier lines follows:

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

Three cross-sectional properties follow from Γ and connect the model directly to observable moments.

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*

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

The average frontier share across firms equals μ^*, so the composition moment used in estimation is simply the cross-sectional mean of s_F.

(ii)

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

---

[9] The normalizing constant $-\log(1-q) = \sum_{n=1}^{\infty} q^n/n$ ensures probabilities sum to one (logarithmic series identity). The distribution is logarithmic rather than geometric because each of the firm's n lines independently generates a new line at rate x^* and loses one at rate τ: both birth and death rates scale linearly with n. A geometric distribution arises when the net gain rate is constant, independent of portfolio size; the proportional scaling here yields the logarithmic form as the unique stationary solution.

As μ^* falls, more firms end up with no frontier lines at all: each line they hold is less likely to be frontier-oriented, so the share of purely incremental firms rises.

(iii) $\mathbb{E}[\pi(G)] = \mu^* \bar{\pi}_F + (1 - \mu^*)\bar{\pi}_I(1 + G)$. *A lower frontier share shifts the portfolio toward incremental lines, which earn lower markups, reducing average per-line profits.*

Higher government demand lowers μ^* and, through the equilibrium entry condition, also adjusts q; both shifts push the cross section toward incremental-only portfolios. The next section estimates $(\kappa, \lambda_I, \lambda_F, h_E)$ by matching five empirical moments that span the composition and size margins.

4 Quantitative Analysis

The quantitative analysis fits the model to observed innovation composition and firm dynamics, then compares policy 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), when frontier-share, firm-size, and new-product variables are jointly available.

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 $\bar{G} = 0.018$ implies that procurement in MIP-covered industries equals approximately 1 percent of those industries' gross output — obtained by scaling total reported German procurement (EUR 132 billion, Vergabestatistik 2022 (BMWK, 2023)) by the share flowing to MIP-covered sectors (27.5 percent from TED supplier-industry composition) and dividing by Destatis Input-Output gross output for those sectors. The model uses this benchmark to translate the compositional pattern in the MIP 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 $m(\boldsymbol{\psi})$ stacks model-implied moments, $\hat{m}$ their empirical counterparts, and W is the inverse of the moment variance-covariance matrix estimated via influence functions. For each candidate $\boldsymbol{\psi}$, the innovation cost scale $c_x(\boldsymbol{\psi})$ is adjusted to match the observed German long-run growth rate exactly, so the estimated parameters govern composition and firm dynamics independently of the aggregate growth level.

4.2.1 Identification and Model-to-Data Mapping

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 revenue share among new products) 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 identifies h_E along a margin independent of firm size and turnover. The procurement coefficients are not targeted moments: their signs discipline the mechanism externally, and the model reproduces the qualitative pattern — less frontier innovation, more unchanged-product revenue, and greater scale as procurement rises.[10] Table 4 presents the estimates.

---

[10] Appendix E.1.1 details the measurement equations for each targeted moment. Appendix Table E2 provides a magnitude check: under a proportional mapping from TED contract counts to G, the model implies a market-novelty coefficient of -0.097, compared with -0.199 in the data.

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	0.158	Implied by Bertrand pricing
$\bar{\pi}_I$	Incremental per-line profit	0.095	Implied by Bertrand pricing
$c_x(\boldsymbol{\psi})$	Innovation cost scale	21.550	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	0.018	(BMWK, 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	1.188 (0.003)	GMM estimate
λ_I	Incremental quality step	1.105 (0.002)	GMM estimate
κ	Direction cost	0.055 (0.004)	GMM estimate
h_E	Setup labor per entrant	1.638 (0.009)	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 $\bar{G} = 0.018$ implies procurement in MIP-covered industries equal to $(1 - \mu^*)\bar{G} \approx 1\%$ of those industries' gross output, calibrated from Vergabestatistik 2022 and Destatis Input-Output accounts (BMWK, 2023). 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, and Destatis.

The parameter estimates are economically sensible. The frontier quality step $\hat{\lambda}_F = 1.188$ implies an 18.8% productivity improvement per creative-destruction event; the incremental step $\hat{\lambda}_I = 1.105$ implies a 10.5% 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.

Under Bertrand competition, per-line profits follow from the quality steps: $\bar{\pi}_F = 15.8\%$ and $\bar{\pi}_I = 9.5\%$, a frontier premium of 6.3 p.p. 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 $\hat{\kappa} = 0.055$ places adjustment frictions well below the profit premium, so composition responds to the procurement-induced gap narrowing. The setup requirement $\hat{h}_E = 1.638$ labor units clears the free-entry condition at the observed 5.1% creative-destruction rate.

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)	56.293	1.956
Large-firm share	0.104	(0.003)	0.096	0.008
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 ($J = 104.4$, one degree of freedom). The R&D personnel target measures employees formally classified as R&D personnel, whereas the model counterpart measures incumbent innovation labor; this mapping is approximate.
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. 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 versus 0.029); part-time and multitasking R&D roles, excluded from the model, likely account for part of the gap.

4.2.2 Comparative Statics

Figure 5 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. Under free entry, aggregate innovation effort adjusts to clear 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.

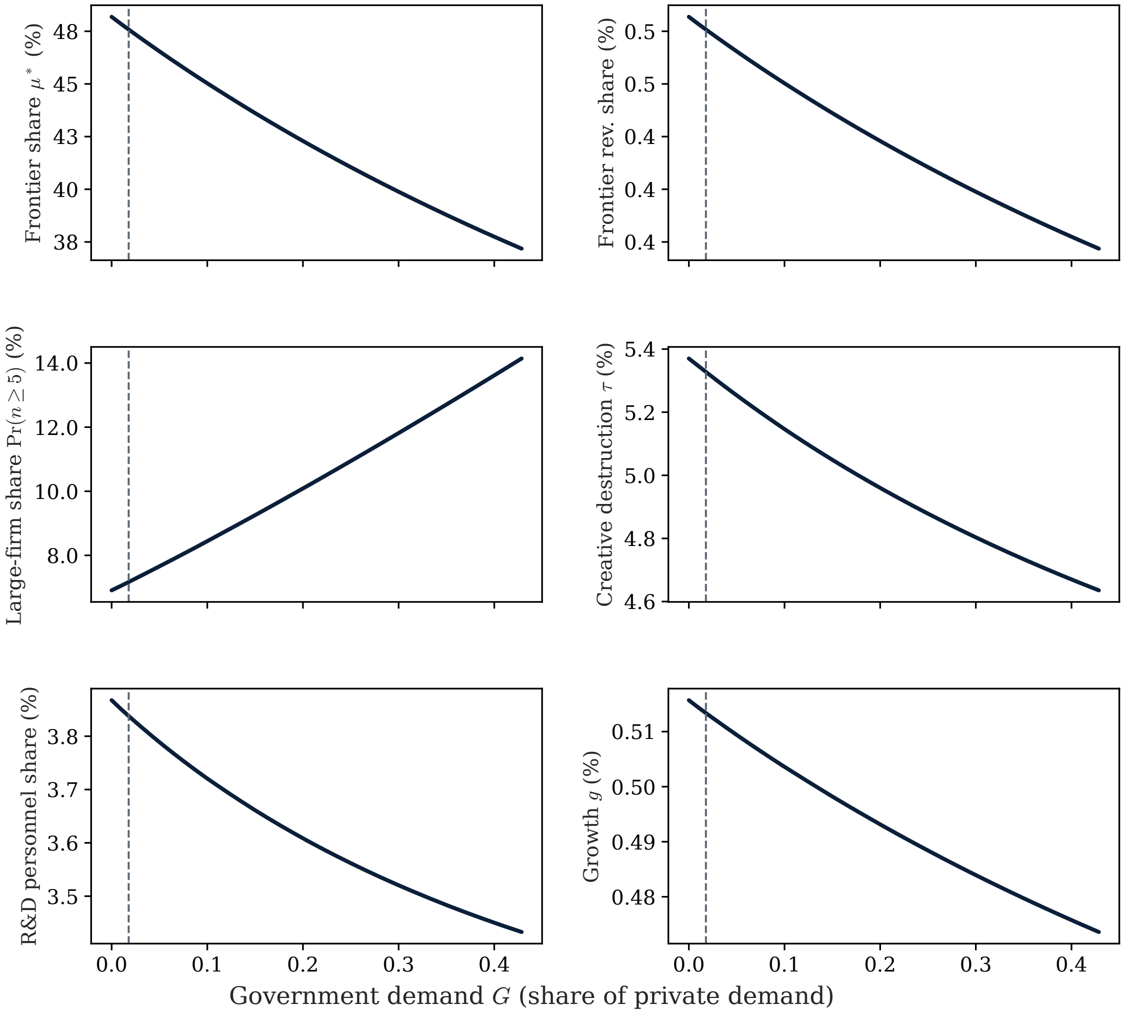

Figure 5: Comparative statics over government demand

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 $\bar{G} = 0.018$. 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

We examine three counterfactual experiments from the estimated model, each varying a single policy instrument while holding calibrated parameters at their benchmark values.

Each counterfactual uses the same fiscal accounting: procurement is financed by a lump-sum household tax $T = (1 - \mu^*)G$, equal to $T_{bm} \approx 0.010$ at the calibrated benchmark. Because a lump-sum tax is neutral in isolation, the welfare effect of each experiment depends on how the freed fiscal resources are redeployed.

Table 6 summarizes three 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. A design reform keeps total procurement spending fixed but changes how awards are allocated: we realign contract specifications toward open awards. A third exercise reduces G to a historically scaled estimate of its 2008 level.

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 permanent-consumption equivalent between policy B and benchmark A is

$$\mathrm{CEV}_{B,A} = 100 \left\{ \exp\left[\Delta \log C_0 + \frac{g_B - g_A}{\rho} \right] - 1 \right\}, \tag{34}$$

with C_0 from the household budget (2) and $\rho = 0.04$. In all scenarios, the household tax adjusts with 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 the full procurement-tax envelope 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 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

$$(r+\tau)v_k = \pi_k(0) - (1-s)\Phi(x) - \Psi(\mu) + x\bar{v}(\mu; 0), \qquad k \in \{I, F\}, \tag{35}$$

where the cost reduction increases equilibrium innovation effort and, through the frontier-choice condition (16), moves the innovation 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 this instrument. The counterfactual redirects the procurement tax $T_{bm} = 0.010$ to a proportional R&D credit at rate $s = 0.141$. The eligible base extends beyond R&D wages to entry costs and a production-implementation component equal to 2.53% of production labor (covering prototyping, process redesign, and adoption costs that credit designs routinely include), making total credit expenditure equal to T_{bm}.[11] As a lower bound, removing procurement without an R&D credit yields a growth gain of 0.28 basis points and CEV of 0.10%. Table 6 reports the main exercise.

Under the full-budget implementation, the frontier share increases from 0.441 to 0.502 ($\Delta\mu^* = 0.061$), and innovation effort is 16.4% higher. R&D labor rises by 18.7% while production labor falls by only 0.3%: because R&D labor is a small share of the labor endowment at the benchmark, the credit draws workers overwhelmingly from a mix of production labor and entry labor rather than displacing production one-for-one, so both the price channel (frontier premium) and the quantity channel (labor available for R&D) move in the same direction. Growth rises by 10.15 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. The net CEV reflects a negative direct consumption effect: the household tax finances the credit rather than falls away, offset by a dominant long-run growth gain.

5.2 Alternative Instruments and Design Reforms

The remaining exercises separate two margins. The first asks how much the composition shifts when procurement design is altered but total spending is held fixed. The second scales aggregate demand to its 2008 level, returning the fiscal savings to households.

Specification realignment. Specification-intensive procurement concentrates government demand on incremental lines by requiring suppliers to meet product requirements rather than performance objectives. Realigned specifications shift toward open, functional awards (Howell et al., 2025). Open awards are associated with higher frontier innovation rates, and functional specifications, rather than product-specific requirements, account for the pattern (Uyarra et al., 2014; Edquist and Zabala-Iturriagagoitia, 2020).

Let government demand decompose into a frontier component G_F and an incremental component G_I. In the benchmark, $G_F = 0$ and all demand flows to incremental lines; realigned specifications allow the government to place orders on both line types. Hence, the budget constraint is $(1-\mu^*)G_I + \mu^* G_F = T_{bm}$. Setting $G_I = G_F$ and imposing the budget constraint yields $G_I = G_F = T_{bm}$: both line types face the same per-line government demand and profits satisfy $\pi_k = \bar{\pi}_k(1 + T_{bm})$ for

---

[11] 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 2.53% component is the model-implied residual, not a separate data target. Appendix E.5 varies both entry eligibility and the implementation share.

$k \in \{I, F\}$. The incremental demand advantage disappears, widening the net frontier premium and raising equilibrium frontier orientation. The frontier share rises by $\Delta\mu^* = 0.014$, growth by 0.38 basis points, and CEV by 0.17%. Firms do not attempt more innovations under this reform ($\Delta x^* = 0.00\%$): the arrival rate of new lines is unchanged. What changes is what those attempts aim for. More of the same fixed number of attempts now target frontier lines, and frontier-oriented projects cost more labor to direct than incremental ones. That alone raises R&D labor by 1.0% and frees up production labor by 0.1%: a labor-market response to procurement reform that runs entirely through the type of innovation firms pursue, not through how often they try. Appendix E.3 formally derives the equilibrium response.

Historical rollback. This exercise sets G to its 2008 level, using the OECD national-accounts procurement share to calibrate $G_0 = 0.013$ against the benchmark $\bar{G} = 0.018$. The rollback also reduces the household tax $T = (1 - \mu^*)G$: households gain from both improved composition and a smaller fiscal burden. Because Germany's broad procurement-to-GDP ratio was relatively stable over this horizon, the absolute reduction in $\bar{G}$ is modest; the frontier share rises by $\Delta\mu^* = 0.003$, growth by 0.07 basis points per year, and CEV by 0.026%. R&D labor rises by 0.2% and production labor falls by 0.02%, an order of magnitude smaller than under either of the other two exercises, consistent with the modest scale of this reform. The direction is consistent with the empirical estimate, but the historical scaling is not a calibration target [12].

The comparison with the design reform is instructive: specification realignment, which holds total spending fixed, generates nearly five times the frontier-share gain ($\Delta\mu^* = 0.014$ versus 0.003) at the same fiscal cost. The composition distortion thus reflects how procurement resources are allocated across line types, not the aggregate scale of government demand. All three exercises compare initial and final steady states; transition dynamics are not modeled, so the welfare estimates approximate the long-run gain once the economy has fully converged.

---

[12] TED contract counts grew far faster than the OECD GDP share over this period, partly reflecting improved reporting compliance rather than real procurement growth.

Table 6: Policy counterfactuals from the GMM benchmark

Counterfactual	Δx^* (%)	$\Delta\mu^*$	ΔL_P (%)	ΔL_R (%)	Δg (bp/yr)	CEV (%)
R&D tax credit	16.4	0.061	-0.32	18.69	10.15	1.18
Specification realignment	0.00	0.014	-0.12	1.00	0.38	0.17
Historical rollback	0.00	0.003	-0.02	0.22	0.07	0.026

Notes: Δx^* is the percent change in equilibrium innovation effort. ΔL_P and ΔL_R are the percent changes in aggregate production and incumbent R&D labor relative to the benchmark, holding the fixed labor endowment $\bar{L}$ fixed; the entry-labor residual absorbs the remainder. CEV is the permanent-consumption equivalent $100\{\exp[\Delta \log C_0 + \Delta g/\rho] - 1\}$ with $\rho = 0.04$. Row 1 removes procurement and redeploys the fixed tax $T_{bm} = 0.00997$ to a broad-base R&D credit with $s = 14.1\%$, $\omega_E = 0.5$, and $\omega_P = 0.0253$. Row 2 holds total procurement spending fixed using $(1-\mu)G_I + \mu G_F = T_{bm}$. Row 3 scales G by the OECD procurement-to-GDP ratio for Germany in 2008 relative to the calibrated benchmark and returns the reduction in the tax to households. *Source:* Authors' calculations from the calibrated model.

The welfare effects are comparable in scale to other quantitative benchmarks. The main R&D-credit counterfactual (1.18% CEV) exceeds the inflation-cost estimate of Lucas (2000) and the short-termism distortion in Terry (2023), falls below estimates for R&D misallocation and monopsony costs (Lehr, 2024, 2025), and lies at the lower end of trade-liberalization gains (Arkolakis et al., 2012; Caliendo and Parro, 2015; Melitz and Redding, 2015; Costinot and Rodríguez-Clare, 2015). It is also below business-cycle costs of around 2% (Krusell et al., 2009). Budget-neutral design reform is smaller because it operates through composition alone; the gap between the two types of reform reflects the fiscal channel: removing procurement eliminates both the composition wedge and the household tax, while design reforms hold spending fixed.

These gains translate into long-run TFP and GDP per capita growth: in quality-ladder models, output expansion comes entirely from quality improvements. Against a German baseline of $\bar{g} \approx 50$ basis points per year (Bergeaud et al., 2016), the R&D-credit counterfactual adds 10.15 basis points and specification realignment adds 0.38 basis points. Both run through composition: procurement reform moves the innovation mix toward creative destruction, the highest-step margin. The 10.15 basis-point gain is of the same order as the composition-driven component of the US productivity slowdown: García-Macià et al. (2019) attribute 16 of the 34 basis points of TFP growth lost between the 1980s and 2000s to a shift away from creative destruction; against Germany's $\bar{g} \approx 50$ basis-point baseline, the procurement composition wedge accounts for a comparable share of forgone growth. By contrast, Atkeson and Burstein (2019) find limited productivity gains from scale-neutral innovation subsidies, reinforcing the importance of composition over scale.

5.3 Robustness Checks

We conduct four robustness exercises. The first explores an alternative policy scenario: a procurement officer risk-preference reform that shifts demand toward frontier lines through reduced career risk, microfounded by a mean-variance government objective. The second evaluates the calibrated model at sector-specific procurement wedges, testing whether the aggregate composition pattern holds within industries. The third checks structural robustness through parameter perturbations and an extended model variant in which entrants also innovate. The fourth checks whether the welfare ranking survives once government purchases are credited with their own value rather than treated as pure transfer.

Procurement officer risk preferences. Procurement rules create career incentives that push buyers toward low-risk, specification-compliant suppliers. When transparency requirements expose shortfalls and career penalties fall on contracting officers who oversee failed frontier projects, officers rationally favor incremental lines whose delivery risk is lower and easier to document (Prendergast, 2007; Bandiera et al., 2009). Rigid award criteria that reward compliance over performance compound the bias. Procurement reforms that reduce career penalties or tighten technical evaluation lower the effective risk premium on frontier awards and shift demand toward frontier lines at zero additional fiscal cost (Uyarra et al., 2014; OECD, 2017; Decarolis et al., 2021; Bosio et al., 2022).

Let the government hold mean-variance preferences over delivered quality with risk-aversion parameter $\gamma_G \geq 0$. Frontier lines offer higher expected quality but also higher delivery variance than incremental lines. For sufficiently risk-averse buyers, the lower delivery variance outweighs the frontier quality advantage, rationalizing the baseline demand asymmetry $G_F = 0$; Appendix E.6 derives the preference threshold and its comparative statics. Redirecting one-quarter of the procurement budget toward frontier lines raises the frontier share by $\Delta\mu^* = 0.008$, growth by 0.21 basis points, and CEV by 0.09%.

Sector-specific wedges. We evaluate the calibrated model at sector-specific procurement wedges G_s instead of re-estimating it separately by sector. We map mean sector-year reported contract counts proportionally into model units, normalizing their mean to the covered-sector benchmark $\bar{G} = 0.018$. Among sectors with positive exposure, the resulting wedges range from $G_s = 0.008$ to $G_s = 0.025$. The frontier share in each sector is inversely related to its procurement exposure, matching the cross-sectoral pattern: the largest declines in μ_s^* occur in the most exposed sectors. Removing the sector wedge raises the frontier share by 0.48–1.44 percentage points and annual growth by 0.13–0.39 basis points. Because the count-to-G conversion is imposed rather than estimated, we treat these results as illustrative rather than as sector-specific calibration. Appendix E.4.1 reports the full mapping.

Model variants. We assess robustness to two alternative specifications. In the first, we perturb each structural parameter by one standard error around the calibrated baseline; annual growth moves by at most 0.72 basis points in either direction and the welfare change ranges from -0.313 to $+0.312$

percent, with κ generating the largest movement in the frontier-innovation share. In the second, we re-estimate an extended specification in which entrants also raise line quality and the policy ranking survives. Appendices E.5 and **??** report the full results.

Government goods in welfare. The baseline CEV treats procurement spending as a pure transfer: households pay the tax that finances it, but the model assigns no utility value to what the government buys, even though procurement in practice pays for infrastructure, IT, medical supplies, and other outputs the government needs to operate. Appendix E.6 extends household utility to include a log-weighted index of government-purchased output, weighted by $\alpha \in [0, 1]$ relative to private consumption, and recomputes CEV for all three main counterfactuals under this alternative aggregator. The design reforms are essentially unaffected (government-goods value rises slightly as μ^* increases, since more government demand reaches frontier lines). The R&D credit, which eliminates procurement entirely, loses the full benchmark government-goods term; even so, the break-even weight is $\alpha^* \approx 1.19$, above the maximum reasonable weight of $\alpha = 1$ (equal value to private and public consumption), so the credit remains welfare-improving even when government purchases are valued at par with private consumption. Only the historical rollback, whose baseline gain is small, turns welfare-reducing once α is close to one. Table E9 reports the full sensitivity.

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. This paper documents public procurement as a channel associated with the observed composition shift. Government contracts specify technical standards, certification requirements, and delivery schedules that favor existing products over frontier novelty; sectors with higher procurement exposure introduce market-novel products less often and generate more revenue from existing products.

Three patterns characterize the evidence. On innovation composition, a one-unit increase in sector procurement exposure is associated with a 0.199 percentage-point lower market-novelty rate and a 0.379 percentage-point higher existing-product revenue share. On firm scale, matched procurement winners expand in scale, with higher employment and turnover, without a corresponding gain in productivity or profitability. On heterogeneity, the composition association concentrates in supplies and services contracts and is larger for central-government buyers; works contracts show no significant association on either margin. The evidence is consistent with a shift away from frontier novelty and toward scale-oriented delivery rather than a reduction in innovation effort overall.

We extend Klette and Kortum (2004) by adding a direction-choice margin: each period, firms allocate innovative effort between frontier and incremental lines, facing separate quality steps and a reallocation cost. Procurement enters as a demand wedge that compresses the frontier premium, shifting the equilibrium mix toward incremental delivery. We estimate four structural parameters by GMM, targeting five empirical moments from the microdata; the fitted model attributes the growth cost to composition rather than to the level of public spending.

The primary counterfactual redirects procurement funds to an R&D tax credit modeled on Germany's *Forschungszulage*. Growth rises by 10.15 basis points per year and welfare by 1.18 percent of permanent consumption; about one-fifth of Germany's trend TFP growth and above standard inflation-cost estimates.

A budget-neutral specification reform raises growth by 0.38 basis points; a scaled volume reduction adds 0.07 basis points. Across counterfactuals, the destination of public funds matters more for growth than their volume or contractual form: redirecting procurement to an innovation subsidy yields a gain an order of magnitude larger than reforming contract design, and design reforms in turn dominate volume reduction.

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 (Uyarra et al., 2014; OECD, 2017; Edquist and Zabala-Iturriagagoitia, 2020). 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 would decompose the composition margin into vertical quality upgrading and horizontal product-line expansion using establishment-level records, following García-Macià et al. (2019).

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.

Bandiera, Oriana, Andrea Prat, and Tommaso Valletti, "Active and Passive Waste in Government Spending: Evidence from a Policy Experiment," *American Economic Review*, 2009, *99* (4), 1278–1308.

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.

BMWK, “Fostering Innovation through Public Procurement,” Technical Report, Berlin 2023.

Bosio, Erica, Simeon Djankov, Edward Glaeser, and Andrei Shleifer, “Public Procurement in Law and Practice,” *American Economic Review*, 2022, *112* (4), 1091–1117.

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.

European Commission, “Economic and Legal Analysis of the Factors Leading to a Low Rate of Publication of Public Procurement Opportunities in Germany,” Report, Directorate-General for Internal Market, Industry, Entrepreneurship and SMEs 2019.

— , “Single Market Scoreboard: Public Procurement,” 2024. Directorate-General for Internal Market, Industry, Entrepreneurship and SMEs.

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.

Jones, Charles I. and Peter J. Klenow, “Beyond GDP: Welfare Across Countries and Time,” *American Economic Review*, 2016, *106* (9), 2426–2457.

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.

Peters, Michael, "Heterogeneous Markups, Growth, and Endogenous Misallocation," *Econometrica*, 2020, *88* (5), 2037–2073.

Prendergast, Canice, "The Motivation and Bias of Bureaucrats," *American Economic Review*, 2007, *97* (1), 180–196.

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 five 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 and numerical solution details. Appendix E presents counterfactual diagnostics, robustness exercises, and sector-level direction-cost results.

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 A3 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 $\text{arcsinh}(\text{contracts}_{s,t})$, 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 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 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 A1: Summary Statistics: MIP–Procurement Merged Panel, 2000–2023 (regression sample: 2009–2023)

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 turnover)	108,632	9.839	20.77
Unchanged-product revenue (% of turnover)	72,999	84.49	27.77
Cost-reduction revenue (% of turnover)	88,431	2.503	6.642
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.

A.3 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 precisely this pattern. If procurement shifts firms from frontier novelty toward market-known improvements (leaving total product-innovation activity approximately unchanged), then a measure that aggregates both types should be near zero. Table A3 confirms the prediction: market novelty falls with procurement exposure, while the Krieger et al.

Table A2: 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^{***}
Unchanged-product revenue (% turnover)	90.00	89.01	-0.99^{***}
Cost-reduction revenue (% turnover)	1.79	2.12	0.33^{***}
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 $\text{arcsinh}(\text{contracts}_{s,t})$ 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. $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

aggregate indicator is near zero. The near-zero aggregate is a consequence of the composition shift, not evidence against it.

Table A3: Variable Definition Comparison: This Paper vs. Krieger, Pruefer, and Strecke (2024)

	This paper		KPS (2024) comparable	
	(1) Mkt. novelty	(2) Existing-prod. rev. (%)	(3) Prod. innov.	(4) New-prod. rev. (%)
arcsinh(proc. lots, MIP sector)	-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
MIP-sector clusters	21	21	21	21

Notes: Firm and year fixed effects. Standard errors are clustered by MIP sector (21 clusters). Columns (1)–(2) use this paper's market-novelty outcomes. Columns (3)–(4) use the broader definitions in Krieger et al. (2024): product innovation includes any new or improved product, and new-product revenue is its turnover share. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.4 Procurement Composition by Contract Value

Figure A1 replicates Figure 3 using contract value rather than contract count. Works contracts account for a larger value share (around 30–50 percent in most years) than count share (7 percent), reflecting the high unit value of construction and infrastructure awards. The composition by buyer category is similar across both measures. The count-based decomposition in the main text is more directly relevant to the treatment variable, which is the arcsinh of contract counts.

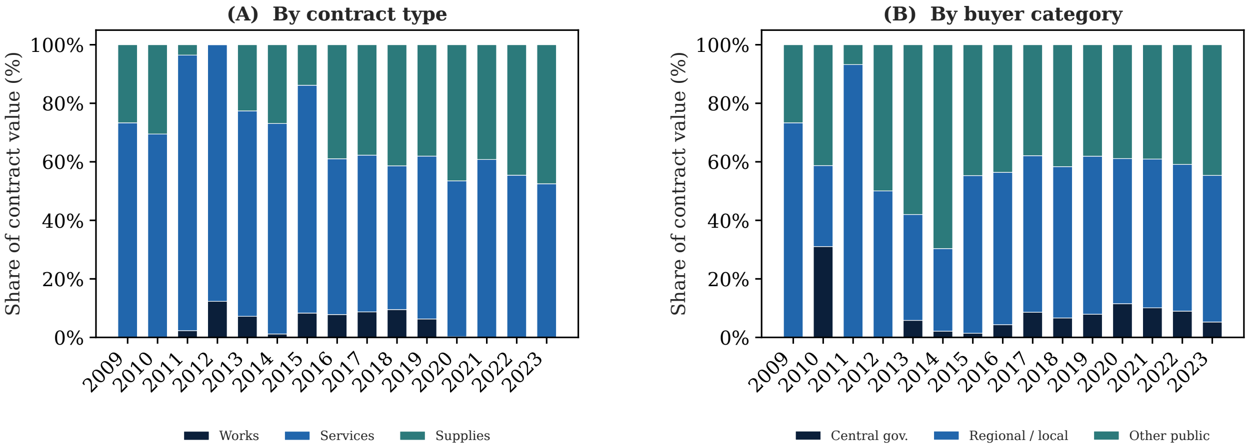

Figure A1: Procurement composition in MIP-covered industries by contract value, 2009–2023

Notes: Share of total above-threshold contract value by type (Panel A: works, services, supplies) and by buyer category (Panel B: central government, regional and local government, other public bodies). Restricted to German industry sectors covered by the Mannheim Innovation Panel. Excludes the top 0.1 percent of contracts by value to remove data-entry outliers.

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

A.5 Procurement Data Coverage, Sector Classification, and Reporting Robustness

German procurement agencies publish only a fraction of their contracts on TED. Published contract value amounts to approximately 1.8 percent of GDP, against an EU average of 5.8 percent (European Commission, 2024), with the gap attributed to fragmented state-level procurement systems and uneven compliance with above-threshold reporting requirements (European Commission, 2019). In April 2016, Germany tightened its award-notice rules and expanded electronic reporting as part of a broader EU directive transposition; TED contract counts more than doubled in the years that followed, a jump far larger than the moderate rise in public spending recorded in national accounts. Pre-2016 exposure is therefore a lower bound on true sector procurement activity rather than a representative measure.

Figure A2 decomposes above-threshold German procurement contracts into three categories for each year 2009–2023. Project and investment contracts account for 69 percent of contracts and 80 percent of contract value over the post-2016 window. Staple and routine contracts account for 24 percent of contracts and 16 percent of value. The composition is stable across the post-2016 years; category shares are unaffected by the April 2016 reporting reform.

Table A4 tests whether the jump in TED contract counts after April 2016 reflects reporting compliance rather than actual procurement growth. Adding sector-specific linear time trends leaves both

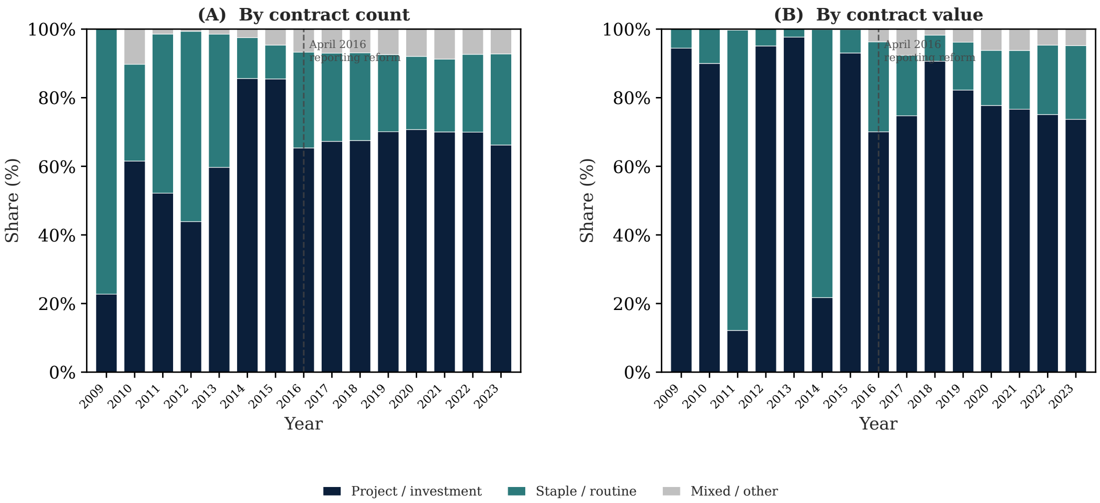

Figure A2: Contract category composition of above-threshold German procurement, 2009–2023
Notes: Share of contracts by count (A) and by value (B), grouped by contract category. Project / investment: construction, engineering, specialized equipment, IT, and R&D services. Staple / routine: food supply, printing, maintenance, financial, and social services. Dashed line: April 2016 EU-mandated reporting reform.
Source: Authors' calculations using TED, 2009–2023.

the market-novelty and unchanged-revenue estimates close to their baseline values (specification B: -0.299 pp, $p = 0.052$, and $+0.366$ pp, $p = 0.027$, respectively). Restricting to 2017–2023, where compliance is more uniform, produces wider confidence intervals rather than sign reversals; the shorter window reduces identifying variation rather than uncovering a different pattern. The main results are not an artifact of the reporting shift.

Table A4: Robustness to TED Coverage Break

	(A) Post-2016 only		(B) Full + sector trends		(C) Post-2016 + sector trends	
	(1) Mkt. nov.	(2) Unchanged rev.	(3) Mkt. nov.	(4) Unchanged rev.	(5) Mkt. nov.	(6) Unchanged rev.
arcsinh(proc. lots, MIP sector)	0.037 (0.207)	0.313 (0.340)	-0.299* (0.144)	0.366** (0.153)	-0.369 (0.989)	1.114 (0.993)
Observations	27,772	8,118	66,980	25,962	27,772	8,118
Clusters	21	21	21	21	21	21
R^2 (within)	0.000	0.000	0.000	0.001	0.000	0.000
Sample	2017–2023	2017–2023	2009–2023	2009–2023	2017–2023	2017–2023
Sector linear trends	No	No	Yes	Yes	Yes	Yes

Firm and year FEs throughout. Specifications (B)–(C) additionally absorb sector-specific linear time trends (sector × year). SE clustered by MIP sector (21 clusters). Outcome (1): market novelty (binary ×100 = pp). Outcome (2): share of turnover from unchanged products (percent). Post-2016 window = 2017–2023, the period of reliable TED reporting compliance. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The treatment variable assigns each procurement contract to an industry sector by mapping its EU product code to the corresponding German industry sector. A natural validation asks whether firms that win contracts in a given product category actually operate in the industry sector to which the mapping assigns them. We assess this by linking each contract winner in the TED records to the firm's primary industry classification in ORBIS, using the matching described in Appendix B. Because German and European industry classification codes share the same letter-based sector scheme, the comparison is direct.

The overall count-weighted agreement rate is 54.4 percent and the value-weighted rate is 66.0 percent. Agreement is highest for product categories that map closely to a single industry: petroleum and gas services (87.8%), transport services (86.1%), and architectural and engineering services (84.3%). Agreement is lowest for product categories whose contracts attract buyers and winners from many industries: defense and public-administration services, IT office equipment (won by hardware resellers whose primary industry is often wholesale trade rather than manufacturing), and landscape and environmental services. These low-agreement categories are precisely those for which the sector assignment is most approximate; the treatment-robustness table excludes them as a check (Table C2).

Within the post-2016 window the agreement rate is stable at 53–57 percent per year, with higher agreement by contract value (60–83 percent) because large contracts tend to go to firms in the sector the contract was written for. Table A5 summarizes the agreement rates, and Figure A3 shows the year-by-year pattern.

Table A5: CPV-to-WZ Crosswalk Validation

	Agreement rate	Observations
Panel A: Overall		
Count-weighted	54.4%	145,952
Value-weighted	66.0%	132,810
Panel B: By supply type		
Services	52.1%	49,273
Works	61.0%	49,240
Supplies	49.7%	46,741
Panel C: Pre- vs. post-2016 reform		
2009–2015	61.3%	419
2016–2023	54.4%	145,533
Difference ($p = 0.004$)	-7.0%	

Notes: Agreement is the share of contracts where the CPV-mapped WZ section of the contract matches the NACE section of the winning firm as recorded in ORBIS. The sample includes contracts with at least one winning bidder matched to ORBIS. WZ section codes correspond directly to NACE Rev. 2 section letters. Pre-2016 agreement is lower because compliance with TED electronic reporting was incomplete; the post-2016 subsample has both higher coverage and higher within-sector agreement. Sources: TED award notices and ORBIS financial data.

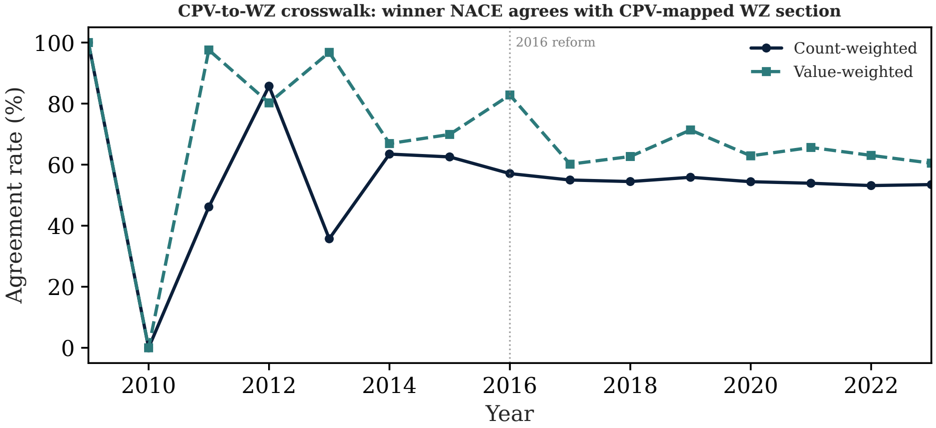

Figure A3: Sector assignment accuracy by year

Notes: Agreement is the share of matched contract rows where the industry sector assigned by the EU product-code mapping matches the primary industry classification of the winning firm as recorded in ORBIS. Count-weighted (solid) and value-weighted (dashed) rates are shown. The vertical dotted line marks April 2016 (EU Procurement Directive transposition). Pre-2016 observations are fewer than 500 rows due to low TED reporting compliance before the reform; the stable post-2016 series is the informative window. Sample: TED 2009–2023, ORBIS-matched winning bidders.
Source: Authors' calculations using TED and ORBIS.

B ORBIS Validation

The MIP results describe the innovation-composition margin. ORBIS provides evidence on the scale margin among matched procurement winners. Figure B1 describes the profile of winners; the regression tables that follow show that procurement exposure is associated with higher employment and turnover, while profitability effects remain small. The full-sample estimates appear in Table 3 in the main text.

B.1 Winner Profile and Scale Dynamics

Figure B1 describes the size distribution and sector composition of procurement winners in the matched ORBIS panel. Winners skew toward small and medium-sized enterprises: 47 percent have 10–49 employees and 16 percent are micro firms; large firms (250 or more employees) account for only 10 percent of winners. Construction and manufacturing together account for 44 percent of contracts won; professional services, wholesale trade, and administrative-support sectors account for a further 35 percent.

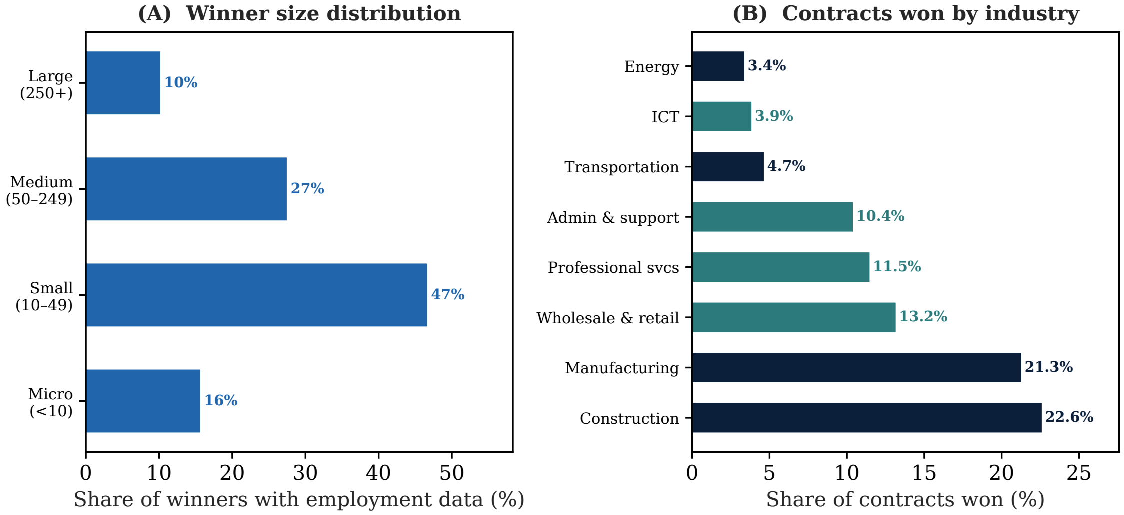

Figure B1: Size distribution and sector composition of procurement winners

Notes: Panel A reports the size distribution of winning firms by employment class among firms with employment data at first win; the median winner has 33 employees. Panel B reports the share of contracts won by the eight largest industry sectors (accounting for 91 percent of all contracts); navy bars indicate production-intensive sectors (construction, manufacturing, energy, transportation), teal bars indicate knowledge and service sectors. Sample: ORBIS-matched winner panel, TED 2009–2023.
Source: Authors' calculations using TED matched to ORBIS.

Table B1 splits winner-panel estimates by firm age at first win. Age is measured as years since the firm first appears in ORBIS. Young entrants (five years or less) and established entrants show similar employment and turnover responses to procurement exposure; return on capital is unchanged in both groups, consistent with scale expansion rather than a profitability gain for either cohort.

Table B1: Scale Outcomes by Firm Age at First Win

	All winners			**Young entrants ($\leq$5 yr)**			**Established entrants ($>$5 yr)**		
	Log emp.	Log turn.	Ret. on capital	Log emp.	Log turn.	Ret. on capital	Log emp.	Log turn.	Ret. on capital
arcsinh(contracts won)	0.025^{***}	0.008	-0.145	0.027^{***}	0.011	0.090	0.019	0.005	-0.761
	(0.007)	(0.019)	(0.386)	(0.007)	(0.017)	(0.296)	(0.019)	(0.073)	(1.327)
Observations	28,746	12,773	8,538	20,461	9,470	5,986	8,285	3,303	2,552

Notes: ORBIS-matched winner panel, 2009–2023. Firm age at first win is proxied by years since the firm first appears in the ORBIS financial database. Young entrants have five or fewer years in ORBIS before their first win; established entrants have more than five. All specifications include firm and year fixed effects. Standard errors, in parentheses, are clustered by firm. Return on capital is profit before tax divided by total assets, in percentage points. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.2 Scale, Profitability, and Contractor Heterogeneity

Table B2 reports the matched-winner estimates for the full sample and for small, medium, and large firms. In the full matched sample, a one-unit increase in $\mathrm{arcsinh}$(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) growth, and profit margins show no statistically significant association, consistent with scale expansion rather than a profitability gain.

Table B2: Procurement and Firm Outcomes: Full Sample and by Firm Size

	Full (1)	Small (< 50) (2)	Medium (50–249) (3)	Large (≥ 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: $\mathrm{arcsinh}$(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. $^{*}p{<}0.1$, $^{**}p{<}0.05$, $^{***}p{<}0.01$.

Table B3 examines ownership heterogeneity. Multinational enterprises (MNEs) have larger employment responses than domestic firms (0.049 versus 0.025 log points). Return on capital is not statistically significant in any ownership subgroup, consistent with capacity growth rather than a rent channel.

Table B3: 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 $\text{arcsinh}($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.

Table B4 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 statistically imprecise among repeat contractors. Return on capital remains unchanged in both groups, and revenue labor productivity is not higher among repeat contractors.

Table B4: 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 $\text{arcsinh}(\text{contracts won})$, 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.

C Mechanism Robustness

This appendix presents three supplementary analyses. The first compares procurement exposure to public R&D subsidies, showing that the composition pattern is specific to procurement rather than common to government funding broadly. The second rules out a denominator artifact and shows the pattern is robust to alternative exposure definitions and common-sample restrictions. The third places the within-firm estimates in aggregate context, examining how much of Germany's market-novelty decline the estimated associations can account for.

C.1 Procurement versus Public R&D Subsidies

Table C1 compares two public-funding instruments 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 exposure is associated with higher existing-product revenue and lower market novelty, consistent with a within-firm composition shift. Public R&D subsidies show the opposite pattern: they are associated with higher market novelty and lower existing-product revenue, consistent with an overall innovation stimulus. The two instruments are associated with opposite patterns of innovation composition.

Table C1: Procurement vs. Public Subsidies: Two Channels

	(1) Market novelty (pp)	(2) Market-novel share (% of new products)	(3) Existing-product revenue (pp)
Panel A. Procurement exposure, 2009–2023			
$\operatorname{arcsinh}$(contracts)	-0.199^{*}	-0.099	0.379^{**}
	(0.101)	(0.198)	(0.137)
N	66,980	15,830	25,962
Panel B. 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 is asked. 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.2 Sample and Exposure Robustness

Ruling out a denominator artifact. One concern is that procurement exposure coincides with higher total sales, which would mechanically raise the existing-product revenue share even if the level of new-product activity were unchanged. Two additional revenue-share outcomes speak against this interpretation. The share of sales from market-novel products falls with exposure, and the share from any newly introduced product also falls directionally. If the pattern were a pure denominator artifact, these new-product shares would remain flat; the fact that all three revenue-mix margins move in the same direction is consistent with a genuine reallocation.

Alternative exposure definitions. Table C2 checks whether the results change when procurement exposure is measured differently. Replacing contract counts with total awarded value in euros produces smaller, imprecise estimates, consistent with contract value being measured with more error in TED records than counts. Scaling exposure by the number of MIP-surveyed firms in each sector-year makes the market-novelty estimate more negative while the unchanged-revenue estimate remains significant. Restricting to project and investment contract categories—which account for 69 percent of contracts and 80 percent of contract value post-2016—gives estimates that are more negative for market novelty and more positive for unchanged revenue, suggesting the baseline understates the association in the most economically meaningful contract types.

Table C2: Robustness: Alternative Treatment Variable Scalings

	(1) arcsinh(euro value)		(2) arcsinh(contracts / MIP firms)		(3) arcsinh(investment lots only)	
	(1) Mkt. nov.	(2) Unchanged rev.	(3) Mkt. nov.	(4) Unchanged rev.	(5) Mkt. nov.	(6) Unchanged rev.
arcsinh(euro value)	-0.039	0.055				
	(0.036)	(0.041)				
arcsinh(contracts / MIP firms)			-0.533	0.849**		
			(0.311)	(0.387)		
arcsinh(investment lots)					-0.282*	0.546***
					(0.149)	(0.133)
Obs.	66,980	25,962	66,980	25,962	102,654	55,930
Clusters	21	21	21	21	23	23

Firm and year FEs throughout. SE clustered by MIP industry. Binary outcomes ×100 = pp.

Treatment (1): arcsinh of total contract value in euros. Treatment (2): arcsinh of contracts divided by firms in sector-year.

Treatment (3): arcsinh of lots in project and investment categories only (construction, engineering, equipment, IT/R&D; 69% of lots and 80% of value) in 2016–2023). Investment-lots columns use the full MIP panel.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

C.3 Aggregate Market-Novelty Decline: Descriptive Accounting

Table C3 asks how much of the aggregate market-novelty decline the within-firm estimates can account for. Multiplying the full-sample coefficients by the observed 2009–11 to 2021–23 increase in exposure predicts a decline of 0.98–1.47 percentage points, or 5.4–8.1 percent of the 18 percentage-point national decline. This is not a historical procurement contribution because the exposure increase includes the TED reporting expansion; in the higher-compliance period, mean exposure is essentially unchanged between 2017–19 and 2021–23 and the predicted contribution is approximately zero. Procurement exposure is one correlate of the composition trend; the estimates do not pin down its contribution to Germany's historical novelty decline.

Table C3: Partial accounting of the market-novelty decline

Specification	Window	$\hat{\beta}$	Δ exposure	Predicted Δ novelty (pp)	Share of 18 pp (%)
Firm and year FE	2009–11 to 2021–23	-0.199	4.909	-0.977	5.4
Add sector trends	2009–11 to 2021–23	-0.299	4.909	-1.466	8.1
Post-2016 FE	2017–19 to 2021–23	0.037	-0.009	-0.000	0.0
Post-2016 plus trends	2017–19 to 2021–23	-0.369	-0.009	0.003	-0.0

Notes: The predicted change is $\hat{\beta}$ times the firm-observation-weighted change in sector-year arcsinh lot exposure. The final column divides that change by the 18 percentage-point national decline shown in Figure 1. The corresponding change in the MIP estimation sample is -9.59 pp over 2009–11 to 2021–23 and -2.29 pp over 2017–19 to 2021–23. The full-period exposure change includes the expansion of TED reporting coverage and cannot be interpreted as a change in actual procurement alone. Post-2016 coefficients are imprecise. This is descriptive partial accounting, not a causal decomposition.
Source: Authors' calculations using the Mannheim Innovation Panel and TED.

D Technical Derivations for the Klette–Kortum Model

The appendix collects the technical material behind the model: 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 Fixed-Point Reduction

The incumbent block is governed by

$$(r+\tau)v_I = \pi_I(G) - \Phi(x) - \Psi(\mu) + x\bar{v}(\mu;G), \tag{36}$$

$$(r+\tau)v_F = \pi_F(G) - \Phi(x) - \Psi(\mu) + x\bar{v}(\mu;G), \tag{37}$$

where $\bar{v}(\mu;G) \equiv (1-\mu)v_I + \mu v_F$ and $r+\tau-x>0$ ensures finite values. The value gap $v_F - v_I = [\bar{\pi}_F - \bar{\pi}_I(1+G)]/(r+\tau)$ follows from subtracting these two equations; its derivation appears in the main text. Defining

$$v_F - v_I = \frac{\bar{\pi}_F - \bar{\pi}_I(1+G)}{r+\tau}, \tag{38}$$

premultiplying (36) by $(1-\mu)$ and (37) by μ, then adding, gives the average line value:

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

where

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

The effort and direction first-order conditions are

$$c_x \zeta x^{\zeta-1} w = \bar{v}(\mu;G), \tag{40}$$

$$\kappa\mu^{\eta-1} = x(v_F - v_I). \tag{41}$$

Using (38), the interior direction condition becomes

$$\mu^{\text{int}} = \left[\frac{x}{\kappa}\frac{\bar{\pi}_F - \bar{\pi}_I(1+G)}{r+\tau}\right]^{1/(\eta-1)}. \tag{42}$$

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

$$\mu^* = \min \left\{1, \max \left\{0, \mu^{\text{int}}\right\}\right\}. \tag{43}$$

Equivalently, for $D(\mu) = x(v_F - v_I) - \kappa\mu^{\eta-1}$, the KKT conditions are $D(\mu^*) = 0$ when $0 < \mu^* < 1$, $D(0) \leq 0$ when $\mu^* = 0$, and $D(1) \geq 0$ when $\mu^* = 1$. For $\eta > 1$, the two boundary conditions are $x(v_F - v_I) \leq 0$ and $x(v_F - v_I) \geq \kappa$, respectively. Counterfactuals apply these inequalities through the projection in (43).

Equations (39), (40), and (43) 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 determines innovation effort, and the third determines the frontier share.

D.2 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

$$F(\mathbf{z}; G, \boldsymbol{\psi}) = 0, \tag{44}$$

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

$$F_1(\mathbf{z}; \cdot) \equiv c_x \zeta x^{\zeta-1} w - \bar{v}(\mathbf{z}; G) = 0, \tag{45}$$

$$F_2(\mathbf{z}; \cdot) \equiv \kappa\mu^{\eta-1} - \frac{x[\bar{\pi}_F - \bar{\pi}_I(1 + G)]}{r + \tau} = 0. \tag{46}$$

At the boundaries $\mu = 0$ or $\mu = 1$, the algorithm applies the projection in (43).

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.

Optimization and Inference. The GMM objective is minimized by a multistart global search. At each start, we draw candidate parameter vectors uniformly from the admissible region, evaluate $Q(\boldsymbol{\psi})$ at each candidate, and retain the best. The search spans tens of thousands of function evaluations across multiple random seeds, avoiding reliance on gradient information unreliable given the discrete choice in μ^*.

Inference uses the delta method. Let $\hat{\mathbf{J}} = \partial \mathbf{m}(\boldsymbol{\psi})/\partial \boldsymbol{\psi}^\top \big|_{\hat{\boldsymbol{\psi}}}$ be the Jacobian of model moments at the estimates, computed by symmetric finite differences. The asymptotic covariance of the parameter vector is

$$\mathrm{Var}(\hat{\boldsymbol{\psi}}) = \left(\hat{\mathbf{J}}^\top W^* \hat{\mathbf{J}}\right)^{-1} \hat{\mathbf{J}}^\top W^* \hat{\Sigma} W^* \hat{\mathbf{J}} \left(\hat{\mathbf{J}}^\top W^* \hat{\mathbf{J}}\right)^{-1}, \tag{47}$$

where $\hat{\Sigma}$ is the empirical moment covariance and $W^* = \hat{\Sigma}^{-1}$ is the efficient weighting matrix estimated via influence functions, with MIP moments clustered at the firm level and the creative-destruction moment using the cross-year standard deviation from BHP. Standard errors are square roots of the diagonal elements. Appendix E.5 replaces the influence-function covariance with a firm-block bootstrap over 399 replications; parameter uncertainty changes little.

D.3 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 $\bar{G} = 0.018$. 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.

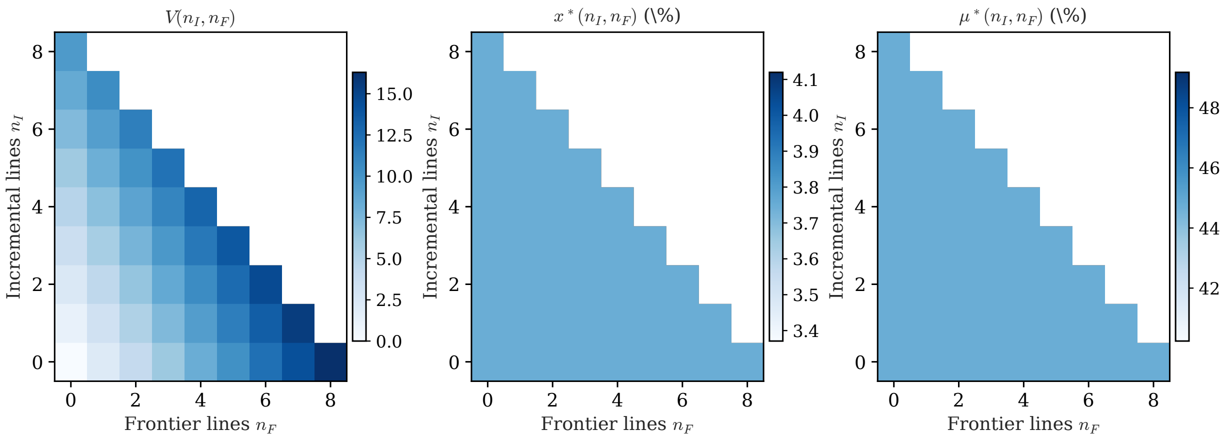

Figure D1: Benchmark value and policy functions in the firm state

Notes: The left panel reports firm value $V(n_I, n_F)$ at the benchmark $\bar{G} = 0.018$. 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.4 Entry, Creative Destruction, and Stationarity

The baseline has two sources of ownership changes but only one source of quality growth. Incumbent innovation improves a rival line and transfers it to the innovator at rate x^* per existing line. Entrants use setup labor to acquire existing lines and create new firms with one line, without generating an additional quality step. Existing lines change owner at total hazard τ. With a unit mass of product lines, stationarity requires

$$x^* + e = \tau, \tag{48}$$

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

$$e = \tau - x^*. \tag{49}$$

The entry-labor requirement is

$$L_E = h_E e = h_E(\tau - x^*), \tag{50}$$

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

Free entry determines the expected value of an entrant line:

$$\bar{v}(\mu^*; G) = w h_E. \tag{51}$$

Here h_E is the fixed setup-labor requirement per entrant and wh_E is the corresponding monetary cost. Labor-market clearing uses production labor, incumbent R&D labor, and entrant labor:

$$\bar{L} = L_P(\mu, G, w) + \ell_R(x, \mu, w) + h_E(\tau - x). \tag{52}$$

D.5 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

$$0 = x^*(n-1)m_{n-1} + \tau(n+1)m_{n+1} - (x^* + \tau)nm_n. \tag{53}$$

The solution is

$$m_n = \frac{e}{x^*}\frac{1}{n}\left(\frac{x^*}{\tau}\right)^n, \qquad n \geq 1. \tag{54}$$

Let

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

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

$$\Pr(N = n) = \frac{q^n}{n[-\log(1-q)]}, \qquad n \geq 1. \tag{55}$$

Equation (55) 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^*:

$$\Pr(n_F = b \mid N = n) = \binom{n}{b}(\mu^*)^b(1-\mu^*)^{n-b}. \tag{56}$$

The joint distribution is

$$\Gamma(n_I, n_F; G) = \Pr(N = n_I + n_F)\binom{n_I + n_F}{n_F}(\mu^*)^{n_F}(1-\mu^*)^{n_I}. \tag{57}$$

The product-innovation probability over a horizon T is

$$p_{\text{innov}} = 1 - \sum_{n=1}^{\infty} \Pr(N = n) e^{-Tx^* n}. \tag{58}$$

Using (55),

$$p_{\text{innov}} = 1 - \frac{-\log\left(1 - qe^{-Tx^*}\right)}{-\log(1-q)}. \tag{59}$$

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 $\operatorname{Var}(s_F) = \mu^*(1-\mu^*)\mathbb{E}[1/N]$. Under (55),

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

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 (55) and recognizing 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.

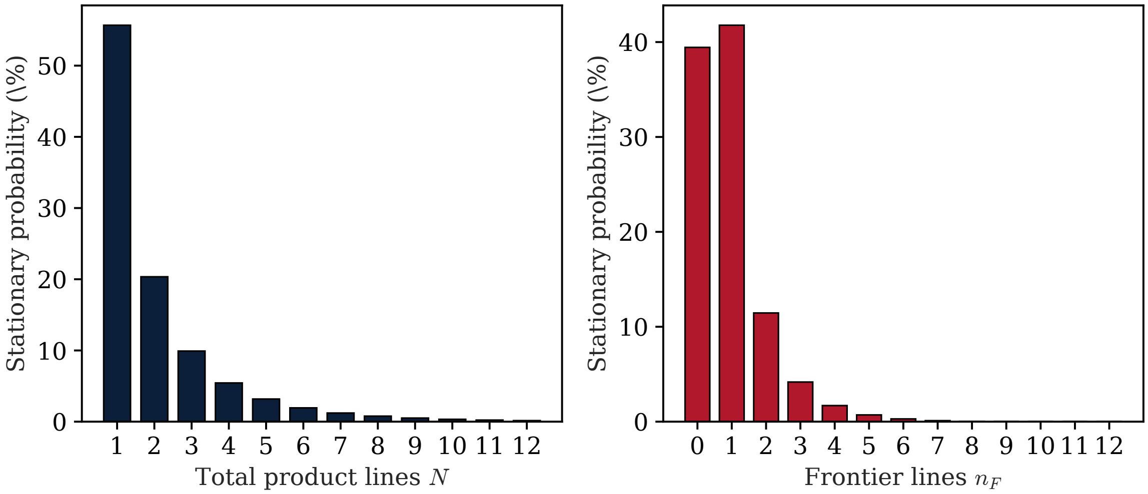

Figure D2: Model-implied stationary distributions at the GMM benchmark

Notes: Panel A reports the firm-scope distribution in (55). Panel B reports the marginal frontier-line distribution implied by binomial mixing under the benchmark equilibrium.
Source: Authors' calculations from the calibrated model.

E Quantitative Analysis

The appendix collects the quantitative exercises supporting the model and data mapping.

E.1 Identification and Model Fit

E.1.1 Detailed Measurement Mapping and External Validation

Table E1 reports the corresponding external-validation exercise. Because the empirical and model treatments are measured in different units, the table does not treat the magnitudes as a formal fit criterion. It asks whether a transparent increase in model procurement reproduces the joint empirical signs for innovation composition and incumbent scale. To make the unit conversion explicit, Panel B derives the model derivative of the market-novelty counterpart by the numerical product of market novelty to procurement times a proportional mapping from sector-year contract counts to sector procurement shares.

Table E1: External validation: empirical and model responses to procurement

Panel A. Outcome	Data effect	Model effect	Sign match	Role
Market novelty	-0.709 pp	-0.024 pp	Yes	Composition validation
Unchanged-product revenue	$+1.352$ pp	$+0.139$ pp	Yes	Composition validation
Employment / firm scope	$+2.300$ percent	$+0.339$ percent	Yes	Scale validation
Panel B.	Data value	Model value	Sign match	Unit
Procurement coefficient	$\hat{\beta} = -0.199$	-0.097	Yes	pp / arcsinh unit

Notes: Data effects correspond to a one-standard-deviation increase in the relevant procurement exposure. Model effects compare the benchmark $\bar{G} = 0.01783$ with a 25% increase, $G = 1.25\bar{G}$, holding structural parameters and the recovered innovation-cost scale fixed while re-solving general equilibrium. We compute the model coefficient from the change of market novelty to procurement exposure and a mapping between sector-year procurement shares and contract counts.

To make the unit conversion explicit, Table E2 combines the model derivative of the market-novelty counterpart with a proportional mapping from sector-year lot counts to sector procurement shares.

Table E2: Model-implied procurement coefficient

Object	Value	Unit
$\partial \mathrm{MarketNovelty}/\partial G$	-5.411	pp / unit-G
$\partial G/\partial \operatorname{arcsinh}(\ell)$	0.01791	unit-G / arcsinh unit
$\hat{\beta}_{\mathrm{model}}$	-0.097	pp / arcsinh unit
$\hat{\beta}_{\mathrm{empirical}}$	-0.199	pp / arcsinh unit

*Notes:*The first derivative is computed by symmetric finite differences at $\bar{G} = 0.01783$, holding other parametes fixed. $\partial G/\partial \operatorname{arcsinh}(\ell)$ assumes that sector-year procurement shares are proportional to sector-year contract counts. The model-implied coefficient is the product of these two components.

E.2 Moment Sensitivity

Figure E1 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.

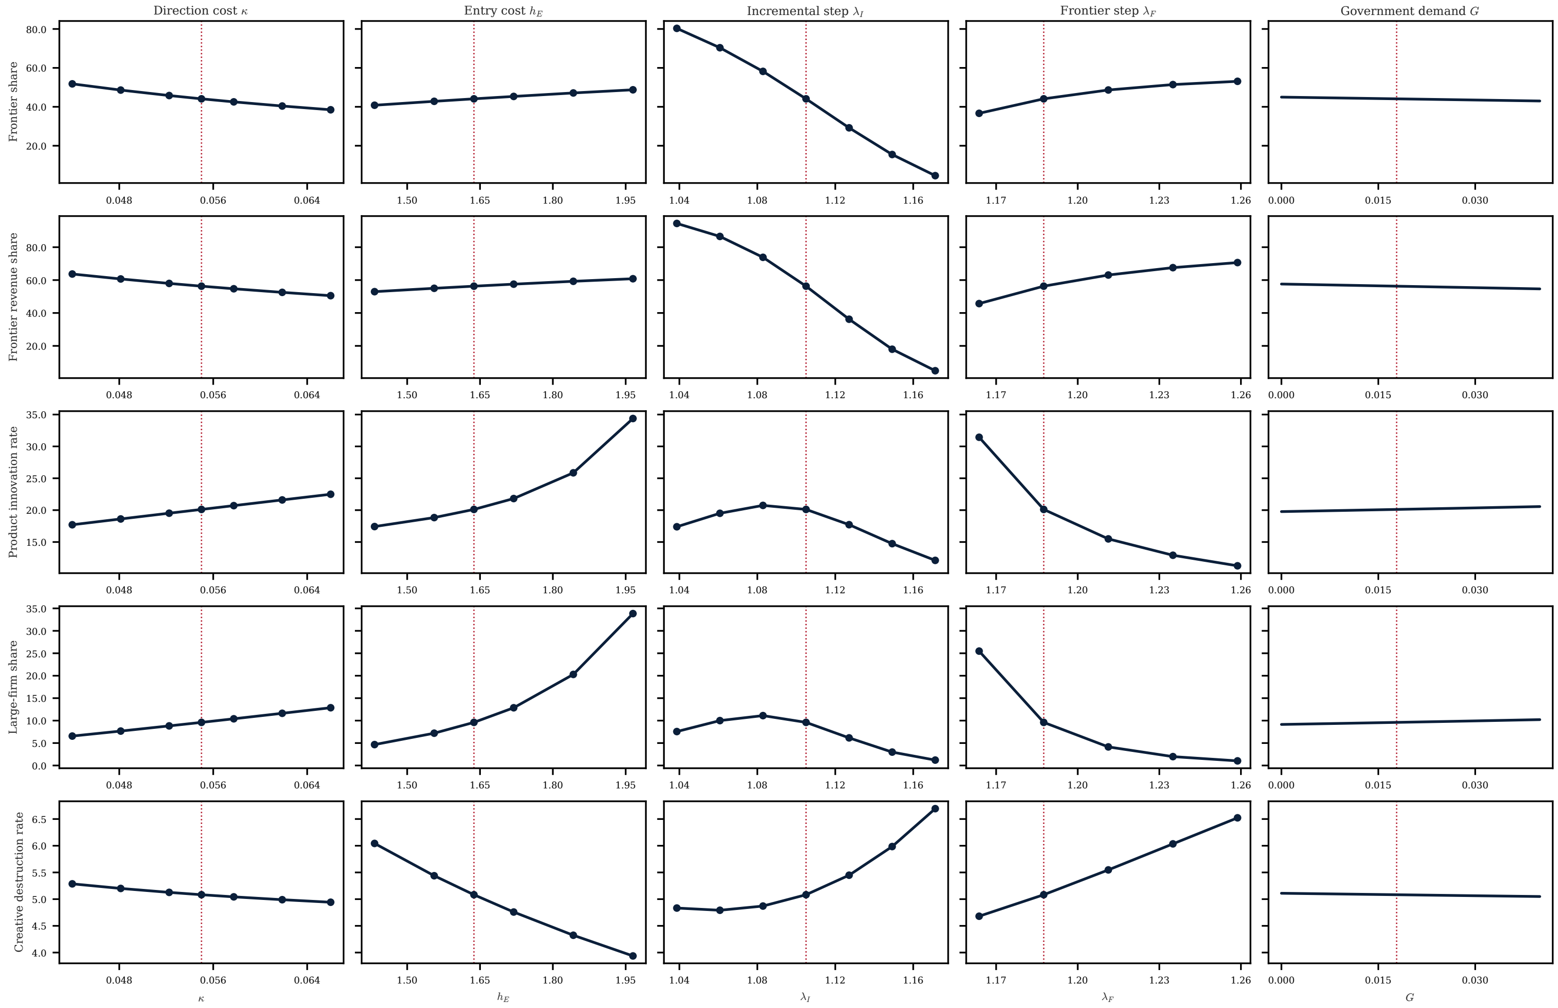

Figure E1: Moment comparative statics: identification

Notes: Rows are targeted moments; columns are parameters. For each structural parameter (κ, h_E, λ_I, λ_F), the curve evaluates up to 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 $G = \bar{G} = 0.018$; dots mark converged equilibria. The G column (right) traces $G \in [0, 0.04]$ at benchmark structural parameters. The dotted red vertical line marks the calibrated benchmark of each parameter.
Source: Authors' calculations from the calibrated model.

E.3 Specification Design Reform: Formal Derivations

The specification design reform reallocates a fixed procurement budget T_{bm} across line types. Let $\alpha \in [0,1]$ denote the share of the budget assigned to frontier lines. Given the equilibrium frontier mass μ, the per-line demand intensities are

$$G_I(\mu,\alpha) = \frac{(1-\alpha)T_{bm}}{1-\mu}, \qquad G_F(\mu,\alpha) = \frac{\alpha T_{bm}}{\mu}, \tag{60}$$

so $(1-\mu)G_I + \mu G_F = T_{bm}$ in every counterfactual equilibrium. The modified profit expressions are

$$\pi_I(\alpha) = \bar{\pi}_I\Big[1 + G_I(\mu,\alpha)\Big], \tag{61}$$

$$\pi_F(\alpha) = \bar{\pi}_F\Big[1 + G_F(\mu,\alpha)\Big]. \tag{62}$$

At $\alpha = 0$ all demand flows to incremental lines. At the benchmark, $T_{bm} = (1-\mu_{bm})\bar{G}$, so $G_I = \bar{G}$ and $G_F = 0$. The value gap under the reform is

$$v_F(\alpha) - v_I(\alpha) = \frac{\bar{\pi}_F[1 + G_F(\mu,\alpha)] - \bar{\pi}_I[1 + G_I(\mu,\alpha)]}{r+\tau}. \tag{63}$$

The direction condition gives

$$\mu^*(\alpha) = \frac{x^*}{\kappa} \cdot \frac{\bar{\pi}_F[1 + G_F(\mu^*,\alpha)] - \bar{\pi}_I[1 + G_I(\mu^*,\alpha)]}{r+\tau}. \tag{64}$$

Proposition E.1 (Budget-neutral specification realignment)**.** *Suppose the equilibrium is interior. Reallocating a marginal unit of the fixed procurement budget from incremental to frontier lines raises the current profit gap by*

$$\frac{\partial(\pi_F - \pi_I)}{\partial\alpha} = T_{bm}\left(\frac{\bar{\pi}_F}{\mu} + \frac{\bar{\pi}_I}{1-\mu}\right) > 0.$$

Holding general-equilibrium prices fixed, this raises frontier orientation. Quantitatively, the model jointly resolves the induced change in μ, wages, and creative destruction while imposing equation (60)*.*

Proof. Raising α by a unit shifts government demand away from incremental lines and toward frontier lines. From (61)–(62), 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 (63), so the direction first-order condition moves μ^* toward frontier orientation. Formally, differentiating (64) 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. □

Table E3 summarizes observable contract-design frequencies in the post-coverage-break sample. Roughly one-third of covered lots contain technical-requirement text above the empirical long-specification threshold. These frequencies establish that design reform concerns a quantitatively relevant set of contracts.

Table E3: Observed procurement-design frequencies, 2017–2023

Contract-design feature	Share of covered lots
Technical-requirement text observed	0.648
Technical-requirement text above 912 characters	0.316
Long text, conditional on positive text length	0.487
Unique tender–lot records	482,840

Notes: The sample contains winning tender–lot records mapped to MIP-covered sectors. The 912-character threshold is the median used in the empirical heterogeneity analysis; the conditional median in this tender–lot sample is 881 characters because the regression moderator is constructed at the sector–year level. Procedure type and text length are observable design proxies, not measures of the model's frontier budget share α.

E.4 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

$$\kappa_s \mu_s^* = x^*(v_F - v_I) = x^* \frac{\bar{\pi}_F - \bar{\pi}_I(1+G)}{r+\tau}, \tag{65}$$

so the equilibrium frontier share is

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

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 E.2 (Conditional invariance of proportional compression)**.** *Let direction costs be sector-specific, $\kappa_s > 0$, and suppose sectors take the common equilibrium responses $(x(G), r(G), \tau(G))$ as given. For $\eta = 2$, proportional compression satisfies*

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

which is independent of κ_s. This invariance is conditional: it need not survive if sectoral direction costs change aggregate equilibrium responses or if sectors face different wages or destruction hazards.

Proof. Taking logs of (66):

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

Differentiating with respect to G along the common equilibrium path gives

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

which does not depend directly on κ_s under the maintained common-equilibrium-response assumption. ☐

Proposition E.2 isolates a partial-equilibrium comparison across sectors. Conditional on common responses of effort, growth, and creative destruction, sectors with high κ_s have lower frontier shares and therefore smaller absolute changes, while their proportional response is the same. We do not use the conditional result to claim invariance in a multisector general equilibrium.

Empirical range of direction costs. Table E4 recovers $\hat{\kappa}_s$ for each sector and for the 21 MIP industries by inverting (66) 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 0.047 (Manufacturing) to 0.099 (Transport). 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 E2 plots the full subsector distribution.

Table E4: Sector-level direction costs

WZ	Sector	$\bar{\mu}_s$	$\hat{\kappa}_s$ (SE)
G	Trade/Retail	0.513	0.047 (0.001)
C	Manufacturing	0.497	0.049 (0.000)
M	Prof. Services	0.440	0.055 (0.001)
K	Finance	0.374	0.065 (0.002)
D	Energy/Water	0.367	0.066 (0.003)
J	ICT/Media	0.334	0.073 (0.002)
N	Other Services	0.286	0.085 (0.005)
B	Mining	0.274	0.088 (0.008)
H	Transport	0.244	0.099 (0.007)
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 E.2 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 E5: Direction costs at the 21-subsector level

Subsector	$\bar{\mu}_s$	N	$\hat{\kappa}_s$ (SE)
Mining (WZ B)			
Mining	0.274	339	0.088 (0.008)
Manufacturing (WZ C)			
Chemicals & Pharma	0.572	1637	0.042 (0.001)
Glass & Ceramics	0.545	721	0.044 (0.002)
Machinery & Vehicles	0.539	2291	0.045 (0.001)
Electrical Equip.	0.531	3077	0.046 (0.001)
Textiles	0.498	1000	0.049 (0.002)
Plastics	0.494	1103	0.049 (0.001)
Metals	0.472	1862	0.051 (0.001)
Furniture & Med. Tech.	0.460	1401	0.053 (0.002)
Wood & Paper	0.399	1141	0.061 (0.002)
Food & Tobacco	0.381	1202	0.064 (0.002)
Energy/Water (WZ D)			
Energy/Water	0.367	632	0.066 (0.003)
Trade/Retail (WZ G)			
Retail	0.576	1772	0.042 (0.001)
Wholesale	0.323	588	0.075 (0.004)
Transport (WZ H)			
Transport	0.244	704	0.099 (0.007)
ICT/Media (WZ J)			
IT/Telecom	0.348	1570	0.070 (0.002)
Media Services	0.313	1068	0.078 (0.004)
Finance (WZ K)			
Finance	0.374	1413	0.065 (0.002)
Prof. Services (WZ M)			
Prof. Services I	0.480	2044	0.051 (0.001)
Prof. Services II	0.375	1230	0.065 (0.002)
Other Services (WZ N)			
Other Services	0.286	766	0.085 (0.005)
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.

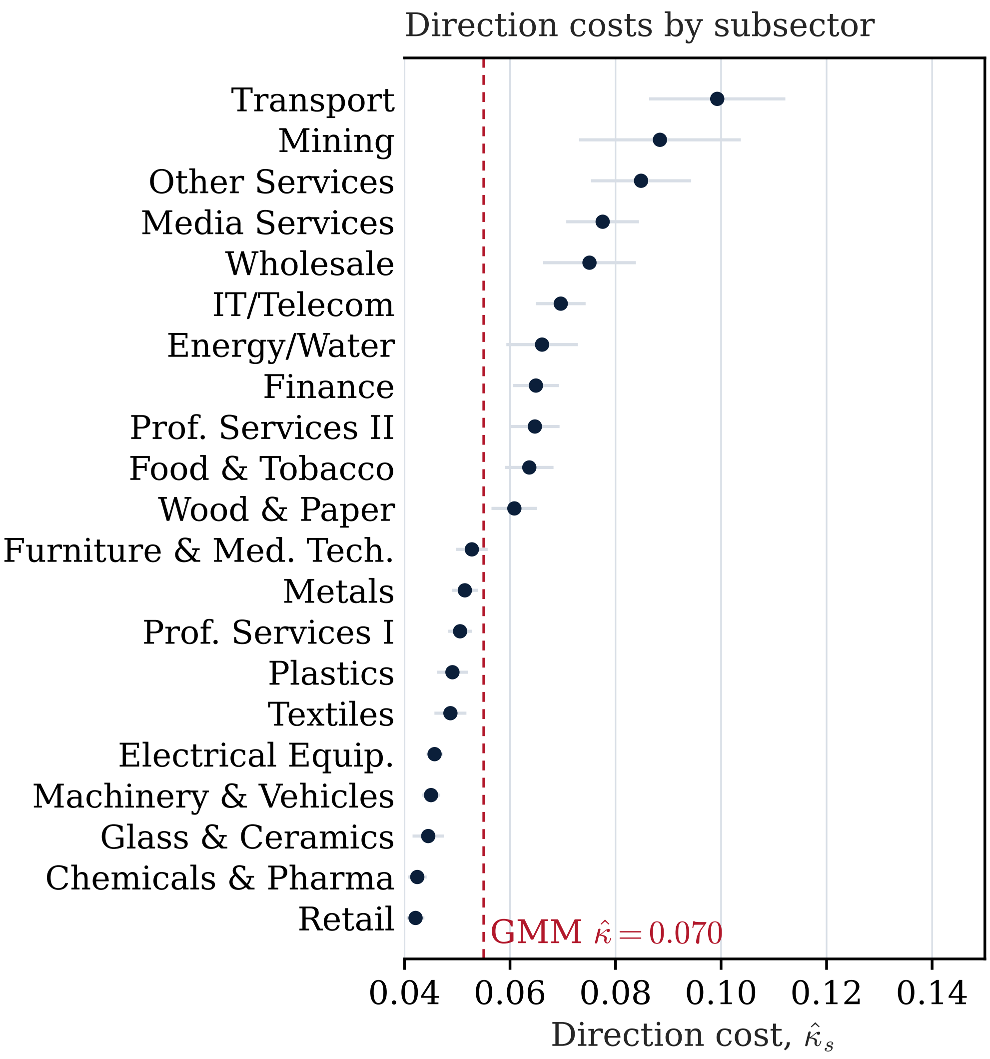

Figure E2: Sector-level direction costs

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 $\hat{\kappa} = 0.055$. Proposition E.2 establishes conditional proportional invariance when sectors share the same equilibrium responses of effort, interest rates, and creative destruction.
Source: Authors' calculations using the Mannheim Innovation Panel, 2003–2018.

E.4.1 Sector Exposure Heterogeneity

The composition wedge scales with the government demand parameter G. We map mean sector-year reported contract counts to model units proportionally, normalizing the mean across sectors with positive exposure to the benchmark $\bar{G} = 0.018$. Manufacturing and professional services receive the largest mapped wedges: without procurement, their frontier-innovation shares are 1.44 and 1.35 percentage points higher, respectively. Transport and IT services lie in the middle range (1.16 and 1.11 pp), while finance and mining have smaller gains (0.67 and 0.48 pp). The corresponding annual growth gains range from 0.13 to 0.39 basis points. The pattern is the composition-channel counterpart to Proposition E.1: sectors assigned larger G_s face larger composition wedges.

Table E6: Frontier-Innovation Distortion by Sector

Sector	$\bar{G}_s$	$\hat{\mu}_s^*$	$\Delta\mu^*$ (pp)	Δg (bp)
Manufacturing	0.025	0.437	$+1.4$	$+0.4$
Professional services	0.023	0.438	$+1.3$	$+0.4$
Transport	0.020	0.440	$+1.2$	$+0.3$
IT and information	0.019	0.440	$+1.1$	$+0.3$
Administrative services	0.017	0.441	$+1.0$	$+0.3$
Energy	0.017	0.441	$+1.0$	$+0.3$
Finance	0.012	0.445	$+0.7$	$+0.2$
Mining	0.008	0.446	$+0.5$	$+0.1$
Wholesale and retail	0	0.451	0	0
Economy-wide benchmark	0.018	0.441	$+1.0$	$+0.3$

Notes: $\bar{G}_s$ is the sector G equivalent, computed by mapping mean sector-year arcsinh contract counts to model units proportionally: the economy-wide mean corresponds to $\bar{G} = 0.018$. This proportional mapping is a descriptive incidence exercise rather than a sector-specific calibration. $\hat{\mu}_s^*$ is the equilibrium frontier-innovation share at that G, read from the estimated model's CF1 G-sweep. $\Delta\mu^*$ and Δg report the gain in frontier share and annual growth if procurement were removed ($G_s \to 0$), holding the R&D tax credit fixed. The R&D credit contributes ≈ 10 additional basis points economy-wide in the main counterfactual, independent of sector exposure.

E.5 Robustness Exercises

Three exercises assess robustness of the quantitative results. We perturb structural parameters locally, and vary exposure intensity across sectors. None of these exercises removes the model-specification rejection; they instead characterize numerical reliability and conditional sensitivity.

E.5.1 Parameter Sensitivity

Table E7 perturbs each structural parameter by one conventional GMM standard error, holding the remaining parameters at their 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 0.72 basis points in either direction and the full CEV ranges from -0.313 to $+0.312$ percent. The frontier quality step λ_F and the frontier-orientation cost κ generate comparable welfare movements; κ generates the largest movement in the frontier-innovation share.

Table E7: 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	1.26	-0.01	-0.041	-1.26	0.03	0.047
λ_F	-1.11	-0.71	-0.313	1.08	0.72	0.310
κ	2.60	0.71	0.312	-2.32	-0.64	-0.285
h_E	-0.35	-0.39	-0.041	0.35	0.38	0.041

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: $\text{SE}(\lambda_I) = 0.0019$, $\text{SE}(\lambda_F) = 0.0029$, $\text{SE}(\kappa) = 0.0035$, $\text{SE}(h_E) = 0.0095$.

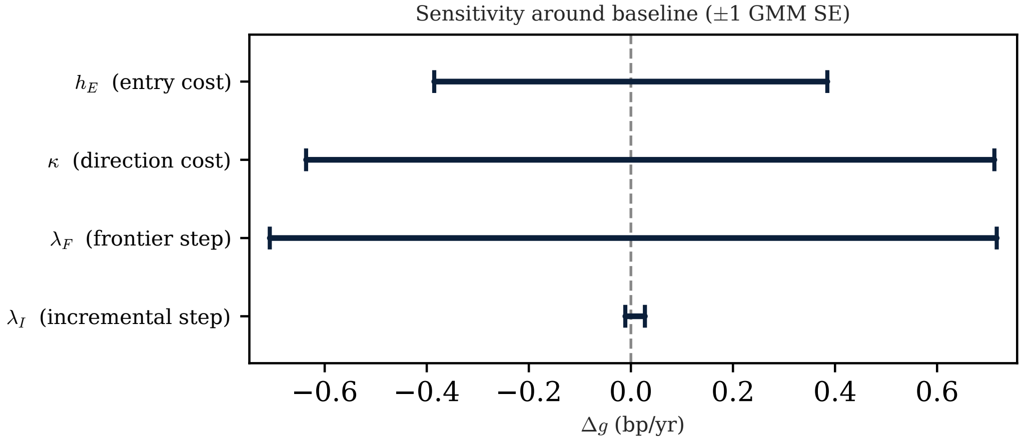

Figure E3: Baseline sensitivity (±1 GMM SE per parameter)

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.

E.6 Government Extensions

Risk aversion and the demand wedge. 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.

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 E.3 (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 E.3 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. 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.

Remark E.4 (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 E8 evaluates the threshold using the GMM estimates $\lambda_F = 1.188$ and $\lambda_I = 1.105$. Setting $\sigma_I = 0$ and varying σ_F gives the following values:

Table E8: 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 $\lambda_F - \lambda_I = 0.083$ 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 E8 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 $\lambda_F = 1.188$ and $\lambda_I = 1.105$.

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.

Government goods in the welfare aggregator. The baseline CEV captures gains from higher private consumption and faster growth but assigns no household utility to government-purchased output itself — infrastructure maintenance, IT systems, medical supplies, and other goods government buys. This extension adds that value back and checks whether the main welfare ranking survives. The concern is most acute for the R&D credit, which eliminates procurement entirely: households gain from the composition improvement but lose the output government was buying. The break-even weight $\alpha^* \approx 1.19$ exceeds the maximum of $\alpha = 1$ (equal weight on private and public consumption), so the R&D credit remains welfare-improving throughout $\alpha \in [0, 1]$.

Environment. Extend household utility to $U = \log C^{priv} + \alpha \log Q^{gov}$, where

$$\log Q^{gov} = (1 - \mu^*) \log(1 + G_I) + \mu^* \log(1 + G_F)$$

is a log-weighted index of government demand across line types, and $\alpha \geq 0$ governs the relative value of public services. At the benchmark $G_I = \bar{G}$ and $G_F = 0$, so only incremental lines generate public-good value. With log aggregation, total output is

$$Y^{total} = Y^{private} \cdot (1 + G)^{1-\mu^*}, \tag{67}$$

and no equilibrium condition changes. The parameter α enters only the welfare calculation, adding one term to the main-text CEV:

$$\mathrm{CEV}^{gov}_{B,A}(\alpha) = 100\left\{\exp\left[\Delta \log C_0 + \alpha\, \Delta[(1 - \mu^*) \log(1 + G_I) + \mu^* \log(1 + G_F)] + \frac{\Delta g}{\rho}\right] - 1\right\}. \tag{68}$$

Setting $\alpha = 0$ recovers the main-text formula exactly.

Mechanism. Each counterfactual type has a distinct effect on Q^{gov}. Budget-neutral reforms (specification realignment, monitoring) hold G fixed; Q^{gov} changes only because μ^* rises—more frontier lines, fewer incremental lines receiving government demand. The effect is small and positive: the public-good term adds $+0.004\%$ under specification realignment and $+0.003\%$ under monitoring. The historical rollback reduces G directly, partially destroying Q^{gov} and generating a large offset (-0.248 p.p. at $\alpha = 1$) that turns the net CEV negative once α is close to one. The R&D credit eliminates procurement entirely, losing the full benchmark public-good term $(1 - \mu^*_{bm}) \log(1 + \bar{G}) = 0.988\%$. The sensitivity as a function of α is

$$\mathrm{CEV}^{gov}_{CF1}(\alpha) = 100\Big\{\exp\Big[\log(1 + 0.011805) - 0.009881\,\alpha\Big] - 1\Big\}. \tag{69}$$

At $\alpha = 0$ this recovers the main-text CEV of 1.18%. The break-even weight $\alpha^* = \log(1.011805)/0.009881 \approx 1.19$ exceeds one, so the R&D credit dominates for every $\alpha \in [0, 1]$.[13] At $\alpha = 0.8$, the exact CEV is

---

[13] The R&D credit assumes full fiscal redeployment of the procurement budget. Above-threshold TED records already exclude routine below-threshold purchases (framework agreements, direct negotiation), and project and investment

0.38%. Table E9 reports the full set of counterfactuals. Jones and Klenow (2016) motivate valuing government services; Bandiera et al. (2009) motivate allowing procurement inefficiency. We report the full range rather than selecting one welfare weight: the point of the exercise is that the main-text welfare ranking is not an artifact of treating government purchases as pure waste, since the ranking survives even under the maximal weight on public consumption.

Table E9: Welfare effects of all counterfactuals with and without government goods in the aggregator

Counterfactual	$\Delta\mu^*$	Δg (bp/yr)	CEV	Pub. good Δ	CEV with gov. goods
Design reforms (G fixed, budget neutral)					
Specification realignment	$+0.014$	$+0.38$	$+0.166\%$	$+0.004\%$	$+0.169\%$
Monitoring reform	$+0.008$	$+0.21$	$+0.090\%$	$+0.003\%$	$+0.093\%$
Volume reduction					
Historical rollback	$+0.003$	$+0.07$	$+0.026\%$	-0.248%	-0.222% at $\alpha = 1$
Fiscal redeployment to R&D credit					
R&D credit ($s = 14.1\%$)†	$+0.061$	$+10.15$	$+1.18\%$	0.38% at $\alpha = 0.8$; $\alpha^* \approx 1.19$	

Notes: CEV is the exact permanent-consumption equivalent $100\{\exp[\Delta \log C_0 + \Delta g/\rho] - 1\}$ with $\rho = 0.04$. Government services enter inside the same exponential through $\alpha\Delta[(1-\mu)\log(1+G_I) + \mu \log(1+G_F)]$. The benchmark service term is 0.988%. Specification and monitoring hold procurement spending fixed; the historical rollback lowers G to 0.01339; and the R&D credit sets procurement demand to zero while retaining the benchmark tax to finance the credit. $\alpha = 0$ recovers the main-text CEV.

Source: Authors' calculations from the calibrated model.

---

contracts account for 69 percent of lots and 80 percent of value in 2016–2023. The $G = 0$ scenario is a hypothetical upper bound, not a feasible policy prescription.