Internal and External R&D: An analysis of costs and benefits *

Regi Kusumaatmadja †

August 2026

Abstract

This paper analyzes the costs and benefits of internal and external R&D activities. Using Dutch production and innovation surveys between 2000 and 2020, focusing on the ICT industry, I document an increasing share of firms performing R&D and the four-way mix of in-house research, contracted-out research, both, and neither. To rationalize these findings, I build and estimate a dynamic discrete-choice model of R&D with mode-specific investment costs. The cost of contracted-out R&D is much higher than the cost of in-house R&D, which explains the small share of firms that contract out alone. Doing both is cheaper than the sum of the two costs: in-house R&D offsets almost all of the extra cost of going outside. I read that extra cost as a hold-up problem with the partner. The offset is absorptive capacity: a firm that already does in-house R&D keeps more of the collaboration, so the two activities cost less together than apart. Productivity gains are similar across modes, so the mix of activities is driven by these costs. A uniform subsidy that does not favor a mode raises the share of R&D-active firms and firm value by more than a subsidy aimed only at in-house R&D.

Keywords: Research and Development, Complementarity, Innovation, Productivity, Internal R&D, External R&D, Outsourcing

JEL Codes: O31, O32, L25, C61, H25

*I would like to thank my supervisors, Sabien Dobbelaere and José-Luis Moraga-González, for their constant guidance and support. I would also like to thank Ulrich Doraszelski, Paul Grieco, Kenneth Judd, Claire Lelarge, Davide Luparello, Bettina Peters, Joris Pinkse, Mark Roberts, and Kazuma Takakura for their helpful comments and suggestions. I am grateful to seminar discussants and audiences at JEI-Gran Canaria (2022), IO/Applied Micro Brownbag - Penn State (2023), VU Amsterdam (2023), Tinbergen Institute (2023), EARIE (2023), DKIIS-Valencia (2023), CAED-Penn State (2024), SED Meeting-Barcelona (2024), CEPR Workshop-Rotterdam (2024), APIOC-Seoul (2024), CEPR Symposium-Paris (2024), and EEA-Bordeaux (2025) for their comments. The views expressed in this paper are mine and do not reflect the views of Statistics Netherlands (CBS). All errors are mine.

†VU Amsterdam and Tinbergen Institute. Mail: r.kusumaatmadja@vu.nl

1 Introduction

Firms typically conduct R&D either in-house or by contracting it out to strategic partners within or outside their home country. This paper develops a structural model to capture the costs and benefits of these R&D activities. Using Dutch microdata from 2000 to 2020, with a focus on the Dutch ICT industry, I quantify the complementarity between internal and external R&D and simulate its role in R&D subsidization programs.

I first document several stylized facts. The share of firms engaging in at least one type of R&D activity has increased over time. Most of these firms conduct R&D in-house, either exclusively or in combination with external partners. Specifically, the share of firms performing only internal R&D has risen steadily. In contrast, the share of firms relying solely on external R&D has remained constant. The proportion of firms conducting no R&D has declined overall but increased after 2016. I also report a reduced-form check of whether combining the two modes raises value added or the share of new products by more than adding them separately.

To rationalize these facts, I construct a model that captures the costs and benefits of internal R&D (conducted in-house) and external R&D (contracted out to other entities). The model comprises a static component, where demand and supply follow standard monopolistic competition assumptions, and a dynamic component, which adopts a single-agent dynamic discrete choice framework for R&D decisions. The dynamic part incorporates specific investment costs for each R&D activity and a cost-complementarity between internal and external R&D. When the firm contracts research out, a hold-up problem with the partner generates those two costs: the firm keeps only a share of the collaboration, and in-house capability raises that share. In the model, firms' R&D choices serve as inputs that influence productivity evolution, which in turn affects profitability.

Model estimation proceeds in stages. In the static part, I first estimate demand elasticity, followed by semiparametric productivity estimation to recover parameters governing productivity evolution. In the dynamic part, using estimates from the static part and a nested fixed-point algorithm (Rust, 1987), I recover the structural parameters: internal R&D investment costs, external R&D investment costs, and the cost-complementarity parameter.

Three patterns emerge. First, capital, labor, and the R&D mode the firm pursues are significant productivity determinants, and the productivity gains from in-house R&D, contracted-out R&D, and both are similar. Second, the average cost of in-house R&D is 0.13 million EUR, while the average cost of contracted-out R&D is 2.45 million EUR. Before the Innovation Box expansion, contracted-out R&D was about four times as costly as in-house R&D; in the full sample the gap is larger. That cost gap explains why so few firms rely on external R&D alone. Third, doing both is cheaper than the sum of the two costs. The extra cost of contracting out is largely waived when the firm also does in-house R&D. The mix of activities is therefore driven by costs, including a relationship-specific friction that in-house capability undoes. I read that friction as a hold-up problem on the external contract.

These estimates enable counterfactual exercises that mimic Dutch R&D subsidization as reductions in associated R&D costs. A uniform 5% subsidy across all costs raises the share of R&D-active firms and firm value (producer surplus) by the most. A 5% subsidy on contracted-out R&D costs is next. A 20% subsidy on in-house R&D costs alone produces the smallest change in participation and the smallest change in firm value. When cost-complementarity is shut off, the share of firms that do both falls sharply and the value gains from the subsidies shrink. The waiver of the external cost is what sustains the joint mode.

Related Literature.—This paper contributes to several strands of literature. The static part of my model aligns with Aw et al. (2011), Boler et al. (2015), and Peters et al. (2017), though it differs in its dynamic structure and distributional assumptions. The dynamic part draws on Igami (2017) and Igami and Uetake (2019), employing specific investment costs to represent general and activity-specific R&D expenses. Because prices are not observed, I model the dynamic part as a single-agent problem, and I explicitly incorporate complementarity. In related work, Akcigit and Kerr (2018) develop an endogenous growth model distinguishing internal (in-house improvements) and external (acquiring new lines) innovations, with external being costlier and scaling less with firm size, leading small firms to specialize in major breakthroughs. For this, I adapt insights from Miravete and Pernias (2006), who examine complementarity between product and process innovations, to the cost side of my dynamic framework.

My reduced-form evidence builds on Cassiman and Veugelers (2006), who test complementarity between internal and external R&D using insights from Arora (1996) and Athey and Stern (1998). Veugelers (1997) and Veugelers and Cassiman (1999) show that external technology sourcing raises in-house R&D when the firm already has a lab, and that larger firms combine the two. I extend the analysis by leveraging panel data. Cassiman and Veugelers (2006) focus solely on reduced-form results from a single CIS survey wave. Related work includes Mohnen and Roeller (2005), though I limit analysis to internal and external R&D, whereas they consider multiple activities. Grimpe and Kaiser (2010) document an inverted-U return to R&D outsourcing, moderated by internal R&D, which is the cost-side reading I take to the structural model. My use of a panel random-effects Tobit model to handle censored variables and panel structure echoes Hagedoorn and Wang (2012) and Love et al. (2014).

For productivity estimation, mainstream approaches often use value-added output (e.g., De Loecker and Warzynski, 2012; Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015), but these can face identification challenges with gross output. Solutions include nonparametric methods (Gandhi et al., 2021). I follow semiparametric techniques from Aw et al. (2011), Doraszelski and Jaumandreu (2013), and Peters et al. (2017).

Theoretically, this paper relates to hybrid R&D in Goyal et al. (2008), who derive conditions for complementarity in firm projects arising from cost advantages, even without external spillovers. That condition can hold from costs when the extra productivity from combining modes is small. My findings confirm cost-side complementarity between internal and external R&D. A leading microfoundation of that cost pattern is a hold-up problem on the external contract (Grossman and Hart, 1986; Hart and Moore, 1990), applied to research contracts by Aghion and Tirole (1994), together with absorptive capacity from in-house R&D (Cohen and Levinthal, 1989, 1990). Haggling over incomplete contracts produces the same qualitative pattern (Bajari and Tadelis, 2001; Tadelis, 2002). Jungbauer et al. (2026) study product-level outsourcing when an innovator risks cannibalizing existing products. I study contracted-out research in the Community Innovation Survey, and the extra cost of that contract when the firm also has a lab. I draw on Coase (1937) and Williamson (1975) for the broader transaction-cost view of the firm's boundary. I use Grossman–Hart with absorptive capacity because it is closed form and it nests in the four-mode Bellman. The estimated objects remain the three dynamic costs. Section 3 records the subperiod as part of the model.

On R&D subsidies, Galaasen and Irarrazabal (2021) show uniform subsidies boost investment, growth, and welfare for Norwegian firms, while size-dependent ones increase aggregate R&D but reduce growth and welfare. My results similarly favor uniform, non-discriminatory subsidies for increasing the share of R&D-active firms. The fiscal-incentive literature finds that R&D tax credits lower the user cost of research (Hall and Van Reenen, 2000; Bloom et al., 2002). For the Netherlands, Lokshin and Mohnen (2012) find that the WBSO level-based credit raises firms' R&D capital. Bilir and Morales (2020) find that for US multinationals, parent R&D spills over to foreign affiliates with strong complementarity from affiliate R&D, leading 20% of returns to domestic subsidies to be realized abroad and highlighting spatial disconnects in policy design. Gonzalez et al. (2005) find subsidies induce non-R&D Spanish firms to invest. I find that a uniform 5% cost cut raises R&D participation by more than a 20% cut in in-house costs alone.

Additional relevant papers on complementarity and outsourcing include Ceccagnoli and Jiang (2013); Ceccagnoli et al. (2014), who explore firm-level drivers of complementarity between internal R&D and external knowledge acquisition, and Berchicci (2013), who examines the innovative performance of R&D outsourcing. For firm boundaries in R&D, Howells (2008) discusses blurring boundaries in global outsourcing, and Amoroso (2017) analyzes complementarities among internal, external, and cooperative R&D.

Structure of the Paper.—This paper has the following structure. The next section describes the data, information on the Dutch Tax Incentives for Innovation scheme, and reduced-form evidence. In Section 3, I present the static profit, the four-mode dynamic discrete choice, and the contracting subperiod that generates the external cost and the complementarity. Section 4 outlines the empirical implementation of the structural model and briefly discusses sources of identification. The results of the model are presented in Section 5. I then proceed with some counterfactual exercises in Section 6. Section 7 concludes.

2 Institutional Setting, Data, and Motivating Evidence

This section outlines the Dutch Tax Incentives for Innovation, describes the data used, and presents stylized facts as motivating evidence.

2.1 Dutch Tax Incentives for Innovation

To promote research and development in the Netherlands, the Dutch government introduced a comprehensive innovation program that provides tax incentives for Dutch companies. This program has two components: cost-side and profit-side incentives. On the cost side, eligible firms receive significant reductions in wage taxes for R&D performed in the Netherlands. On the profit side, eligible firms benefit from a special corporate income tax regime for profits generated from associated R&D activities. These incentives result in a lower effective tax rate for participating firms.

The cost-side program is the WBSO (NL: *Wet Bevordering Speur- en Ontwikkelingswerk*, EN: The Promoting Research and Development Act), administered by the RVO (NL: *Rijksdienst voor Ondernemend Nederland*, EN: the Dutch Enterprise Agency). The WBSO provides a remittance reduction on wage taxes and national insurance contributions for R&D activities. To qualify, firms must apply for a WBSO statement from the RVO, demonstrating: (1) development of a technically new product, production process, software, or scientific research; and (2) resolution of technical bottlenecks that cannot be addressed by known techniques, necessitating R&D[1]. Upon approval, firms receive the WBSO statement and allocated R&D hours to calculate their "R&D base," enabling the remittance reduction.

In 2022, the reduction was 32% of the R&D base up to €350,000 and 16% thereafter, with a 40% rate up to €350,000 for new entrants (firms employing staff for less than five years and granted WBSO statements in fewer than three calendar years). The R&D base can be calculated based on R&D hours or actual costs and expenses[2].

The profit-side program is the Dutch Innovation Box policy, which offers a reduced corporate income tax rate. Originally the Patent Box (introduced in 2007), it expanded in 2010 to the Innovation Box (Mohnen et al., 2017). As of 2022, qualifying profits are taxed at 9% instead of the standard 25.8%[3]. Only profits attributable to R&D qualify, requiring a WBSO statement. Large firms also need a patent or plant breeders' right (NL: *Kwekersrecht*). Firms are "small" if their five-year gross margin from intangibles is under €37.5 million and turnover under €250 million.

[1]For example, a project using existing libraries like TensorFlow to build a machine learning model would likely be rejected, as it relies on established technologies. However, developing a model from scratch to solve novel technical issues might qualify. The key criterion is whether the project addresses unsolved technical problems.

[2]For the hours-based method, use an average hourly wage of €29 per granted R&D hour, plus €10 per hour for the first 1,800 hours and €4 per additional hour.

[3]The rate has increased from 5% since 2010. The Netherlands has two brackets: 15% on profits up to €395,000 and 25.8% thereafter. For example, a firm with €500,000 in profits pays 15% on the first €395,000 and 25.8% on the remainder.

Four methods determine the taxable profit under the Dutch Innovation Box policy, with the Dutch Tax Authority selecting the appropriate one to allocate profits attributable to R&D: (1) The peel-off method (NL: *Afpelmethode*), the most common, allocates earnings before interest and tax (EBIT) to entrepreneurship, sales, production, and R&D[4]. (2) The cost-plus method adds an 8-15% markup to R&D costs. (3) The single intangible asset method, similar to peel-off, focuses on R&D expenses and benefits from a specific asset[5]. (4) The flat-rate method attributes 25% of profits to the Box if other methods are infeasible[6].

The Dutch government has adopted the OECD "Nexus Approach" for the Innovation Box Policy, under which only qualifying income relating to intangible assets developed by firms "in-house" will be eligible for the application. This decision effectively prevents firms from benefiting from the tax benefits of the policy if they do not have a substantial economic presence in the Netherlands or if they are not engaged in any research or innovative activities in the country. Furthermore, the decision implies that only innovation developed in-house by firms will be eligible for the tax benefits[7].

In counterfactual simulations, I approximate the WBSO and the Innovation Box as reductions in the costs of R&D.

2.2 Data

I use datasets from multiple sources. The primary dataset is the annual Production Statistics (PS) survey from Statistics Netherlands (CBS), providing firm characteristics like capital, intermediate inputs, and revenue. It covers all large firms; smaller firms are sampled randomly[8].

The second dataset is the biennial Community Innovation Survey (CIS). This dataset contains information about the mode of R&D. Let us first discuss the mode of R&D captured in this survey. The definition of the mode of R&D in this paper closely follows the definition outlined by Cassiman and Veugelers (2006), building on Veugelers (1997) and Veugelers and Cassiman (1999). For example, suppose a firm performs in-house R&D with its own staff. In that case, I define this R&D activity as internal R&D. If a firm decides to contract out its R&D to other firms, universities, or other strategic partners, I define this activity as external R&D. The survey only contains information on the extensive margin, i.e., whether firms perform a certain type of R&D or not in a given

[4]The Tax Authority assesses shares (e.g., 20% entrepreneurship, 10% sales, 30% production, 40% R&D) and ensures R&D is essential to operations, with costs compensated by attributable profits.

[5]Used in royalty structures for clear revenue attribution, e.g., from a single patent.

[6]Capped at €25,000 annually; typically for small firms.

[7]Technically, R&D outsourcing to an affiliated entity is feasible as long as the R&D costs and risks are owned by the firm in which it has a significant domestic presence. However, the benefit of doing such a strategy is greatly reduced. Only a maximum of 30% of the R&D expenses can be outsourced to an affiliated entity to be eligible for the Innovation Box benefits. In practice, such eligibility is difficult to attain and, once attained, to demonstrate to the tax authority, so contracted-out R&D receives little effective support under the scheme.

[8]Details at `https://www.cbs.nl/en-gb/onze-diensten/methods/surveys/korte-onderzoeksbeschrijvingen/production-statistics`. For firms with fewer than 10 workers, tax data supplements; those with under 50 are sampled; all with 50+ are included.

year[9].

I focus on the Dutch Information, Communication, and Technology (ICT) industry, including subsectors like software, IT, data processing, telecommunications, and related activities[10]. The ICT sector is growing and vital to the Dutch economy, outpacing overall growth and central to innovation policies like the Digital Agenda (van Economische Zaken, 2017). It ranks high in EU digital metrics and benefits from policies favoring tech innovation. In my dataset, ICT contributes the largest number of observations (750 firms). Evaluations indicate high WBSO take-up in knowledge-intensive sectors like ICT, with €1.5 billion in annual benefits economy-wide (Dialogic et al., 2025), and a positive effect of the credit on firms' R&D capital (Lokshin and Mohnen, 2012).

I use OECD STAN deflators for output, capital, and intermediates. The final unbalanced panel merges PS and CIS from 2000-2020 (even years).

Table 1: Firm Characteristics of Dutch ICT Industry

Variables	**Average**
Output/Revenue	59822.95
Capital	7899.39
Intermediate Inputs	28810.87
Value Added	30353.23
FTE Workers	229.55
Energy	247.57
Share of New Products	11.01%
No R&D	37%
Internal-only R&D	32%
External-only R&D	3%
Both R&D	28%
Observations	2883
Number of Firms	750

Notes.—The unit for firms' output, capital, intermediate inputs, value-added, and energy is thousands of EUR (real). The unit for FTE workers is the number of people employed by a firm.

Table 1 summarizes key variables. Among 750 firms (2,883 observations), average revenue is €59.8 million, capital €7.9 million, intermediates €28.8 million, workers 230 FTE, and energy €248,000. On average, 37% perform no R&D at all, 32% choose to perform only internal R&D, 3% choose to perform only external R&D, and 28% choose to perform both internal and external R&D.

2.3 Motivating Evidence

I present facts on Dutch R&D activities: time trends, persistence, and complementarity evidence.

[9]From the CIS, the indicator of R&D type must be filled in by the firm. However, the information on how much R&D expenses for each activity is not mandatory and is not always available or asked in the survey questionnaire in every wave of the survey.

[10]Firms in SBI codes: 61 (telecommunications), 62 (IT support), 63 (information services), 582 (software publishing).

Facts 1 (Time-trend of R&D: Internal v. External R&D) *The share of firms performing Internal R&D is larger than External R&D. The share of firms performing Internal R&D is increasing, especially after the expansion of the Dutch Innovation Box Policy. The share of firms performing External R&D remains constant throughout the sample period.*

Figure 1: Time trend of R&D Activities

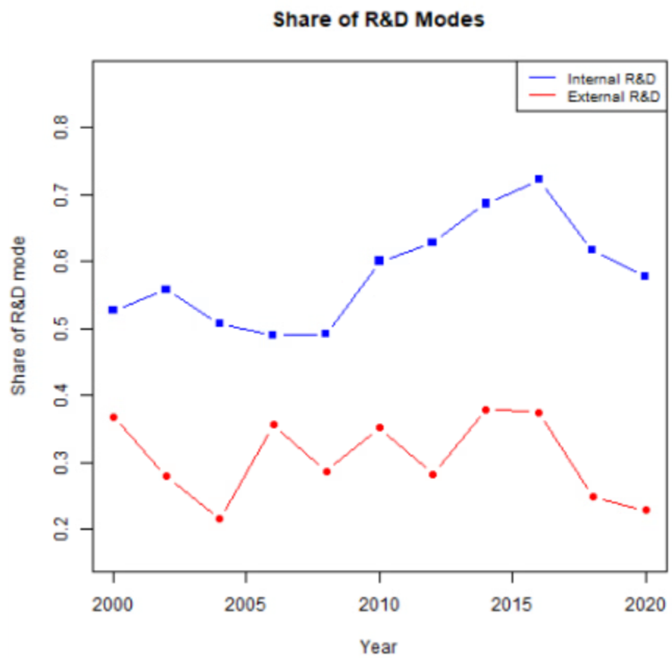

Notes.—The figure is constructed from an unbalanced panel of Dutch Production Statistics (PS) and Community Innovation Survey (CIS) datasets by taking the mean of R&D modes in each year. *Internal R&D* indicates that a firm performs Internal R&D in a given year. *External R&D* indicates that a firm performs External R&D in a given year.

Fact 1 is summarized by Figure 1. The share of firms performing internal R&D is larger than external R&D. The trend of internal R&D is also increasing, especially after the expansion of the Dutch Innovation Box policy in 2010; that is, from around 50% to a peak of around 70% in 2016. The share of firms performing external R&D is also constant in the range of around 20% to 40% throughout the sample period.

Facts 2 (Time-trend of R&D: None, Internal-only, External-only, Both R&D) *A decrease in the share of firms performing no R&D at all, especially after the expansion of the Dutch Innovation Box Policy. An increase in the share of firms performing Internal-only R&D and Both R&D. A persistently small share of firms performing External-only R&D.*

Figure 2: Time trend of R&D Modes

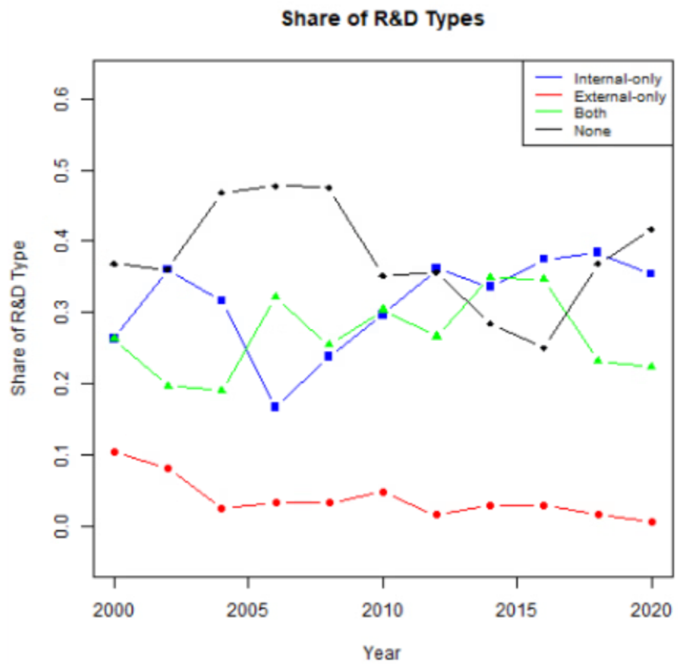

Notes.—The figure is constructed from an unbalanced panel of Dutch Production Statistics (PS) and Community Innovation Survey (CIS) datasets by taking the mean of R&D modes in each year. *No R&D* indicates that a firm performs no R&D in a given year. *Internal-only R&D* denotes that a firm performs only internal R&D, while *External-only R&D* means that a firm only contracts out R&D to other strategic partners. *Both R&D* marks that a firm conducts both Internal and External R&D.

Fact 2 is summarized by Figure 2. The share of firms performing internal-only R&D (the blue line) has been increasing over the years, especially after the expansion of the Dutch Innovation Box policy in 2010. The share of firms performing no R&D at all decreases after the expansion, then increases again in the later years (from 2018 to 2020). The share performing both R&D rises after 2010 but falls in 2018 and 2020. The share performing only external R&D is persistently small over the sample period.

Facts 3 (R&D activities persistence) *Past R&D activities influence current R&D activities.*

Table 2: OLS: Persistence of R&D modes

OLS	*Dep. var*			
	No R&D (t)	**Internal-only (t)**	**External-only (t)**	**Both (t)**
No R&D (t-1)	0.4202*** (0.0296)			
Internal-only (t-1)		0.2788*** (0.0366)		
External-only (t-1)			0.1780** (0.0561)	
Both (t-1)				0.3538*** (0.0351)
Year FE	Yes	Yes	Yes	Yes
Sub-Industry FE	Yes	Yes	Yes	Yes

Notes.—OLS of the current period of R&D modes on the previous period of R&D modes. Standard errors are in parentheses. I use heteroskedasticity and autocorrelation-consistent (HAC) standard errors. Significance level: *** : $p < 0.001$, ** : $p < 0.01$, * : $p < 0.05$, $^{\cdot}$: $p < 0.1$.

Fact 3 is summarized by Table 2[11]. Table 2 reports an OLS regression of current R&D modes on lagged R&D modes. The coefficients are all positive and statistically significant. Conditional on year and sub-industry fixed effects, last period's no-R&D choice is associated with a 42.02 percentage-point higher probability of current no R&D (linear probability). The analogous associations are 27.88 percentage points for internal-only, 17.80 for external-only, and 35.38 for both.

Facts 4 (Complementarity of Internal and External R&D) *Reduced-form check of incremental returns to combining internal and external R&D.*

Fact 4 is summarized by Table 3. The table follows Arora (1996), Athey and Stern (1998), and Cassiman and Veugelers (2006), based on Milgrom and Roberts (1990, 1995).[12] The four coefficients are positive and significant. Complementarity requires $\theta_{\text{both}} - \theta_{\text{internal}} \geq \theta_{\text{external}} - \theta_{\text{no}}$ (Equation (23)). Define Δ as the difference between those two increments. The last row of the table reports Wald p-values of H_0: $\Delta = 0$: 0.38, 0.19, and 0.44. I do not reject equal incremental returns. The share-of-new-products OLS point satisfies the inequality; the value-added and Tobit points lie slightly the other way. Cost-complementarity is identified later from the share of firms that choose both.

Discussion on Identification. The reduced-form results have caveats. Ignoring censoring in pooled OLS yields biased estimates, as the censored sample is unrepresentative (more positive disturbances). The panel random-effects Tobit mitigates this. Endogeneity persists: unobserved firm-specific factors may correlate with R&D choices and outcomes. Controls do not fully address this, and Athey and Stern (1998) notes risks from unobserved heterogeneity. I interpret this result

[11] A probit specification confirms the same sign pattern. See Appendix A.
[12] See Appendix B for details.

Table 3: Reduced-form of Complementarity of Internal and External R&D

	Dependent Variable		
	$log(VA)$	Share New Product	
	OLS	**OLS**	**Tobit (RE)**
No R&D	9.822^{***} (0.2313)	15.92^{***} (3.678)	12.81^{***} (2.31)
Internal-only R&D	9.892^{***} (0.2326)	15.36^{**} (3.605)	25.56^{***} (2.09)
External-only R&D	10.46^{***} (0.3031)	13.82^{**} (4.759)	18.75^{***} (3.55)
Both R&D	10.33^{***} (0.2360)	18.00^{***} (3.612)	30.93^{***} (2.14)
Year F.E	Yes	Yes	Yes
Sub-Industry F.E	Yes	Yes	Yes
Wald p-value of $\Delta = 0$	0.38	0.19	0.44

Notes.—Significance level: $^{***}: p < 0.001$, $^{**}: p < 0.01$, $^{*}: p < 0.05$, $^{\cdot}: p < 0.1$. The dependent variable is the share of new products in revenue $\in [0, 100]$. Standard errors are in parentheses. I use heteroskedasticity and autocorrelation-consistent (HAC) standard errors. The last row reports p-values for a Wald test of equal incremental returns, $\Delta = 0$ (Equation (23)). Large p-values mean the extra return to combining the two modes is not distinguishable from adding them separately.

cautiously. These caveats are my primary motivation to develop the dynamic structural model discussed in the next section.

3 Model

In this section, I present my model. The model has a static part, a dynamic part, and, when the chosen mode uses an outside partner, a contracting subperiod. In the static part, I discuss the supply and demand sides of the model, and the firm's revenue and profit under monopolistic competition. In the dynamic part, I outline a single-agent discrete choice over four R&D modes, with mode-specific investment costs and a cost-complementarity parameter. For modes that contract research out, I embed a Grossman–Hart problem with one-sided investment (Grossman and Hart, 1986). The equilibrium nets of that problem locate the cost of contracted-out R&D and the complementarity in the Bellman. I estimate those three costs. I do not estimate the bargaining share.

3.1 Static Decisions: Supply and Demand

The general structure of the static part follows Aw et al. (2011), Boler et al. (2015), and Peters et al. (2017)[13]. In particular, I assume monopolistic competition in the firm's static decisions. Monop-

[13] ***Alternative Demand System***. A related discrete-choice representation with log prices yields a similar revenue equation (Appendix E). Own-price and cross-price elasticities coincide by that restriction, not as a finding from a Berry (1994) or Berry et al. (1995) demand system. I use the CES-like demand because prices and quantities are not observed at the firm level.

olistic competition approximates a market in which strategic interactions among firms are weak (see Zhelobodko et al. (2012) and Thisse and Uschev (2018) for the discussion of the general case of monopolistic competition).

3.1.1 Demand

The representative consumer has a CES-like utility function. By implication, the demand curve faced by firm i is then assumed to follow the Dixit-Stiglitz form. That is,

$$Q_{it} = Q_t \left(\frac{P_{it}}{P_t} \right)^{\eta} \tag{1}$$

where P_{it} is the firm's output price, P_t is the average price in the period t, and η is the elasticity of demand, which is assumed to be constant for all firms in the industry.

3.1.2 Supply

Now, suppose firm i's (log of) marginal cost is defined as follows

$$c_{it} = c(k_{it}, l_{it}, w_t) - \omega_{it} = \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_w ln(w_t) - \omega_{it} \tag{2}$$

where k_{it} is the log of firm's capital stock, l_{it} is the log of the number of FTE workers, w_t is a vector of variable input prices common to all firms, and ω_{it} is the firm's productivity. The firm is assumed to produce a single output. There are several sources of cost heterogeneity in this model[14]. The first one is coming from the firm's capital stock which is directly observable in the data. The second source is from the number of workers the firms have, which is also directly observable in the data. The third source is the firm's productivity which is observed by the firm, but not observable by the econometrician. Notice here that the more productive the firm is, the lower the marginal cost it faces. The marginal cost does not vary with the firm's output level, implying that demand shocks in one market do not affect the static output decision in other markets. Therefore, I can compute revenue and profits in each market independently of the output levels in other markets.

3.1.3 Revenue

Using information from the demand-side and the marginal cost, I can get the log of firm's revenue as follows

Results 1 (Revenue equation) *The (log) revenue equation is derived from both the demand and supply*

[14]***Alternative Supply-Side*** .—The lack of information on prices, the number of products each firm produces, and the quality of each product precludes an analysis of the supply-side as in the standard empirical IO literature. For example, the absence of firm's output price makes it hard to model the supply-side as an oligopoly. In this model, I also assume a single-product firm throughout the economy since I do not observe the number of products.

sides, and follows the following expression

$$r_{it} = (\eta+1)ln\left(\frac{\eta}{\eta+1}\right) + ln(\Phi_t) + (\eta+1)\Big[\beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_w ln(w_t) - \omega_{it}\Big] \tag{3}$$

where $\Phi_t := Q_t / P_t^{\eta}$.

Proof. See Appendix C for the full derivation. ■

From the above expression, the firm's revenue provides information on its marginal cost, in particular the productivity level ω_{it}. In the empirical implementation of the model, I will estimate the revenue functions and can interpret the source of the unobserved heterogeneity, ω_{it}. Note that while I describe ω_{it} as the firm's productivity, it could also include the quality of the product which would affect the demand for the firm's product, as well as the cost.

3.1.4 Static Profit

I can now recover the profit of the firm using the expression of the firm's revenue above as follows

Results 2 (Static profit equation) *The static profit equation is derived from both the revenue equation and the (log) marginal cost equation, and follows the following expression*

$$\pi_{it} = -\left(\frac{1}{\eta}\right) exp(r_{it}(k_{it}, l_{it}, \omega_{it}; \eta, \beta_0, \beta_k, \beta_l)) \tag{4}$$

Proof. The revenue expression follows from taking the exponent of the log revenue equation from Results 1. The cost expression follows from taking the exponent of the log marginal cost with the quantity from the demand side. To retrieve the static profit, I subtract cost from revenue. See Appendix C for the full derivation. ■

The above expression allows me to recover the profit of the firm using estimates of the model and observables in the data. Note that, this profit will be an important ingredient of the firm's decision to choose modes of R&D in the dynamic discrete choice model developed in the next subsections.

3.2 Dynamic Decisions: Dynamic Discrete Choice of R&D modes

As mentioned in Section 2, I try to capture the WBSO through some form of subsidy from the government. In particular, if the firm receives this subsidy, it experiences a decrease in its costs of investing in certain modes of R&D. As such, the following single-agent dynamic discrete choice model will incorporate investment costs as the main structural elements.

3.2.1 Setup

I assume that time is discrete with infinite horizon. A finite number of firms are indexed by i. The firm's state space is defined as $s_{it} = (\omega_{it}, k_{it}, l_{it})$, which endogenously evolves as the firm decides which mode of R&D to choose. ω_{it} is the (discretized) productivity of firm i in a year t. In each year, firm i can choose to do no R&D at all, only internal R&D, only external R&D, or both internal and external R&D.

Each mode of R&D entails specific investment costs. For example, suppose firm i decides to perform internal R&D. In that case, it needs to invest a certain cost of $\kappa^{internal}$ to procure the internal R&D. If the firm decides to do external R&D, then, it needs to invest a certain cost of $\kappa^{external}$ in order to perform the external R&D. If the firm decides to do both R&D, it has to pay both types of investment costs. These costs are mode-specific investments that firms need to pay to perform a particular R&D. For firms that choose not to do R&D at all, there is no specific investment cost.

There are several possible reasons why a firm decides to do both internal and external R&D. First, both types of R&D may benefit the firm in terms of cost savings or better product quality. For example, in the Dutch ICT industry, a firm focusing on online booking technology and services can decide to build a dynamic pricing algorithm from scratch (i.e., an internal R&D). At the same time, this firm can also consult with economists and computer scientists at several universities on the technical know-how on issues and state-of-the-art solutions related to the dynamic pricing algorithm (i.e., an external R&D). Doing both types of R&D allows the firm to gain an additional benefit from faster technological adoption of the dynamic pricing algorithm, even though it incurs an additional investment cost of doing both R&D types. I call this phenomenon the cost-complementarity of internal and external R&D. It is also possible that doing both internal and external R&D yields no meaningful benefit for the firm. Following Miravete and Pernias (2006), I capture the cost-complementarity of internal and external R&D in a structural parameter $\kappa^{complement}$. In particular, I assume that this parameter $\kappa^{complement}$ can take on the following values.

$$\kappa^{complement} \begin{cases} > 0, & [Complement] \\ = 0, & [Independent] \\ < 0, & [Substitute] \end{cases}$$

In the above expression, three possible values exist for the cost-complementarity parameter, $\kappa^{complement}$. If it is positive, then both types of R&D complement each other. It is also possible that doing both types of R&D does not benefit the firm, i.e., both are independent. I also include in the parameter $\kappa^{complement}$ the possibility that both types of R&D are substitutes.

Each firm draws an IID private cost shock $\epsilon_{it}^{a_{it}}$ which follows a Type-1 Extreme Value distribution. After observing the four shocks, firm i takes an action a_{it} from the choice set $a_{it} \in$

$\{\text{none}, \text{internal-only}, \text{external-only}, \text{both}\}$. Private shocks reflect each firm's informational, managerial, and organizational conditions. If the chosen action uses an outside partner, the firm then sinks a relationship-specific investment and bargains, as described in the contracting subperiod below. Bargaining therefore takes place inside a mode that has already been selected. The logit formula for the four choice probabilities is the integral of that discrete problem. If the firm bargained before drawing the four shocks, that closed form need not apply.

3.2.2 State Transitions

I assume that ω_{it} follows a controlled Markov process. In particular, the distribution of ω_{it} can be written as $G_\omega(\omega_{it} \mid \mathcal{I}_{it-1}) = G_\omega(\omega_{it} \mid \omega_{it-1}, \text{none}_{it-1}, \text{int}_{it-1}, \text{ext}_{it-1}, \text{both}_{it-1})$, where $\mathcal{I}_{it}$ is an information set which contains the information that the firm can use to solve its decision problem in period t. I can express the persistent productivity ω_{it} as $\omega_{it} = g_\omega(\omega_{it-1}, \text{none}_{it-1}, \text{int}_{it-1}, \text{ext}_{it-1}, \text{both}_{it-1}) + \xi_{it}$, where $E[\xi_{it} \mid \mathcal{I}_{it-1}] = 0$. Variables $\text{none}_{it-1}, \text{int}_{it-1}, \text{ext}_{it-1}, \text{both}_{it-1}$ are indicator variables (in period $t-1$) that refer to the decision of the firm to choose no R&D at all, only internal R&D, only external R&D, and both R&D, respectively. The random variable ξ_{it} captures the (unanticipated at period $t-1$) stochastic nature of improvement on the firm's persistent productivity ω_{it} in period t. By construction, ξ_{it} does not correlate with ω_{it-1}. Note that $\text{none}_{it-1}, \text{int}_{it-1}, \text{ext}_{it-1}, \text{both}_{it-1}$ are assumed to be predetermined variables; i.e., these variables are realized before ω_{it} is realized and functions of previous period information set; that is, $\{\text{none}_{it-1}, \text{int}_{it-1}, \text{ext}_{it-1}, \text{both}_{it-1}\} = \mathbb{X}(\mathcal{I}_{it-1}) \in \mathcal{I}_{it}$. Therefore, those variables also do not correlate with ξ_{it}. The transition therefore depends on which of the four modes the firm chose. It does not depend on the size of the relationship-specific investment inside an external contract.

3.2.3 Dynamic Programming Problem

Firms make their dynamic discrete choices to maximize their expected values. The future stream of profits is discounted by a factor $\beta \in (0, 1)$. The following Bellman equation characterizes the dynamic programming problem of each firm:

$$
\begin{aligned}
V(s_{it}) &= \pi_{it}(s_{it}) \\
&\quad + \max \Big\{ \underbrace{\beta EV^{\text{no R\&D}}(s_{it}) + \epsilon_{it}^{\text{no R\&D}}}_{\text{No R\&D}}; \\
&\quad \underbrace{-\kappa^{\text{internal}} + \beta EV^{\text{internal}}(s_{it}) + \epsilon_{it}^{\text{internal}}}_{\text{Internal-only}}; \\
&\quad \underbrace{-\kappa^{\text{external}} + \beta EV^{\text{external}}(s_{it}) + \epsilon_{it}^{\text{external}}}_{\text{External-only}}; \\
&\quad \underbrace{-\kappa^{\text{internal}} - \kappa^{\text{external}} + \kappa^{\text{complement}} + \beta EV^{\text{both}}(s_{it}) + \epsilon_{it}^{\text{both}}}_{\text{Both}} \Big\}
\end{aligned}
\tag{5}
$$

As explained, the key parameters of this dynamic discrete choice R&D model are the investment cost of internal R&D ($\kappa^{internal}$), the investment cost of external R&D ($\kappa^{external}$), and the cost-complementarity parameter ($\kappa^{complement}$). Here, I assume that unobservables ϵ_{it}^{a} are additive to the key parameters and the expected future value, $EV(\cdot)$. I also assume that the scale parameter is normalized to one, $\sigma = 1$.

To illustrate how these choices manifest in the Bellman equation (5), consider the example of a firm in the Dutch ICT industry developing a dynamic pricing algorithm. The firm weighs four options to maximize its long-term value: doing no R&D at all, which provides only the current profit plus a discounted expected future value adjusted by some random shock; pursuing internal R&D only, which involves paying an internal cost but gaining a potentially higher future value from in-house development like building the algorithm from scratch; opting for external R&D only, which subtracts an external cost while benefiting from outside expertise, such as university consultations; or choosing both, which combines the internal and external costs but includes a complementarity adjustment that could reduce the net expense if positive, alongside an even greater expected future value from the synergies. Even if the immediate net cost of doing both remains higher than internal R&D alone (for instance, when the complementarity savings do not fully offset the added external expense), firms may still select this option because the long-term benefits, captured in the higher expected future value from faster innovation and improved outcomes, can more than compensate for the upfront investment, ultimately leading to greater overall profits when discounted back to the present.

I follow Rust (1987, 1994) in exploiting the property of the logit error and the assumption of the conditional independence over time. I get a closed-form expression for the expected value *before* observing private shocks $\epsilon_{it}^{a_{it}}$ as follows (i.e., the McFadden's social surplus)

Results 3 (McFadden's Social Surplus) *The McFadden's social surplus can be expressed as follows*

$$\begin{aligned} EV(s_{it}) = {} & \pi_{it}(s_{it}) + \gamma^{Euler} \\ & + ln\left[exp(\overline{V}_{it}^{no}) + exp(\overline{V}_{it}^{internal}) + exp(\overline{V}_{it}^{external}) + exp(\overline{V}_{it}^{both})\right] \end{aligned} \tag{6}$$

where γ^{Euler} is the Euler's constant. $\overline{V}_{it}^{no}, \overline{V}_{it}^{internal}, \overline{V}_{it}^{external}, \overline{V}_{it}^{both}$ represent the deterministic part of the conditional (or "alternative-specific") values in Equation (5)*.*

Proof. The derivation of McFadden's social surplus is in the Appendix F ■

Without McFadden's social surplus expression, computing the expected maximum payoff would require costly high-dimensional integration over shock distributions for each state and period. This log-sum-exp formulation simplifies solving the Bellman equation, deriving choice probabilities, and estimating model parameters empirically.

As for the expected future value $(EV(\cdot))$, I can define it as follows

Results 4 (Expected future value) *The expected future value can be expressed as follows*

$$EV^a(s_{it}) = \sum_{\omega'} V(s_{it+1})\, P(\omega' \mid s_{it}, a_{it}) \qquad , \qquad \forall a_{it} \in \{no\ R\&D, internal, external, both\} \tag{7}$$

where $P(\omega' \mid s_{it}, a_{it})$ can be thought of as the transition probability or, equivalently, the state transition.

Proof. The derivation of $EV(\cdot)$ can be found in the Appendix G. ■

I can now define the conditional choice probabilities as follows

Results 5 (Conditional choice probabilities) *The conditional choice probabilities can be expressed as follows*

$$Pr(a_{it} = action) = \frac{exp(\overline{V}_{it}^{action})}{exp(\overline{V}_{it}^{no}) + exp(\overline{V}_{it}^{in}) + exp(\overline{V}_{it}^{out}) + exp(\overline{V}_{it}^{both})} \qquad , \qquad \forall a_{it} \in \{no\ R\&D, internal, external, both\} \tag{8}$$

Proof. The full derivation of the probability can be found in the Appendix H. ■

I use these optimal choice probabilities to construct a likelihood function for the empirical implementation.

3.2.4 Hold-up and the cost of contracted-out R&D

The objects I estimate remain the three costs in the Bellman equation (5). This subsection supplies the theory of two of those costs. When the firm chooses external-only or both, it enters a one-period Grossman–Hart problem with a partner (Grossman and Hart, 1986; Hart and Moore, 1990), in the tradition of Aghion and Tirole (1994) on the organization of research. The equilibrium nets of that problem locate $\kappa^{external}$ and $\kappa^{complement}$ in the Bellman. I do not estimate the bargaining share, and I do not add a separate hold-up payoff on top of the Bellman.

I use this pairing because it is closed form and it sits inside the four-mode problem already written down. In-house R&D raises the firm's claim on the collaboration, as in the absorptive-capacity argument of Cohen and Levinthal (1989, 1990). Haggling over an incomplete contract (Bajari and Tadelis, 2001; Tadelis, 2002) produces the same two intercepts: a lab makes the external agreement cheaper to use. The Community Innovation Survey records the mode, not the contract, so those stories are not separately distinguished. Grossman–Hart with absorptive capacity is the illustration I carry through the algebra.

Primitives. Consider a choice that involves external R&D, $a_{it} \in \{\text{external-only}, \text{both}\}$. The firm sinks a relationship-specific investment $I \geq 0$. That investment costs I^2 and raises the collaboration's pie by θI, with $\theta > 0$ a surplus scale in the same units as the Bellman (million EUR). Because the contract is incomplete, the partner can hold the firm up after I is sunk. Disagreement payoffs after that investment is sunk are $(0, 0)$: the investment has no scrap outside this match. The firm then keeps a share $\tau \in (0, 1]$ of the pie, and the partner keeps $1 - \tau$. When disagreement payoffs are $(0, 0)$ and the pie is transferable, that split is the asymmetric Nash solution (Nash, 1950) and the proportional solution of Kalai (1977). I treat τ as a primitive. I do not estimate a nested bargaining model.

The firm anticipates the split and chooses

$$\max_{I \geq 0} \ \tau\theta I - I^2. \tag{9}$$

The first-order condition is $\tau\theta - 2I = 0$, so

$$I^* = \frac{\tau\theta}{2}, \qquad \text{net}(\tau) := \tau\theta I^* - (I^*)^2 = \frac{\tau^2\theta^2}{4}. \tag{10}$$

The second derivative is $-2 < 0$, so I^* is the unique maximum. Appendix I records the steps.

If the firm kept the entire pie, the same program would be run at $\tau = 1$. Call that the efficient investment. Then $I^{\mathrm{e}} = \theta/2$ and

$$S := \text{net}(1) = \frac{\theta^2}{4}. \tag{11}$$

S is the surplus the firm would collect if it owned the collaboration outright. The equilibrium net is the fraction τ^2 of that surplus: $\text{net}(\tau) = S\tau^2$. A low τ produces underinvestment and a small net from contracting out.

If one normalizes $\theta = 1$, then $\text{net}(\tau) = \tau^2/4$ and $S = 1/4$. Connecting the subperiod to an external cost measured in millions of euro requires keeping θ (or S) free. Appendix I also records an isomorphic technology with surplus $\theta\sqrt{\tilde{I}}$ and linear cost $\tilde{I}$. Both deliver $\text{net}(\tau) = \tau^2\theta^2/4$.

Location in the Bellman. Measure the cost of contracting out against the efficient net S. Then

$$\kappa^{external} = \text{net}(1) - \text{net}(\tau) = S\big(1 - \tau^2\big), \tag{12}$$
$$\kappa^{complement} = \text{net}(\tau') - \text{net}(\tau) = S\big((\tau')^2 - \tau^2\big), \tag{13}$$

where $\tau' = \tau + \delta$ is the firm's share when it also conducts internal R&D, with $\delta \geq 0$ the extra protection the lab buys (Cohen and Levinthal, 1989, 1990). Those two identities are definitions of the intercepts, not extra parameters.

The current-period payoffs in the Bellman are therefore

$$v^{external} = \text{net}(\tau) - \text{net}(1) = -\kappa^{external}, \tag{14}$$
$$v^{both} = -\kappa^{internal} + \text{net}(\tau') - \text{net}(1) = -\kappa^{internal} - \kappa^{external} + \kappa^{complement}. \tag{15}$$

The second line uses

$$\text{net}(\tau') - \text{net}(1) = \big(\text{net}(\tau') - \text{net}(\tau)\big) + \big(\text{net}(\tau) - \text{net}(1)\big) = \kappa^{complement} - \kappa^{external}.$$

The Bellman already contains those nets as the three estimated costs. Adding $\text{net}(\tau)$ again as a separate term would count the collaboration twice.

What the two costs do and do not identify. Equations (12)–(13) map three primitives (θ, τ, τ') into two intercepts. The Jacobian of that map has rank 2. Any surplus scale $S > \kappa^{external}$ delivers a different τ and a matching τ':

$$\tau = \sqrt{1 - \kappa^{external}/S}, \qquad \tau' = \sqrt{1 - \left(\kappa^{external} - \kappa^{complement}\right)/S}.$$

The likelihood sees the intercepts. I do not report a fitted bargaining share.

The ratio of the two costs does not depend on S:

$$r := \frac{\kappa^{complement}}{\kappa^{external}} = \frac{(\tau')^2 - \tau^2}{1 - \tau^2}. \tag{16}$$

On the whole set consistent with a given r, $\tau' \geq \sqrt{r}$. Internal R&D therefore buys a waiver of the hold-up tax, relative to whatever unprotected τ one is on, and the size of that waiver is a property of the two estimated costs. Section 5 reports the sample value. Section 5.4 plots the inversion over a range of S.

The estimated external cost is the firm's private gap relative to owning the collaboration. It splits into underinvestment, $S(1-\tau)^2$, plus a transfer to the partner, $2S\tau(1-\tau)$. The split needs S. The sum is already $\kappa^{external}$.

Proposition 1 (Observational equivalence) *Suppose (i) next period's productivity depends on the R&D mode, not on the equilibrium investment or the bargaining share; (ii) after the relationship-specific investment is sunk, disagreement payoffs are* $(0, 0)$*; (iii) the firm's equilibrium net from the collaboration is a mode-specific intercept; (iv) the Type-I extreme-value shocks sit on the four modes, and bargaining takes place only after the mode has been chosen. Then the alternative-specific values in* (5) *coincide with those of the Grossman–Hart subperiod once* (12) *and* (13) *hold. The choice probabilities therefore coincide. The likelihood recovers the three costs, not* τ*.*

Proof. See Appendix I. ■

I maintain those four restrictions, which are the ones already implicit in the estimated Bellman: productivity evolves on mode dummies, the three costs are constants, and bargaining takes place after the mode is chosen.

4 Empirical Implementation

In the previous section, I outlined the structural model, which comprises static decisions on production and dynamic decisions on R&D modes. To estimate the model, I proceed in stages. First, from the static part, I recover the demand elasticity η and estimate firm productivity ω_{it}, which allows me to compute static profits $\pi_{it}(s_{it})$. Next, I estimate the parameters governing productivity evolution. Finally, using these estimates, I recover the dynamic parameters: the investment costs for internal R&D ($\kappa^{internal}$), external R&D ($\kappa^{external}$), and the cost-complementarity parameter ($\kappa^{complement}$).

4.1 Retrieving Elasticity of Demand

The demand elasticity η is recovered directly from the data using the implication of the monopolistic competition model that the ratio of total variable costs (TVC) to revenue equals $1 + 1/\eta$.

Results 6 (Elasticity of demand) *The ratio of total variable cost (TVC) to the revenue (R) is defined as follows*

$$\frac{TVC}{R} = 1 + \frac{1}{\eta}$$

Proof. From the static model, revenue is $R_{it} = P_{it}Q_{it}$ and total variable costs are $\frac{1+\eta}{\eta}R_{it}$, yielding the ratio directly. ■

I compute the mean cost-revenue ratio across firms in the Dutch ICT industry to estimate $1 + 1/\eta$, and invert to obtain η. TVC is the sum of nominal intermediate inputs and energy costs, while revenue is nominal gross output.

Source of identification.—This parameter is identified from the markup implied by the monopolistic competition assumption, directly observable in the cost and revenue data. Alternatively, η can be estimated via regression of TVC on revenue with fixed effects, as in Aw et al. (2011); Boler et al. (2015); Peters et al. (2017):

$$TVC_{it} = R_{it}\left(1 + \frac{1}{\eta}\right) + \gamma_j + \gamma_t + \varepsilon_{it},$$

where γ_j and γ_t are sub-industry and time fixed effects.

4.2 Productivity Estimation

To recover productivity ω_{it}, I employ a semiparametric control function approach following Olley and Pakes (1996), Aw et al. (2011), Doraszelski and Jaumandreu (2013), and Peters et al. (2017). This method addresses the endogeneity of inputs correlated with unobserved productivity while incorporating the role of R&D modes in productivity evolution. The approach proceeds in two stages. In the first stage, I control for unobserved productivity in the revenue equation by using flexible inputs (intermediate materials m_{it} and energy n_{it}) as proxies, assuming they are chosen after ω_{it} and are strictly monotonic in it. This yields the fitted value $\hat{\phi}_{it}$, which combines the effects of capital, labor, and productivity scaled by the demand elasticity. In the second stage, I model productivity evolution as a controlled Markov process depending on lagged productivity and lagged R&D modes. This yields the following estimating equation of productivity.

Results 7 (Estimating equation of productivity) *The estimating equation to retrieve productivity ω_{it} takes the following form:*

$$
\begin{aligned}
\hat{\phi}_{it} = \beta_k^* k_{it} + \beta_l^* l_{it} + \alpha_1 \left(\hat{\phi}_{it-1} - \beta_k^* k_{it-1} - \beta_l^* l_{it-1}\right) \\
- \alpha_2^* \left(\hat{\phi}_{it-1} - \beta_k^* k_{it-1} - \beta_l^* l_{it-1}\right)^2 \\
+ \alpha_3^* \left(\hat{\phi}_{it-1} - \beta_k^* k_{it-1} - \beta_l^* l_{it-1}\right)^3 \\
- \alpha_4^* no_{it-1} - \alpha_5^* internal_{it-1} - \alpha_6^* external_{it-1} - \alpha_7^* both_{it-1} - \xi_{it}^*
\end{aligned}
\tag{17}
$$

where $\xi_{it}^* \sim \mathcal{N}(0, \sigma_{\xi^*})$. The above estimating equation will be estimated using non-linear least squares.

Proof. See Appendix D. ■

Using Results 6 and 7, I retrieve the productivity as follows

Results 8 (Estimated productivity) *The estimate of productivity can be expressed as the following*

$$
\hat{\omega}_{it} = -(1/(\hat{\eta}+1))\hat{\phi}_{it} + \hat{\beta}_k k_{it} + \hat{\beta}_l l_{it} \tag{18}
$$

Proof. See Appendix D. ■

Source of identification.—Identification follows the control function literature (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015; Gandhi et al., 2021). In the first stage, monotonicity ensures invertibility of ω_{it} from proxies m_{it}, n_{it}, given fixed k_{it}, l_{it}. Timing assumes k_{it}, l_{it} are predetermined (chosen before ω_{it} is realized), while m_{it}, n_{it} are flexible (chosen after ω_{it} is realized). In Dutch ICT, materials (e.g., software licenses, hardware components) and energy (e.g., server electricity) respond quickly to productivity shocks, validating the proxies. Scalar unobservability assumes no other unobservables drive inputs, which is reasonable in this homogeneous sector.

4.3 Costs of R&D and Complementarity

To recover $\kappa = (\kappa^{internal}, \kappa^{external}, \kappa^{complement})$, I use maximum likelihood on the choice probabilities (Equation 8):

$$
L(a_{it}|s_{it};\kappa) = \prod_i^N \prod_t^{T_i} Pr(a_{it} = action)^{I_{\{a_{it}=action\}}} \tag{19}
$$

There are two loops in estimating the dynamic discrete choice model. The maximum likelihood estimation serves as the outer loop in estimating Equation (19). Within each loop of the MLE, there is an inner loop involving the calculation of value functions. Essentially, it is a contraction mapping of the value function described in Equation (6). The algorithm to recover the structural parameters follows the successive-approximations variant of the nested fixed-point algorithm (NFXP-SA), as in Rust (1987)[15].

Source of identification.—As a single-agent model (monopolistic competition implies no strategic interactions), identification follows Rust (1987, 1994) and Aguirregabiria et al. (2021): additive Type-1 extreme-value shocks, conditional independence ($\omega_{it} \perp \epsilon_{it}$), normalized scale, fixed $\beta = 0.95$. Costs $\kappa^{internal}$ and $\kappa^{external}$ are identified from choice frequencies: the small share of external-only R&D identifies a high $\kappa^{external}$. Complementarity $\kappa^{complement}$ is identified from excess both-mode choices beyond additivity of the two costs. The Grossman–Hart subperiod in Section 3.2.4 is the theory of $\kappa^{external}$ and $\kappa^{complement}$ (Proposition 1). The objects I estimate remain the three costs. The bargaining share is a translation of those two intercepts once a surplus scale is chosen; I do not estimate it.

5 Results

5.1 Elasticity of Demand

I first estimate the demand elasticity η, a crucial input for recovering productivity and for profit calculations. As derived in Results 6, η is identified from the markup implied by monopolistic competition, observable in the ratio of total variable costs (TVC) to revenue. Table 4 reports estimates for the Dutch ICT industry, computed as the mean cost-revenue ratio inverted to yield $\eta \approx -1.8$ (standard error 0.16). This magnitude aligns with prior studies in innovation-intensive sectors (e.g., Aw et al. (2011); Peters et al. (2017)), indicating moderate price sensitivity consistent with differentiated ICT products like software and data services. By the CES demand, own-price and cross-price elasticities coincide. That is an implication of the demand system, not a finding that separates the two elasticities.

[15] There are several ways to improve the speed of NFXP-SA; the first one is to perform mathematical programming with equilibrium constraints (MPEC), developed by Su and Judd (2012). Another approach is to change the successive approach algorithm to the Newton-Kantorovich algorithm (NFXP-NK) developed by Iskhakov et al. (2016).

Table 4: Estimated Demand Elasticity

	Estimates
η	-1.794^{***} (0.164)

Notes.—The elasticity of demand, η, is derived from $TotalVariableCost/Revenue = 1 + 1/\eta$. The total variable cost consists of the sum of the firm's nominal intermediate inputs and energy. The revenue of the firm is nominal gross output. Standard errors of the demand elasticity are derived using the delta method.

5.2 Productivity Parameters

Productivity ω_{it} evolves as a controlled Markov process, shaped by lagged R&D modes, as estimated via nonlinear least squares in Results 7. Table 5 presents the parameters, revealing strong persistence ($\hat{\alpha}_1 = 0.704$, significant at 1%) and positive but similar impacts across R&D modes ($\hat{\alpha}_5$ to $\hat{\alpha}_8$ ranging from 0.345 to 0.390, all significant at 1%). Notably, both R&D ($\hat{\alpha}_8 = 0.387$) lies between internal-only ($\hat{\alpha}_6 = 0.384$) and external-only ($\hat{\alpha}_7 = 0.390$), and the differences are small, suggesting limited differential effects from combining modes.

Table 5: Estimated Productivity Parameters

Parameters	**Productivity Estimation Dependent var: ϕ**	**Parameters**	**Productivity Estimation Dependent var: ϕ**
$\beta_k(k)$	-0.193^{***} (0.012)	$\alpha_5(no\_rd)$	0.345^{***} (0.084)
$\beta_l(l)$	-0.650^{***} (0.025)	$\alpha_6(internal)$	0.384^{***} (0.084)
$\alpha_1(\omega)$	0.704^{***} (0.077)	$\alpha_7(external)$	0.390^{***} (0.102)
$\alpha_2(\omega^2)$	0.051^{*} (0.024)	$\alpha_8(both)$	0.387^{***} (0.085)
$\alpha_3(\omega^3)$	-0.002 (0.002)		
Subindustry FE	Yes		
Year FE	Yes		
$SD(\xi)$	0.347		
Wald p-value of $\Delta = 0$	0.54		

Notes.—Significance level: $^{***}: p < 0.001$, $^{**}: p < 0.01$, $^{*}: p < 0.05$, $^{\cdot}: p < 0.1$. Standard errors are in parentheses. I use heteroskedasticity and autocorrelation-consistent (HAC) standard errors. $SD(\xi)$ indicates the standard deviation of the residuals from the estimating Equation (38). The last row is a Wald p-value of equal incremental returns on $\alpha_5, \ldots, \alpha_8$.

The supermodularity contrast $(\alpha_8 - \alpha_6) - (\alpha_7 - \alpha_5)$ equals -0.042. A Wald test of H_0: that contrast is zero has $p = 0.54$. The return to R&D is similar across modes. The mode intercepts in Table 5 come from the $\hat{\phi}$ equation. Equation (18) recovers productivity as $\hat{\omega}_{it} = -(1/(\hat{\eta}+1))\hat{\phi}_{it} + \hat{\beta}_k k_{it} + \hat{\beta}_l l_{it}$, so a difference in those intercepts maps into ω by the factor $-1/(\eta+1)$. Shifting from no R&D raises ω by $(\alpha_5 - \alpha_6)/(\eta+1) \approx 0.049$ for internal-only, $(\alpha_5 - \alpha_7)/(\eta+1) \approx 0.057$ for external-only, and $(\alpha_5 - \alpha_8)/(\eta+1) \approx 0.053$ for both, given $\eta = -1.794$ in Table 4. Because ω is in logs, those shifts are about 5.0%, 5.8%, and 5.4% on average [$\exp(0.049) - 1$, $\exp(0.057) - 1$, and $\exp(0.053) - 1$], aligning with literature on R&D returns (Hall and Lerner, 2010; Mohnen et al., 2017). Combining the two modes does not add a distinguishable extra productivity increment beyond those similar level shifts. That finding matches the contracting subperiod in Section 3.2.4, where productivity continues to evolve on the mode dummies and the friction lives in the current-period cost of contracting out. The mode mix is then left to be explained by those costs. The complementarity I recover in the next subsection is a cost object: doing both is cheaper than the sum of the two modes. Goyal et al. (2008) show that complementarity in the firm's objective can arise from cost advantages alone. That is the structure estimated here: the productivity intercepts are similar, and $\kappa^{complement}$ makes doing both cheaper. In terms of magnitudes of β_k and β_l, the estimated parameters are quite similar to the estimates of the elasticity of output to capital and labor in recent literature[16].

Returning to the previous example of a Dutch ICT firm: Internal R&D might build the algorithm's core logic, boosting ω_{it} via firm-specific knowledge, while external consultations resolve novel bottlenecks, adding incremental gains. Combining both can accelerate adoption. The estimated productivity differences across modes remain small.

The estimated parameters shown in Table 5 will be used to construct the (discretized) productivity of the dynamic structural model. The standard deviation from the (predicted) residuals, shown in the second to the last row of Table 5, will also be used to construct the transition probability.

5.3 Costs of Innovation and Complementarity

Dynamic parameters are recovered by maximum likelihood on choice probabilities (Equation 8), using nested fixed-point iteration (Rust, 1987). Table 6 reports estimates for internal ($\kappa^{internal}$), external ($\kappa^{external}$), and complementarity ($\kappa^{complement}$) costs, across the full sample (2000–2020), pre-expansion (2000–2008), and post-expansion (2012–2020) periods[17].

[16] See for example De Loecker and Warzynski (2012), De Loecker et al. (2020), and Dobbelaere and Mairesse (2013). The negative signs on β_k and β_l arise because my model starts from the cost side (Equation 2), where higher capital and labor inputs reduce marginal costs, yielding negative coefficients. In contrast, revenue- or production-based models (e.g., De Loecker (2011); De Loecker and Warzynski (2012); De Loecker et al. (2020)) estimate positive output elasticities, as inputs directly increase revenue or output. These approaches are dual: Under cost minimization and production theory duality (e.g., Varian (1992)), the absolute magnitudes reflect equivalent input contributions. For instance, $|\hat{\beta}_k| \approx 0.193$ and $|\hat{\beta}_l| \approx 0.650$ align with positive elasticities of 0.1–0.3 for capital and 0.5–0.7 for labor in those studies, capturing diminishing returns in ICT after accounting for R&D-driven efficiencies.

[17] I make the decision that the post-expansion period begins in 2012 to account for policy diffusion and adoption lags following the 2010 Innovation Box expansion. During the initial years, awareness and utilization varied: some

Table 6: Costs of Innovation and Cost-Complementarity

Parameters	Costs of R&D and Complementarity		
	Full Period: 2000 - 2020	**Before Innovation Box: 2000 - 2008**	**After Innovation Box: 2012 - 2020**
$\kappa^{internal}$	0.127^{***} (0.049)	0.569^{***} (0.096)	-0.063 (0.062)
$\kappa^{external}$	2.453^{***} (0.120)	2.239^{***} (0.186)	2.695^{***} (0.177)
$\kappa^{complement}$	2.295^{***} (0.132)	2.153^{***} (0.217)	2.502^{***} (0.188)
log-likelihood	-2836.652	-786.940	-1784.994

Notes.—The estimates are in millions of EUR (real). The gradients and Hessian are obtained by finite-difference approximation evaluated at the estimated parameters. Gradients are below $1e-4$. Tilt and radius of curvature for these parameters are below $1e-3$, computed by axial search at each parameter. Standard errors are obtained by computing the square root of the diagonal of the inverse Hessian.

The estimates reveal a large cost asymmetry. External R&D costs range from 2.239 to 2.695 million EUR, compared with 0.127 to 0.569 million EUR for internal R&D in the full sample and the pre-expansion period. Before the Innovation Box expansion, contracted-out R&D is about four times as costly as in-house R&D (2.239/0.569). In the full sample the ratio is larger because in-house costs are already low. That gap is why the share of external-only firms is about 3 percent.

The estimates map onto the contracting subperiod of Section 3.2.4 through the Bellman, without a second estimation. A firm that contracts out alone keeps share τ and earns $\text{net}(\tau) = S\tau^2$ (Equations 9–11). The cost of that mode is the gap to the efficient net,

$$\kappa^{external} = \text{net}(1) - \text{net}(\tau) = S\big(1 - \tau^2\big).$$

When the firm also runs a lab, its share rises to τ' and

$$\kappa^{complement} = \text{net}(\tau') - \text{net}(\tau) = S\big((\tau')^2 - \tau^2\big).$$

The current-period payoffs are then $v^{external} = -\kappa^{external}$ and $v^{both} = -\kappa^{internal} + \text{net}(\tau') - \text{net}(1) = -\kappa^{internal} - \kappa^{external} + \kappa^{complement}$, which are the terms already in (5). The lab itself still costs $\kappa^{internal}$. I do not recover τ or S as parameters. I read the estimated costs as those two gaps.

That reading fits the numbers. Complementarity is positive and of the same order as the external cost ($\kappa^{complement}$ from 2.153 to 2.502 million EUR). The net cost of both is $\kappa^{internal} + \kappa^{external} - \kappa^{complement} = 0.285$ million EUR in the full sample, 0.655 million EUR before the expansion, and 0.130 million EUR after. Treating the statistically insignificant post-expansion $\kappa^{internal}$ as zero, as I

firms quickly adapted, while others faced delays in information dissemination, application processing, or internal adjustments. By 2012, widespread knowledge and implementation enabled fuller participation, as evidenced by gradual increases in WBSO and Patent Box policy take-up among small firms from 2006 to 2010 (Verhoeven and Stel, 2013), allowing a cleaner assessment of long-term effects (Mohnen et al., 2017).

do when describing policy, that net cost is 0.193 million EUR. Those net costs sit next to in-house R&D alone. In the full sample, complementarity offsets $2.295/2.453 \approx 0.94$ of the extra cost of contracting out, so doing both is almost as cheap as doing in-house R&D by itself. Contracting out alone remains expensive, so only about 3 percent of firms choose external-only. Contracting out together with in-house R&D is cheap enough that about 28 percent choose both. The 2.45 million EUR is the firm's private gap relative to owning the collaboration. Part of it is underinvestment and part is a transfer to the partner; the split depends on S, which the sample does not identify. Haggling over an incomplete contract (Bajari and Tadelis, 2001; Tadelis, 2002) produces the same pair of costs: in-house R&D is what makes the external agreement cheap enough to use. The three columns of Table 6 give three versions of this reading, which I compare in Section 5.4.

The post-expansion internal-cost estimate is -0.063 million EUR and is statistically indistinguishable from zero.[18] I treat that estimate as about zero when describing policy.

Related work on absorptive capacity finds that a lab raises the return to external knowledge (Cassiman and Veugelers, 2006; Cohen and Levinthal, 1990; Schmiedeberg, 2008; Lokshin et al., 2008). In the Dutch ICT example, a firm can build a dynamic pricing algorithm in-house and consult university researchers on the remaining bottlenecks. In the contracting subperiod, the algorithm is $\kappa^{internal}$; the consultation is the relationship-specific I; the lab raises the firm's share from τ to τ' and thereby creates $\kappa^{complement}$. After the Innovation Box expansion, in-house costs fall to about zero, while external costs rise to 2.695 million EUR. That is WBSO on the lab's wage bill and Nexus on the external contract, sitting on top of the hold-up pattern. Complementarity remains of the same order as the external cost, so the lab continues to undo most of the friction once it is in place.

5.4 Implied bargaining shares given a surplus scale

Section 3.2.4 defined S as the surplus the firm would collect if it owned the collaboration outright. The sample does not identify S. For any S larger than the estimated external cost, the two identities in that section invert to bargaining shares

$$\tau = \sqrt{1 - \frac{\kappa^{external}}{S}}, \qquad \tau' = \sqrt{1 - \frac{\kappa^{external} - \kappa^{complement}}{S}}. \tag{20}$$

Here τ is the firm's share of the collaboration when it contracts out without doing in-house R&D, and τ' is its share when it does both.

One comparison does not require a value of S. The ratio $r = \kappa^{complement}/\kappa^{external}$ is the fraction of the extra cost of contracting out that is offset when the firm also does in-house R&D. Table 7

[18]The statistically insignificant $\kappa^{internal}$ partly results from reduced variation in R&D choices post-expansion. As shown in Fact 2, the share of firms choosing no R&D drops significantly after 2010, leaving the model to rely on subtler shifts among R&D-performing firms. This lower variation weakens identification, attenuating the estimate and inflating standard errors.

reports that ratio in each column of Table 6: 0.962 before the Innovation Box expansion, 0.936 in the full sample, and 0.928 after. In every window, doing both undoes almost all of the extra cost of going outside. In the hold-up algebra, the same ratio also bounds τ' from below by $\sqrt{r}$. That lower bound is 0.981 before the expansion, 0.967 in the full sample, and 0.964 after. If the estimated costs come from hold-up, a firm that already does in-house R&D keeps almost the entire surplus of the outside contract, in every sample.

What did change across columns is the cost of contracting out *alone*. That cost rose from 2.239 million EUR before the expansion to 2.695 million EUR after. Complementarity rose with it, which is why the ratio r stayed high: going outside without in-house R&D became more expensive, but going outside *with* in-house R&D remained cheap. The cost of in-house R&D itself, $\kappa^{internal}$, is a different parameter. It falls to about zero after 2012, which I already read as WBSO on the wage bill of the firm's own researchers. The formulas for τ and τ' use only the external cost and complementarity. They are silent on that in-house cost.

Table 7: How much of the external cost is offset by also doing in-house R&D

Sample window	$\kappa^{external}$	$\kappa^{complement}$	r	Lowest τ'
Before Innovation Box, 2000–2008	2.239	2.153	0.962	0.981
Full period, 2000–2020	2.453	2.295	0.936	0.967
After Innovation Box, 2012–2020	2.695	2.502	0.928	0.964

Notes. Costs are the point estimates from Table 6, in million EUR. $r = \kappa^{complement}/\kappa^{external}$ is the share of the extra cost of contracting out that complementarity offsets. The last column is $\sqrt{r}$, the lowest bargaining share when the firm does both that is consistent with that r (Section 3.2.4). These numbers are transformations of the estimated costs. They are not a second estimation, and they have no standard error of their own.

Saying anything about τ, the share when the firm contracts out alone, does require a value of S. Figure 3 traces (20) for each sample window. Each curve starts at that window's $\kappa^{external}$, because a surplus smaller than the estimated external cost cannot generate it. Putting all three curves on the same plot therefore requires $S > 2.695$. At $S = 5$, implied τ is 0.74 before the expansion, 0.71 in the full sample, and 0.68 after. The share of the external cost that is underinvestment is then 0.15, 0.17, and 0.19; the remainder is a payment to the partner. Those numbers treat S as the same in all three windows. If it was, the higher external cost after 2012 would mean a smaller share for a firm that contracts out without in-house R&D. The sample does not say whether S itself moved with the Innovation Box, so I do not read that tilt as an estimate of how policy changed bargaining power. The ratio in Table 7 does not need that assumption, and it is stable.

I do not pick a preferred S, and I do not put a standard error on τ in the cost table.

The next section shuts off $\kappa^{complement}$ and re-solves the model. In the contracting subperiod that restriction is $\tau' = \tau$: the firm keeps only τ even when it also does in-house R&D, so doing both loses its cost advantage. The joint mode should then collapse.

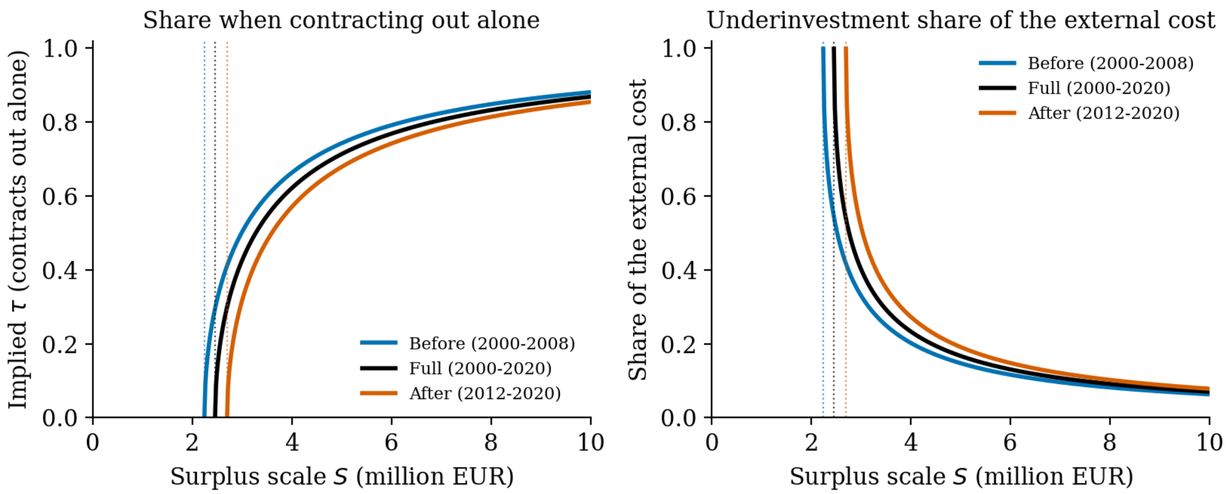

Figure 3: Left: bargaining share when the firm contracts out without in-house R&D. Right: share of the estimated external cost that is underinvestment. Each curve uses the $(\kappa^{external}, \kappa^{complement})$ pair from Table 6 for that window. Vertical dotted lines mark the three $\kappa^{external}$ thresholds. The inversion is undefined for $S \leq \kappa^{external}$.

6 Simulation Analysis

This section uses the estimated model to evaluate counterfactual R&D subsidy policies inspired by the Dutch Tax Incentives for Innovation scheme. I focus on cost-side interventions, modeling subsidies as reductions in R&D investment costs ($\kappa^{internal}$, $\kappa^{external}$). The questions are how a uniform subsidy and a mode-specific subsidy affect R&D participation and firm value, and what role cost-complementarity plays. The Dutch scheme tilts toward in-house R&D under the OECD Nexus Approach, which limits benefits for outsourced activities (OECD, 2015).

I compute counterfactuals by re-solving the model under modified costs, iterating to a new steady-state distribution of states and choices. I report changes in R&D shares and in firm value, $\Delta V^c = (V^c - V)/V$, where V is the ex-ante expected value (McFadden's social surplus) from Equation (6). This is producer surplus in partial-equilibrium monopolistic competition.

I consider three subsidy scenarios, calibrated so that fiscal bite is comparable: (1) a uniform 5% reduction across all R&D costs; (2) a 20% reduction in internal costs only, in the spirit of the scheme's in-house tilt; and (3) a 5% reduction in external costs only. The Dutch scheme does not subsidize contracted-out R&D. The external-only case is a symmetric comparison. External costs are much larger than internal costs, so a 5% external cut is larger in euros than a 20% internal cut under the full-sample estimates. The percentages follow the policy analogy: a uniform cut, an in-house tilt, and a contracted-out alternative. Through the contracting subperiod of Section 3.2.4, a cut in $\kappa^{external}$ is a cut in $S(1-\tau^2)$, a cut in $\kappa^{internal}$ is a cheaper lab, and shutting off $\kappa^{complement}$ is $\tau' = \tau$ (the lab no longer raises the firm's share of the collaboration). I re-solve each policy under that last restriction as the mechanism check from Section 5. The reported experiments are operations on the estimated intercepts. They do not require a value of τ.

6.1 Change in the Share of R&D activities

Table 8 reports changes in steady-state R&D shares relative to baseline (actual shares: 36.8% no R&D, 32.4% internal-only, 3.2% external-only, 27.7% both). The uniform 5% subsidy (Case 1) yields the largest increase in R&D-active firms, reducing the no-R&D share by 1.6 percentage points (pp) while boosting both (2.4 pp) and external-only (0.2 pp). Through the contracting subperiod, that policy cheapens the lab and the external contract together, so firms can pay $\kappa^{internal}$ and still capture τ' of the collaboration.

Table 8: Change in the share of R&D activities

Change in the Share of R&D Activities	**Actual Share**	$\Delta^c = share^c - share$					
		5% uniform		**20% internal**		**5% external**	
		Base	***No Comp.***	***Base***	***No Comp.***	***Base***	***No Comp.***
No R&D	0.368	-0.016	0.115	-0.006	0.116	-0.015	0.116
Internal-only	0.324	-0.012	0.104	0.003	0.113	-0.012	0.103
External-only	0.032	0.002	0.015	-0.001	0.01	0.002	0.015
Both	0.277	0.024	-0.235	0.003	-0.239	0.024	-0.235

Notes.—Base: changes only to the costs. No Comp: Set the cost-complementarity parameter to zero, i.e., $\kappa^{complement} = 0$. Negative sign indicates ”decreasing”.

The internal-only 20% subsidy (Case 2) has small effects: the no-R&D share falls by 0.6 pp, with modest gains in internal-only (0.3 pp) and both (0.3 pp), and a slight decline in external-only (-0.1 pp). Baseline $\kappa^{internal}$ is already low (0.127 million EUR), so cheapening the lab does little to the hold-up tax on the contract. The external-only 5% subsidy (Case 3) ranks second, after the uniform policy, cutting the no-R&D share by 1.5 pp and raising both by 2.4 pp. That subsidy acts on $\kappa^{external} = S(1-\tau^2)$, and thereby makes the contract usable for firms that already have, or can add, a lab.

When cost-complementarity is shut off, the uniform subsidy raises the no-R&D share by 11.5 percentage points and the internal-only share by 10.4 percentage points, while the joint-mode share falls by 23.5 percentage points. The same collapse of the joint mode appears under the other two subsidies. In the contracting subperiod, this is $\tau' = \tau$: in-house R&D no longer raises the firm's share of the collaboration, so $\kappa^{complement}$ no longer offsets $\kappa^{external}$, and doing both is once again the sum of a lab and an expensive contract. Firms drop the joint mode. That is the mechanism check of Section 5.

6.2 Change in Firm's Valuation

Table 9 shows average firm value changes (ΔV^c), capturing producer surplus gains from higher productivity and profits. The uniform subsidy yields the highest gains (2.74%). It cheapens both the lab and the contract, so firms can use τ' at a lower total cost. The external-only subsidy follows at 2.49%, because it cuts the hold-up tax that keeps contracting out rare. The internal-only subsidy lags at 0.97%, because it cheapens an already inexpensive lab and leaves τ unchanged.

Table 9: Change in Value

	$\Delta V^c = (V^c - V)/V$					
	5% uniform		**20% internal**		**5% external**	
Change in	***Base***	***No Comp.***	***Base***	***No Comp.***	***Base***	***No Comp.***
Value	2.735%	1.022%	0.968%	0.923%	2.49%	0.792%

Notes.—Base: changes only to the costs. No Comp: Set the cost-complementarity to zero. Negative sign indicates "decreasing".

With cost-complementarity shut off, the value gains fall to 1.02%, 0.92%, and 0.79% under the uniform, internal-only, and external-only subsidies. Matching the baseline uniform gain would then require a much larger subsidy. When $\tau' = \tau$, the extra term $\text{net}(\tau') - \text{net}(\tau)$ vanishes, so a given cut in $\kappa^{external}$ no longer opens the joint mode and the productivity transition that the baseline assigns to doing both.

A uniform subsidy therefore raises participation and firm value by more than a subsidy aimed only at in-house R&D, in line with evidence on neutral R&D policies (Galaasen and Irarrazabal, 2021; Akcigit et al., 2021, 2022). The Dutch scheme's in-house tilt leaves the hold-up tax on contracted-out R&D in place, which is consistent with the flattening of R&D shares after 2016 (Fact 2). When $\tau' = \tau$, those policy gains shrink.

A policy that instead raised the firm's bargaining share would, given an S, move $\kappa^{external} = S(1 - \tau^2)$ and $\kappa^{complement} = S((\tau')^2 - \tau^2)$ together (Section 5.4). That mapping is well defined once S is chosen, and reporting new mode shares or firm value would require re-solving the dynamic program. I do not run that experiment.

7 Conclusion

This paper analyzes the costs and benefits of in-house and contracted-out R&D. I first document stylized facts on the four-way R&D menu in the Dutch ICT industry. I then build and estimate a dynamic discrete-choice model that recovers the cost of each mode, including the extra cost of contracting out and the saving from doing both.

The cost of contracted-out R&D is much higher than the cost of in-house R&D, which is why so few firms contract out alone. Doing both is cheaper than the sum of the two costs: the extra cost of contracting out is largely waived when the firm also has a lab. Productivity gains are similar across modes, so the mix of activities is driven by these costs. A hold-up problem on the external contract, together with absorptive capacity from in-house R&D, is a leading microfoundation of that pattern. Haggling over an incomplete contract produces the same cost shape. After the Innovation Box expansion, in-house costs fall and contracted-out costs rise, which is the policy layer sitting on top of the transaction-cost pattern.

A uniform cost subsidy raises R&D participation and firm value by more than a subsidy aimed only at in-house R&D. When cost-complementarity is shut off, the share of firms that do both collapses and the value gains from the subsidies shrink. The waiver of the external cost is what sustains these responses.

I would like to complement this ex-ante evaluation with an ex-post evaluation of the Dutch Tax Incentives for Innovation scheme, and to study whether the scheme creates incentives for firms to relabel other spending as innovation-related.

References

J. Abbring and O. Daljord. Identifying the discount factor in dynamic discrete choice models. *Quantitative Economics*, 2020.

D. A. Ackerberg, K. Caves, and G. Frazer. Identification properties of recent production function estimators. *Econometrica*, 83(6), 2015.

P. Aghion and J. Tirole. The management of innovation. *The Quarterly Journal of Economics*, 109(4): 1185–1209, 1994.

V. Aguirregabiria, A. Collard-Wexler, and S. P. Ryan. Dynamic games in empirical industrial organization. *Handbook of Industrial Organization, Vol IV*, 2021.

U. Akcigit and W. R. Kerr. Growth through heterogeneous innovations. *The Journal of Political Economy*, 2018.

U. Akcigit, D. Hanley, and N. Serrano-Velarde. Back to basics: Basic research spillovers, innovation policy, and growth. *Review of Economic Studies*, 2021.

U. Akcigit, D. Hanley, and S. Stantcheva. Optimal taxation and r&d policies. *Econometrica*, 2022.

S. Amoroso. Multilevel heterogeneity of r&d cooperation and innovation determinants. *Eurasian Business Review*, 2017.

S. P. Anderson, A. De Palma, and J.-F. Thisse. *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, MA, 1992.

A. Arora. Testing for complementarities in reduced-form regressions: A note. *Economics Letters*, 1996.

S. Athey and S. Stern. An empirical framework for testing theories about complementarity in organizational design. *NBER Working Paper*, 1998.

B. Y. Aw, M. J. Roberts, and D. Y. Xu. R&d investment, exporting, and productivity dynamics. *American Economic Review*, 101(4), 2011.

P. Bajari and S. Tadelis. Incentives versus transaction costs: A theory of procurement contracts. *The RAND Journal of Economics*, 2001.

L. Berchicci. Towards an open r&d system: Internal r&d investment, external knowledge acquisition and innovative performance. *Research Policy*, 2013.

S. Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 1994.

S. Berry, J. Levinsohn, and A. Pakes. Automobile prices in market equilibrium. *Econometrica*, 1995.

L. K. Bilir and E. Morales. Innovation in the global firm. *The Journal of Political Economy*, 2020.

N. Bloom, R. Griffith, and J. Van Reenen. Do r&d tax credits work? evidence from a panel of countries 1979–1997. *Journal of Public Economics*, 85(1):1–31, 2002.

E. A. Boler, A. Moxnes, and K. H. Ulltveit-Moe. R&d, international sourcing, and the joint impact on firm performance. *American Economic Review*, 2015.

B. Cassiman and R. Veugelers. In search of complementarity in innovation strategy: Internal r&d and external knowledge acquisition. *Management Science*, 2006.

M. Ceccagnoli and L. Jiang. The cost of integrating external technologies: Supply and demand drivers of value creation in the markets for technology. *Strategic Management Journal*, 2013.

M. Ceccagnoli, M. J. Higgins, and P. Vincenzo. Behind the scenes: Sources of complementarity in r&d. *Journal of Economics & Management Strategy*, 2014.

R. H. Coase. The nature of the firm. *Economica*, 1937.

W. M. Cohen and D. A. Levinthal. Innovation and learning: The two faces of R&D. *The Economic Journal*, 99(397):569–596, 1989.

W. M. Cohen and D. A. Levinthal. Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 1990.

J. De Loecker. Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica*, 2011.

J. De Loecker and F. Warzynski. Markups and firm-level export status. *American Economic Review*, 102(6), 2012.

J. De Loecker, J. Eeckhout, and G. Unger. The Rise of Market Power and the Macroeconomic Implications*. *The Quarterly Journal of Economics*, 135(2):561–644, 01 2020. ISSN 0033-5533. doi: 10.1093/qje/qjz041. URL `https://doi.org/10.1093/qje/qjz041`.

Dialogic, PwC, and Berenschot. Evaluatie wet bevordering speur- en ontwikkelingswerk (wbso) 2018-2022. Technical report, Ministerie van Economische Zaken en Klimaat, 2025. URL `https://open.overheid.nl/documenten/06f98f9c-f172-4a41-88ae-38390ff74f12/file`.

S. Dobbelaere and J. Mairesse. Panel data estimates of the production function and product and labor market imperfections. *Journal of Applied Econometrics*, 2013.

U. Doraszelski and J. Jaumandreu. R&D and Productivity: Estimating Endogenous Productivity. *The Review of Economic Studies*, 80(4), 2013.

S. Galaasen and A. Irarrazabal. R&d heterogeneity and the impact of r&d subsidies. *The Economic Journal*, 2021.

A. Gandhi, S. Navarro, and D. A. Rivers. On the identification of gross output production functions. *The Journal of Political Economy*, 2021.

X. Gonzalez, J. Jaumandreu, and C. Pazo. Barriers to innovation and subsidy barriers. *The RAND Journal of Economics*, 2005.

S. Goyal, J. L. Moraga-Gonzalez, and A. Konovalov. Hybrid r&d. *Journal of the European Economic Association*, 2008.

C. Grimpe and U. Kaiser. Balancing internal and external knowledge acquisition: The gains and pains of r&d outsourcing. *Journal of Management Studies*, 47(8):1483–1509, 2010.

S. Grossman and O. Hart. The costs and benefits of ownership: A theory of vertical and lateral integration. *The Journal of Political Economy*, 1986.

J. Hagedoorn and N. Wang. Is there complementarity or substitutability between internal and external r&d strategies? *Research Policy*, 2012.

B. H. Hall and J. Lerner. The financing of r&d and innovation. *Handbook of the Economics of Innovation*, 2010.

B. H. Hall and J. Van Reenen. How effective are fiscal incentives for r&d? a review of the evidence. *Research Policy*, 29(4-5):449–469, 2000.

O. Hart and J. Moore. Property rights and the nature of the firm. *The Journal of Political Economy*, 98(6):1119–1158, 1990.

J. Howells. New directions in r&d: current and prospective challenges. *R&D Management*, 2008.

M. Igami. Estimating the innovator's dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981-1988. *The Journal of Political Economy*, 2017.

M. Igami and K. Uetake. Mergers, innovation, and entry-exit dynamics: Consolidation of the hard disk drive industry, 1996-2016. *The Review of Economic Studies*, 2019.

F. Iskhakov, J. Lee, J. Rust, B. Schjerning, and K. Seo. Comment on 'constrained optimization approaches to estimation of structural models'. *Econometrica*, 84(1), 2016.

T. Jungbauer, S. Nicholson, J. Pan, M. Waldman, and L. X. Wang. The organization of innovation: Incomplete contracts and the outsourcing decision. *American Economic Journal: Microeconomics*, 18(2):56–109, 2026.

E. Kalai. Proportional solutions to bargaining situations: Interpersonal utility comparisons. *Econometrica*, 45(7):1623–1630, 1977.

J. Levinsohn and A. Petrin. Estimating Production Functions Using Inputs to Control for Unobservables. *The Review of Economic Studies*, 70(2), 2003.

B. Lokshin and P. Mohnen. How effective are level-based r&d tax credits? evidence from the netherlands. *Applied Economics*, 44(12):1527–1538, 2012.

B. Lokshin, R. Belderbos, and M. Carree. The productivity effects of internal and external r&d: Evidence from a dynamic panel data model. *Oxford Bulletin of Economics and Statistics*, 70(3): 399–413, 2008.

J. H. Love, S. Roper, and P. Vahter. Dynamic complementarities in innovation strategies. *Research Policy*, 2014.

P. Milgrom and J. Roberts. The economics of modern manufacturing: Technology, strategy, and organization. *American Economic Review*, 1990.

P. Milgrom and J. Roberts. Complementarities and fit: Strategy, structure and organizational change in manufacturing. *Journal of Accounting and Economics*, 1995.

E. J. Miravete and J. C. Pernias. Innovation compelementarity and scale of production. *The Journal of Industrial Economics*, 2006.

P. Mohnen and L.-H. Roeller. Complementarities in innovation policy. *European Economic Review*, 2005.

P. Mohnen, A. Vankan, and B. Verspage. Evaluating the innovation box tax policy instrument in the netherlands, 2007–13. *Oxford Review of Economic Policy*, 2017.

J. F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

OECD. Countering harmful tax practices more effectively, taking into account transparency and substance, action 5 - 2015 final report. *OECD Report*, 2015.

G. S. Olley and A. Pakes. The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6), 1996.

B. Peters, M. J. Roberts, V. A. Vuong, and H. Fryges. Estimating dynamic r&d choice: an analysis of costs and long-run benefits. *The RAND Journal of Economics*, 48(2), 2017.

J. Rust. Optimal replacement of GMC buses: An empirical model of Harold Zurcher. *Econometrica*, 55(5), 1987.

J. Rust. Structural estimation of markov decision processes. *Handbook of Econometrics*, 1994.

C. Schmiedeberg. Complementarities of innovation activities: An empirical analysis of the german manufacturing sector. *Research Policy*, 2008.

C.-L. Su and K. L. Judd. Constrained optimization approaches to estimation of structural models. *Econometrica*, 80(5), 2012.

S. Tadelis. Complexity, flexibility, and the make-or-buy decision. *American Economic Association: Papers and Proceedings*, 2002.

J.-F. Thisse and P. Uschev. Monopolistic competition without apology. *Handbook of Game Theory and Industrial Organization, Vol I*, 2018.

M. van Economische Zaken. Beleidsnota digitale agenda 2020. Technical report, Rijksoverheid, 2017. URL `https://open.overheid.nl/documenten/ronl-71c5e1c9-04ba-49d7-9058-b480a1447be0/pdf`.

H. R. Varian. *Microeconomic Analysis*. W. W. Norton & Company, New York, 3 edition, 1992. ISBN 978-0393957358.

W. Verhoeven and A. Stel. Hoe effectief is de wbso-regeling door de jaren heen? *Beleidsonderzoek Online*, 2013.

R. Veugelers. Internal r&d expenditures and external technology sourcing. *Research Policy*, 26(3): 303–315, 1997.

R. Veugelers and B. Cassiman. Make and buy in innovation strategies: evidence from belgian manufacturing firms. *Research Policy*, 28(1):63–80, 1999.

O. E. Williamson. The economics of organization: The transaction cost approach. *American Journal of Sociology*, 1975.

E. Zhelobodko, S. Kokovin, M. Parenti, and J.-F. Thisse. Monopolistic competition: Beyond the constant elasticity of substitution. *Econometrica*, 2012.

A Probit: Persistence of R&D modes

Table 10: Probit: Persistence of R&D modes

Probit	*Dep. var*			
	No R&D (t)	**Internal-only (t)**	**External-only (t)**	**Both (t)**
No R&D (t-1)	1.171*** (0.0895)			
Internal-only (t-1)		0.7635*** (0.0950)		
External-only (t-1)			1.198*** (0.2105)	
Both (t-1)				1.017*** (0.0821)
Year FE	Yes	Yes	Yes	Yes
Sub-Industry FE	Yes	Yes	Yes	Yes

B Details on the Reduced-form Testing of Complementarity between Internal and External R&D

I now provide a simple, reduced-form test for complementarity between $InternalR\&D$ and $ExternalR\&D$. To do so, I follow similar approaches used by Arora (1996), Athey and Stern (1998), and Cassiman and Veugelers (2006). First, let us define Π as a firm's profit function that is assumed to be supermodular. The arguments of this profit function are $InternalR\&D$ and $ExternalR\&D$. Then, as shown by Milgrom and Roberts (1990, 1995), $InternalR\&D$ and $ExternalR\&D$ are complementary only if

$$\Pi(\text{both R\&D}) - \Pi(\text{external}) \geq \Pi(\text{internal}) - \Pi(\text{no R\&D}) \tag{21}$$

The main intuition of Equation (21) is that adding an activity while the other activity is already being performed has a higher incremental effect on profits than adding the activity in isolation. Here, performing external R&D on top of internal R&D can only be deemed more profitable (compared to profits from performing *only* internal R&D) if the above condition holds.

In the empirical implementation, I run the following specification

$$\begin{aligned} y_{ijt} &= \theta_{\text{both}} BothR\&D_{ijt} + \theta_{\text{internal}} InternalR\&D_{ijt} \\ &+ \theta_{\text{external}} ExternalR\&D_{ijt} + \theta_{\text{no}} NoR\&D_{ijt} + \zeta_t + \zeta_j + \varepsilon_{ijt} \end{aligned} \tag{22}$$

where the dependent variable y_{ijt} is either value-added or the share of new products produced by each firm in a given year. ζ_t indicates year fixed effects. Since the share of new products in revenue ranges from 0 to 100, a standard OLS will yield biased estimates. Given this concern, and the ability to exploit the panel structure of the model, I estimate Equation (22) using a Panel (Random Effect) Tobit model.

Testing for the complementarity boils down to checking if

$$\theta_{\text{both}} - \theta_{\text{internal}} \geq \theta_{\text{external}} - \theta_{\text{no}} \tag{23}$$

Complementarity is $\Delta \equiv (\theta_{\text{both}} - \theta_{\text{internal}}) - (\theta_{\text{external}} - \theta_{\text{no}}) \geq 0$. The last row of Table 3 is a Wald test of H_0: $\Delta = 0$, using the χ^2 statistic of that linear restriction and the sign of $\widehat{\Delta}$. The point estimates are $\widehat{\Delta} = -0.200$ for log value added, $\widehat{\Delta} = 4.74$ for the share of new products under OLS, and $\widehat{\Delta} = -0.57$ under the random-effects Tobit. None of the three tests rejects $\Delta = 0$. The same contrast on the productivity intercepts in Table 5 is $\widehat{\Delta}^{\omega} = (\alpha_8 - \alpha_6) - (\alpha_7 - \alpha_5) = -0.042$, with $p = 0.54$. Complementarity in the dynamic model is a cost object, identified from choice frequencies, not from these payoff-side contrasts.

C Details on the Derivation of the Static Part

Proof of Result 1: Revenue Equation.. First, recall the demand function is the following:

$$\begin{aligned} Q_{it} &= Q_t \left(\frac{P_{it}}{P_t} \right)^{\eta} \\ &= \frac{Q_t}{P_t^{\eta}} P_{it}^{\eta} \\ &= \Phi_t P_{it}^{\eta} \end{aligned} \tag{24}$$

Each firm maximizes its profit,

$$\pi_{it} = P_{it} Q_{it} - C_{it} Q_{it} \tag{25}$$

We can substitute Q_{it} using the expression we have in the demand function, we have the following.

$$\pi_{it} = \Phi_t P_{it}^{1+\eta} - C_{it} \Phi_t P_{it}^{\eta} \tag{26}$$

Take the first-order condition with respect to P_{it}, we have the following.

$$\frac{\partial \pi_{it}}{\partial P_{it}} = (1+\eta) \Phi_t P_{it}^{\eta} - \eta C_{it} \Phi_t P_{it}^{\eta - 1} = 0 \tag{27}$$

Simple algebra of the above first-order condition provides us with the following relation between price P_{it} and marginal cost C_{it}.

$$P_{it} = \frac{\eta}{\eta + 1} C_{it} \tag{28}$$

Now, revenue can be expressed as $R_{it} = P_{it} Q_{it}$. We can further re-write the revenue equation, using both the simplified expression of demand function and the relationship between price and marginal cost, as follows.

$$
\begin{aligned}
R_{it} &= P_{it}Q_{it} \\
&= \Phi_t P_{it}^{1+\eta} \\
&= \Phi_t \left(\frac{\eta}{\eta+1}C_{it}\right)^{1+\eta}
\end{aligned}
\tag{29}
$$

Take the log of the above expression and recall the functional form of log of marginal cost, $ln(C_{it}) = c_{it}$, we have the following.

$$
\begin{aligned}
r_{it} &= ln(R_{it}) \\
&= (1+\eta)ln\left(\frac{\eta}{\eta+1}\right) + ln(\Phi_t) \\
&+ (1+\eta)\left[\beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_w ln(w_t) - \omega_{it}\right]
\end{aligned}
\tag{30}
$$

We now have the final expression of Results 1. ∎

Proof of Result 2: Static Profit Equation.. First, we define profit as follows.

$$
\pi_{it} = P_{it}Q_{it} - C_{it}Q_{it} \tag{31}
$$

using the relation between price and marginal cost, we can rewrite the above expression as follows.

$$
\begin{aligned}
\pi_{it} &= P_{it}Q_{it} - \frac{\eta+1}{\eta}P_{it}Q_{it} \\
&= R_{it} - \frac{\eta+1}{\eta}R_{it} \\
&= -\frac{1}{\eta}R_{it}
\end{aligned}
\tag{32}
$$

Then, using the expression derived in Results 1, we have the following.

$$
\pi_{it} = -\frac{1}{\eta}exp(r_{it}(k_{it}, l_{it}, \omega_{it}, w_t, \Phi_t; \eta, \beta_0, \beta_k, \beta_l, \beta_w)) \tag{33}
$$

As we will see later, r_{it} can be reduced further to only depend only on k_{it}, l_{it}, ω_{it}, and parameters $\eta, \beta_0, \beta_k, \beta_l$. In other words, w_t and Φ_t will be absorbed as time-specific effects. Therefore, we can now have the final expression of π_{it} as follows.

$$
\pi_{it} = -\frac{1}{\eta}exp(r_{it}(k_{it}, l_{it}, \omega_{it}; \eta, \beta_0, \beta_k, \beta_l)) \tag{34}
$$

We now have the final expression of Results 2.

∎

D Details on the Productivity Estimation

First, the expression in Equation (3) is appended with an IID error term u_{it}. That is,

$$\begin{aligned} r_{it} = (\eta+1)ln\left(\frac{\eta}{\eta+1}\right) + ln(\Phi_t) \\ + (\eta+1)\left[\beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_w ln(w_t)\right] \underbrace{-(\eta+1)\omega_{it} + u_{it}}_{\text{Composite Error Term}} \end{aligned} \tag{35}$$

where $\Phi_t := \frac{Q_t}{P_t^{\eta}}$.

I borrow insights from Olley and Pakes (1996) to re-write the unobserved productivity, ω_{it}, in terms of some observables which are correlated with it. In particular, as in Levinsohn and Petrin (2003); Ackerberg et al. (2015), and Gandhi et al. (2021), I use the firm's intermediate input, m_{it}, and the firm's energy input, n_{it}, as firm i's choice of this variable input depends on the level of productivity. The essence of Olley and Pakes (1996) is that this observable, i.e., intermediate input expenditures of the firm, contains information on its productivity level[19]. Then, I can write the level of productivity, conditional on the capital stock, as a function of the variable input levels, $\omega_{it}(k_{it}, l_{it}, m_{it}, n_{it})$. This method allows me to use the intermediate inputs of the firm to control for the productivity in Equation (35). I then combine the demand elasticity terms into an intercept γ_0, and the time-varying aggregate demand shock and market-level factor prices into a set of time and industry dummies ($D_t^{(1)}$ and $D_j^{(2)}$), I can re-write Equation (35) as follows

$$\begin{aligned} r_{it} &= \gamma_0 + \sum_{t=1}^{T} \gamma_t^{(1)} D_t^{(1)} + \sum_{j=1}^{J} \gamma_j^{(2)} D_j^{(2)} + (\eta+1)\left[\beta_k k_{it} + \beta_l l_{it} - \omega_{it}\right] + u_{it} \\ &= \gamma_0 + \sum_{t=1}^{T} \gamma_t^{(1)} D_t^{(1)} + \sum_{j=1}^{J} \gamma_j^{(2)} D_j^{(2)} + h(k_{it}, l_{it}, m_{it}, n_{it}) + \nu_{it} \end{aligned} \tag{36}$$

The function $h(\cdot)$ represents the combined effect of capital and productivity. In particular, I approximate this function with third-order polynomials. I estimate Equation (36) with a simple OLS. From here, I can get a fitted value of the $h(\cdot)$ where I denote it as $\hat{\phi}_{it}$. This $\hat{\phi}_{it}$ is an estimate of $(\eta+1)\left[\beta_k k_{it} + \beta_l l_{it} - \omega_{it}\right]$.

As discussed in the previous section, I assume that ω_{it} follows a controlled Markov process. I assume the distribution of ω_{it} can be written as $G_{\omega}(\cdot)$. For the empirical implementation of my model, I model the evolution of ω_{it} as follows

$$\begin{aligned} \omega_{it} &= g_{\omega}(\omega_{it-1}, \text{no}_{it-1}, \text{int}_{it-1}, \text{ext}_{it-1}, \text{both}_{it-1}) + \xi_{it} \\ &= \alpha_0 + \alpha_1\omega_{it-1} + \alpha_2(\omega_{it-1})^2 + \alpha_3(\omega_{it-1})^3 \\ &+ \alpha_4\text{no}_{it-1} + \alpha_5\text{int}_{it-1} + \alpha_6\text{ext}_{it-1} + \alpha_7\text{both}_{it-1} + \xi_{it} \end{aligned} \tag{37}$$

[19]There are two crucial assumptions for this insight to work. The first one is the scalar unobservability assumption. The intuition behind this assumption is that in choosing the input, firms only require one unobservable (in our case ω_{it}). The second assumption is the strictly monotone assumption. This assumption dictates that the intermediate input must be strictly monotone in ω. Levinsohn and Petrin (2003) show that the assumption is not too restrictive in the case of manufacturing plants in Chile. These assumptions allow us to invert the intermediate input, $m_{it} = \mathbb{M}_t(k_{it}, l_{it}, \omega_{it})$.

Here, I assume that ξ_{it} is IID across time and firms and is drawn from a normal distribution with zero mean and variance σ_{ξ}^2.

I can now substitute the productivity ω_{it} using $\hat{\phi}_{it}$. That is, substituting $\omega_{it} = -(1/(\eta+1))\hat{\phi}_{it} + \beta_k k_{it} + \beta_l l_{it}$ to the Equation (37). After some algebra, I obtain the following estimating equation

$$\begin{aligned}\hat{\phi}_{it} &= \beta_k^* k_{it} + \beta_l^* l_{it} + \alpha_1\left(\hat{\phi}_{it-1} - \beta_k^* k_{it-1} - \beta_l^* l_{it-1}\right) \\ &- \alpha_2^*\left(\hat{\phi}_{it-1} - \beta_k^* k_{it-1} - \beta_l^* l_{it-1}\right)^2 \\ &+ \alpha_3^*\left(\hat{\phi}_{it-1} - \beta_k^* k_{it-1} - \beta_l^* l_{it-1}\right)^3 \\ &- \alpha_4^* \text{no}_{it-1} - \alpha_5^* \text{internal}_{it-1} - \alpha_6^* \text{external}_{it-1} - \alpha_7^* \text{both}_{it-1} - \xi_{it}^*\end{aligned} \tag{38}$$

where the star, $(*)$, represents that the coefficients α and β are multiplied by $(\eta+1)$[20]. The above expression will be estimated with nonlinear least squares. From the estimated parameters, I can recover an estimate of productivity for each observation as follows. The factor $-1/(\eta+1)$ is also how the mode intercepts in Table 5 are mapped into a productivity difference in Section 5.

$$\hat{\omega}_{it} = -(1/(\hat{\eta}+1))\hat{\phi}_{it} + \hat{\beta}_k k_{it} + \hat{\beta}_l l_{it} \tag{39}$$

Source of identification.—This approach is essentially a semiparametric one. As noted by Doraszelski and Jaumandreu (2013), the main ingredient for identification is that variables in the parametric part of the model are not perfectly predictable by the variables in the non-parametric part (in a simple regression sense). Apart from a similar identification as in Olley and Pakes (1996) and Doraszelski and Jaumandreu (2013), the estimating equation in Equation (38) makes use of information on the demand elasticity. This demand elasticity is identified, and the source of its identification has been discussed in the previous subsection.

E Alternative Demand System

I follow a standard discrete choice model of consumer choice a la Berry (1994) and Berry et al. (1995). The indirect utility function can be described as follows

$$V_{ijt} = \eta p_{jt} + \xi_{ijt} + \varepsilon_{jt} \tag{40}$$

where p_{jt} is the logarithmic price, ξ_{ijt} is the idiosyncratic preference shock which follows a Type-1 Extreme Value distribution, and ε_{jt} is an unobserved demand shock.

Using the Berry inversion as in Berry (1994) for the above expression, I can get a well-known expression for the market share of good j relative to the outside option as follows

$$ln(ms_{jt}) - ln(ms_{ot}) = \eta p_{jt} + \xi_{ijt} + \varepsilon_{jt} \tag{41}$$

Consider only the case of single-product firm ($i = j$). From the log of market share, I can rearrange the above expression using that $ln(ms_{it}) = ln(q_{it}) - ln(Q_t)$, as in the demand expressed in Equation (1). Then, I obtain the expression for log price $ln(p_{jt})$. Note that $ln(r_{it}) = ln(p_{it}q_{it})$.

[20]With some exceptions for $\alpha_2^* = \alpha_2(\eta+1)^{-1}$ and $\alpha_3^* = \alpha_3(\eta+1)^{-2}$

Rearranging this expression with information on the expression of $ln(p_{it})$ and Equation (1), this yields the following expression for log revenue

$$ln(r_{it}) = \kappa_0 + \kappa_k ln(k_{it}) + \omega^*_{it} + u_{it} \tag{42}$$

where $\kappa_0 = \frac{1}{|\eta|} ln(ms_{ot})$, $\kappa_k = (\eta + 1)\beta_k$, and $\omega^*_{it} = -(\eta + 1)\omega_{it}$. The above expression is similar to the estimating equation denoted in Equation (36) in the absence of the year fixed effects. As noted by De Loecker (2011), the total output enters in exactly the same way, leading to identification of η in the production function framework. Anderson et al. (1992) note that in the usual logit demand structure, the estimated parameter η is used to compute own and cross-price elasticities. However, in this setup with log prices in the indirect utility function, they are identical.

F Derivation of the McFadden's Social Surplus

Under the Type 1 Extreme Value distribution assumption for ε^a_{it}, there is a closed-form expression for the expected value of the **max** over options. This is a standard result from McFadden (1978) for static logit models, extended to dynamic settings by Rust (1987),

Let us denote the deterministic part of each action's payoff, excluding the shock, as follows:

$$\begin{aligned}
\bar{V}^{no}_{it} &= \beta EV^{no} \\
\bar{V}^{internal}_{it} &= -\kappa^{internal} + \beta EV^{internal} \\
\bar{V}^{external}_{it} &= -\kappa^{external} + \beta EV^{external} \\
\bar{V}^{both}_{it} &= -\kappa^{internal} - \kappa^{external} + \kappa^{complement} + \beta EV^{both}
\end{aligned}$$

The realized value is $V(s_{it}) = \pi_{it}(s_{it}) + \max_a\{\bar{V}^a_{it} + \varepsilon^a_{it}\}$. Then, the ex-ante expectation is:

$$EV(s_{it}) = \pi_{it}(s_{it}) + E_\varepsilon \left[\max_a (\bar{V}^a_{it} + \varepsilon^a_{it})\right]$$

where $\gamma^{Euler} \approx 0.57721$ is Euler's constant from the Type 1 Extreme Value distribution's properties.

For an IID Type 1 EV shock with scale 1, the expected max has a closed-form solution:

$$E_\varepsilon \left[\max_a (v_a + \epsilon_a)\right] = ln\left(\sum_a exp(v_a)\right) + \gamma^{Euler}$$

Then, we have that the ex-ante expectation is the following

$$\begin{aligned}
EV(s_{it}) = \pi_{it}(s_{it}) &+ \gamma^{\text{Euler}} \\
&+ ln\left[exp(\overline{V}^{no}_{it}) + exp(\overline{V}^{internal}_{it}) + exp(\overline{V}^{external}_{it}) + exp(\overline{V}^{both}_{it})\right]
\end{aligned}$$

G Derivation of the Expected Future Value

The value function $V(s_{it})$ represents the maximum expected discounted lifetime profits starting from state s_{it}, following optimal behavior. After choosing action a, the firm receives current profit $\pi(s_{it})$ minus costs, plus the discounted expected value from $t+1$ onward.

The expected future value conditional on action a is defined as the expectation of $V(s_{it+1})$ over the distribution of next states:

$$EV^a(s_{it}) = \mathbb{E}[V(s_{it+1} \mid s_{it}, a]$$

This is the standard definition in the typical dynamic discrete choice model (see Rust (1987), Abbring and Daljord (2020)). It captures how the choice of a influences the evolution of the state and thus future values.

Assuming that the state space is continuous, we can express the expectation using the transition density $p(s_{it+1}|s_{it}, a)$:

$$EV^a(s_{it}) = \int_s V(s_{it+1}) p(s_{it+1} \mid s_{it}, a) ds$$

Then, in our case,

$$\begin{aligned} EV^a(s_{it}) &= \int_{\omega'} V(s_{it+1}) dG(\omega_{it+1} \mid s_{it}, a_{it}) \\ &, \qquad \forall a_{it} \in \{\text{no R\&D}, \text{internal}, \text{external}, \text{both}\} \end{aligned}$$

Since, I binned the state, we need to model the state space as a discrete one. In that case, the expected future value has the following expression.

$$\begin{aligned} EV^a(s_{it}) &= \int_s V(s_{it+1}) p(s_{it+1} \mid s_{it}, a) ds \\ &\approx \sum_s V(s_{it+1}) p(s_{it+1} \mid s_{it}, a) \end{aligned}$$

Similarly, if the continuous state space (e.g., productivity ω) is discretized into a finite grid for computational tractability, the expected future value can be approximated as a discrete sum over possible next states. This is computationally efficient, as it can be represented as a matrix multiplication $EV^a = P^a \cdot V$, where P^a is the action-specific transition matrix and V is the vector of value functions. It preserves action-dependence: Different actions (e.g., internal R&D vs. no R&D) lead to different transition probabilities P^a, reflecting how R&D choices affect state evolution (e.g., higher probability of jumping to a high-innovation state with both types of R&D).

$$\begin{aligned} EV^a(s_{it}) &= \int_{\omega'} V(s_{it+1})\, dG(\omega' \mid s_{it}, a_{it}) \\ &\approx \sum_{\omega'} V(s_{it+1})\, P(\omega' \mid s_{it}, a_{it}) \\ &, \qquad \forall a_{it} \in \{\text{no R\&D}, \text{internal}, \text{external}, \text{both}\} \end{aligned}$$

This is computationally efficient since it can be represented simply as a matrix multiplication. It preserves action-dependence: Different actions (e.g., "internal R&D" vs. "no R&D") lead to different transition probabilities P^a, reflecting how R&D choices affect state evolution (e.g., higher probability of jumping to a "high innovation" state with "both" R&D).

H Derivation of the Probability

As I have outlined in the main text, suppose $\varepsilon_{it}(a_{it})$ is IID that follows a Type 1 Extreme Value distribution. The density of each private shock is $f(\varepsilon_{it}) = e^{-\varepsilon_{it}}e^{-e^{-\varepsilon_{it}}}$. The cumulative distribution is $F(\varepsilon_{it}) = e^{-e^{-\varepsilon_{it}}}$.

Suppose firm i chooses action *both*. In choosing this particular alternative, it must be that the value of doing *both* is higher than other alternatives than itself. We can express it as follows

$$\begin{aligned} Pr(a_{it} = both) &= Prob\left(\overline{V}_{it}^{both} + \varepsilon_{it}^{both} > \overline{V}_{it}^{j} + \varepsilon_{it}^{j}, \forall j \neq both\right) \\ &= Prob\left(\varepsilon_{it}^{j} < \varepsilon_{it}^{both} + \overline{V}_{it}^{both} - \overline{V}_{it}^{j}, \forall j \neq both\right) \end{aligned}$$

where $\overline{V}_{it}^{both}$ and $\overline{V}_{it}^{j}, \forall j \neq both$ are the deterministic parts of the conditional (alternative-specific) values.

Since ε's are independent, we can express the conditional probability as follows

$$Pr(a_{it} = both|\varepsilon_{it}^{both}) = \prod_{j \neq both} e^{-e^{-\left(\varepsilon_{it}^{both} + \overline{V}_{it}^{both} - \overline{V}_{it}^{j}\right)}}$$

Then,

$$\begin{aligned} Pr(a_{it} = both) &= \int \left(Pr(a_{it} = both|\varepsilon_{it}^{both})\right) f(\varepsilon_{it}^{both}) d\varepsilon_{it}^{both} \\ &= \int \left(\prod_{j \neq both} e^{-e^{-\left(\varepsilon_{it}^{both} + \overline{V}_{it}^{both} - \overline{V}_{it}^{j}\right)}}\right) e^{-\varepsilon_{it}^{both}} e^{-e^{-\varepsilon_{it}^{both}}} d\varepsilon_{it}^{both} \\ &= \int \left(\prod_{j} e^{-e^{-\left(\varepsilon_{it}^{both} + \overline{V}_{it}^{both} - \overline{V}_{it}^{j}\right)}}\right) e^{-\varepsilon_{it}^{both}} d\varepsilon_{it}^{both} \\ &= \int exp\left(-\sum_{j} e^{-\left(\varepsilon_{it}^{both} + \overline{V}_{it}^{both} - \overline{V}_{it}^{j}\right)}\right) e^{-\varepsilon_{it}^{both}} d\varepsilon_{it}^{both} \\ &= \int exp\left(-e^{-\varepsilon_{it}^{both}} \sum_{j} e^{-\left(\overline{V}_{it}^{both} - \overline{V}_{it}^{j}\right)}\right) e^{-\varepsilon_{it}^{both}} d\varepsilon_{it}^{both} \end{aligned}$$

Now suppose $x = exp(-\varepsilon_{it}^{both})$. We also know that $-exp(-\varepsilon_{it}^{both})d\varepsilon_{it}^{both} = dx$. As ε_{it}^{both} approaches infinity, x approaches zero. Similarly, as ε_{it}^{both} approaches negative infinity, x becomes infinitely large. Replacing the above expression with the new term, we have the following

$$
\begin{aligned}
Pr(a_{it} = both) &= \int_{\infty}^{0} exp\left(-x\sum_j e^{-\left(\overline{V}_{it}^{both} - \overline{V}_{it}^{j}\right)}\right)(-dx) \\
&= \int_{0}^{\infty} exp\left(-x\sum_j e^{-\left(\overline{V}_{it}^{both} - \overline{V}_{it}^{j}\right)}\right) dx
\end{aligned}
$$

Evaluating the integral in the last line, we arrive at the following expression

$$
\begin{aligned}
Pr(a_{it} = both) &= \left.\frac{exp\left(-x\sum_j e^{-\left(\overline{V}_{it}^{both} - \overline{V}_{it}^{j}\right)}\right)}{-\sum_j e^{-\left(\overline{V}_{it}^{both} - \overline{V}_{it}^{j}\right)}}\right|_0^{\infty} \\
&= \frac{1}{\sum_j e^{-\left(\overline{V}_{it}^{both} - \overline{V}_{it}^{j}\right)}} \\
&= \frac{e^{\overline{V}_{it}^{both}}}{\sum_j e^{\overline{V}_{it}^{j}}}
\end{aligned}
$$

Similarly, we can derive the optimal choice probability for other actions as well.

I Grossman–Hart algebra and observational equivalence

This appendix fills in the contracting subperiod of Section 3.2.4. The first block derives the equilibrium net, the efficient benchmark, and the location of $\kappa^{external}$ and $\kappa^{complement}$ in the Bellman. The second block proves Proposition 1. Section 5.4 inverts the estimated costs at each surplus scale S as a reading of those two intercepts, not as a re-solved experiment.

I.1 Equilibrium net

The firm solves

$$\max_{I \geq 0} \ \tau\theta I - I^2, \qquad \tau \in (0,1],\ \theta > 0. \tag{43}$$

Differentiate with respect to I:

$$\frac{\partial}{\partial I}\left(\tau\theta I - I^2\right) = \tau\theta - 2I. \tag{44}$$

Set the derivative to zero:

$$\tau\theta - 2I = 0 \quad \Longrightarrow \quad I^* = \frac{\tau\theta}{2}. \tag{45}$$

The second derivative is

$$\frac{\partial^2}{\partial I^2}\left(\tau\theta I - I^2\right) = -2 < 0, \tag{46}$$

so the critical point is a maximum. At $I = 0$ the first derivative equals $\tau\theta > 0$, so the firm strictly prefers a positive investment. Substitute I^* into the objective:

$$\text{net}(\tau) = \tau\theta \cdot \frac{\tau\theta}{2} - \left(\frac{\tau\theta}{2}\right)^2 = \frac{\tau^2\theta^2}{2} - \frac{\tau^2\theta^2}{4} = \frac{\tau^2\theta^2}{4}. \tag{47}$$

That is equation (10) in the text.

Efficient investment. If the firm kept the entire pie, the same program would be run at $\tau = 1$. Then $I^{e} = \theta/2$ and

$$S := \mathrm{net}(1) = \frac{\theta^2}{4}. \tag{48}$$

The equilibrium net is the fraction τ^2 of that surplus:

$$\mathrm{net}(\tau) = S\tau^2. \tag{49}$$

The special case $\theta = 1$ yields $\mathrm{net}(\tau) = \tau^2/4$ and $S = 1/4$. An external cost measured in millions of euro cannot be generated from $S = 1/4$, so the paper keeps θ (equivalently S) free.

An isomorphic technology. The same net is obtained from a concave surplus $\theta\sqrt{\tilde{I}}$ with linear cost $\tilde{I}$. The program is $\max_{\tilde{I}\geq 0} \tau\theta\sqrt{\tilde{I}} - \tilde{I}$. The first-order condition is $\tau\theta/(2\sqrt{\tilde{I}}) = 1$, so $\tilde{I}^* = (\tau\theta/2)^2$ and

$$\mathrm{net}(\tau) = \tau\theta \cdot \frac{\tau\theta}{2} - \left(\frac{\tau\theta}{2}\right)^2 = \frac{\tau^2\theta^2}{4}. \tag{50}$$

I use the quadratic-cost version in the main text because the first-order condition is linear in I. Nothing in the Bellman depends on which of the two technologies one writes down.

Location identities. Measure costs against S:

$$\begin{aligned} \kappa^{external} &= \mathrm{net}(1) - \mathrm{net}(\tau) = S - S\tau^2 = S\big(1 - \tau^2\big), & (51)\\ \kappa^{complement} &= \mathrm{net}(\tau') - \mathrm{net}(\tau) = S\big((\tau')^2 - \tau^2\big). & (52)\end{aligned}$$

The current-period payoffs relative to the efficient net are

$$\begin{aligned} v^{external} &= \mathrm{net}(\tau) - \mathrm{net}(1) = -\kappa^{external}, & (53)\\ v^{both} &= -\kappa^{internal} + \mathrm{net}(\tau') - \mathrm{net}(1). & (54)\end{aligned}$$

The second line expands as

$$\begin{aligned} \mathrm{net}(\tau') - \mathrm{net}(1) &= \big(\mathrm{net}(\tau') - \mathrm{net}(\tau)\big) + \big(\mathrm{net}(\tau) - \mathrm{net}(1)\big) \\ &= \kappa^{complement} - \kappa^{external}, & (55)\end{aligned}$$

and therefore

$$v^{both} = -\kappa^{internal} - \kappa^{external} + \kappa^{complement}, \tag{56}$$

which is the current-period term already in (5). Replacing $\mathrm{net}(\tau') - \mathrm{net}(1)$ by $\mathrm{net}(\tau') - \mathrm{net}(\tau)$ would drop $-\kappa^{external}$ and would not match the Bellman.

Underinvestment and transfer. Social surplus at investment I is $\theta I - I^2$. At the efficient investment this equals S. At $I^* = \tau\theta/2$ it equals $S(2\tau - \tau^2)$. The deadweight loss is therefore

$$\mathrm{DWL}(\tau) = S - S(2\tau - \tau^2) = S(1-\tau)^2. \tag{57}$$

The partner collects share $1-\tau$ of the realized pie θI^* and invests nothing, so the transfer is

$$(1-\tau)\theta I^* = (1-\tau)\tau \cdot 2S = 2S\tau(1-\tau). \tag{58}$$

Adding the two pieces recovers the firm's private gap:

$$S(1-\tau)^2 + 2S\tau(1-\tau) = S(1-\tau^2) = \kappa^{external}. \tag{59}$$

The split between deadweight loss and transfer depends on S. The sum does not.

I.2 Proof of Proposition 1

Write $\overline{V}^a$ for the deterministic alternative-specific value of mode a in (5), omitting the Type-I shock. Under the estimated Bellman,

$$\begin{aligned}
\overline{V}^{\text{no}} &= \beta E V^{\text{no}}, \\
\overline{V}^{\text{internal}} &= -\kappa^{internal} + \beta E V^{\text{internal}}, \\
\overline{V}^{\text{external}} &= -\kappa^{external} + \beta E V^{\text{external}}, \\
\overline{V}^{\text{both}} &= -\kappa^{internal} - \kappa^{external} + \kappa^{complement} + \beta E V^{\text{both}}.
\end{aligned}$$

Under the Grossman–Hart subperiod, the current-period nets for external-only and both are $v^{external}$ and v^{both} from (14)–(15). Condition (iii) says those nets are intercepts. Conditions (12)–(13) then make the intercepts coincide with $-\kappa^{external}$ and $-\kappa^{internal} - \kappa^{external} + \kappa^{complement}$. Condition (i) says $P(\omega' \mid s, a)$ depends on the mode dummy, not on I^* or τ, so the four continuation values EV^a coincide with those of the estimated Bellman. Condition (ii) says disagreement after I is sunk does not bring continuation values into the bargaining problem, so the net does not depend on V. Condition (iv) says the Type-I shocks sit on modes, with bargaining inside the chosen mode, so the McFadden inclusive value is the log-sum-exp of the same four $\overline{V}^a$. Therefore the choice probabilities (8) coincide, and so does the likelihood.

The map $(\theta, \tau, \tau') \mapsto (\kappa^{external}, \kappa^{complement})$ has Jacobian rank 2, so the fibre over a given pair of costs is one-dimensional. The likelihood therefore recovers the three costs, not τ.

The claim is a sufficiency statement. If productivity depended on $I^*(\tau)$, the nested continuation would move with τ and would no longer match the estimated Bellman. I also maintain that disagreement payoffs after investment is sunk are $(0,0)$, so that the net does not depend on V.